Digital Video & Digital Signal Processing

IC Handbook

December 1993

GEC PLESSEY
SEMICONDUCTORS

DIGITAL VIDEO & DIGITAL SIGNAL PROCESSING

IC Handbook

GEC PLESSEY
SEMICONDUCTORS

Foreword

GEC Plessey Semiconductors (GPS) has been supplying high speed Data Converters and highly integrated Digital Signal Processing (DSP) building block components since the mid-1980s. Later DSP devices offered a significant increase in integration over the traditional Multipliers, ALUs and Address Generators, while the Data Converters offered an accurate flash conversion from video speed to sampling rates beyond 100 MHz.

GPS has recently launched a new range of Algorithm Specific DSP devices. These components are dedicated high performance (sampling rates up to 40MHz) solutions to common DSP algorithms such as FIR Filtering, Co-ordinate Conversion, Fast Fourier Transforms and 2D Convolution. The high level of functional integration offered by the Algorithm Specific components allows high performance DSP functions to be implemented with reduced component count and improved power consumption. In addition, GPS continues to offer basic building-block DSP devices that handle Complex Arithmetic as standard. This means that real time data handling, involving the manipulation of complex-values, can be performed with ease. In general, this provides the system designer with four-fold improvements in speed and substantial reduction in board area and power consumption.

The GPS Data Converter range of products divides into those designed to support traditional Digital Video applications and others geared more towards the high speed requirements of digital instrumentation and back-end graphics systems. The devices offer excellent linearity and 8-bit resolution. As an example of the continuous improvement in this area, the VP8708 is a 30MHz flash ADC which contains a 3:1 input multiplexing system plus signal conditioning stages which greatly reduces the normal amount of external components required on a typical video capture board.

The latest developments in Digital Video Compression from GPS take advantage of our strengths in both DSP and Data Conversion. The H.261 standard for video compression is addressed by the VP26xx chip-set. The level of integration of these devices is such that they require the absolute minimum of external components in a typical compression system. In support of this chip-set, GPS has also developed the VP510 RGB-YUV Colour Space Converter and VP520 Video Filter and will soon be announcing the VP53X series of NTSC-PAL Encoders and the VP1394 Video Clock Generator.

In summary, the application areas addressed by the PDSP family (building block and application specific) plus the SP/VP Signal Conversion products include:

Digital Filtering
Pulse & Signal Compression/ Decompression
Digital Modulation/ Demodulation
Correlation
Convolution
Image Processing
Digital Waveform Synthesis
Video Clock Generation
Analog-to-Digital /Digital-to-Analog Conversion
Video Signal Conditioning

Contents

Product index - Video Compression

Type	Function	Clock	Page
VP2611	H.261 encoder	27MHz	9
VP2612	Video multiplexer	27MHz	21
VP2614	H.261 video de-multiplexer	27MHz	30
VP2615	H.261 decoder	27MHz	38
VP520	PAL/NTSC to CIF/QCIF converter	27MHz	58

Product index - Digital Video

Type	Function	Clock	Page
VP510	Bi-directional colour space converter	27MHz	47
VP8708	30MHz 8-bit analog video input interface	30MHz	71

Product index - Building block DSP ICs

Arithmetic logic units

Type	Function	Speed	Page
PDSP1601A	ALU and barrel shifter	20MHz	83

Complex arithmetic

Type	Function	Speed	Page
PDSP16112/A	Complex multiplier (16x12 bits)	20MHz	99
PDSP16116/A	Complex multiplier (16x16 bits)	20MHz	106
PDSP16318/A	Complex accumulator	20MHz	138

Product index - Algorithm specific DSP ICs

Image processing

Type	Function	Speed	Page
PDSP16488	Single-chip 2-D convolver	40MHz	171

Digital filtering

Type	Function	Speed	Page
PDSP16256/A	Programmable FIR filter	25MHz	121

Co-ordinate conversion

Type	Function	Speed	Page
PDSP16330/A/B	Pythagoras processor	25MHz	145
PDSP16340	Polar to Cartesian converter	20MHz	152

Waveform generation and modulation

Type	Function	Speed	Page
PDSP16350	I/Q splitter/NCO	20MHz	160

Frequency domain processing

Type	Function	Speed	Page
PDSP16510A	Stand alone FFT processor	40MHz	200
PDSP16540	32K bucket buffer	40MHz	221

Product List - alpha numeric

Digital Video

GEC PLESSEY
SEMICONDUCTORS

VP 2611

H.261 ENCODER

IMPORTANT

At the time of publication of this handbook the VP2611 datasheet was being extensively revised. Please check with your nearest Customer Service Centre for the latest information.

FEATURES

- Fully integrated H261 video encoder
- Up to full CIF resolution and 30 Hz frame rates
- Inputs YUV data in 8 x 8 sub block format
- Outputs run length coded coefficients
- On chip motion vector estimator with +/-7 pixel search window
- Addresses and control generated internally for DRAM frame store
- Packaged in132 pin quad flat pack

ASSOCIATED PRODUCTS

- VP510 Colour Space Converter
- VP520CIF/QCIF Converter
- VP530 NTSC/PAL Encoder
- VP2612 Video Multiplexer
- VP2614 Video Demultiplexer
- VP2615 H.261 Decoder

DESCRIPTION

The VP2611 Video Compression Source Coder forms part of a chip set used in video conferencing, video telephony and multimedia applications. It produces data which conforms to the H261 standard for video compression with rates between 64K and 2M bits per second. With a 27 MHz clock the device will accept data produced to full CIF resolution at 30 Hz frame rates. The pipeline latency through the device is only 3 macro block periods.

The VP2611 contains all the elements necessary for the compression algorithm. It incorporates a Motion Vector Estimator which performs a +/- 7 pixel search. The decision to use inter or intra frame compression is made by the device, and the selected data blocks are read from the frame store. New or difference data is then passed through a Discrete Cosine Transformer and quantized. Data from the quantizer is also inverse quantized and passed through an Inverse Discrete Cosine Transformer. This re-constructed data is then written to the frame store for use in the next frame period.This frame store is managed by an internal DRAM controller, and no external logic is needed.

The input data must be in YUV space, and must also conform to the six sub blocks per macro block format defined by H261. Any conversion from RGB format is performed by the VP510 Colour Space Converter. Any reduction in spatial resolution, down to CIF or QCIF requirements, is done by the VP520 Three Channel Video Filter.

The quantized data is zig-zag scanned and run length coded before being output, together with block information and motion vectors.

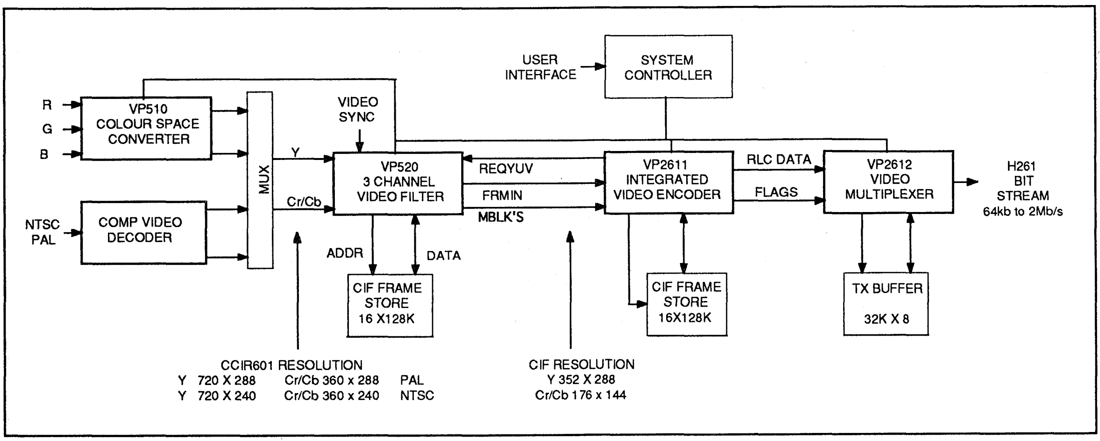

Fig 1 : Typical Video Conferencing Transmission System

PIN DESCRIPTIONS

YUV7:0 — This input bus accepts YUV data one pixel at a time from the preprocessor, clocked in on the rising edge of PCLK.

PCLK — This signal is used to strobe in data at the YUV port and must be derived by dividing SYSCLK with an integer greater than one.

FRMIN — This input should be pulled high to prepare the VP2611 to code a new frame. It must be held high for at least one SYSCLK cycle and then must be pulled low again before the next frame begins. The VP2611 will respond to the rising edge of FRMIN by asserting REQYUV appproximately 184 SYSCLK cycles later.

REQYUV — This output is pulled high to request that YUV data be input for a new MacroBlock. It is pulled low again 1871 SYSCLK cycles later. It remains low during Dummy MacroBlocks and during the lay period between frames.

DBUS7:0 — This output bus serves several functions as defined by DMODE3:0. In addition to providing the quantized coefficients and motion vectors, it is used to output control information.

DMODE3:0 — Output flag port for DBUS7:0 bus. The value at this port identifies the data type appearing on DBUS7:0 during the same period.

DCLK — This output pulses high for a minimum of 37ns each time new data is output on DBUS or DMODE. It can be used as an edge sensitive strobe signal or a level sensitive "valid" signal.

SW15:0 — This bidirectional port is connected to the frame store.

$\overline{\text{RAS}}$ — Row Address Strobe output for the external DRAMs.

$\overline{\text{CAS}}$ — Column Address Strobe output for the external DRAMs.

R/$\overline{\text{W}}$1 — Read/Write control for external DRAM 1.

R/$\overline{\text{W}}$2 — Read/Write control for external DRAM 2. N/C if 256k DRAMs.

$\overline{\text{OE1}}$ — Output Enable control for external DRAM 1 or ADR8.

$\overline{\text{OE2}}$ — Output Enable control for external DRAM 2. N/C if 256k DRAMs.

ADR7:0 — Address output for the external DRAMs.

CBUS7:0 — Bi-directional data bus for use by a Microprocessor. Data and insructions are clocked on and off the chip on the rising edge of CSTR .

$\overline{\text{CSTR}}$ — Data strobe for the CBUS port.

$\overline{\text{CEN}}$ — An enabling signal for the CBUS port.

CADR — When high, this signal defines CBUS as a data bus, and when low as an instruction input.

SYSCLK — System clock, run at 27MHz maximum. The clock must be high for between 35% and 65% of each clock cycle. This clock is used for all internal operations.

$\overline{\text{RESET}}$ — Active low power on reset which must be held low for at least 2064 cycles.

TCK — Test clock for JTAG.

TMS — Test Mode Select for JTAG.

TDI — Input JTAG test data.

TDO — Output JTAG test data.

$\overline{\text{TRST}}$ — Reset JTAG controller (active low).

NOTE:
"Barred" active low signals do not appear with a bar in the main body of the text.

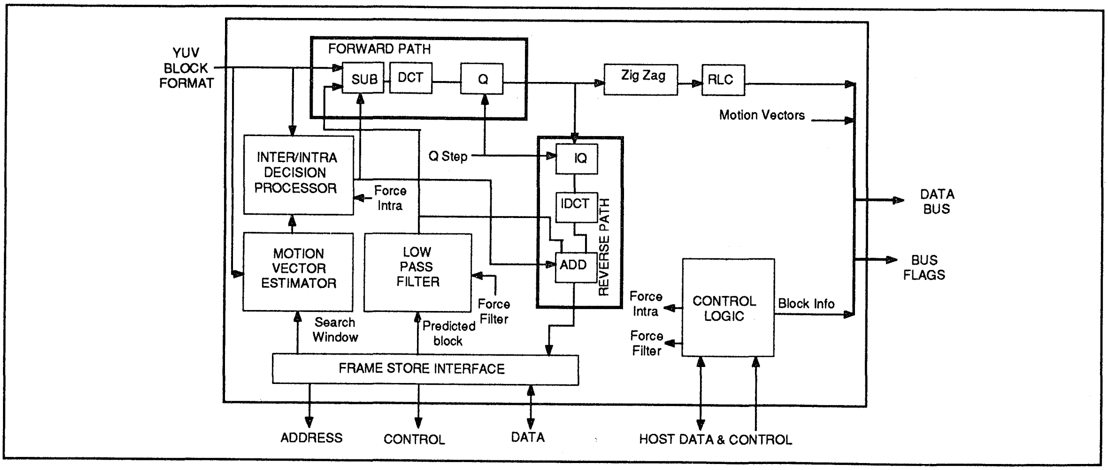

Fig 2 : Simplified Block Diagram

OPERATION OF MAJOR BLOCKS

Motion Vector Estimator

The motion estimator calculates the mean absolute error (MAE) for each possible position of the combined luminance block in a search window from the previous frame. The combined luminance block consists of 16 x 16 pixels, and in the search window this is displaced between -7 to +7 vertically, and -8 to +7 horizontally. The two lsb's of each pixel are discarded and the MAE value is contained within 14 bits.

The minimum MAE value, representing the best match between the previous and current block, is passed to the motion compensation decision block, together with the position of this best fit in the search window. The zero displacement MAE value is also passed to this block, which then decides whether the best fit is sufficiently better than the zero displacement fit. It uses the characteristic shown in Figure 3, and in the area to the right of the line all points defined by the two MAE values will cause motion compensation to be applied. In this case the best fit MAE value is used by the inter/intra decision processor, otherwise the zero displacement value is used.

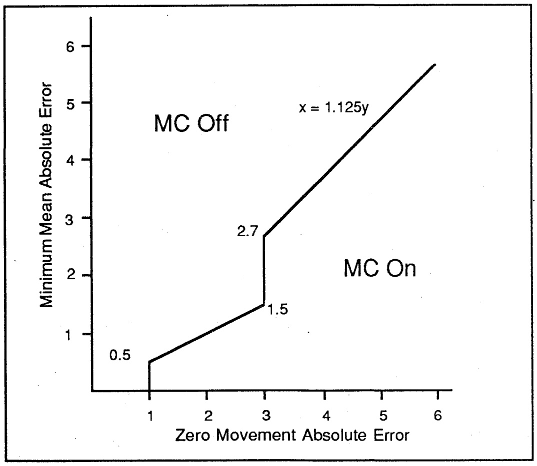

Fig 3 : MC Decision Slope

The user may overide the internal Inter/Intra decision at any time using the CBUS control port. A user generated forced inter mode will overide an internally generated `Force Intra'.

Inter/Intra Decision Processor

The MAE value passed by the motion compensation decision block is compared to the simplified variance of the current block. This simplified variance is calculated by summing the moduli of the differences between each luminance pixel and the mean luminance value over the whole macroblock. Eight bit pixels are used, and the variance value is expressed in 14 bits by discarding the two lsb's from the actual 16 bit result. It can then be directly compared to the 14 bit MAE value.

If the MAE value is below a user defined threshold inter mode coding is always selected. The default threshold is 3, on a scale from 0 to 15 using the 4 msb's from the 14 bit value. Above this threshold inter mode is only selected if the variance of the current block is greater than or equal to the MAE value in use.

In order to avoid gradual picture degredation, every 61st Macroblock input to the VP2611 is coded in intra mode regardless of the above decision. As 61 is a prime number, this will ensure that each macroblock will be transmitted in intra mode at least once in every 61 transmissions. If FIX MACROBLOCK or SKIP PICTURE is invoked this Force Intra' counter will be disabled.

Low Pass Filter

The macroblock selected from the previous frame in motion compensated inter mode coding, will be filtered before it is subtracted from the current block. This decision can be overidden externally by the system controller. The Filter uses a simple [1 2 1] characteristic in both vertical and horizontal dimensions as specified in H.261.

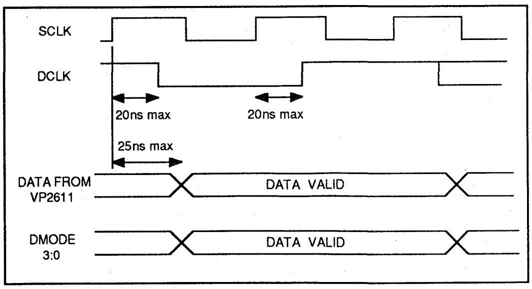

Fig 4 : Timing diagram

SYMBOL	PARAMETER	MINIMUM	MAXIMUM
t RAC	Access time from RAS	-	105ns or under
t CAC	Access time from CAS	-	25ns or under
t RP	RAS precharge time	50ns or under	-
t CP	CAS precharge time	15ns or under	-
t RAS	RAS pulse width	90ns or under	-
t CAS	CAS pulse width	50ns or under	-
t REF	Time to refresh 256 rows	-	0.25ms or over

N.B. All times are quoted assuming 27MHz operation. For lower clock frequencies increase the above values proportionately.

Table 1 : External DRAM timing requirements

Frame Store Manager

The previous picture is stored in an external CIF DRAM frame store, which is connected by a glueless interface. The internal Frame Store Manager controls all read, write, and refresh operations to these DRAMs. No provision is made to allow the use of smaller DRAM's, if only QCIF operation is required.

During the coding of each macroblock columns of the search window are read from these DRAMs, and finally the "best fit" macroBlock is obtained. At the completion of coding the fully processed new macroblock is written to the DRAM's, after it has been decoded again. In this way the frame store maintains a bit-accurate duplicate of the image seen by the Decoder (excepting transmission errors).

Several configurations are possible to make the required 128Kx16 store. Two 64K x 16 DRAMs could be employed; in this case use the default 1M DRAM mode when setting up the chip. Otherwise, a single 256K x 16 DRAM or four 256K x 4 DRAMs could be used. In these last two cases use OE1 as ADR8, RW1 as R/W, and do not connect RW2 and OE2. Also, use the Setup instruction at the CPORT to put the device into 4M DRAM mode.

Table 1 details the critical timing parameters which the external DRAM must meet with SYSCLK running at 27MHz. Note that, if used at slower speeds, the requirements on the DRAM timing are relaxed with the exception of refresh. The number of refresh cycles the VP2611 produces is directly proportional to the SYSCLK frequency.

Discrete Cosine Transform

This circuit performs a Discrete Cosine Transform on each 8x8 sub block, whether in inter or intra mode. In intra mode, eight bit pixel data is used, with a ninth implied sign bit (all pixel data is positive). In inter mode the difference between the current and best fit previous block is used. This will be a two's complement number. Twelve bit coefficients are produced by the DCT, and passed on to the quantizer.

Quantize

This section quantizes the results of the DCT by dividing the 12 bit output from the DCT with a host supplied value. The 5 bit quantization value supplied corresponds to division of the 12 bit coefficients (range $\pm$ 2048) by values from 2 to 62, but in steps of 2. This variable quantization strategy allows the volume of data generated by the encoder to be adjusted dynamically, depending on the fullness of the transmission buffer. For H.261 applications it uses the quantisation value provided at the control port during the previous Macroblock period (or at some earlier time). An option is provided which allows two quantisation values to be used, one for use with inter coded macroblocks, and the other for use with intra coded macroblocks.

As specified in H.261, the DC coefficient of an Intra coded Block is treated differently ; a quantizer value of 8 is always used in this case.

When the quantization value is small, and the DCT coefficient is large, there is a danger of overflow in the eight bit output. To avoid this a clipping circuit is included at the output of the quantizer, which saturates at the maximum values.

Zig Zag Scan

This is essentially an address generator which reorders the DCT coefficients according to the standard zig-zag scan pattern. This has the effect of concentrating the significant coefficients at the beginning of the sub-block, improving the efficiency of the Run Length Coder.

Run Length Coder

Each coefficient output from the zig zag scan is examined. If it is non-zero, then the Run Length Coding circuit will pass the coefficient magnitude to the output port along with its zero count i.e. the number of zero magnitude coefficients preceding it within the same 8x8 sub-block.

Inverse Quantize

This circuit replicates the operation of the inverse quantizer in the decoder. It reconstructs the 12 bit DCT coefficients from the 8 bit quantized inputs, using the 5 bit quantization value. This is achieved using the following formulae.

If QUANT is odd :
REC = QUANT*(2*LEVEL+1) : LEVEL > 0
REC = QUANT*(2*LEVEL-1) : LEVEL < 0

If QUANT is even :
REC = QUANT*(2*LEVEL+1)-1 : LEVEL > 0
REC = QUANT*(2*LEVEL-1)+1 : LEVEL < 0

For Intra Coded DC Coefficients :
REC = 8*LEVEL
except if LEVEL=255 when REC=1024

If LEVEL=0 then REC=0 in all cases.

The reconstructed values (REC) are passed through a Clipping Circuit in case of arithmetic overflow.

Thus, the Inverse Quantizer restores the DCT coefficients to their original value but with quantisation error.

Inverse DCT

This circuit replicates the operation of the Inverse Cosine Transform in the Decoder, and outputs 9 bit signed pixel data (intra mode) or pixel difference data (inter mode). The IDCT fully meets the CCITT specification.

Reconstruction Adder

In Inter Mode, the IDCT data is added to the best fit block from the previous frame store. In Intra mode, the IDCT data is simply added to zero. After the adder, the sign bit is removed from the result to give 8 bit pixels. Clipping circuits ensure that any pixels with values exceeding 255 are clipped to 255, and any with negative values are clipped to zero (such values are possible due to quantization noise).

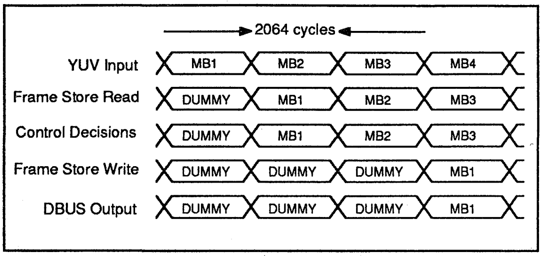

Fig 5: MacroBlock Pipelining

OPERATION OF INTERFACES

Macroblock Delays

The VP2611 has a three macroblock pipeline delay between pixel inputs and run length coded outputs. This is illustrated in Figure 4. Whilst the second macroblock is being input, the best fit macroblock from the previous frame is being identified and then read from the frame store. At this time any Control Decisions which are to effect the first macroblock must be supplied by the host controller. The run length coded outputs for the first macroblock are not available until the fourth macroblock is supplied at the input pins.

YUV Input Port

The YUV port accepts pixel data from the preprocessor in block format as illustrated in Figure 5. Within a complete system the VP2611 is always the master device, and must be supplied with macroblock data when it makes a demand. The order in which pixels are supplied is pre-determined, and must be strictly maintained. There are 64 pixels per sub-block and 4 luminance and 2 chrominance sub-blocks per Macroblock. The macroblocks themselves are divided into groups of blocks (GOB's), and the sequence specified in H.261 must also be maintained. Note that, since the chrominance resolution is half the luminance resolution both vertically and horizontally, then the two chrominance blocks cover the same picture area as the four luminance blocks.

The pre-processor producing macroblock data must produce a frame start signal (FRMIN) when it has a complete frame of data available. This resets the input controller within the VP2611, which will then generate sequential GOB and macroblock numbers for the coded outputs referenced to this input.

FRMIN must go high for at least one system clock period, and must go low before the next frame is available. The VP2611 responds to FRMIN with a request for macroblock data (REQYUV), which occurs approximately 184 SYSCLK periods after FRMIN. It must then receive a complete macroblock within 1871 SYSCLK periods, and at the end of this time REQYUV will go inactive. The VP2611 must be provided with a PCLK signal to strobe in the data. This must be derived from SYSCLK, and must only be present when there is valid data at the input. Data must meet the set up and hold times with respect to PCLK as specified in Figure 6.

The maximum peak rate for PCLK is the SYSCLK rate divided by two, but since there are 384 bytes per macroblock

SUBBLOCK ORDER WITHIN MACROBLOCK

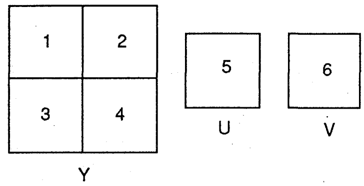

PIXEL ORDER WITHIN SUBBLOCK

00	01	02	03	04	05	06	07
08	09	10	11	12	13	14	15
16	17	18	19	20	21	22	23
24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39
40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55
56	57	58	59	60	61	62	63

Fig 6 : Ordering of Pixels

then theoretically the average rate need only be 384/1871 times the SYSCLK rate. Note that PCLK must always be obtained by driving SYSCLK by an integer greater than one. When the VP520 CIF/QCIF Converter is supplying the VP2611 with data, it provides a peak PCLK rate equivalent to SYSCLK divided by two, and an average rate of SYSCLK divided by four.

The mimimum gap between REQYUV going active is 2064 SYSCLK periods. In full CIF mode "dummy" macroblocks are internally inserted between rows, in order to give the chip sufficient time to load a new search window. No new YUV data must be loaded during these dummy macroblocks, and REQYUV will remain inactive. No dummy macroblocks are required in QCIF mode. With a 27MHz SYSCLK all macroblocks will be coded in less than a 30Hz frame rate period, and there will be a period of inactivity before FRMIN goes active again. During this period the output bus will remain static at all ones, and no output strobe (DCLK) will be produced.

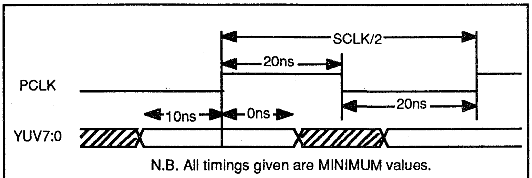

Fig 7 : Timing at YUV Port

DBUS Output Port

The DBUS port is used to pass data and control information directly to the VP2612 Video Multiplexer. The type of data on the output pins is identified by the DMODE 3:0 outputs, using the codes shown in Table 2. An output strobe is also produced (DCLK) which always goes high one system clock period after the data defined by DMODE 3:0 becomes valid. This edge is used to strobe the data into the Video Multiplexer, and thus the data set up time is always one SYSCLK period minus differential output delays.

The number of SYSCLK periods during which data remains valid is dependent on the type of data, and DCLK remains high for this same period. It goes low as the result of the same SYSCLK rising edge which produces a change in DMODE 3:0. The output delays with respect to SYSCLK are illustrated in Figure 8, and Figure 9 shows a typical output sequence during which DCLK remains high for several cycles as the sub-block number (code 7) is produced. During a Wait State (code 15) no DCLK transitions are produced. The actual sequence of output events which occur for each macroblock, and the duration of each event, are illustrated in Figure 7.

The output events are defined in more detail below;

GOB Number : At the start of each new macroblock, the current GOB Number is output on DBUS3:0. (DBUS3 is MSB).

DMODE3:0	FUNCTION
0000	GOB Number
0001	MB Number
0010	Control Decisions
0011	Quant Value
0100	Horizontal MV
0101	Vertical MV
0110	Coded Blk Pattern
0111	Sub-Block No
1000	Zero Run Count
1001	RLC Coefficient
1010	Not used
1011	Not used
1100	Not used
1101	Not used
1110	Not used
1111	Wait State

Table 2 : DBUS Functions

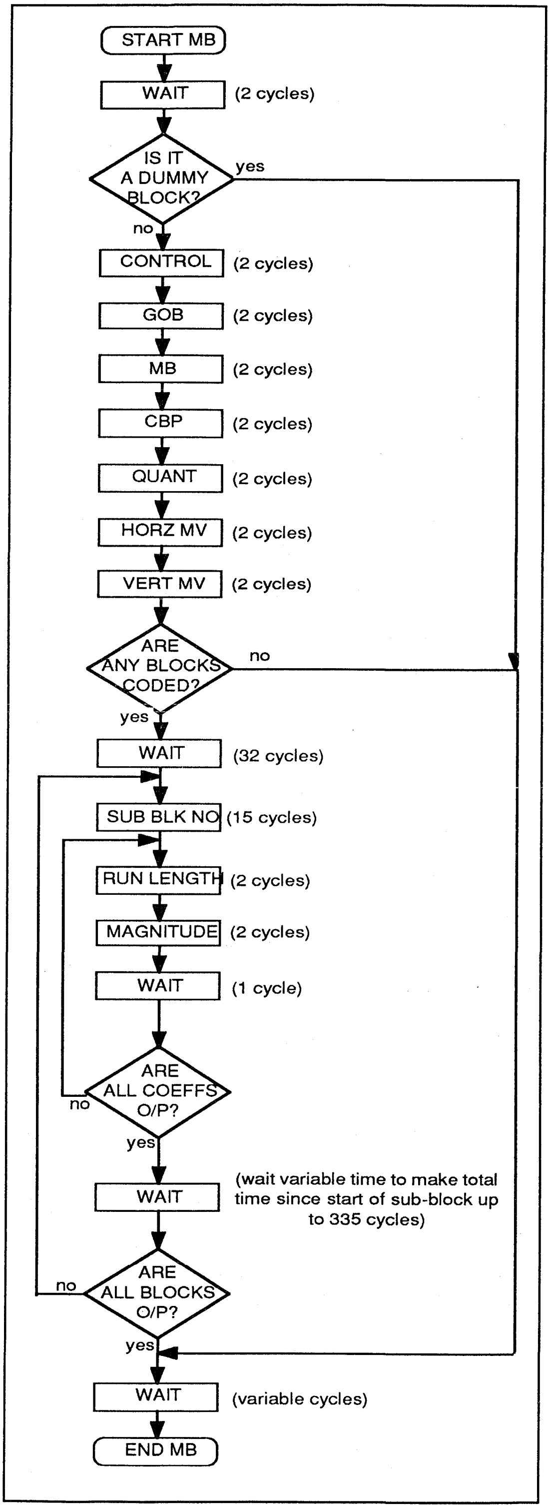

Fig 8 : DBUS Port Flow Chart

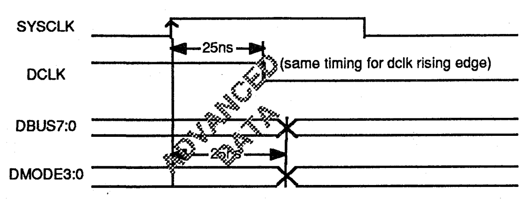

N.B. All timings given are maximum values assuming a 30pF load.

Fig 9 : Output Delays at DPORT

MB Number : After the GOB Number, the macroblock Number is output on DBUS5:0 (DBUS5 is MSB).

Control Decisions : This byte shows which control decisions have been taken for the forthcoming macroblock. DBUS0 will be high if a Fixed Macroblock (FIX MB) was enforced i.e. no new data will be transmitted this macroblock. DBUS1 indicates whether Inter (high) or Intra (low) coding was used for the macroblock. DBUS2 will be high if the macroblock was filtered, and DBUS3 will be high if motion compensation was used. DBUS5 will be high if the current frame is being coded in FAST UPDATE mode. In this mode the complete frame will be intra coded. DBUS6 will be high if the current frame is a SKIP FRAME i.e. not being coded - so no coefficients will be transmitted. DBUS4 and DBUS7 are not used.

Quant Value : The quantisation value used in processing the current macroblock is output on DBUS4:0 (DBUS4 is MSB). This represents an actual quantisation level between 2 and 62 as defined in H.261.

Horizontal MV : If motion compensation is used, the horizontal component of the motion vector will be output on DBUS4:0 (DBUS4 is MSB). This 5 bit value represents a two's complement number in the range +/-15 (although only vectors in the range -8 to +7 are currently possible with the VP2611).

Vertical MV : If motion compensation is used, the vertical component of the motion vector will be output on DBUS4:0 (DBUS4 is MSB). This 5 bit value represents a two's complement number in the range ±15 (although only vectors in the range ±7 are currently possible with the VP2611).

Coded Block Pattern : This byte contains a 6 bit linear code that indicates which of the sub-blocks actually contain

CBUS3:0	INSTRUCTION
0000	Input VAR Threshold
0001	Reserved
0010	Input Inter Quantiser
0011	Input Intra Quantiser
0100	Input Setup Data
0101	Input Control Functions
0110	Reserved
0111	Reserved
1000	Output GOB Number
1001	Output MB Number
1010	Reserved
1011	Output Control Decisions
1100	Output Setup Data
1101	Reserved
1110	Reserved
1111	Reserved

Table 3 : CBUS Instruction Codes

coded data. DBUS6 will be high if sub-block 1 contains coded data, through to DBUS 1 being high if sub-block 6 contains coded data. DBUS7 and DBUS0 are not used.

Sub-block Number : An identifier for the run length coded coefficients which are about to be made available. DBUS 2:0 contain the coded sub-block number from 1 to 6. All zero sub-blocks will not be produced at the outputs, and their corresponding numbers will not appear.

Zero Run Count : The number of zero valued coefficents preceding the next non zero coefficient is available on DBUS5:0 (DBUS5 is MSB). Normally, DBUS7:6 are low, except to signify the end of a Sub-block, when they will both be high. Zero Run Count is always followed by a coefficient, even at the end of a sub-block.

RLC Coefficient : This byte contains the 8 bit coefficient value. It will always be a non-zero value, except when the previous Zero Run Count signalled the end of sub-Block. A zero value is then possible since, as stated above, the run count is always followed by a coefficient byte, which may have be zero if the last coefficient is zero.

Wait State : This indicates that no valid data is being output from the DBUS port during this cycle. No DCLK is produced for this state.

Pins which are "not used" for certain functions will be forced low.

This diagram shows a typical Sub-block being output from the VP2611.

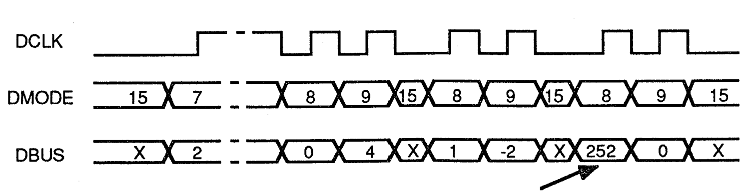

Fig 10 : DBUS Timing

CBUS Control Port

The CBUS control port is used to input control and setup information and also to output status information. In order to save on pin count, a microprocessor driving this port is required to execute two I/O instructions in order to transfer a single byte of information to or from the device. The first transfer is always a write operation, with a low level on the single address line which is used by the interface. Data on the bus then defines the instructions listed in Table 3. The second transfer can be a read or write operation as necessary, but the address line must then be high with the set up time given in Figure 10.

In addition to the single addresss line (CADR), data transfers use a control strobe (CSTR) which is only effective when a chip enable is present (CEN). Detailed timing information is given in Figure 10, and when writing data or instructions to the VP2615 the set up and hold times which are referenced to the rising edge of CSTR must be maintained.

When a write instruction has been defined CADR should be pulled high, valid data presented to CBUS7:0 and then strobed in using CSTR. Other system I/O transfers can occur between defining a write operation and supplying the data to be written, assuming CEN is not active during those other transfers. If CSTR does not go active because of I/O transfers to other devices, then CEN can remain active low between the instruction and data.

When a read instruction has been specified the requested data will then be output on CBUS7:0 after the access time specified from CEN going low, assuming that CADR was already high. Otherwise the data will become valid after the access time specified from CADR going high after CEN was

WRITING DATA FROM THE CBUS:

This diagram shows a typical instruction and associated data field being written to the device.

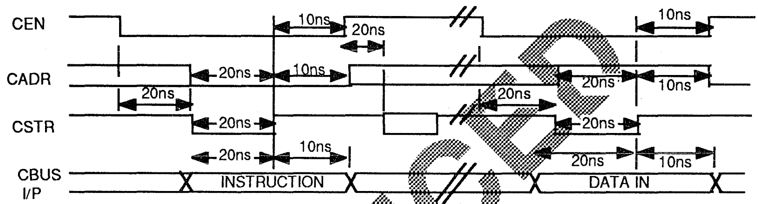

READING INFORMATION ON CBUS :

This diagram shows a typical instruction and associated data field being read from the device.

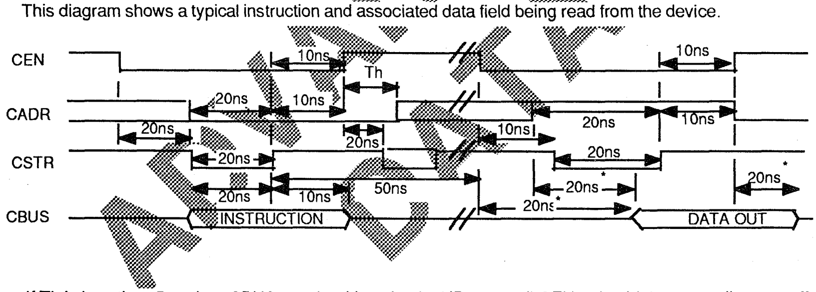

If Th is less than 5 ns then CBUS may be driven by the VP2611until CEN going high eventually turns off the drivers. It will not prevent correct data being read when CEN again goes active

N.B. All timings shown are minimum values except those marked * which are maximums.

Fig 11 : Use of the Control Port

low. Note that in the data read phase CADR must always go high before CSTR goes high, with the set up time specified. When CEN goes high, or CADR goes low, the CBUS will go high impedance after the delay specified.

Note that the access times under the conditions given above are only true when the gap between CSTR going high in the instruction phase, and CEN going low in the data phase, is greater than the minimum specified in figure 10.

Only the four LSBs, CBUS3:0, are used when writing instructions to the VP2611. The remaining bits, CBUS7:4, should be pulled low while the instruction is strobed into the VP2611.

The instructions listed in Table 3 are described below in greater detail;

Input VAR Threshold: VAR is the difference between the best fit MAE value and the variance of the current macroblock. The VAR Threshold is the best fit MAE value below which Inter Frame Prediction is always used, no matter what the variance of the current block. Above this threshold inter mode coding is only used if the best fit value is less than the current block variance. The default value is 3, within a range of 0 - 15 using the four most significant bits of the 14 bit value. This instruction allows the default value to be changed within the range 0 - 15 using bits 3:0 in the following data byte.

Input Inter Quantiser: Coefficents of inter coded macroblocks will be quantized using the value on CBUS4:0 following this instruction. Internally this represents a 6 bit number with the lsb always zero, giving a value between 0 and 62 in steps of two. Where only one quantization value is to be used for both inter and intra cases, this instruction should be used. On reset the value will default to the maximum allowed.

Input Intra Quantiser: This instruction is similar to the above, except that it defines the quantization level for intra mode coding when it is to be different to that of inter mode coding.

Input Setup Data: This instruction allows several user defined options to be specified, using individual bits in the following data word. If CBUS0 is LOW the device will work in full CIF mode , if HIGH it uses the QCIF mode. If CBUS3 is HIGH both inter and intra quantization values will be used, otherwise a common value will be used. If CBUS5 is high then the motion compensation circuits will be disabled. If CBUS6 is high, then the device will be configured to use 256K x 16 or 256K x 4 DRAM's, otherwise it will assume the use of two 64K x 16 DRAM's. The default conditions after RESET are those selected by the Low level. CBUS1, CBUS2, CBUS4 and CBUS7 are not used but must be low during the definition phase. This instruction may be used any time after RESET has gone high.

Input Control Functions: This instruction specifies several control options using individual bits in the following data word. If CBUS0 is HIGH then the on board Inter/Intra Decision circuitry will be overidden according to CBUS1; if CBUS1 is HIGH then all subsequent macroblocks will be intra coded, if it is LOW they will be inter coded. When CBUS2 is HIGH the on-board Filter Decision circuitry is overidden according to CBUS3; if CBUS3 is HIGH then the filter will be forced on, if it is LOW the filter will be forced off. If CBUS4 is HIGH then FIX MB will be implemented, and no new data from the current macroblock will be coded. A two macroblock delay exists between defining the Force Inter/Intra, Force Filter or FIX MB decisions through the control bus and data being affected at the outputs. These decisions will stand for all subsequent macroblocks until they are again changed. If CBUS5 is HIGH a FAST UPDATE will be performed on the next frame and all blocks will be coded in intra mode. If CBUS6 is HIGH then the next frame will not be transmitted (SKIP FRAME). Note that these two global frame bits do not take effect until the start of the next frame, and stay in effect for all frames until they are removed. If CBUS7 is HIGH, then the on-board Force Update Controller will be overidden, and the user will have to enforce their own Force Update policy using the Force Intra command. RESET will cause the options to default to those defined by the LOW state. Note that SKIP FRAME has priority over any other bits and that FIXMB has priority over all bits bar SKIP FRAME.

Output GOB Number: This instruction will output the GOB Number on CBUS3:0, for the data currently being output on DBUS. CBUS7:4 are not used (always low).

Output MB Number: This instruction will output the macroblock number on CBUS5:0, for the data currently being output on DBUS. If CBUS6 is low it indicates that the macroblock number has just changed, or is about to change. New Quantization Value or Control Function words should not be written at this time since it is uncertain which macroblock they will effect. CBUS7 is not used (always low).

Output Control Decisions: This instruction will output the details of several control decisions on the CBUS. CBUS0 shows whether the MacroBlock currently being output on DBUS was inter or intra coded (0=Intra). CBUS1 shows whether motion compensation was used (1=MC used). CBUS3 shows whether the macroblock was passed through the loop filter or not (1=Filtered). CBUS4 will be high if the FIX MB instruction was enforced. CBUS5 will be high if FAST UPDATE is currently being undertaken. CBUS6 will be high if SKIP FRAME is in force. CBUS2 and CBUS7 are not used.

Output Setup Data: This instruction allows the user to verify the internal setup previously selected. The bits have the same significance as in the Input Setup Data Instruction.

Initialising the VP2611

On power-up, RESET should be low and must remain low for at least 2064 cycles of SYSCLK. After RESET is pulled high, FRMIN may be activated to start the first frame. Before activating FRMIN for the first time, it is advisable to use the CBUS to implement a FAST UPDATE for the first frame (i.e. all blocks Intra coded).

JTAG Test Interface

The VP2611 includes a test interface consisting of a boundary scan loop of test registers placed between the pads and the core of the chip. The control of this loop is fully JTAG/ IEEE 1149-1 1990 compatible. Please refer to this document for a full description of the standard.

The interface has five dedicated pins: TMS, TDI, TDO, TCK and TRST. The TRST pin is an independent reset for the interface controller and should be pulsed low, soon after power up; if the JTAG interface is not to be used it can be tied low permanently. The TDI pin is the input for shifting in serial instruction and test data; TDO the output for test data. The TCK pin is the independent clock for the test interface and registers, and TMS the mode select signal.

TDI and TMS are clocked in on the rising edge of TCK, and all output transitions on TDO happen on its falling edge.

Instructions are clocked into the 8 bit instruction register (no parity bit) and the following are available.

Instruction Register (MSB first)	Name
11111111	BYPASS
00000000	EXTEST
01000000	INTEST
XX001011	SAMPLE/PRELOAD

Timing details for the JTAG control signals are shown in fig 11.The maximum TCK frequency is 5 MHz.

The test registers, their positions in the boundary loop and the corresponding i/o pad are detailed in Table 4. Note that the three state control signals also have test registers associated with them. The order given in Table 4 determines the serial data stream needed for JTAG testing.

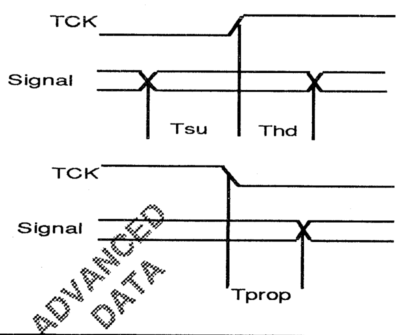

	Tsu	Thd	Tprop
TMS toTCK timing	15	5	
TDI to TCK timing	15	5	
Chip i/p to TCK timing	15	5	
TCK to TDO timing			20

Fig 12 : JTAG Interface timing

Pad	Type	Reg No	Pad	Type	Reg No	Pad	Type	Reg No
RESET	IN	1	DBUS5	IN	32		IN	63
CADR	IN	2	DBUS6	IN	33	SW1	OP	64
CSTR	IN	3	DBUS7	IN	34		IN	65
CEN	IN	4	SW15	OP	35	SW0	OP	66
CBUS0	OP	5		TRI	36		IN	67
	TRI	6		IN	37	DHZ	TRI	68
	IN	7	SW14	OP	38	RAS	OP	69
CBUS1	OP	8		IN	39	CAS	OP	70
	IN	9	SW13	OP	40	RW1	OP	71
CBUS2	OP	10		IN	41	RW2	OP	72
	IN	11	SW12	OP	42	OE1	OP	73
CBUS3	OP	12		IN	43	OE2	OP	74
	IN	13	SW11	OP	44	ADR0	OP	75
CBUS4	OP	14		IN	45	ADR1	OP	76
	IN	15	SW10	OP	46	ADR2	OP	77
CBUS5	OP	16		IN	47	ADR3	OP	78
	IN	17	SW9	OP	48	ADR4	OP	79
CBUS6	OP	18		IN	49	ADR5	OP	80
	IN	19	SW8	OP	50	ADR6	OP	81
CBUS7	OP	20		IN	51	ADR7	OP	82
	IN	21	SW7	OP	52	PCLK	IN	83
DCLK	IN	22		IN	53	YUV7	IN	84
DMODE0	IN	23	SW6	OP	54	YUV6	IN	85
DMODE1	IN	24		IN	55	YUV5	IN	86
DMODE2	IN	25	SW5	OP	56	YUV4	IN	87
DMODE3	IN	26		IN	57	YUV3	IN	88
DBUS0	IN	27	SW4	OP	58	YUV2	IN	89
DBUS1	IN	28		IN	59	YUV1	IN	90
DBUS2	IN	29	SW3	OP	60	YUV0	IN	91
DBUS3	IN	30		IN	61	SYSCLK	IN	92
DBUS4	IN	31	SW2	OP	62	FRMIN	IN	93
						REQYUV	OP	94

Table 4 Pin and JTAG test registers

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage VDD	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to VDD+ 0.5V
Output voltage V_{OUT}	-0.5V to VDD + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	0°C to 70°C
Junction temperature	100°C
Package power dissipation	3000mW

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as negative into the device.

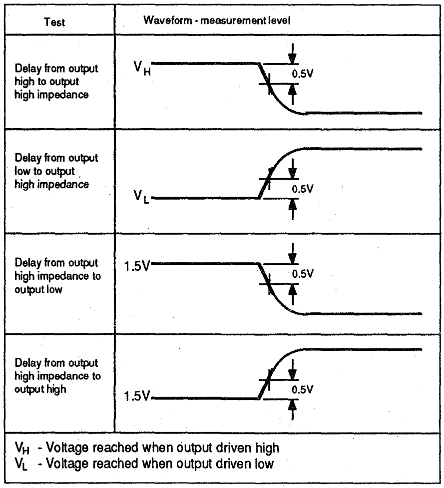

STATIC ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)
Tamb = 0 C to +70°C VDD = 5.0v ± 5%

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	3.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.0		-	V	3.0V for SYSCLK and PCLK
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{DD}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{DD}
Output S/C current	I_{SC}	10		300	mA	V_{DD} = Max

Pin	Function	Pin	Function	Pin	Function	Pin	Function
1	SW3	34	NC	67	CSTR	100	NC
2	SW4	35	NC	68	CADR	101	NC
3	NC	36	NC	69	NC	102	NC
4	SW5	37	DBUS1	70	RESET	103	ADR7
5	GND	38	DBUS0	71	VDD	104	ADR6
6	VDD	39	DMODE3	72	GND	105	ADR5
7	SW6	40	GND	73	TCK	106	VDD
8	SW7	41	VDD	74	NC	107	GND
9	NC	42	DMODE2	75	TMS	108	NC
10	SW8	43	DMODE1	76	TDI	109	ADR4
11	SW9	44	NC	77	TDO	110	ADR3
12	SW10	45	DMODE0	78	(CLK54)	111	ADR2
13	SW11	46	DCLK	79	REQYUV	112	ADR1
14	NC	47	VDD	80	FRMIN	113	GND
15	GND	48	CBUS7	81	NC	114	ADR0
16	SW12	49	CBUS6	82	VDD	115	VDD
17	VDD	50	GND	83	SYSCLK	116	GND
18	NC	51	VDD	84	NC	117	OE2
19	SW13	52	CBUS5	85	GND	118	OE1
20	SW14	53	GND	86	NC	119	VDD
21	NC	54	CBUS4	87	YUV0	120	RW2
22	SW15	55	CBUS3	88	YUV1	121	RW1
23	NC	56	CBUS2	89	YUV2	122	CAS
24	DBUS7	57	CBUS1	90	YUV3	123	RAS
25	DBUS6	58	NC	91	NC	124	VDD
26	NC	59	GND	92	YUV4	125	GND
27	DBUS5	60	VDD	93	YUV5	126	NC
28	GND	61	CBUS0	94	NC	127	SW0
29	VDD	62	TRST	95	VDD	128	SW1
30	DBUS4	63	CEN	96	GND	129	SW2
31	NC	64	NC	97	YUV6	130	NC
32	DBUS3	65	NC	98	YUV7	131	NC
33	DBUS2	66	NC	99	PCLK	132	NC

Pin out table for 132 pin power ceramic QFP package (GC132)

ORDERING INFORMATION

VP2611 CG GCAR (Commercial - GG Pcakage)

NOTE: Prior to completion of full device characterisation, pre-production parts will be designated VP2611 PR GCAR

VP2612

VIDEO MULTIPLEXER

FEATURES

- Fully integrated H261 video multiplexer
- Inputs data direct from VP2611 source coder
- Output to X21 line buffers
- Line rates from 64kbits/s up to 2Mbits/s
- 100 Pin Quad Flatpack

ASSOCIATED PRODUCTS

- VP2611 H.261 Encoder
- VP2615 H.261 Decoder
- VP2614 Video Demultiplexer
- VP520 CIF/QCIF Converter
- VP530 PAL/ NTSC Encoder

The VP2612 Video Multiplexer forms part of the GPS chip-set for video conferencing, video telephony, and multimedia applications. This chip set implements the H261 standard for video compression for line rates of between 64K and 2M bits per second. With a 27MHz clock rate full CIF resolution images can be coded at a frame rate of up to 30Hz.

The device contains all the elements necessary to convert the run length coded data from the VP2611 source coder into an H261 compatible bit stream. It also calculates the differential motion vectors and macroblock addresses from the absolute values received from the VP2611. These values are variable length coded, and bit packed for temporary storage in the transmission buffer. The size of this buffer can be either 256Kbits or 512Kbits. Data from the transmission buffer is output through an X21 compatible serial interface, and consists of frames containing framing bits, data, and the BCH (511,493) forward correction code.

The system processor interface is used to write data for PTYPE, PSPARE, GSPARE, and to select the source of temporal reference. The interface can also be used to monitor the pointers into the transmission buffer, so that the buffer fullness can be controlled using proprietary software algorithms. In addition to the bus interface, flags are supplied which indicate the start of each macroblock, each FEC stuffed frame, the number of bits per picture is reaching the allowable maximum, and impending buffer overflow.

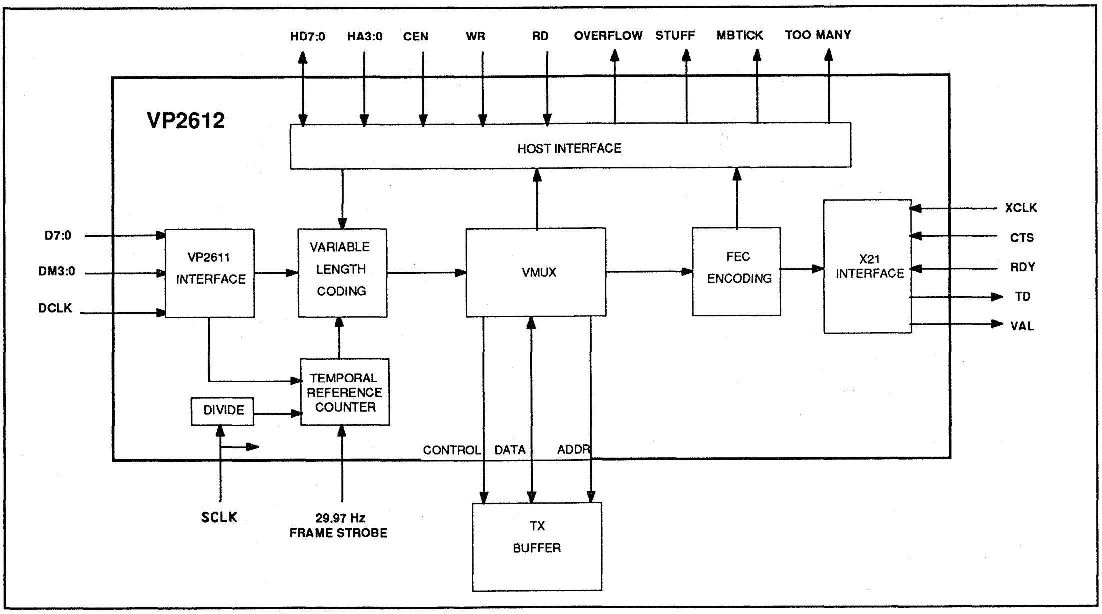

Fig. 1. VP2612 Video Multiplexer

PIN DESCRIPTIONS

DBUS7:0 The input data bus from VP2611. The data type is defined by the value present on DMODE3:0

DMODE3:0 These inputs define the data type present on the data bus D7:0. Polarities are given in Table 1.

DCLK A strobe for DM3:0 and DBUS7:0.The high going edge latches data into the VMUX.

HD7:0 A bidirectional tri-state data bus connecting the VMUX to the system processor.

HA3:0 Four system processor address bits used to address internal registers.

$\overline{\text{WR}}$ An active low write strobe from the system processor.

$\overline{\text{RD}}$ An active low read strobe from the system processor.

$\overline{\text{CEN}}$ An active low chip select input from the system processor.

OVR An active high output which signals impending buffer overflow.

STUFF An active high output that signals that FEC stuffing is occuring.

MTICK An output which pulses high for every macroblock received.

TOOM This active high output indicates that the picture is likely to exceed the allowable number of bits per picture.

VAL This line is taken low to indicate that the VMUX is ready to transmit valid data. The C line in an X21 system.

TD This is the serial data output from the VMUX.

CTS Indicates that the receiver can accept data. The I line in an X21 system.

RDY Indicates that the receiver can accept data. The R line in an X21 system.

XCLK X21 line clock input. 0 to 2.048MHz.

SCLK System clock input. Only the high going edge is used internally, apart from TXWE generation.

FS A 29.97 Hz frame strobe for the temporal reference counter. Must be high for at least 4 SCLK periods.

$\overline{\text{RES}}$ Active low reset signal. Must be low for at least 16 SCLK periods.

TXA14:0 Address output to Transmission buffer.

TXD7:0 Bidirectional data interface to Transmission buffer.

TXE1 Active low chip enable for the Transmission buffer. If a 256kBit buffer is being used this Chip Enable should be used.

$\overline{\text{TXE2}}$ Active low chip enable for the Transmission buffer. This is used for the optional second memory chip, if a 512kBit buffer is being used.

$\overline{\text{TXWE}}$ Active low write enable for the Transmission buffer.

$\overline{\text{TXOE}}$ Active low O/P enable for the Transmission buffer.

TCK Test clock for JTAG.

TMS Test mode select.

TDI Test data I/P.

TDO Test data O/P.

$\overline{\text{TRST}}$ JTAG reset.

TOE When low ALL O/P pins are high impedance.

NOTE: "Barred" active low signals do not appear with a bar in the main body of the text.

OPERATIONS OF MAJOR BLOCKS

Variable Length Coding

This block is responsible for ordering the data from the VP2611 Encoder into the correct sequence for the H261 bit stream, and for performing the variable length coding. It also uses data supplied by the system controller and the Temporal Reference Counter.

Data for PTYPE, PSPARE, GSPARE is only obtained from the system controller, and only 8 bits of PSPARE and GSPARE information can be transmitted per picture or GOB respectively. The temporal reference can either be obtained from an internal counter, from the VP2611 outputs, or can be written by the system controller. The actual source is determined by bits in a control register as described later. The internal counter is clocked from either a frame clock with a maximum frequency of 29.97Hz, or a 29.97Hz clock derived from the 27MHz system clock, or it simply counts H.261 frames from the encoder.

There is no support provided for macroblock stuffing, however FEC stuffing is implemented, and can be used to provide bit stuffing.

This block is also responsible for converting the absolute values that are output from the V2611 into the relative values that are required in parts of the H261 bitstream. The VMUX has been designed so that it can accept ±15 motion vectors, rather than the +7/-8 motion vectors produced by the VP2611. Thus it will be compatible with any future upgrades to the VP2611 that increase the size of the motion estimator search window.

VMUX Block

The VMUX section performs the bit packing on the data coming from the variable length coder. This data is in the form of a delimiter and a variable number of valid bits. The VMUX section packs these variable length fields into bytes for storage in the transmission buffer.

The transmission buffer is controlled by this block. It thus generates read and write pointers, and performs the arbitration between read and write operations. Buffer level monitoring is, however, done by the FEC block as described later.

The two address pointers can be read by the system processor, thus allowing the level of the buffer to be monitored. These are provided as 16 bit words with no truncation, and thus require two bytes. The 16 bit value is internally frozen when the most significant byte is requested by the system processor, and for accuracy the write pointer should be read first. There is also a control register bit which selects a buffer size of either 256kbits or 512kbits.

FEC Block

The FEC section performs the framing, and adds the error correction parity bits. If sufficient data for a frame is not available in the transmission buffer, then the frame will be stuffed automatically. There is no absolute threshold at which the FEC will start to stuff, as the buffer level monitor in the FEC only works to a resolution of ±128bits. FEC stuffing can also be forced by setting the "Force FEC stuffing" bit in the VMUX/FEC control register.

If the buffer level reaches a threshold, internally set to 512 short of the buffer being full, the OVERFLOW output is asserted.

DMODE3:0	FUNCTION
0000	GOB Number
0001	MB Number
0010	Control Decisions
0011	Quant Value
0100	Horizontal MV
0101	Vertical MV
0110	Coded Blk Pattern
0111	Sub-Block No.
1000	Zero Run Count
1001	RLC Coefficient
1010	Not Used
1011	Not Used
1100	Not Used
1101	Not Used
1110	Not Used
1111	Wait State

Table 1

This is to warn the system processor that drastic action is needed to avert a buffer overflow, which will result in corruption and loss of data. Since the buffer level monitor only works to resolution of ±128bits, then the overflow detection can only be accurate to ±128bits.

VP2611 Interface

The VMUX has been designed to interface directly to the VP2611 encoder, with no buffering. The interface consists of two buses DBUS7:0 and DMODE3:0, and a strobe signal DCLK. The value on DMODE3:0 identifies the data type on DBUS7:0 during the same period (see Table 1).

The output of the VP2611 is structured such that the data on DBUS7:0 and DMODE3:0 is always valid for at least two cycles, and DCLK is high for minimum of one cycle. The rising DCLK edge occurs one cycle after DBUS7:0 and DMODE3:0 are valid, as shown in Figure 2.

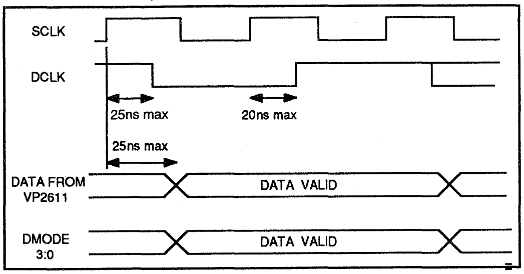

Figure 2. DBUS Timing

The sequence of events, and the duration of each event, is shown in Figure 3. These duration times have been chosen to satisfy the internal requirements of the VP2612, and Wait States are inserted such that the time to transfer a macroblock is always 2064 SCLK periods.

The parameters used by the VP2612 are described in more detail below;

GOB Number : The current GOB Number is provided on DBUS3:0 after the Control Decisions byte. (DBUS3 is MSB).

MB Number : After the GOB Number, the macroblock Number is provided on DBUS5:0 (DBUS5 is MSB).

Control Decisions : This byte shows which control decisions have been taken for the forthcoming macroblock, and is the first in the sequence. DBUS0 will be high if a Fixed Macroblock (FIX MB) was enforced i.e. no new data will be transmitted this macroblock. DBUS1 indicates whether Inter (high) or Intra (low) coding was used for the macroblock. DBUS2 will be high if the macroblock was filtered, and DBUS3 will be high if motion compensation was used. DBUS5 will be high if the current frame is being coded in FAST UPDATE mode. In this mode the complete frame will be intra coded. DBUS6 will be high if the current frame is a SKIP FRAME i.e. not being coded - so no coefficients will be transmitted. DBUS4 and DBUS7 are not used.

Quant Value : The quantisation value used in processing the current macroblock is provided on DBUS4:0 (DBUS4 is MSB). This represents an actual quantisation level between 2 and 62 as defined in H.261.

Horizontal MV : If motion compensation was used the horizontal component of the motion vector will be provided on DBUS4:0 (DBUS4 is MSB). (This 5 bit value represents a two's complement number in the range (-15 to +15) (although only vectors in the range +7/-8 are currently possible with the VP2611). If motion compensation was not used this is a don't care value.

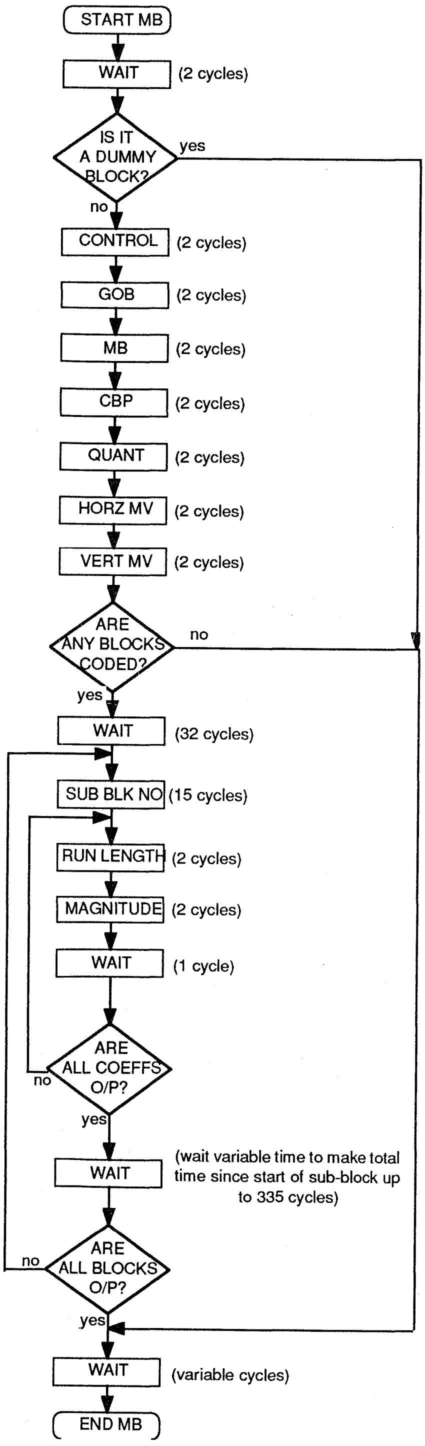

Fig 3 : DBUS Port Flow Chart

Vertical MV : If motion compensation was used the vertical component of the motion vector will be provided on DBUS4:0 (DBUS4 is MSB). (This 5 bit value represents a two's complement number in the range ±15 (although only vectors in the range (±7) are currently possible with the VP2611). If motion compensation was not used this is a don't care value.

Coded Block Pattern : This byte contains a 6 bit linear code that indicates which of the sub-blocks actually contain coded data. DBUS6 will be high if sub-block 1 contains coded data, through to DBUS1 being high if sub-block 6 contains coded data. DBUS7 and DBUS0 are not used.

Sub-block Number : An identifier for the run length coded coefficients which are about to be made available. DBUS2:0 contain the coded sub-block number from 1 to 6. All zero sub-blocks will not be produced, and their corresponding numbers will not appear.

Zero Run Count : The number of zero valued coefficents preceding the next non zero coefficient is provided on DBUS5:0 (DBUS5 is MSB). Normally, DBUS7:6 are low, except to signify the end of a Sub-block, when they will both be high. Zero Run Count is always followed by a coefficient, even at the end of a sub-block.

RLC Coefficient : This byte contains the 8 bit coefficient value. It will always be a non-zero value, except when the previous Zero Run Count signalled the end of sub-Block. A zero value is then possible since, as stated above, the run count is always followed by a coefficient byte, which may be zero if the last coefficient is zero.

Wait State : This indicates that no valid data is being output from the DBUS port during this cycle. No DCLK is produced for this state.

SYSTEM PROCESSOR INTERFACE

The system processor interface is a memory mapped microprocessor compatible interface. It has been designed for use with any system processor, and consists of the following buses and signals:

HD7:0	Processor Data Bus
HA3:0	LSBs of addresses bus
WR	Active Low Write strobe
RD	Active Low Read strobe
CEN	Decoded Active Low chip select

Detailed interface timing is shown in Figure 4. Since there are several internal pipeline registers which are clocked by SCLK, then access times and strobe widths are dependent on the period of SCLK.

Table 2 shows the addresses used for each of the user accessible registers, and the function of each register is described in detail below;

PTYPE This is the picture type as defined in H261. The bits are assigned as follows:

Address	Function	Read / Write
0	PTYPE	W
1	Temporal Reference	W
2	PSPARE	W
3	TR Source	W
4	GSPARE	W
5	Not Used	
6	Not Used	
7	Not Used	
8	TX-buffer Write Address MSB	R
9	TX-buffer Write Address LSB*	R
A	TX-buffer Read Address MSB	R
B	TX-buffer Read Address LSB*	R
C	FEC / VMUX status word	W
D	Bits per Picture Threshold	W
E	Not Used	
F	Not Used	
* N.B. The LSB must be read after the appropriate MSB.		

Table 2. Addresss Locations

Bit 0	Split screen indicator, "0" off, "1" on.
Bit 1	Document camera indicator, "0" off, "1" on.
Bit 2	Freeze picture release, "0" off, "1" on.
Bit 3	Source Format, "0" QCIF, "1" CIF.
Bit 5:4	Both are set to one as presently defined in the H261 specification.

These values can be changed at will by the system processor, and will be transmitted at the start of each picture.

Temporal Reference If the temporal reference is being written from the system processor, then the 5 LSB's in this register are used to define the next temporal reference value to be transmitted.

PSPARE This register holds 8 bits of PSPARE information which may be transmitted for each picture. The data in the PSPARE register will be transmitted at the start of the next picture after it has been written. Once an item of data has been transmitted, it will not be re-transmitted until data is written from the system processor. It is the responsibility of the system processor to ensure that it does not rewrite to this register before the previous value has been transmitted. This can be done by utilizing a frame interrupt from the video source in conjunction with the MBTICK output from the VMUX.

TR Source The 3 LSB's in this register define the source for the strobe used by the 5 bit temporal reference counter. When the sytem processor is selected, the counter value is replaced by the contents of the Temporal Reference Register.

VALUE	SOURCE
0XX	System Processor
100	Actual coded frames from the VP2611 are counted
101	SCLK is divided down to provide a 29.97 Hz frame strobe
110	The strobe is provided by the frame strobe input pin (FS)
111	Illegal

GSPARE This register holds 8 bits of GSPARE information which may be transmitted every GOB. Once written the data is transmitted at the start of the next GOB, but will not be re-transmitted until the system processor again writes to this address. The system processor must ensure that data is not overwritten before it is used.

TX Buffer Addresses These allow the system processor to monitor the level of the buffer. The write pointer should be read first to minimize the error between the the two

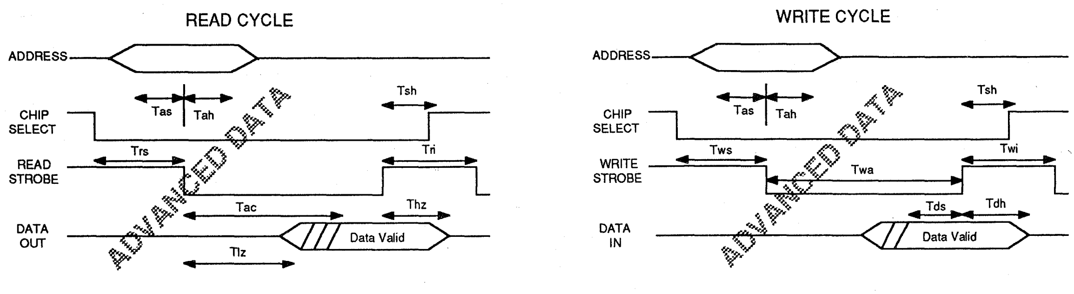

CHARACTERISTIC	SYMBOL	MIN	MAX
Addresss Set Up Time	Tas	10ns	
Address Hold Time	Tah	10ns	
Cip Select Set Up Time	Trs	10ns	
Chip Select Hold Time	Tsh	2ns	
Strobe In active Time	Tri	Øns	
Data Access Time	Tac	4Øns	10 +5Øns
Delay to O/P's low Z	Tlz		25ns
Delay to O/P's high Z	Thz		25ns

NOTE

Ø is the period of the input clock

CHARACTERISTIC	SYMBOL	MIN	MAX
Addresss Set Up Time	Tas	10ns	
Address Hold Time	Tah	10ns	
Chip Select Set Up Time	Tws	10ns	
Chip Select Hold Time	Tsh	2ns	
Strobe In active Time	Twi	2Øns	
Strobe Active Time	Twa	2Øns	
Data Set Up Time	Tds	10ns	
Data Hold Time	Tdh	10ns	

Figure 4. Host Controller Timing

values. With a 2Mbits/sec line the error will increase at a rate of 0.25 bytes per microsecond. Reading the most significant bytes will trigger the internal latching of the least significant bytes.

FEC / VMUX Control This register controls the operation of the transmission buffer and the FEC block. Actions taken when bits are set are given below;

BIT	FUNCTION
0	Select 512K buffer. The buffer size must not be changed during normal operation and must be defined within 2.4 ms of reset.
1	Enable FEC framing.The option to disable FEC framing is only provided as a test mode.
2	Force FEC stuffing. If force FEC stuffing is selected it will start at the beginning of the next frame and only stop at the start of subsequent frames. The system processor must ensure that the transmission buffer does not overflow with forced stuffing. In normal operation FEC stuffing only occurs when there is insufficient data in the transmission buffer.

Bits Per Picture Register When the number of bits which have been coded has been subtracted from the maximum possible (as defined by H.261), and the result reaches the value in this register, then the TOO MANY interrupt will be generated. The programmed value thus defines in Kbits the number of bits which may still be generated before reaching the maximum allowed. The default value is 8 Kbits, and the maximum number used internally changes between CIF and QCIF.

INTERRUPT OUTPUTS

The special signals listed below are provided to drive timers and interrupt inputs on the system processor.

OVERFLOW	**(OVR)**
FEC STUFF	**(STUFF)**
MACROBLOCK TICK	**(MTICK)**
TOO MANY	**(TOOM)**

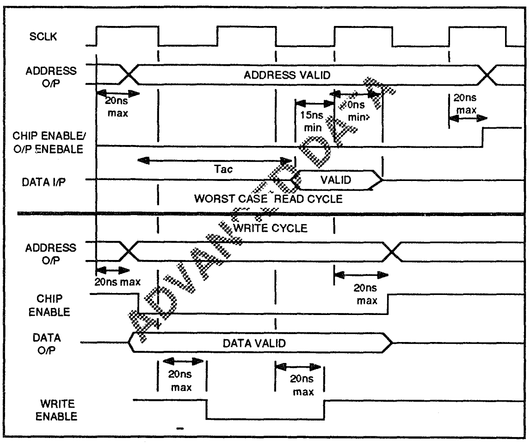

Figure 5. Transmission Buffer Timing

They perform the following functions:

OVR This line signals an impending buffer overflow. When the buffer reaches 512±128 bits from being full this line will be taken high, and will remain high until the buffer level falls below the threshold. It is intended that this line be used as a processor interrupt, to signal that drastic action must be taken.

STUFF This line signals that FEC stuffing is occuring, and can be used to monitor the amount of stuffing being performed. It will pulse high once at the start of each FEC stuffed frame, the length of the pulse being one line clock period. It is intended that this should be used to clock a system processor counter, to keep a running total of the number of FEC stuffed frames.

MTICK This output pulses high once for every Macroblock received from the VP2611. The pulse is 3 clock cycles in duration, and the leading edge will occur 6 SCLK cycles after the Macroblock address was received from the VP2611. It is anticipated that this should be used to clock a counter in the system processor, so that the system processor can keep track of which MB is being processed. In conjunction with the frame pulse this will enable the system processor to write information to the VP2611 at appropriate times.

TOOM This signal indicates that the present picture has reached a threshold relative to the maximum number of bits per picture allowed by H.261 (256k if CIF, 64k if QCIF). It is set when the number of bits remaining before the maximum will be exceeded reaches the value in the Bits Per Picture Register, and stays high until the end of the current picture.

TRANSMISSION BUFFER INTERFACE

The transmission buffer can consist of either one or two 32K x 8 bit static RAMs. Fifteen address outputs are provided for direct connection to the memory devices, and two RAM select pins are provided to define the device in use. If only a single device is being used then CE2 is redundant.
An internal FIFO is provided to average out high speed bursts of transmission buffer cycles. This allows the external SRAM read cycles to occupy at least three SCLK periods. Detailed timing for the buffer is given in Figure 5, and shows that with a 27 MHz clock the RAM must have an access time of less than 39 nanoseconds. Figure 5 illustrates the worst case read access time, which occurs when a second read cycle follows the first without an intermediate write cycle. Chip enable and output enable remain low from the first read cycle.
The write cycle uses two SCLK periods and requires the use of both the falling and rising edges of SCLK. The Write Enable output thus remains active for one SCLK period minus differential rising and falling edge delays. These are limited to two nanoseconds. Note that when consecutive read or write operations take place then Chip Enable will remain active, and not go inactive between cycles.

LINE INTERFACE

A serial interface is provided which facilitates the operation of the encoder and decoder in a back to back configuration. It is similar in operation to an X21 interface but does not support balanced lines. Alternatively the interface can be used in a simple serial manner by tying the control lines to fixed logic levels. It uses the following signals:

XCLK Line rate clock
VAL Ready to send
TD Transmitted Data
CTS Clear To Send
RDY Receiver ready

Of these signals XCLK, CTS and RDY are supplied by the receiving device, the latter two indicating that the receiver is ready to accept data. The VAL line is used to signal that the VMUX is ready to start transmitted valid data, and the TD line provides the data. The signaling convention is as follows:

CTS = 1 RDY = 0	Receiving device not ready
CTS = 0 RDY = X	Receiving device ready to accept data
CTS =1 RDY =1	Receiving device ready to accept data

The VAL line is taken high by the reset input, and when the receiving device signals that it is ready to accept data then the VP2612 takes the VAL line low on a falling edge of an XCLK. The data is then clocked out on subsequent falling edges of the XCLK signal, so that it can be sampled by the receiver on the rising edge of the clock.

If a simple serial interface is required, the CTS input should be tied low and the RDY input tied high. It is possible to use a variable rate clock provided the maximum instantaneous bit rate does not exceed 8Mbits/s, and the average clock rate over 32 bits does not exceed 2Mbits/s. Timing delays with respect to the incoming XCLK are shown in Figure 6.

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage VDD	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to VDD + 0.5V
Output voltage V_{OUT}	-0.5V to VDD+ 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	0°C to 70°C
Junction temperature	100°C
Package power dissipation	1000mW

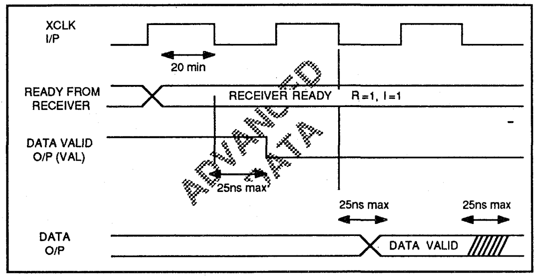

Figure 6. Serial Interface Timing

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as negative into the device.

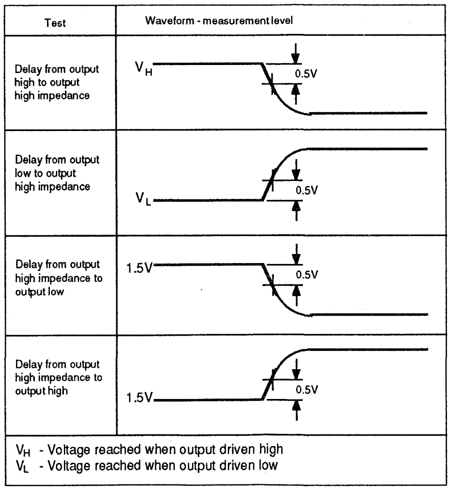

STATIC ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)
Tamb = 0 C to +70°C VDD = 5.0v ± 5%

Characteristic	Symbol	Value Min.	Value Typ.	Value Max.	Units	Conditions
Output high voltage	V_{OH}	3.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.0		-	V	3.0V for SYSCLK and DCLK
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{DD}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{DD}
Output S/C current	I_{SC}	10		300	mA	

PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC
1	N/C	21	GND	41	HD3	61	TXD1	81	TXA1
2	N/C	22	DCLK	42	HD4	62	TXD0	82	TXA0
3	TOE	23	XCLK	43	HD5	63	TXA14	83	TXWE
4	N/C	24	RDY	44	HD6	64	TXA13	84	TXOE
5	OVR	25	CTS	45	HD7	65	TXA12	85	GND
6	DMODE0	26	TD	46	VDD	66	TXA11	86	VDD
7	DMODE1	27	VAL	47	GND	67	TXA10	87	TXE2
8	DMODE2	28	N/C	48	WR	68	TXA9	88	TXE1
9	DMODE3	29	N/C	49	RD	69	TXA8	89	TDI
10	GND	30	N/C	50	CEN	70	TXA7	90	TMS
11	VDD	31	HA0	51	N/C	71	VDD	91	TRST
12	DBUS0	32	HA1	52	N/C	72	GND	92	TCK
13	DBUS1	33	HA2	53	TXD7	73	TXA6	93	TDO
14	DBUS2	34	HA3	54	TXD6	74	TXA5	94	VDD
15	DBUS3	35	SCLK	55	TXD5	75	TXA4	95	GND
16	DBUS4	36	GND	56	TXD4	76	TXA3	96	RES
17	DBUS5	37	VDD	57	TXD3	77	TXA2	97	MTICK
18	DBUS6	38	HD0	58	TXD2	78	N/C	98	STUFF
19	DBUS7	39	HD1	59	GND	79	N/C	99	TOOM
20	VDD	40	HD2	60	VDD	80	N/C	100	FS

Pin Out Diagram

PAD	TYPE	REG No.	PAD	TYPE	REG No.
TXE1	O/P	1	HD5	(input)	46
TXE2	O/P	2	HD4	(output)	47
TX0E	O/P	3	HD4	(input)	48
TXWE	O/P	4	HD3	(output)	49
TXA0	O/P	5	HD3	(input)	50
TXA1	O/P	6	HD2	(output)	51
TXA2	O/P	7	HD2	(input)	52
TXA3	O/P	8	HD1	(output)	53
TXA4	O/P	9	HD1	(input)	54
TXA4	O/P	10	HD0	(output)	55
TXA5	O/P	11	HD0	(input)	56
TXA6	O/P	12	HD	I/P	57
TXA7	O/P	13	SCLK	I/P	58
TXA8	O/P	14	HA3	I/P	59
TXA9	O/P	15	HA2	I/P	60
TXA10	O/P	16	HA1	I/P	61
TXA11	O/P	17	HA0	I/P	62
TXA12	O/P	18	VAL	O/P	63
TXA13	O/P	19	TD	O/P	64
TXA14	O/P	20	CTS	I/P	65
TXD	I/P	21	RDY	I/P	66
TXD0	(input)	22	XCLK	I/P	67
TXD0	(output)	23	DCLK	I/P	68
TXD1	(input)	24	DBUS7	I/P	69
TXD1	(output)	25	DBUS6	I/P	70
TXD2	(input)	26	DBUS5	I/P	71
TXD2	(output)	27	DBUS4	I/P	72
TXD3	(input)	28	DBUS3	I/P	73
TXD3	(output)	29	DBUS2	I/P	74
TXD4	(input)	30	DBUS1	I/P	75
TXD4	(output)	31	DBUS0	I/P	76
TXD5	(input)	32	DMODE3	I/P	77
TXD5	(output)	33	DMODE2	I/P	78
TXD6	(input)	34	DMODE1	I/P	79
TXD6	(output)	35	DMODE0	I/P	80
TXD7	(input)	36	OVR	O/P	81
TXD7	(output)	37	TOE	I/P	82
CEN	I/P	38	TSE	I/P	83
RD	I/P	39	DEN	I/P	84
WR	I/P	40	FS	I/P	85
HD7	(output)	41	TOOM	O/P	86
HD7	(input)	42	STUFF	O/P	87
HD6	(output)	43	MTICK	O/P	88
HD6	(input)	44	RES	I/P	89
HD5	(output)	45			

JTAG Register Allocation

VP2614

H.261 VIDEO DE-MULTIPLEXER

FEATURES

- Fully integrated H.261 video de-multiplexer
- Inputs an H.261 bitstream. Outputs error corrected run length coded coefficients.
- Interfaces directly to the VP2615 H.261 decoder
- Extracts side information and status for transfer to a System Controller
- User definable system level options for proprietary applications
- Average input rates between 40 Kbit /sec and 2Mbit / sec. Maximum peak input rates of 4 Mbit / sec.
- 100 pin quad flatpack

ASSOCIATED PRODUCTS

- VP2611 H.261 Encoder
- VP2612 H.261 Video Multiplexer
- VP2615 H.261 Decoder
- VP520 CIF / QCIF Converter
- VP510 Colour Space Converter
- VP530 PAL/NTSC Encoder

DESCRIPTION

The VP2614 Video De-Multiplexer forms part of the GPS chip set for video conferencing, video telephony, and multimedia applications. It extracts error corrected parameters, and run length coded DCT coefficients, from an H.261 bitstream. Elements of the data which have been variable length coded according to the specification, are decoded within the device. It produces tagged data, aligned to a macroblock timing structure, in the format needed by the VP2615 Decoder. Side information and status bits are separately made available for the system controller.

The VP2614 will accept data up to a peak rate of 4 Mbits per second, but with an average rate up to 2 Mbits per second. The bursty nature of the input, together with the fact that each coded picture does not use the same number of bits, requires the provision of a received data buffer. Since the VP2615 Decoder accepts macroblock data as it becomes available, it is not necessary to provide storage for a complete coded picture. Worst case analysis has shown that a buffer size of 256K bits is adequate in practice.

The incoming sequence is coded with a strict syntax, and the VP2614 must identify and align with this sequence before correct decoding is possible. Storage for this alignment is contained within the external buffer. The device monitors that lock is always valid, and reports to the system controller.

Data is organized to be within 512 bit frames, with 18 bits within this frame used as parity bits for forward error correction. The device analysizes these parity bits and will correct for any two random bits in error, but up to 5 errors can be detected but not corrected. Errors are reported to the system controller, which can also control the response of the Video De-Multiplexer to such errors.

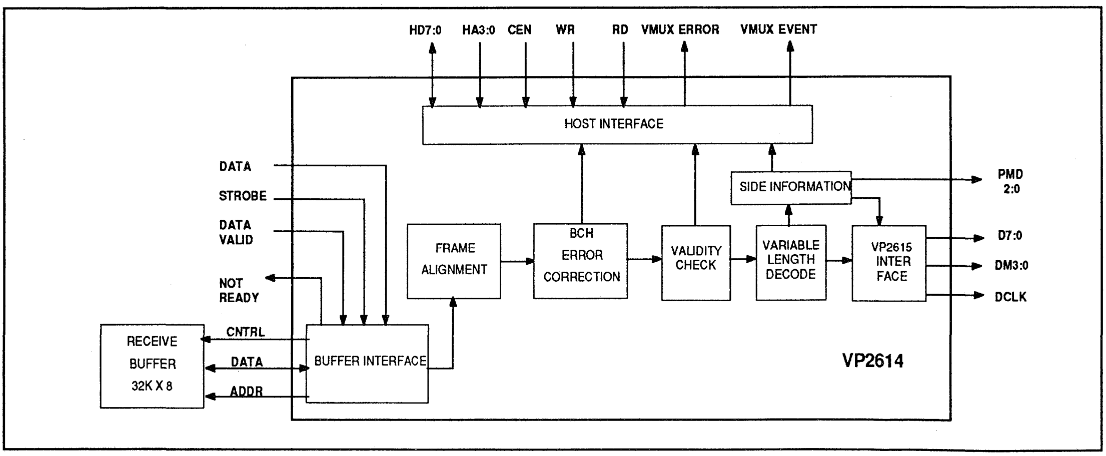

Fig 1 : Simplified Block Diagram

PIN DESCRIPTION

SIG	TYPE	FUNCTION
LD	I	Line input data
$\overline{\text{LEN}}$	I	When low, the line input data is valid.
LCLK	I	Line input strobe
LRED	O	When low, line data cannot be ac cepted.
DBUS7:0	O	Data and control bus to the VP2615.
DMODE3:0	O	These outputs identify the data on DBUS7:0.
PM2:0	O	Identifiers for the additional information on DBUS7:0 .Not used by the VP2615.
DCLK	O	Continuous O/P strobe for the DBUS7:0 bus which is derived from SCLK.
SCLK	I	System clock. Must be 27 MHz for 30 Hz frame rates.
HD7:0	I/O	Bi-directional data bus.
HA3:0	I	Four system controller address bits.
$\overline{\text{WR}}$	I	An active low write strobe from the system controller.
$\overline{\text{RD}}$	I	An active low read strobe from the system controller.
$\overline{\text{CEN}}$	I	An active low chip select from the sys tem controller.
ERR	O	An active low output which Indicates framing and decoding errors.
EVT	O	An active low output which Indicates that new picture status data is available.
B7:0	I/O	Bi-directional data bus to the receive buffer.
A14:0	O	Address bus to the receive buffer.
$\overline{\text{WS}}$	O	An active low write strobe for the receive buffer.
$\overline{\text{BCS}}$	O	An active low select for the receive buffer.
$\overline{\text{BEN}}$	O	An active low output enable for the buffer.
TCK	I	JTAG test clock
TMS	I	JTAG mode select
TDI	I	JTAG I/P data
TDO	O	JTAG O/P data
$\overline{\text{TRST}}$	I	JTAG reset
TOE	I	When low all outputs are high imped- ance
$\overline{\text{RES}}$	I	An active low power on reset

NOTE:
"Barred" active low signals do not appear with a bar in the main body of the text.

OPERATION OF THE MAJOR BLOCKS

FRAME ALIGNMENT

The H.261 continuous bitstream is split into frames of 512 bits (see the section on forward error correction), the first bit in each frame being part of an 8 bit frame alignment pattern. Only the sequence in the pattern is important and detection can start from any bit. To avoid false detection within the actual data, this pattern must be repeated at least three times before " frame lock " can be considered to have been achieved.

The detection of frame lock thus requires data from 24 consecutive 512 bit frames, and a section of the Received Data Buffer is reserved for this purpose. This external RAM is supported by a small internal buffer which allows eight consecutive bits (obtained from reading a byte) to be simultaneously checked for alignment with the corresponding bits in seven other bytes spaced apart by complete frames. The search for alignment over 512 bits takes less than 250 microseconds with a 27 MHz clock, this being less than the time taken to receive 512 bits at the maximum rate of 2Mb/second. Thus the buffer area for frame lock does not overflow.

Once frame lock has been achieved it is continually monitored using the appropriate bit in each frame. If a mismatch occurs then the next four frame alignment bits will be checked for errors. If any one of these four bits is also in error then loss of frame alignment is declared by setting a Status Register Bit, and a search for a new alignment position will commence. If none are in error then a random bit error is assumed and no further action is taken.

The check done on loss of alignment is a compromise between falsely believing that alignment has been lost and not detecting that frame alignment has been lost. The probability of two random bit errors in the five frames used in the check is dependent on the bit rate and also the error rate. With a high error rate of 1:100000, and a bit rate of 2Mb/sec, false detection is possible once per week. The probability of detecting a change in the frame alignment (caused by switching in a new bitstream) is 46.9% in the first five frames, but this rises to 97.4% after 12 frames have been processed.

A Control Bit is provided to enable frame alignment, when this bit is cleared the frame alignment module will be bypassed. Under these conditions Frame Lock will always be indicated and data is still buffered and processed. The error correction module must also be bypassed by clearing both Control Bits.

ERROR CORRECTION

The error correction scheme requires that the bitstream is divided into frames of 512 bits, comprising a first Framing Bit, a second Fill Bit, 492 Data Bits, and 18 Parity Bits. The "Fill" bit is included within the error correction scheme, and it could be corrected if corrupted. This is the second bit in the frame and indicates whether the next 492 bits are coded data, or fill bits used to keep the line busy on every bit clock period. These are generated when the transmission buffer is empty and has no valid data to send, but are disregarded by the de-mux core.

The BCH decoder is designed to either correct up to one error, or up to two errors, or to be bypassed. It will also detect, but not corrrect, three, four, or five errors. When correcting one or two bit errors, the scheme ignores the first framing bit and corrects within the rest of the frame. The second (fill) and the last 18 (parity) bits are stripped to extract the data bits. In the bypass option all 512 bits are treated as data and uncorrected.

The parity bits maintain a Hamming distance of 5 between codewords, so that some three bit errors could be misinterpreted as two bit errors. For this reason the option is provided which restricts the correction to only one bit errors. Separate bits are provided to enable one bit or two bit error correction. Clearing both bits will cause the error correction circuit to be bypassed.

Three counters are provided which indicate to the system controller the number of single, double and multiple bit errors.

A fourth counter is provided to count FEC frames, and thus allows error statistics to be accumulated.These can all be reset by the system controller.

VIDEO LOCK

Once the VP2614 has locked to the error corrected frames, it will begin searching for the 20 bit unique Picture Start Code. Once this has been identified the "Video Lock" status bit will be set, and the bitstream will be translated on a code by code basis. Video lock will be lost and translation process interrupted under the following conditions:

1) A Picture Start Code or Group of Blocks (GOB) Start Code is not present when expected.

2) The codeword is not valid for its context, causing no match to be obtained. Each variable length code in the bitstream is analysed by the VP2614, and invalid patterns will force Video Lock to be lost.

3) Too many coefficients are transferred for the current macroblock because the End of Block code was missing.

Note that only the most frequently occurring coefficients are variable length coded, the others being represented by an escape sequence followed by a fixed length code. The Intra DC coefficient is also a fixed length code. These fixed length codes have bit patterns which are forbidden in the H.261 specification, but they could appear due to bit errors. These invalid codes are trapped by the VP2614, but do not cause Video Lock to be lost. Instead the run length coefficient is replaced by a default value of magnitude 1.

A count is maintained of up 256 occurances of faults 1 - 3, and a status bit is set when lock is lost (the Video Lock Achieved bit is also cleared). An output signal is also provided which can, if required, be used to interrupt the system controller. This indicates any of the above errors which cause Video Lock to be lost and also frame alignment errors; alternatively it can be used to just indicate framing errors.

When Video Lock has been achieved, the detection of a Picture or GOB start code when it is not expected will not cause lock to be lost. Instead the VP2614 will resynchronize to the new start code, and dummy macroblocks will be generated for the missing GOB's. These dummy blocks will be Fixed Macroblocks, and will cause the VP2615 Decoder to use data from the previously decoded picture. Note that Video Lock is actually lost and re-gained under these conditions. The status bit will momentarily be set and then reset, and the Video Lock Lost Counter will be incremented.

Similarly any errors in the actual GOB number will not cause lock to be lost. Since sequential GOB numbers are always produced by the encoder, then the Decoder generates its own GOB numbers and ignores those in the bitstream.

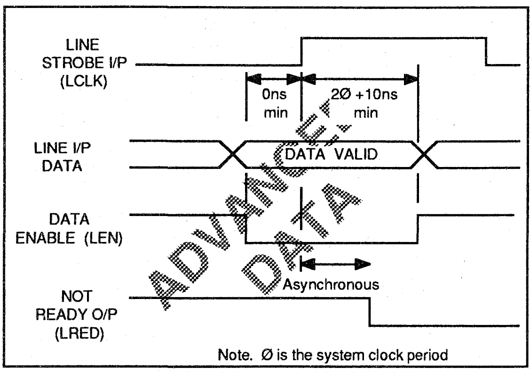

Fig 2 : Line Interface Timing

A Control Bit allows the system controller to take one of two actions when Video Lock is lost. Either the VP2614 can be forced to re-initialize to the next Picture Start Code, or it can abandon the decoding operation until the next GOB Start Code is detected. When lock has been lost, and a new start code has been found, the VP2614 assumes its number to be initially correct and starts its own sequence from that number. If, however, the next number in the bitstream is not in sequence then this new number is used to start a new sequence. This process continues until two sequential numbers are obtained, and then no further checks on the GOB numbers are made until Video Lock is again lost. The VP2614 will generate "Fixed Macroblocks " for the missing GOB numbers since Video Lock was lost, and will output these to the VP2615 decoder. This then uses data from the previous decoded picture for those macroblocks.

A Video Hold bit is provided in one of the System Control Registers which forces Video Lock to be lost immediately. No further data is passed to the VP2615 whilst this bit is set, but the Received Data Buffer is not allowed to fill unnecessarily. Incoming data will be flushed out and lost. When the hold bit is cleared a Picture Start Code must be detected to re-gain Video Lock. The VP2615 will then be provided with any missing GOB's as described above, before GOB's in the new picture are processed.

A Freeze Frame Control Bit is also provided. This has a similar action to the Video Hold Bit, except that it is only actioned when PTYPE has been decoded in the picture layer, and it also sets a Freeze Frame status bit. If Video Lock is lost before the start of a new frame then Freeze Frame will become active and a search will commence for a picture start code. Even though Freeze Frame causes Video Lock to be lost, the VP2614 will still search for picture start codes and will extract PTYPE and Temporal Reference values.

If set, a Release Mode Control Register Bit will allow the freeze condition to be released when the Freeze Bit is cleared, but is only actioned when the next Picture Header is decoded. If the Release Mode Bit is cleared, then the freeze condition is only released when the PTYPE bit in the H.261 stream specifies that this is to occur. Even when automatic release has been selected the system controller can still monitor the length of time that the freeze has been in effect. It can then force a release after a time out period by setting the Release Mode Bit and clearing the Freeze Bit.

DE-MUX CORE

Once Video Lock has been achieved, the core of the VP2615 will convert the H.261 bitstream into video parameters and run length coded coefficients. A state machine, which is a hardware manifestation of the H.261 coding structure, maintains the current position in the bitstream. When necessary variable length de-coding is performed, and side information such as temporal reference and Picture Type Information is stored in registers.

Not all this side information is used by the VP2615 Decoder, but is still made available on the data output Bus DBUS7:0. This is described in the section on Additional Information. In addition the side information can be examined by the system controller.

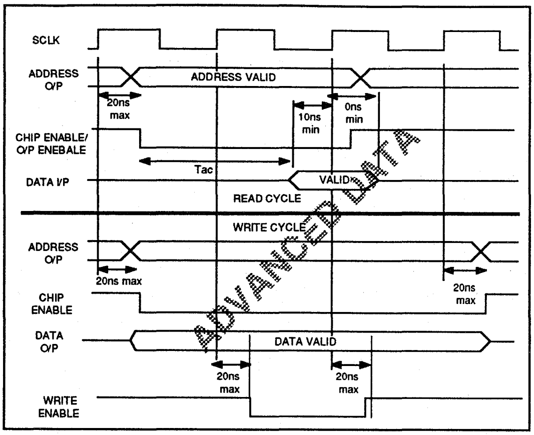

Fig 3 : External Buffer Timing

Requirements for the complete decoder system are such that it is desirable for the VP2614 /15 pair to free run, and to ignore the Temporal References embedded in the video bitstream. The pair then always process the bitstream, whenever code bits are available, using the processing rate needed for the full 30 Hz frame rate. Operating in this manner allows the de-mux core to be closely coupled to the VP2615 Interface circuitry, and no additional buffering is necessary. The de-multiplexing process is then locked to the macroblock timing structure needed by the VP2615.

LINE INTERFACE

Bitstream inputs to the device are controlled by an asynchronous line input strobe, which when data is valid is enabled by a Data Valid signal. Detailed timing information is given in Figure 1.

Maximum input frequency is 4 MHz and the rising edge of the strobe is used to internally latch the data. The VP2614 generates a Ready signal which goes invalid when data cannot be accepted. This ,for example, occurs during system reset or if the Received Data Buffer overflows.

EXTERNAL BUFFER REQUIREMENTS

The external buffer must be a 32K x 8 bit static RAM , and must comply with the timing requirements given in Figure 2. Under normal operating conditions the buffer will not overflow, however it is conceivable that under some unforseen condition the buffer may fill and then overflow. For this reason a Buffer Full Flag is provided in one of the Status Registers. This is asserted when the buffer is 90% full, and is not itself an error condition. If the buffer continues to fill and eventually overflows, then the Ready Signal to the line interface goes invalid. The effect of overflow is to also clear the buffer and the Buffer Empty Flag will be raised. There is no status bit to indicate overflow, but an extended period of Buffer Full followed by Buffer Empty can be used to infer the condition.

VP2615 INTERFACE

The VP2614 provides a glueless interface to the VP2615 Decoder. Run length coded coefficients and control information are transmitted over the DBUS7:0 bus, and are identified by the code on the DMODE3:0 bus given in Table 1. The VP2614 produces a continuous DCLK which is used to strobe data into the VP2615. This is derived by dividing the system clock by two, and when no data is actually available the DMODE3:0 bus will indicate a wait state. Timing is shown in Figure 3.

DMODE3:0	FUNCTION
0000	GOB Number
0001	MB Number
0010	Control Decisions
0011	Quant Value
0100	Horizontal MV
0101	Vertical MV
0110	Coded Blk Pattern
0111	Sub-Block No
1000	Zero Run Count
1001	RLC Coefficient
1010	Not used
1011	Not used
1100	Not used
1101	Not used
1110	Not used
1111	Wait State

Table 1. Output Codes

The VP2615 expects a macroblock and its control information to be transferred over a minimum period, nominally equivalent to 2048 system clock cycles but with allowance for the asynchronous DCLK. Wait states are thus inserted as necessary by the VP2614 in order to enforce this macroblock period. Under normal circumstances the VP2614 will not take longer than 2048 clock periods to produce a macroblock, but some 10% extra time is available for each macroblock before the 30 Hz frame rate becomes impossible to maintain.

The start of a macroblock transfer is identified by the presence of the Control Decisions Byte (DMODE3:0 = 0010). Each macroblock slot must at least consist of this Control Decisions Byte, followed by the GOB number and then the Macroblock number. No further bytes are mandatory.

When high, Bit 0 in the Control Decisions Byte indicates a Fixed Macroblock, and a high on Bit 1 indicates Inter Mode coding. A high on Bit 2 indicates that the macroblock was filtered, and a high on Bit 3 indicates that Motion Compensation was used. When Bit 7 is high this indicates that CIF resolution is in use, but the VP2615 does not use this information. Instead the host controller must supply this information.

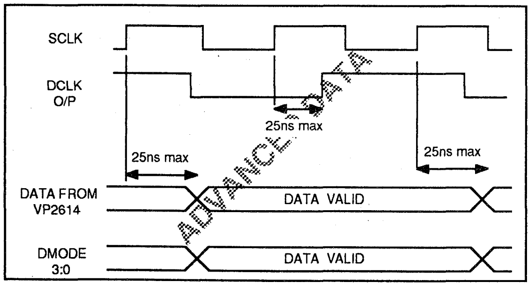

Fig 4 : Output Timing

The VP2615 is essentially a Macroblock Processor which produces decoded data for the position on the screen defined by the GOB and Macroblock number. Since the H.261 specification allows macroblocks to be skipped, then the VP2614 generates dummy Fixed Macroblocks if necessary (see below) which are still separated by 2048 clock cycles. The VP2615 Decoder is not explicitly informed of picture starts, and uses the last macroblock number in the last GOB to indicate that a complete picture has been decoded (after its internal two macroblock pipeline delay). If the data for a particular picture is interrupted because Video Lock is lost, it will simply concatenate the next macroblock which is eventually received onto the partially decoded previous picture. This new macroblock could in reality belong to a new picture and video corruption would result.

This issue is resolved by ensuring that a complete picture, containing dummy data when necessary, is always supplied by the VP2614. The Fixed Macroblock bit in the Control Decisions Byte is set when dummy data is needed, and Intra Mode decoding is specified. This causes the VP2615 to output macroblock data from the previously decoded picture, which was already in the frame store.

ADDITIONAL INFORMATION

Picture Type, PSPARE and GSPARE information is not used by the VP2615 decoder. In future or proprietary uses of H.261 this information could become considerable and be useful to other devices in the system. This can conveniently be supplied by using the DBUS7:0 bus when the DMODE3:0 bus indicates that a wait state is present and there is no useful information for the VP2615. An additional control bus PM2:0 defines the additional information that is present, with the coding given below:

PM2:0	ADDITIONAL PARAMETER
000	Temporal Reference
001	GSPARE transfer
010	PSPARE transfer
011	PTYPE transfer
100	Quantizer step value
111	Data present is that defined by DMODE3:0

SYSTEM CONTROLLER INTERFACE

A conventional microprocessor interface is used consisting of a byte wide bi-directional data bus, four address bits, a chip enable and separate read and write strobes. Detailed timing is given in Figure 4.

In addition two outputs are available which can be used as interrupts if necessary. These can be disabled by control bits. When the Error Interrupt Source Bit is set, the ERROR signal indicates that an error has occurred in the FEC frame alignment module. The Frame Lock Lost Status Bit is also set. The output signal is cleared by reading the status register and will be set again when frame alignment is again achieved. If the host has forced a loss of alignment then ERROR does not go active when lock is lost, but it will still go active when lock is regained.

When the Error Interrupt Source Bit is cleared, then the ERR output also goes active high when Video Lock is lost. Reading the Status Register will detemine the actual cause of the ERR interrupt.

The EVT signal allows the controller to synchronize with picture related parameters extracted from the bitstream. It goes active high when new picture status data is available, as does the Picture Ready bit in Status Register A. This bit and the output signal are cleared when any Status Register is read. The pipeline delay of two macroblock periods through the VP2615 decoder will give the controller time to react to changes in PTYPE affecting the final output of the picture in question. When PTYPE specifies a change between CIF and QCIF, the controller has an amount of time equivalent to that needed to decode the first GOB before it needs inform the VP2615 of the change in operation.

The addresses and functions of the various control and status registers are given below. Setting a Control Bit always performs the function specified, and a high in a Status Register indicates the state is true. All error counters saturate at their maximum values, and are prevented from changing whilst being read.

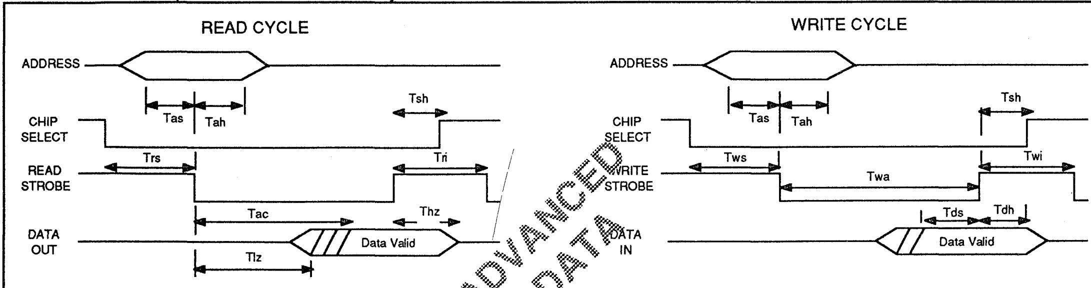

CHARACTERISTIC	SYMBOL	MIN	MAX
Addresss Set Up Time	Tas	10ns	
Address Hold Time	Tah	10ns	
Cip Select Set Up Time	Trs	10ns	
Chip Select Hold Time	Tsh	2ns	
Strobe In active Time	Tri	Øns	
Data Access Time	Tac	6Øns	10 +7Øns
Delay to O/P's low Z	Tlz		25ns
Delay to O/P's high Z	Thz		25ns

NOTE

Ø is the period of the input clock

CHARACTERISTIC	SYMBOL	MIN	MAX
Addresss Set Up Time	Tas	10ns	
Address Hold Time	Tah	10ns	
Chip Select Set Up Time	Tws	10ns	
Chip Select Hold Time	Tsh	2ns	
Strobe In active Time	Twi	3Øns	
Strobe Active Time	Twa	2Øns	
Data Set Up Time	Tds	10ns	
Data Hold Time	Tdh	10ns	

Fig 5 : Host Controller Timing

STATUS REGISTER A (ADDRESS 0)

BIT	FUNCTION
0	GSPARE Byte Available (FIFO not Empty)
1	Freeze Frame
2	Buffer Full
3	Buffer Empty
4	Picture Information Ready
5-7	Unassigned

STATUS REGISTER B (ADDRESS 1)

BIT	FUNCTION
0	Frame Lock Lost
1	Frame Lock Achieved
2	Video Lock Lost
3	Video Lock Achieved
4 -7	Unassigned

CONTROL REGISTER A (ADDRESS 2)

BIT	FUNCTION (when the bit is set)
0	Freeze Frame released by Host
1	Force Freeze Frame
2	Error interrupt only from Frame Lock
3	Enable EVT Interrupt
4	Enable ERR Interrupt
5	Video Hold
6	Clear Buffer
7	System Re-start

CONTROL REGISTER B (ADDRESS 3)

BIT	FUNCTION (when the bit is set)
0	Re-lock to Picture Start Code
1-2	Unassigned
3	Correct 1 bit errors
4	Correct 2 bit errors
5	FEC Framing On
6	Clear Video Lock Lost Counter
7	Clear other Counters apart from above

USER READABLE COUNTERS

ADDRESS	FUNCTION
4	FEC Frame Count
5	Filled Frames Count
6	1 Bit Error Count
7	2 Bit Error Count
8	Greater than 2 Bit Error Count
9	Video Lock Lost Count

PICTURE STATUS REGISTERS

10	Temporal Reference Register
11	Picture Information (see below)
12	First PSPARE Byte
13	Second PSPARE Byte
14	Top of GSPARE Stack

PICTURE INFORMATION REGISTER (ADDRESS 11)

BIT	FUNCTION
0-5	PTYPE from the bitstream
6	PSPARE Byte 1 Valid (cleared by reading byte)
7	PSPARE Byte 2 Valid (cleared by reading byte)

A master - slave arrangement is used for the Picture Status Registers, and the slave is not updated for the duration of the host read operation plus 32 system clock cycles.

Reading any of the counter values (address 4-9) or any Picture Status Register (address 10-13) causes all values in the respective blocks to be frozen for 32 clocks, thus allowing a complete snapshot to be taken of the respective values.

Two bytes of PSPARE data are stored and further bytes will be lost. Note that the VP2612 Video Multiplexer presently only provides one byte of PSPARE information. A FIFO is provided to provide storage for 12 GSPARE bytes, and status bit is provided to indicate that this FIFO is not empty, and that the byte at the top of the stack should be read.

RESET OPERATION

In addition to the hardware reset there are several software reset options which are selective in their action. The hardware reset input will initialize all the internal circuit blocks, and will clear all status registers, error counters, and address pointers. The bits in Control Registers A and B are cleared except that the Video Hold Bit in Register A is set, and Bits 0,4, and 5 are set in Register B. The device will thus re-lock to a Picture Start Code, will correct 2 bit errors, and FEC Framing will be on. The circuit which interfaces to the VP2615 Decoder is reset to the end of picture condition (Macroblock 33 in GOB 12).

The System Re-start bit (Bit 7 in Control Register A) will clear all status bits and will initialize the bitstream decoder, the forward error corrector, and the buffer alignment modules. It should be used if there has been an interruption in the bitstream, and does not affect the circuit producing GOB's and macroblocks for the VP2615 Decoder. Thus, after the re-start, Video Lock can be obtained on a GOB boundary, and Fixed macroblocks can be generated for missing macroblocks within the same picture.

The Clear Buffer bit (Bit 6 in Control Register B) will also cause a System Restart as above, but in addition the read and write address pointers for the external buffer will be reset.

Two bits are also provided in Control Register B for reset operations. One will clear the Video Lock Lost counter, the other clears the FEC frame counter, the Filled Frames counter, and the three error counters in the error detection circuit.

PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC
1	N/C	21	GND	41	HD3	61	B1	81	A1
2	N/C	22	DCLK	42	HD4	62	B0	82	A0
3	N/C	23	LD	43	HD5	63	A14	83	BEN
4	TOE	24	PM0	44	HD6	64	A13	84	BCS
5	N/C	25	PM1	45	HD7	65	A12	85	GND
6	DMODE0	26	PM2	46	VDD	66	A11	86	VDD
7	DMODE1	27	N/C	47	GND	67	A10	87	WS
8	DMODE2	28	N/C	48	WR	68	A9	88	LRED
9	DMODE3	29	N/C	49	RD	69	A8	89	TDI
10	GND	30	N/C	50	CEN	70	A7	90	TMS
11	VDD	31	HA0	51	N/C	71	VDD	91	TRST
12	DBUS0	32	HA1	52	N/C	72	GND	92	TCK
13	DBUS1	33	HA2	53	B7	73	A6	93	TDO
14	DBUS2	34	HA3	54	B6	74	A5	94	VDD
15	DBUS3	35	SCLK	55	B5	75	A4	95	GND
16	DBUS4	36	GND	56	B4	76	A3	96	RES
17	DBUS5	37	VDD	57	B3	77	A2	97	LEN
18	DBUS6	38	HD0	58	B2	78	N/C	98	LCLK
19	DBUS7	39	HD1	59	GND	79	N/C	99	ERR
20	VDD	40	HD2	60	VDD	80	N/C	100	EVT

SIGNAL	DIRECTION	JTAG Bit Number	SIGNAL	DIRECTION	JTAG Bit Number
TOE	IN	84	CEN	IN	41
testoeout	OUT	83	B0	IN	40
DMODE0	OUT	82	B1	IN	39
DMODE1	OUT	81	B2	IN	38
DMODE2	OUT	80	B3	IN	37
DMODE3	OUT	79	B4	IN	36
DBUS0	OUT	78	B5	IN	35
DBUS1	OUT	77	B6	IN	34
DBUS2	OUT	76	B7	IN	33
DBUS3	OUT	75	B0	OUT	32
DBUS4	OUT	74	B1	OUT	31
DBUS5	OUT	73	B2	OUT	30
DBUS6	OUT	72	B3	OUT	29
DBUS7	OUT	71	B4	OUT	28
DCLK	OUT	70	B5	OUT	27
LD	IN	69	B6	OUT	26
PM0	OUT	68	B7	OUT	25
PM1	OUT	67	nroeout	OUT	24
PM2	OUT	66	LRED	OUT	23
HA0	IN	65	WS	OUT	22
HA1	IN	64	BCS	OUT	21
HA2	IN	63	BEN	OUT	20
HA3	IN	62	A0	OUT	19
SCLK	IN	61	A1	OUT	18
HD0	IN	60	A2	OUT	17
HD1	IN	59	A3	OUT	16
HD2	IN	58	A4	OUT	15
HD3	IN	57	A5	OUT	14
HD4	IN	56	A6	OUT	13
HD5	IN	55	A7	OUT	12
HD6	IN	54	A8	OUT	11
HD7	IN	53	A9	OUT	10
HD0	OUT	52	A10	OUT	9
HD1	OUT	51	A11	OUT	8
HD2	OUT	50	A12	OUT	7
HD3	OUT	49	A13	OUT	6
HD4	OUT	48	A14	OUT	5
HD5	OUT	47	RES	IN	4
HD6	OUT	46	LEN	IN	3
HD7	OUT	45	LCLK	IN	2
oeout	OUT	44	ERR	OUT	1
WR	IN	43	EVT	OUT	0
RD	IN	42			

Those signals labelled testoeout, oeout, and nroeout, are not connect to ASIC output pins, but are provided on the JTAG boundary scan to enhance the device testability.

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage VDD	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to VDD + 0.5V
Output voltage V_{OUT}	-0.5V to VDD+ 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	0°C to 70°C
Junction temperature	100°C
Package power dissipation	1000mW

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as negative into the device.

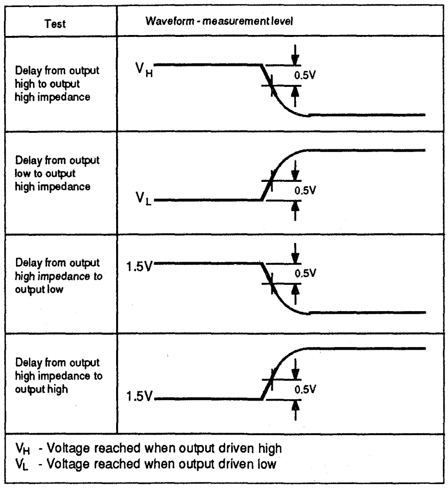

STATIC ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)
Tamb = 0 C to +70°C VDD = 5.0v ± 5%

Characteristic	Symbol	Value Min.	Typ.	Max.	Units	Conditions
Output high voltage	V_{OH}	3.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.0		-	V	3.0V for SYSCLK and LCLK
Input low voltage	V_{IL}	-		0.8	V	0.6V for SYSCLK and LCLK
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{DD}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{DD}
Output S/C current	I_{SC}	10		300	mA	

VP2615

H.261 DECODER

FEATURES

- Inputs run length coded transform data
- Outputs 8 bit pixels in YUV block format
- Up to full CIF resolution and 30 Hz frame rates
- Supports motion compensation with up to 15 pixel movement
- On chip frame store controller
- 100 pin QFP package

ASSOCIATED PRODUCTS

- VP510 Colour Space Converter
- VP520 Three Channel Video Filter
- VP530 NTSC/PAL Encoder
- VP2611 Integrated H261 Encoder
- VP2612 Video Multiplexer
- VP2614 Video Demultiplexer
- VP101 Triple D/A Converter

DESCRIPTION

The VP2615 decoder forms part of a chip set for use in video conferencing and video telephony applications. It conforms to the CCITT H261 standard, and will decode data coded with full or quarter CIF resolution at frame rates up to 30 Hz.

It accepts run length coded coefficients which have already been error corrected and Huffman decoded, and produces multiplexed YUV data in macro block format after a pipeline delay of two MacroBlocks. As shown in Figure 1, other devices in the chip set then convert this data into full resolution, component or composite, video.

The incoming run length coded data is converted to individual coefficient values in the correct order. Data reconstruction is then performed on a block by block basis by multiplying the quantized coefficients with the original quantization value, and then applying the inverse cosine transform. In the inter frame mode this data is then added to the motion compensated block from the previous frame. This block can be passed through a low pass filter when required. A frame store controller produces addresses which allow the best fit block to be read from the frame store, and which also allow the store to be updated with reconstructed data. Refresh cycles are generated when necessary.

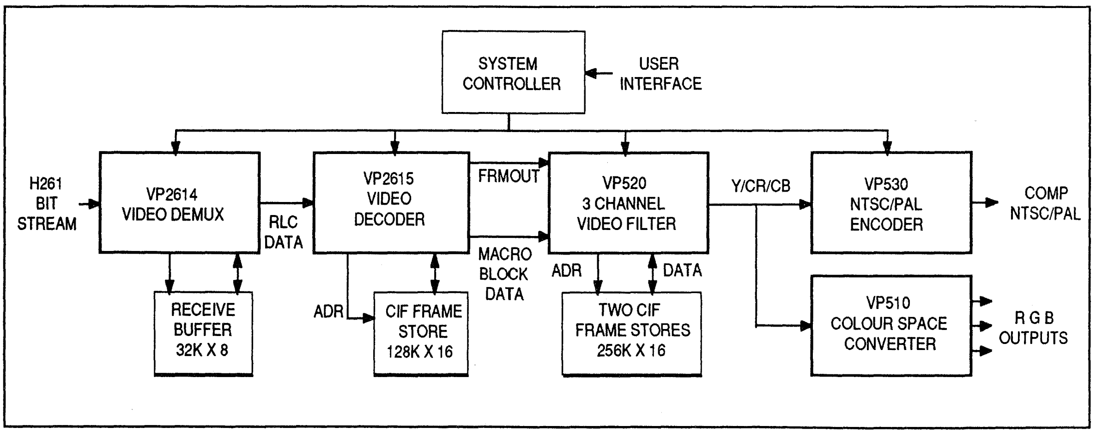

Fig 1 : Typical Video Conferencing Receiver

PIN DESCRIPTIONS

DIN7:0	This port is used to input quantised transform data and control information. Its function is determined by DMODE3:0. Data is clocked in on the rising edge of DCLK.
DMODE3:0	This input controls the function of DIN7:0. Data is clocked in on the rising edge of DCLK.
DCLK	This signal is used to strobe in data at the DIN and DMODE inputs. DCLK can effectively be disabled by inputting a WAIT STATE on DMODE. DCLK must be derived by dividing SYSCLK with an integer greater than one.
YUV7:0	This bus outputs pixel data in YUV block format at quarter SYSCLK frequency.
VPIX	This synchronous output pulses high for two SYSCLK periods when valid pixel data appears at the YUV port. It remains low when inactive.
MBOUT	This synchronous output goes high on the first cycle of a new MacroBlock and stays high until the final pixel of that MacroBlock has been output. At the end of the MacroBlock MBOUT goes low until a new MacroBlock begins.
FRMOUT	This synchronous output goes high to indicate a new Frame is about to begin at the YUV port. It remains high till the last pixel is output. Then, FRMOUT goes low until a new Frame starts.
FS15:0	Data bus for reading and writing to the external DRAM frame store.
ADR7:0	Address bus controlling the external DRAM frame store.
$\overline{\text{RAS}}$	Row Address Strobe controlling the external DRAM frame store.
$\overline{\text{CAS}}$	Column address strobe controlling the external DRAM frame store.

R$\overline{\text{W}}$1	Read/Write control for the external DRAM 1.
R$\overline{\text{W}}$2	Read/Write control for the external DRAM 2.
$\overline{\text{OE1}}$	Output Enable control for external DRAM 1 or ADR8 if 256K DRAM's in use.
$\overline{\text{OE2}}$	Output Enable control for external DRAM 2 N/C if 256k DRAMs in use.
CBUS7:0	Bi-directional data bus for use by a microprocessor. Data and instructions are clocked on and off the chip on the rising edge of CSTR.
$\overline{\text{CSTR}}$	This input strobes the data in and out of the CBUS port.
$\overline{\text{CEN}}$	When this pin is low the CBUS port can be used to input or output data.
CADR	When high this signal defines CBUS as data, and when low as an instruction.
SYSCLK	System clock, run at 27MHz maximum. SYSCLK must remain high for 35% to 65% of each cycle. All internal clocks are derived from this clock.
$\overline{\text{RESET}}$	Active low reset. Must be held low for at least 2048 cycles on power up. If RESET is used during operation, all previous frame data will be lost.
TCK	Test clock for JTAG
TMS	Test mode select for JTAG (Internally pulled low).
$\overline{\text{TRST}}$	JTAG reset pin (Internally pulled low).
TDI	Input JTAG test data (Internally pulled low).
TDO	Output JTAG test data (Internally pulled low).

NOTE:
"Barred" active low signals do not appear with a bar in the main body of the text.

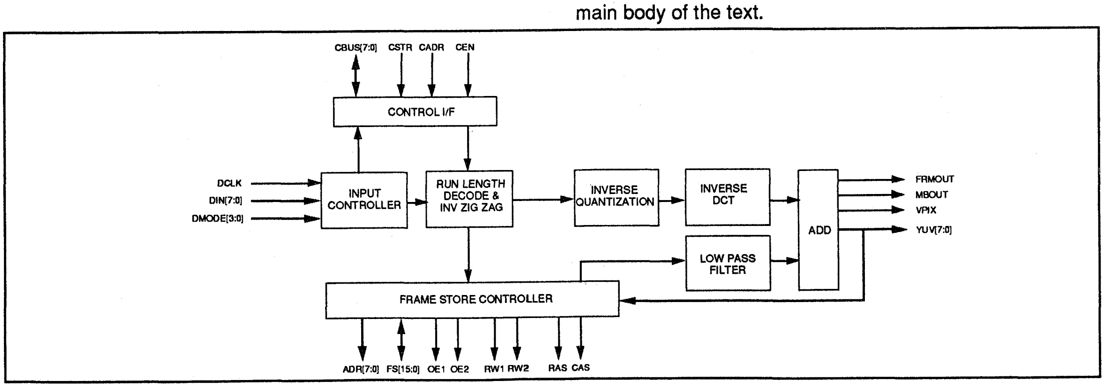

Fig 2 : Simplified Block Diagram

OPERATION OF MAJOR BLOCKS

Run Length Decode

This block converts the run length coded data into 64 individual coefficient values, inserting zero value coefficients where required. It then re-orders these 8 bit quantized DCT coefficients from the zig zag arrangement into normal 8 x 8 format.

Inverse Quantise

This circuit reconstructs the 12 bit DCT coefficients from the 8 bit quantized coefficients using the 5 bit Quantization Value. This is performed using the following formulae.

If QUANT is odd :
REC = QUANT*(2*LEVEL+1) : LEVEL > 0
REC = QUANT*(2*LEVEL-1) : LEVEL < 0

If QUANT is even :
REC = QUANT*(2*LEVEL+1)-1 : LEVEL > 0
REC = QUANT*(2*LEVEL-1)+1 : LEVEL < 0

For Intra coded DC coefficients :
REC = 8*LEVEL
except if LEVEL=255 when REC=1024

If LEVEL=0 then REC=0 in all cases.

The reconstructed values (REC) are passed through a clipping circuit in case of arithmetic overflow.

Inverse DCT

This circuit performs an Inverse Discrete Cosine Transform on an 8x8 block of 12 bit coefficients outputting 9 bit signed pixel data. This IDCT fully meets the CCITT specification.

Frame Store Interface

The whole of the previous picture is stored in either two external 64K x 16 DRAMs, or in a single 256 k x 16 DRAM, or in four 256K x 4 DRAM's. A bit in the user defined Input Set Up Data determines whether 64K or 256K DRAM's are to be used. In the latter case, use OE1 as ADR8, RW1 as R/W and do not connect RW2 and OE2. Table 1 specifies the worst case maximum and minimum times which must be achieved by the DRAM for correct operation with the VP2615. Times in the DRAM specification must be less than or equal to the times stated.

The Frame Store Interface manages all read and write operations to these DRAM's. During the course of each MacroBlock, the "Best Fit" MacroBlock is read from the DRAMs and the fully processed MacroBlock is written back. In this way, the previous frame is continually updated. The DRAM controller also takes care of refresh for the DRAMs.

Figure 3 illustrates the effects of the pipeline delays through the device; whilst macro block 3 is being input the previous macroblock (2) is being decoded and needs the equivalent macroblock from the previous frame to be read from the frame store. At the same time macroblock 1, which has already been decoded, is being written to the frame store for use in the next frame and is also available on the output pins.

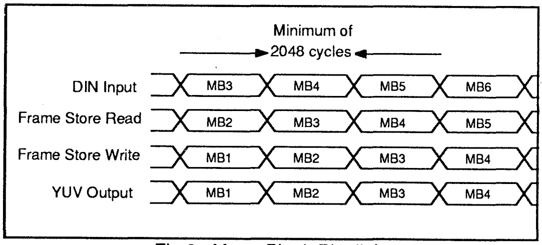

Fig 3 : MacroBlock Pipelining

SYMBOL	PARAMETER	MINIMUM	MAXIMUM
t RAC	Access time from RAS	-	105ns or under
t CAC	Access time from CAS	-	25ns or under
t RP	RAS precharge time	50ns or under	-
t CP	CAS precharge time	15ns or under	-
t RAS	RAS pulse width	90ns or under	-
t CAS	CAS pulse width	50ns or under	-
t REF	Time to refresh 256 rows	-	0.25ms or over

N.B. All times are quoted assuming 27MHz operation. For lower clock frequencies increase the above values proportionately.

Table 1. External DRAM Timing Requirements

Loop Filter

The best matched block from the search window in the previous frame can be passed through a low pass filter to reduce block boundary effects. The filter uses a simple [1 2 1] characteristic in both horizontal and vertical dimensions as laid down in the H261 Specification. An instruction input at the DIN port defines whether the filter should be used or not.

Reconstruction Adder

In Inter Mode, the IDCT data is added to the best fit block from the previous frame store. In Intra Mode, the IDCT data is added to zero. After the adder, the sign bit is removed from the result to give 8 bit pixels. Clipping circuits ensure that any pixels with values exceeding 255 are clipped to 255 and any with negative values are clipped to zero (such values are possible due to quantization effects).

OPERATION OF INTERFACES

DIN Input Port

The DIN port provides a glueless interface to the VP2614 Video Demultiplexer, from which it will accept run length coded transform data and control information. The general purpose nature of the interface will, however, allow other sources of macroblock data to be used.

Data on the input bus is defined by means of the signals DMODE3:0, and is strobed in with the DCLK signal which is provided by the VP2614 and derived from SYSCLK. Set up and hold times with respect to the rising edge of DCLK are given in Figure 4. If DCLK is a continuous strobe, then the WAIT state defined by DMODE 3:0 should be used to disable any clocking actions. If preferred DCLK can alternatively be used as a strobe which is only present when data is valid and action is needed. In this case WAIT states are not strictly necessary.

The VP2615 always expects to receive a complete video frame of data, even if error conditions have occurred in the de-

DMODE3:0	FUNCTION
0000	GOB Number
0001	MB Number
0010	Control Decisions
0011	Quant Value
0100	Horizontal MV
0101	Vertical MV
0110	Coded Blk Pattern
0111	Sub-Block No
1000	Zero Run Count
1001	RLC Coefficient
1010	Not used
1011	Not used
1100	Not used
1101	Not used
1110	Not used
1111	Wait State

Table 2 . DIN Mode Functions

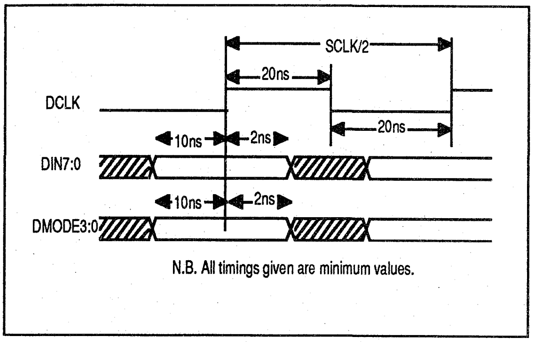

Fig 4 . DIN Port Timing

multiplexer. Skip Picture or Fixed Macroblocks should be supplied if necessary once a frame has started. With the latter, decoded data from the previously stored frame will be produced by the VP2615.

The asynchronous interface will allow the use of other video de-multiplexers, as long as the protocol defined by DMODE3:0 is observed. This protocol is defined below, and summerized in Table 2.

GOB Number: The correct GOB number is required for every macro block in that group. (DIN3 is MSB).

MB Number: Each macroblock in a group requires an identification number. (DIN5 is MSB).

Control Decisions : This byte must always be the first in the sequence since it resets the internal control logic. It defines which control decisions were taken when coding the forthcoming MacroBlock. A high on DIN 0 indicates a Fixed Macro Block (ie no change since the previous frame), and a high on DIN1 indicates that Inter coding was used. Similarly a high on DIN2 indicates that the MacroBlock was filtered, a high on DIN3 indicates that Motion Compensation was used. and a high on DIN6 indicates that SKIP PICTURE is in effect. In the latter case the VP2615 will cease processing until SKIP PICTURE is reversed by writing a new Control Decisions byte. Whilst SKIP PICTURE is active, no further data will be output from the YUV port. SKIP PICTURE effectively resets the VP2615, and the next MacroBlock input should be the first of a new frame. Since the frame store will not be updated then the system should ensure that an Intra coded picture is sent as soon as possible.

Quant Value: This input represents the quantization value (between 2 and 62 with DIN4 as MSB), which has been used for this macroblock. If no new value is provided for a macroblock then the old value is re-used.

Horizontal MV: This input (on DIN4:0) represents the horizontal component of the motion vector. It must always be provided when motion compensated Inter coding is in use.

Vertical MV: This input (On DIN4:0) represents the vertical component of the the motion vector. It must always be provided when motion compensated Inter coding is in use.

Coded Blk Pattern: This byte is defined in the H.261 Specification and is used to indicate which sub blocks contain non zero coefficients. It is produced by the encoder but is not used by the VP2615, and if provided will be ignored. The sub block numbering sequence is actually used to indicate blocks with zero coefficients.

Sub Blk No: Each macroBlock contains 6 Sub-blocks, numbered 1 through 6. The corresponding binary value should be provided on DIN2:0, before the RLC coefficients of that Sub-Block appear. If a Sub-Block contains no coefficients, then its number need not be provided at all, or it can be immediately followed by the next sub block number without any intermediate coefficient values. Even though zero valued sub blocks can simply be ignored in this way, a 2048 clock delay between new macroblocks must still be maintained by the video de-multiplexer.

Zero Run Count: The number of zero valued coefficients preceding the (non-zero) RLC coefficient is defined by this input. DIN 6 and 7 are not used, with the value between 0 and 63 defined by DIN5:0.

RLC Coefficient: This input defines the value of the run length coded coefficient. It will always be a non-zero value

Wait State: This mode should be used on any cycle where no data is being input at the DIN port. Wait States can be inserted between any other instructions as required.

Any undefined bits in the above descriptions may be made high or low as desired.

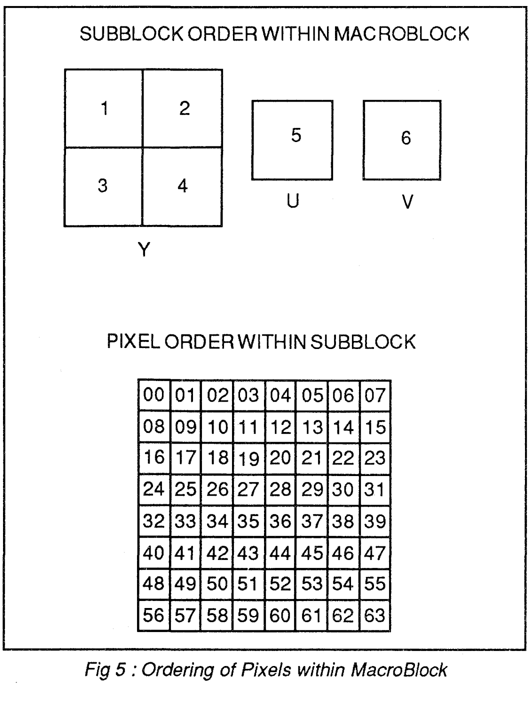

Fig 5 : Ordering of Pixels within MacroBlock

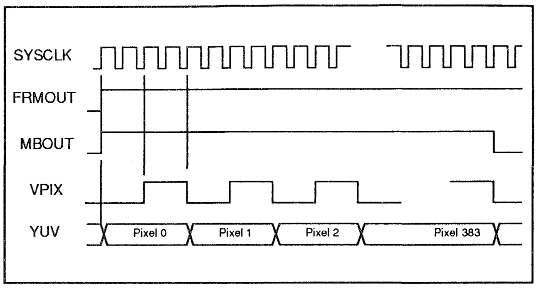

Fig 6 : YUV Port Timing

The first information supplied for a macroblock should be that contained within the Control Decisions byte. Receipt of this instruction resets the internal cycle counter for that MacroBlock. Although some Macro Blocks may contain no data, the VP2615 requires that at least the Control Decisions, GOB Number and Macro Block Number be supplied by the de-multiplexer (in that order). All other side information, which is to be provided for a non zero block, must then be supplied before any sub block data can be accepted. GOB's and Macroblocks must be supplied in the correct sequence, but sub blocks within a macroblock can be in any order. The VP2615 does not need to be explicitly informed that the last coefficient has been received within a sub-block. It will wait for a new sub-block number, or a new Macroblock Control Decision Byte, before processing the previous sub-block since it then knows that the sub block is complete.

At least 2048 SYSCLK cycles must separate the start of one Macro Block (identified by receipt of the Control Decisions byte) from the start of the following Macro Block. There are , however, no specific restrictions on the timing of Sub-Blocks within the MacroBlock. The minimum gap between incoming macroblocks is needed for internal processing and also for the time to output 384 decoded values at one quarter the SYSCLK frequency.

The VP2615 contains two complete macro block buffers in its input circuitry, which swap on the completion of the processing and outputting of the results. Whilst one is used internally the other can be loaded with a new macroblock. It essentially is a macroblock processor and produces the decoded outputs for a macroblock after two macroblock pipeline delays. When it is no longer supplied with macroblock inputs then the pipeline stalls and does not flush out. Thus two macroblocks from a new picture are needed to produce the decoded outputs from the last two macroblocks in a previous picture.

YUV Output Port

Decoded pixel data is presented at the YUV port in standard macroblock format at quarter SYSCLK frequency (6.75MHz max), and in the macroblock order presented at the input. Since the VP2615 always expects a complete picture's worth of GOB's and macroblocks (unless Skip Picture is sent by the video de-mux), then it will always produce a complete coded picture. As explained in the previous section, however, it requires to be supplied with two macroblocks from the next picture before a complete frame is fully decoded. The stand-

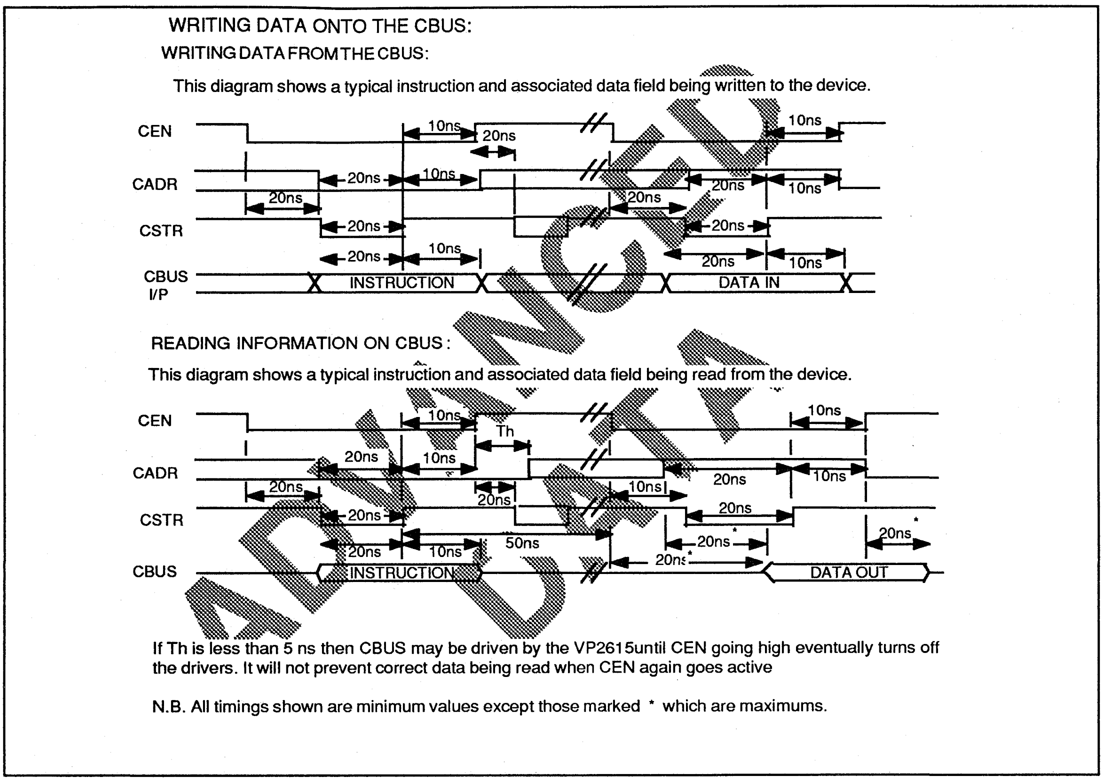

Fig 7 : CBUS Timing

ard macroblock internal configuration is shown in Figure 5.

Output timing is shown in Figure 6. VPIX is toggled high each time a valid pixel is available at the output pins, and remains low when no pixel data is output. MBOUT is used to define the boundaries between MacroBlocks, but is not used when the device is directly connected to the VP520. The Frame Ready Output nominally goes high on the same SYSCLK edge as the first MBOUT goes high, and returns low when the last MBOUT goes low. This will actually be after two macroblocks from the next frame have been supplied as inputs, but this gap will not effect the operation of the VP520 which converts macro block data to full resolution line data. The first VPIX strobe produced after MBOUT goes high, will go high after two SYSCLK periods, with the data being valid for two SYSCLK periods either side of this edge. These delays are subject to internal differential delays and will not be precise clock period delays.

CBUS Control Port

The CBUS control port is used to input control and setup information and also to output status information. In order to save on pin count, a microprocessor driving this port is required to execute two I/O instructions in order to transfer a single byte of information to or from the device. The first transfer is always a write operation, with a low level on the single address line which is used by the interface. Data on the bus then defines the instructions listed in Table 3. The second transfer can be a read or write operation as necessary, but the address line must then be high with the set up time given in Figure 7.

In addition to the single addresss line (CADR), data transfers use a control strobe (CSTR) which is only effective when a chip enable is present (CEN). Detailed timing information is given in Figure 7, and when writing data or instructions to the VP2615 the set up and hold times which are referenced

CBUS3:0	INSTRUCTION
0000	Unassigned
0001	Unassigned
0010	Unassigned
0011	Unassigned
0100	Input Setup Data
0101	Unassigned
0110	Reserved
0111	Reserved
1000	Output GOB Number
1001	Output MB Number
1010	Unassigned
1011	Output Control Decisions
1100	Unassigned
1101	Unassigned
1110	Unassigned
1111	Unassigned

Table 3: CBUS Instructions

to the rising edge of CSTR must be maintained.

When a write instruction has been defined CADR should be pulled high, valid data presented to CBUS7:0 and then strobed in using CSTR. Other system I/O transfers can occur between defining a write operation and supplying the data to be written, assuming CEN is not active during those other transfers. If CSTR does not go active because of I/O transfers to other devices, then CEN can remain active low between the instruction and data.

When a read instruction has been specified the requested data will then be output on CBUS7:0 after the access time specified from CEN going low, assuming that CADR was already high. Otherwise the data will become valid after the access time specified from CADR going high after CEN was low. Note that in the data read phase CADR must always go high before CSTR goes high, with the set up time specified. When CEN goes high, or CADR goes low, the CBUS will go high impedance after the delay specified.

Note that the access times under the conditions given above are only true when the gap between CSTR going high in the instruction phase, and CEN going low in the data phase, is greater than the minimum specified in figure 7.

Only CBUS3:0 are used to define an instruction. The remaining bits, CBUS7:4, should be pulled low. The instructions are listed in Table 3 but are described below in greater detail;

Input Setup Data: This instruction performs several functions, the details being specified in the data field following this instruction. If CBUS0 is high, the device will operate in QCIF mode, otherwise in full CIF mode. If CBUS6 is high, then the device will be configured to use 256K word DRAM's, otherwise it will assume two 64K word DRAM's. All CBUS inputs not defined above must be pulled low during the set up definition phase. On reset the defaults are 64k DRAMs and full CIF mode.

Output GOB number: This instruction will make the VP2615 output the GOB Number associated with the data currently being output at the YUV port. The number will appear on CBUS3:0. CBUS7:4 are not used (always low).

Output MB Number: This instruction will make the VP2615 output the Macroblock Number associated with the data currently being output at the YUV port. The number will appear on CBUS5:0. If CBUS6 is low, this indicates that the MacroBlock number has just changed or is about to change, and is thus not reliable.

Output Control Decisions : This instruction will make the VP2615 output control information received through the DIN port. CBUS0 shows whether the MacroBlock currently being output was Inter or Intra coded (0=Intra). CBUS1 shows whether Motion Compensation was used (1=MC used). CBUS3 will be high if the MacroBlock was passed through the Loop Filter. If CBUS6 is high, this indicates that SKIP PICTURE is currently active.

JTAG Test Interface

The VP2615 includes a test interface consisting of a boundary scan loop of test registers placed between the pads and the core of the chip. The control of this loop is fully JTAG/ IEEE 1149-1 1990 compatible. Please refer to this document for a full description of the standard.

The interface has five dedicated pins: TMS, TDI, TDO, TCK and TRST. The TRST pin is an independent reset for the interface controller and should be pulsed low, soon after power up; if the JTAG interface is not to be used it can be tied low permanently. The TDI pin is the input for shifting in serial instruction and test data; TDO the output for test data. The TCK pin is the independent clock for the test interface and registers, and TMS the mode select signal.

TDI and TMS are clocked in on the rising edge of TCK, and all output transitions on TDO happen on its falling edge.

Instructions are clocked into the 8 bit instruction register (no parity bit) and the following instructions are available.

Instruction Register (MSB first)	**Name**
11111111	BYPASS
00000000	EXTEST
01000000	INTEST
XX001011	SAMPLE/PRELOAD

Timing details (minimums) for the JTAG control signals are shown in Figure 8. The maximum TCK frequency is 5 MHz.

The positions of the test registers in the boundary loop, and their corresponding functional names, are detailed in Table 4. Note that any internal signals controlling the impedance of a bus also have associated registers, even though they are not normally available to the user. This register order will determine the serial data stream for JTAG testing. The signal DHZ will, if loaded with a logic '1', force all the outputs to a high impedance state.

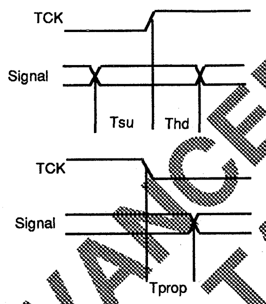

	Tsu	Thd	Tprop
TMS toTCK timing	15	5	
TDI to TCK timing	15	5	
Chip i/p to TCK timing	15	5	
TCK to TDO timing			20

Fig 8. Typical JTAG Interface timing

PAD	TYPE	REG NO	PAD	TYPE	REG NO
DHZ	TRI	1	FS10	IN	48
CADR	IN	2	FS10	OUT	49
CEN	IN	3	FS9	IN	50
CSTR	IN	4	FS9	OUT	51
CBUS0	OP	5	FS8	IN	52
CBUS	TRI	6	FS8	OUT	53
CBUS0	IP	7	FS7	IN	54
CBUS1	OUT	8	FS7	OUT	55
CBUS1	IN	9	FS6	IN	56
CBUS2	OUT	10	FS6	OUT	57
CBUS2	IN	11	FS5	IN	58
CBUS3	OUT	12	FS5	OUT	59
CBUS3	IN	13	FS4	IN	60
SYSCLK	IN	14	FS4	OUT	61
CBUS4	OUT	15	FS3	IN	62
CBUS4	IN	16	FS3	OUT	63
CBUS5	OUT	17	FS2	IN	64
CBUS5	IN	18	FS2	OUT	65
CBUS6	OUT	19	FS1	IN	66
CBUS6	IN	20	FS1	OUT	67
CBUS7	OUT	21	FS0	IN	68
CBUS7	IN	22	FS0	OUT	69
DMODE0	IN	23	ADR7	OUT	70
DMODE1	IN	24	ADR6	OUT	71
RESET	IN	25	ADR5	OUT	72
DCLK	IN	26	ADR4	OUT	73
DMODE2	IN	27	ADR3	OUT	74
DMODE3	IN	28	ADR2	OUT	75
DIN0	IN	29	ADR1	OUT	76
DIN1	IN	30	ADR0	OUT	77
DIN2	IN	31	RW1	OUT	78
DIN3	IN	32	RW2	OUT	79
DIN4	IN	33	DE1	OUT	80
DIN5	IN	34	DE2	OUT	81
DIN6	IN	35	RAS	OUT	82
DIN7	IN	36	CAS	OUT	83
FS15	IN	37	MBOUT	OUT	84
FS	TRI	38	FRMOUT	OUT	85
FS15	OUT	39	VPIX	OUT	86
FS14	IN	40	YUV0	OUT	87
FS14	OUT	41	YUV1	OUT	88
FS13	IN	42	YUV2	OUT	89
FS13	OUT	43	YUV3	OUT	90
FS12	IN	44	YUV4	OUT	91
FS12	OUT	45	YUV5	OUT	92
FS11	IN	46	YUV6	OUT	93
FS11	OUT	47	YUV7	OUT	94

Table 4. Pin and JTAG Test Registers

1	GND	21	DIN7	41	CBUS6	61	TRST	81	RAS
2	N/C	22	DIN6	42	CBUS5	62	TD0	82	OE2
3	FS3	23	DIN5	43	CBUS4	63	YUV7	83	OE1
4	FS4	24	DIN4	44	VDD	64	YUV6	84	GND
5	GND	25	VDD	45	SYSCLK	65	VDD	85	RW2
6	FS5	26	DIN3	46	GND	66	YUV5	86	VDD
7	FS6	27	DIN2	47	CBUS3	67	GND	87	RW1
8	VDD	28	DIN1	48	CBUS2	68	YUV4	88	ADR0
9	FS7	29	N/C	49	CBUS1	69	YUV3	89	ADR1
10	FS8	30	GND	50	CBUS0	70	YUV2	90	ADR2
11	FS9	31	DIN0	51	GND	71	YUV1	91	ADR3
12	GND	32	DMODE3	52	N/C	72	YUV0	92	ADR4
13	FS10	33	DMODE2	53	CSTR	73	VDD	93	ADR5
14	VDD	34	VDD	54	VDD	74	VPIX	94	GND
15	FS11	35	DCLK	55	CEN	75	FRMOUT	95	ADR6
16	FS12	36	GND	56	CADR	76	GND	96	VDD
17	FS13	37	RESET	57	GND	77	MBOUT	97	ADR7
18	FS14	38	DMODE1	58	TD1	78	CAS	98	FS0
19	FS15	39	DMODE0	59	TMS	79	N/C	99	FS1
20	GND	40	CBUS7	60	TCLK	80	GND	100	FS2

Table 5. 100 Pin QFP Pin Assignment

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage VDD	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to VDD + 0.5V
Output voltage V_{OUT}	-0.5V to VDD + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	0°C to 70°C
Junction temperature	100°C
Package power dissipation	1000mW

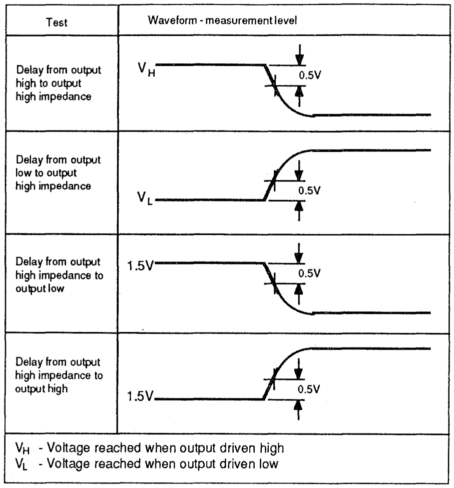

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as negative into the device.

STATIC ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)
Tamb = 0°C to +70°C VDD = 5.0v ± 5%

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	3.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.0		-	V	V_{DD}-IV for SYSCLK, DCLK
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{DD}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{DD}
Output S/C current	I_{SC}	10		300	mA	V_{DD} = Max

VP510

BI DIRECTIONAL COLOUR SPACE CONVERTER

FEATURES

- User definable colour space conversion
- Sampling rates up to 27 MHz
- On chip decimating or interpolating FIR filters
- Conversion from 24 bit inputs to 16 bit outputs or vice versa
- RAM based look up tables for gamma correction
- 100 pin Quad Flat Pack

ASSOCIATED PRODUCTS

- VP2611 Integrated H.261 Video Encoder
- VP2615 H.261 Video Decoder
- VP520 Two dimensional Video Filter
- VP530 NTSC/PAL Encoder

DESCRIPTION

The VP510 converts three channels of RGB data into two channels of decimated chrominance and luminance data. Alternatively it converts two channels of luminance and chrominance data into three channels of interpolated RGB data. Each channel has its own RAM based look up table, which can be loaded from a host system and then used for gamma correction and/or ranging.

The direction of the data flow is controlled by a bit in a Control Register, and causes previous outputs to become inputs and vice versa. The filters change from the decimating to the interpolating mode, and correspondingly follow or precede the colour space conversion.

The 3 x 3 conversion matrix is provided with user definable 12 bit coefficients which have a range from -4.0 to +4.0. The luminance channel is provided with a 23 tap low pass filter which can decimate or interpolate by two. The chrominance channels each have two 11 tap filters in series which can decimate or interpolate by four. This arrangement allows the device to accept or produce RGB data which has been 2x oversampled, thus avoiding the need for external analog anti-aliasing filters. If necessary the device will still accept or produce video data which has not been oversampled.

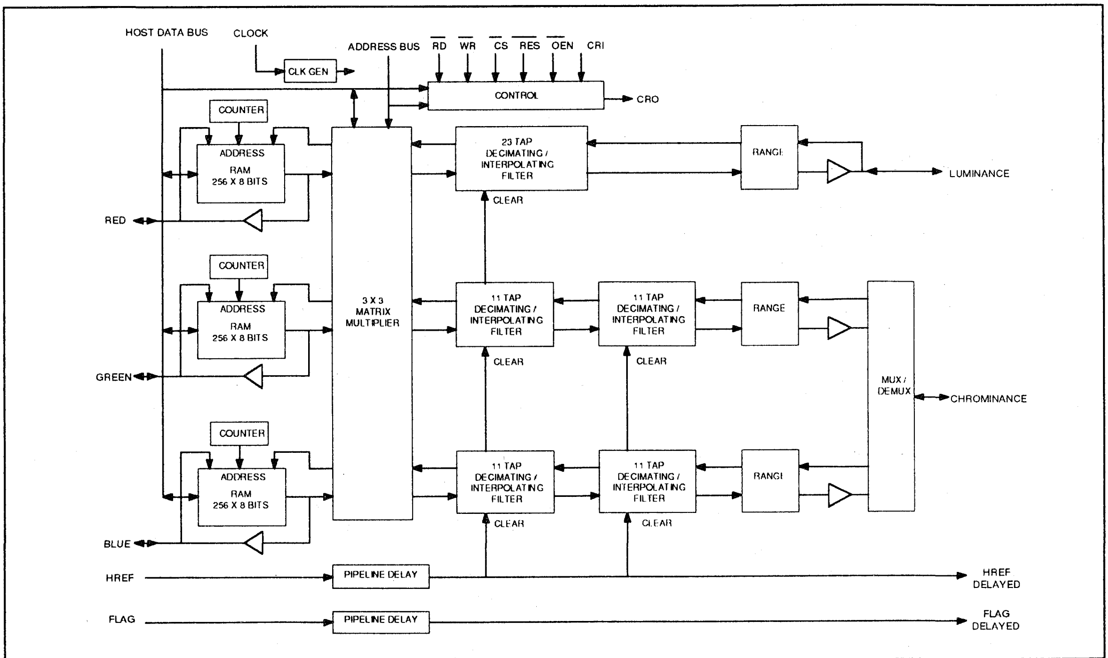

Figure 1. Simplified Block Diagram

PIN DESCRIPTION

PIN	TYPE	DESCRIPTION
R7:0	I/O	Unsigned Red data. Range may be changed by the RAM look up table
G7:0	I/O	Unsigned Green data. Range may be changed by the RAM look up table
B7:0	I/O	Unsigned Blue data. Range may be changed by the RAM look up table
Y7:0	I/O	Unsigned Luminance data in or out. Range is user definable
C7:0	I/O	Two's complement or offset binary multiplexed chrominance data. Range is user definable
D7:0	I/O	Host data bus used for reading or writing
A4:0	I	Host Address Bus. Matrix coefficients and the control register are directly addressable
CLK	I	External line locked clock. All inputs and outputs are referenced to the rising edge
HREF	I	Horizontal or Composite reference used as a start of line indicator and to clear the FIR filters
HDLY	O	HREF input delayed by the 39 clock delay to a correctly filtered output
FI	I	Input Flag as defined by the user. No internal operation.
FO	O	FI delayed by the 39 clock delay to a correctly filtered output
CRI	I	An input which indicates that valid luminance and chrominance data is present
$\overline{\text{CRO}}$	O	An output which indicates that valid luminance and chrominance data is on the output pins
$\overline{\text{OEN}}$	I	Active low output enable for the tristate bus. Used in conjunction with a Control Register bit
$\overline{\text{CS}}$	I	Active low Chip Select from the host system
$\overline{\text{RD}}$	I	Active low request from the host to read the matrix coefficients and RAM contents
$\overline{\text{WR}}$	I	Active low request from the host to write to the device
$\overline{\text{RES}}$	I	Asynchronous low reset used to initialise the device. Must be present for at least 1024 clock periods

LOOK UP TABLES

When the device is configured to produce chrominance and luminance outputs from RGB inputs, each of the three look up tables is addressed by its appropriate colour bus. Any changes to the data thus occur before the colour space conversion. Typically the look up tables are used to provide gamma correction to linear RGB inputs, and / or to limit the range of the inputs. The coefficients in the conversion matrix are usually defined to expect either a range of 1 - 254 or 16 - 235, when converting to Cr and Cb chrominance values.

When the device is configured to produce RGB outputs, the look up tables are positioned just before the output buses. If linear outputs are required the tables can then be used to remove the gamma correction which is produced by the coefficients in the conversion matrix. They can also be used to expand the range produced by the conversion matrix.

The RAM's are not dual ported and use by the host system takes priority over pixel accessing. The RAM's are not directly addressable from the host since the device only uses a 5 bit address bus. Instead each RAM has an internal address counter which must be cleared by writing to address decimal 27. Data is then sequentially written to the Red RAM by supplying 256 bytes of data and address 28. Similarly using address 29 will cause write operations to the green RAM, and address 30 will cause write operations to the blue RAM. The counters do not wrap around and must be reset by using address 27 before further write or read operations are required. Read operations are mechanized in a similar manner to write operations, except that a read strobe must be supplied instead of a write strobe. Since each RAM has its own address counter the red, green, and blue operations can be intermingled on a byte by byte basis, rather than completing one colour before starting the next.

Although host operations are asynchronous to the device clock, this clock must be present to internally effect a read or write operation. The read and write strobes are internally synchronized to the clock, and the read strobe must be active for at least five clock periods, and the write strobe for two clock periods.

CONVERSION MATRIX

The 3 x 3 matrix multiplier performs the following basic operation on three channels with identical sampling rates;

$$\begin{vmatrix} \text{O/PA} \\ \text{O/PB} \\ \text{O/PC} \end{vmatrix} = \begin{vmatrix} c1 & c2 & c3 \\ c4 & c5 & c6 \\ c7 & c8 & c9 \end{vmatrix} \times \begin{vmatrix} \text{I/PA} \\ \text{I/PB} \\ \text{I/PC} \end{vmatrix}$$

When converting from RGB to colour difference information, any decimation of the chrominance channels must be done after the above operation. Conversely when producing RGB data the chrominance channels must be interpolated before the matrix operation. The configuration bit in the Control Register takes care of this reorganization.

The coefficients C9:1 are loaded from the host system, and are directly addressable using the 5 bits provided (see Table 1). Each coefficient must be loaded as two bytes since it uses a total of 12 bits. The upper 4 bits in the most significant byte are don't care values. If the loaded values are read back by the host, these four bits will always be zero's, and are not sign bits.

The 12 coefficient bits are comprised of 3 signed integer and 9 fractional bits. This gives a decimal range of -4.00 to approximately +3.998, with the fractional bits actually giving a decimal resolution of 0.001953.

Pixel data going into the matrix multiplier uses a total of 13 bits; 10 signed integer bits plus 3 fractional bits. This additional pixel accuracy is only obtained from the output of the interpolating filters, where 10 integer bits are necessary to accommodate signed data with undershoot and overshoot beyond the nominal gain.

In the RGB to chrominance and luminance mode, when pre interpolation does not occur, only 8 unsigned integer bits are available from the look up table. Thus, within the 13 bit total, the top 2 bits plus the bottom 3 bits will be made into zero's.

Intermediate precision within the matrix multiplier grows to 15 signed integer bits plus 6 fractional bits. The least significant 9 or 10 of the integer bits are selected at the output, and the fractional bits are rounded to 3 bits. Ten integer bits are used when the matrix is producing RGB from interpolated chrominance and luminance. This allows for undershoot and overshoot beyond the nominal 8 bit unsigned value.

Only 9 integer bits are necessary when the matrix is producing chrominance, and the three fractional bits provide additional precision into the decimating filter. In fact, if the matrix is producing normalized chrominance, the coefficients will have been chosen to produce an output in the range ±127. This range only requires 8 integer bits, and the ninth bit will be a repeated sign bit. Note that ±127 is actually representing ±0.5 in this context. When the NORM bit in the Control Register is reset, the chrominance outputs lie in the range ±1, or ±256 in our internal representation. The full 9 integer bits are then needed.

LUMINANCE FILTER

The luminance channel contains a 23 tap low pass filter with internally defined 10 bit signed coefficients. When the MODE bit in the Control Register is reset the filter will decimate the sampling rate by two. When the MODE bit is set the filter will interpolate the incoming data to produce outputs at twice the incoming sampling rate. The filter coefficients remain the same in both cases, but the gain is adjusted to preserve the energy content..

When the filter is producing decimated luminance it accepts data from the matrix converter with 9 signed integer bits plus 3 fractional bits (9.3). Since luminance is always positive, however, the most significant bit will be zero. Words within the filter calculation are allowed to grow to 15 integer bits plus 6 fractional bits. This is then rounded to 15 bits plus 3 fractional bits, and finally the 10 least significant integer bits are chosen to give a 10.3 result. The 10 bit integer component allows for any undershoot or overshoot in the nominal 0 to 255 luminance range. The three fractional bits are used to round the integer component to a 10 bit value. This is then clipped to a value between 0 and 255. Negative values become zero, and positive values greater then 255 will saturate at 255. Outputs will not saturate under normal operating conditions, and the circuit is only necessary to prevent overflow when the input swings between the maximum and minimum values. Figure 2 illustrates the bit significance at various points in the data path.

When the filter is used to interpolate incoming luminance data, the 8 bit input is padded to the 9.3 format used previously. The 13 bit output from the filter is applied to the matrix converter without further rounding.

The response given by the filter is shown in Figure 3. Stop band attenuation is approximately 45 dB, and the maximum pass band ripple is 0.07 dB. These figures were obtained with 10 bit quantized coefficients and unquantized data. The effects of the various quantization steps within the filter, plus the reduction to 10 bits, is superimposed upon Figure 3. Also shown is the CCIR601 specification for a luminance or RGB

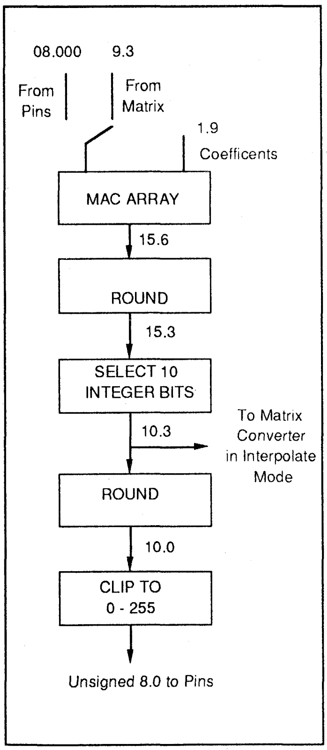

Fig 2. Bit significance in the Y Filter

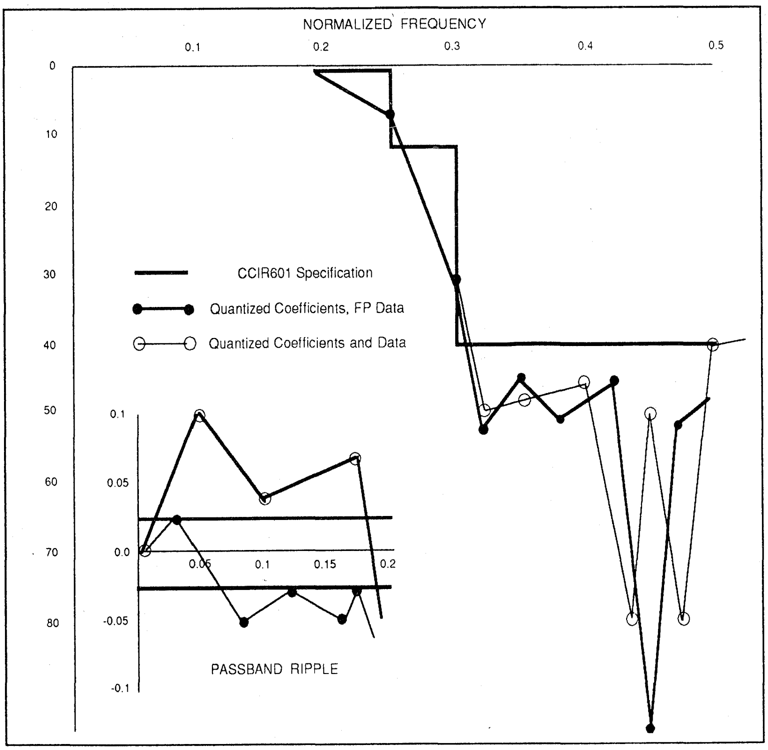

Figure 3. Response of the Luminance Filter

filter with 13.5 MHz output sampling.

CHROMINANCE FILTERS

Each chrominance channel has two 11 tap filters in series and each pair can decimate or interpolate by four. The MODE bit defines whether the filters interpolate or decimate. The coefficients are 10 bit internally defined values, and are the same in both modes. Figure 4 illustrates the bit significance at various points in the calculation.

When the filters are used to decimate chrominance produced by the matrix converter, the inputs are represents by either 8 or 9 signed integer bits plus 3 fractional bits. When the matrix coefficients have been chosen to produce normalized chrominance, the range can be represented by 8 integer bits. Otherwise 9 integer bits are needed. When the inputs are chrominance from the pins, the 3 fractional bits are set to zero, and the ninth bit is sign extended. Words within the filter calculation are allowed to grow to 15 integer bits plus 6 fractional bits. This is then rounded to 15 bits plus 3 fractional bits.

When the filter is used to supply interpolated data to the matrix converter, the least significant 10 integer bits are selected out of the 15 outputs. Only 9 integer bits are actually needed to represent the filtered chrominance with undershoot and overshoot, but the hardware multiplier expects a 10 bit number.

When the filter is producing decimated chrominance, the NORM bit in the Control Register is used to select which 12 integer and fractional bits will be used by the rounding and clipping circuit. For a full description of this operation see the section on Chrominance Outputs.

The response of the filters is given in Figure 5. These results were obtained with 10 bit quantized coefficients and unquantized data. The effects of the various quantization steps within the filter, and then finally rounding down to a 9 bit value are superimposed onto Figure 5. Also shown is the CCIR601 specification for sample rate conversion down to 4:2:2 resolution.

RGB INPUTS

The 24 bit RGB data must meet the set up and hold requirements, with respect to the rising edge of the clock, which are specified in Figure 6. The first edge after HREF has gone inactive (i.e. high) must strobe in the first samples if the delay to the first correctly filtered output is to match the fixed pipline delay of 39 clock to the HDLY and FO outputs. The maximum range is 0 to 255 for each component. If the coefficients in the matrix converter are defined for a restricted input range then this must be guaranteed by the user. Alternatively the look up tables can be used to limit the range. When HREF goes active low the outputs will go low after 39 clocks.

The VP510 has been designed to accept two times oversampled RGB data from an A/D converter. This avoids the need for analog anti aliasing filters before the A/D converters. For this reason the clock used by the VP510 is expected to be twice the sampling clock needed to produce a given number of RGB pixels per line. If the RGB inputs have not been oversampled this double rate clock should still be used. Each incoming sample will then be internally used twice, but the decimating filters will still produce the correct luminance and chrominance values.

Each input directly addresses its own RAM, which has been pre-loaded to meet the system requirements. Linear

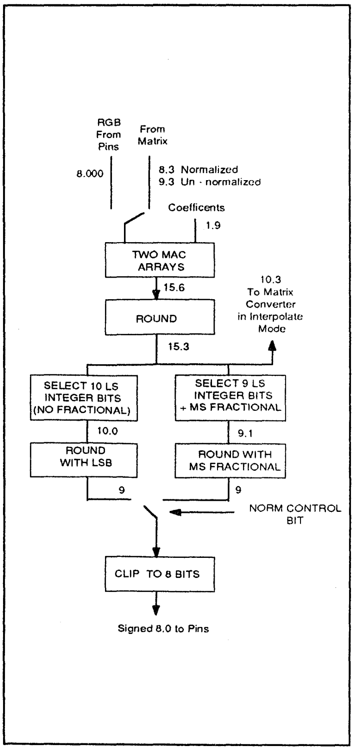

Figure 4. Bit significance

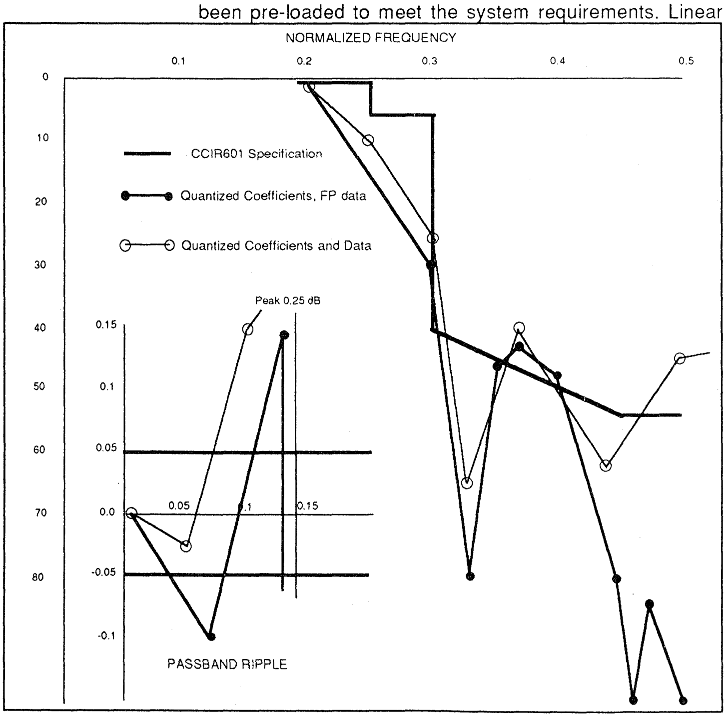

Figure 5. Response of the Chrominance Filters

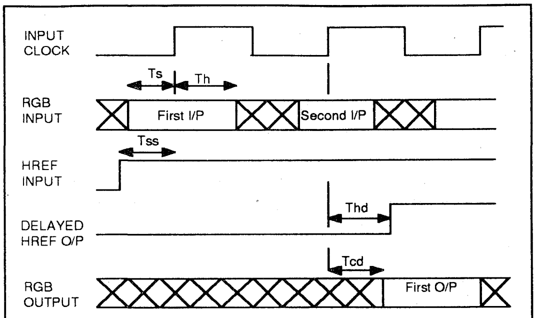

CHARACTERISTIC	SYMBOL	MIN	MAX
I/P Clock rate	Ø	DC	27MHz
I/P Set Up Time	Ts	10ns	
I/P Hold Time	Th	0ns	
HREF Set Up Time	Tss	10ns	
Delayed HREF O/P Delay	Thd		20ns
RGB O/P Delay	Tcd		20ns

Figure 6. RGB I/O Timing (Advanced Data)

RGB data must normally be gamma corrected by the RAM's before colour space conversion.

LUMINANCE AND CHROMINANCE INPUTS

The 16 bit luminance and chrominance values must meet the set up and hold times, with respect to the rising edge of the clock, which are specified in Figure 7. Since the input rate will be half the clock rate an additional signal is required to indicate alternate clock periods. This signal (CRI) must also meet the set up and hold requirements given in Figure 7. On the first occurrence of CRI after HREF goes inactive (High), the 16 bit input bus must contain the first 8 bit luminance component plus the first 8 bit U, I, or Cr component, if the delay to the first correctly filtered output is to match the fixed pipeline delay to the HDLY and FO outputs. On the second occurrence it must contain the second luminance component plus the first V, Q, or Cb component. When HREF goes low the outputs will be forced low after the 39 clock pipeline delay.

YUV or YIQ data is directly applied to the interpolating filters by setting the BYPASS Bit in the Control Register. When Y Cr Cb data is to be used this bit should be reset, and the inputs will then be applied to the ranging and offset circuitry. The SEL bit in the Control Register is used to determine the ranging options. If this bit is reset then the Y input will be

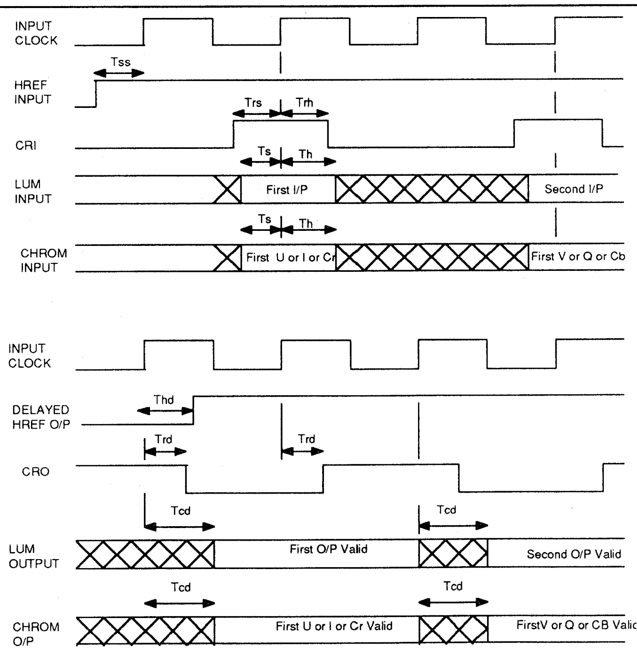

CHARACTERISTIC	SYMBOL	MIN	MAX
I/P Clock Rate	Ø	DC	27 MHz
I/P Set Up Time	Ts	10ns	
I/P Hold Time	Th	0ns	
HREF Set Up Time	Tss	10ns	
CREF Set Up Time	Trs	10ns	
CREF Hold Time	Trh	0ns	
Delayed HREF O/P Delay	Thd		20ns
CREF O/P Delay	Trd		20ns
Data O/P Delay	Tcd		20ns

Figure 7. Chrominance I/O Timing (Advanced Data)

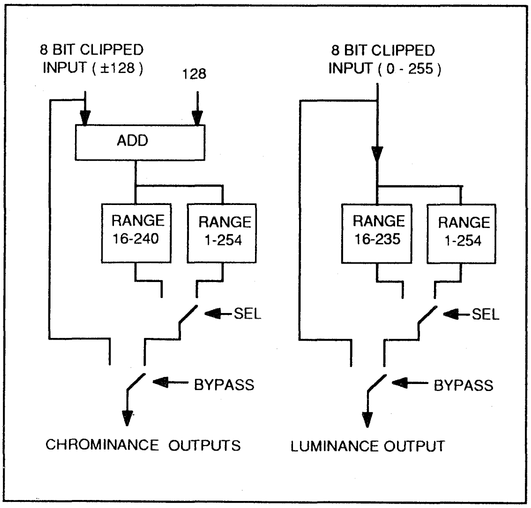

Figure 8. Chrominance and Luminance Output Options

adjusted to have a range of 16 - 235, and the Cr and Cb inputs will be adjusted to 16 - 240. If the SEL bit is set the range will be 1 - 254 for all three inputs. After either ranging option 128 is subtracted from the Cr and Cb channels before they are applied to the matrix converter. Note that if the incoming Y Cr Cb data is already correctly ranged then the range circuit will have no further action. The BYPASS pin must, however, still be reset or the offset of 128 will not be subtracted from the chrominance channels.

RGB OUTPUTS

RGB outputs will be valid after the delay from the rising edge of the clock given in Figure 6. A version of the HREF input is provided (HDLY), which has been delayed by the same number of clock periods as the data. This indicates when the first converted samples are available from each line.

In normal operation of the VP510 the clock input will be two times the sampling clock required to produce a given number of pixels per line. The device then produces RGB outputs at this double rate, and thus avoids the needed for analog anti aliasing filters after the D/A converters. Incoming luminance data is interpolated by two, and chrominance data by four, to achieve these output rates.

For standard CCIR601 video with 720 RGB pixels per line the clock needed would thus be 27 MHz. For square pixel NTSC a clock of 24.54 is needed, and square pixel PAL needs a clock of 29.5 MHz.

If the RGB outputs are connected to a frame store rather than driving a D/A converter, then these oversampled outputs are probably not needed. Since the RGB data will not contain any frequencies above one quarter the clock rate used by the VP510, then the user can simply just use every other output sample without causing aliasing effects.

Each 8 bit output value is obtained from the output of the matrix converter, which is internally represented by 13 bits. This comprises 10 signed integer bits plus three fractional bits. At this point the RGB values have a range of -512 to +511, which is sufficient to accommodate any overshoot or undershoot produced by the filters. If the most significant fractional bit is set, then the integer bits are incremented by one, and the result is then clipped. Negative values will be forced to zero, and values greater than +255 will be forced to saturate at +255. The resulting unsigned 8 bit number is made available on the output pins, as shown in Figure 2.

LUMINANCE AND CHROMINANCE OUTPUTS

The 16 bit output bus changes on alternate rising edges of the clock, with the delay specified in Figure 7. Each output remains valid for two clock periods and is either comprised of a luminance byte plus a U, I, or Cr component, or another luminance byte plus a V, Q, or Cb component. The sequence of events following the HREF delayed output is shown in Figure 7. The CRO signal can be used as a clock enable or a half rate clock for the next component in the system.

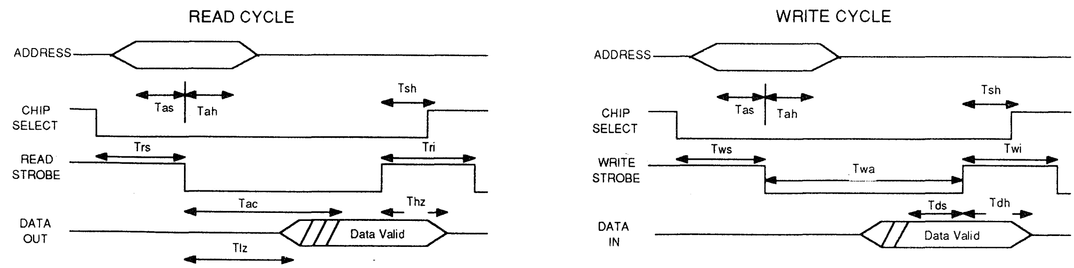

CHARACTERISTIC	SYMBOL	MIN	MAX
Addresss Set Up Time	Tas	10ns	
Address Hold Time	Tah	10ns	
Chip Select Set Up Time	Trs	0ns	
Chip Select Hold Time	Tsh	0ns	
Strobe In active Time	Tri	Øns	
Data Access Time	Tac		10 +5Øns
Delay to O/P's low Z	Tlz	4Øns	
Delay to O/P's high Z	Thz		25ns

NOTE

Ø is the period of the input clock

CHARACTERISTIC	SYMBOL	MIN	MAX
Addresss Set Up Time	Tas	10ns	
Address Hold Time	Tah	10ns	
Chip Select Set Up Time	Tws	0ns	
Chip Select Hold Time	Tsh	0ns	
Strobe In active Time	Twi	3Øns	
Strobe Active Time	Twa	2Øns	
Data Set Up Time	Tds	10ns	
Data Hold Time	Tdh	10ns	

Figure 9. Host Interface Timing (Advanced Data)

ADDR	FUNCTION	ADDR	FUNCTION
0	C1 L Byte	1	C1 H Byte
2	C2 L Byte	3	C2 H Byte
4	C3 L Byte	5	C3 H Byte
6	C4 L Byte	7	C4 H Byte
8	C5 L Byte	9	C5 H Byte
10	C6 L Byte	11	C6 H Byte
12	C7 L Byte	13	C7 H Byte
14	C8 L Byte	15	C8 H Byte
16	C9 L Byte	17	C9 H Byte
27	RAM Address Reset		
28	R/W Red RAM		
29	R/W Green RAM		
30	R/W Blue RAM		
31	Control Register		
18 - 26	Not Used		

Table 1. Internal Address Map

Internally the luminance component obtained from the decimating filter is represented by the 10 least significant integer bits plus 3 fractional bits. The 10 integer bits accommodate any undershoot or overshoot caused by the filter. If the most significant fractional bit is set, then the integer bits are incremented by one. The resulting 10 bit signed integer value ,representing ± 512, is then clipped to provide an 8 bit, positive only, number. Negative values become zero, and values greater than 255 will saturate at 255.

The NORM bit in the Control Register determines which bits out of the 15.3 available are selected from the outputs of the chrominance filters. The choice is illustrated in Figure 4. If the user is working with normalized chrominance, then the matrix coefficients will have been chosen to produce outputs in the range of ±128 (representing ±0.5). This range only requires 8 signed integer bits, and the ninth bit going into the filter will be a repeated sign bit. The 9 least significant integer bits are then selected out of the 15 available from the output of the filter. These are then sufficient to accommodate any undershoot and overshoot beyond the 8 bit input, and are rounded with the most significant fractional bit. The resulting 9 bit signed value is clipped to an 8 bit signed number with a range of ±128, representing ±0.5. Values outside the range are clipped to the maximum values allowed.

When chrominance is not normalized the range becomes ±1, or ±256 in our internal notation. This range needs all 9 bits of the integer component going into the filter, and requires 10 integer bits coming out of the filter to allow for undershoot and overshoot. The 9 bit value expected by the clipping circuit is now produced by using the least significant integer bit to round the next 9 integer bits. This word is then clipped to an 8 bit signed value with a range of ±128, but now representing ±1 since higher order bits were selected at the output of the filter.

If the BYPASS bit is set in the Control Register, these values are passed directly to the output pins. If this bit is reset they are further modified in a manner determined by the SEL bit in the Control Register. This is illustrated in Figure 8.

If the SEL bit is set, then zero luminance values become 1 and value 255 is clipped at 254. If the SEL bit is reset, then values below 16 will be forced to decimal 16 and values greater than 235 will be forced to 235.

When the BYPASS bit is reset decimal 128 will be added to each chrominance channel, to provide a positive only number. The SEL bit then either limits the range to 1 to 254 or to 16 to 240. Values outside those ranges are respectively forced to the minimum or maximum values. Note that if the BYPASS pin is reset then the NORM bit must be set.

HOST INTERFACING

The VP510 utilizes a conventional microprocessor interface except that the RAM based look up tables are not directly addressable. The address inputs must meet set up and hold times with respect to the front edge of the read and write strobes. These are given in Figure 9. Note that the address inputs are internally latched, and need not stay valid for the whole of the strobe times. Chip select, however, must stay active for the whole of the strobe times.

Data, which is to be written to the RAM or Control Register, must meet set up and hold times with respect to the back edge of the write strobe. These are also given in Figure 9. The device clock must be present for the write operation to occur, and internal synchronization takes place. For this reason the write strobe must be active for at least 2 clock periods.

Reading data from the VP510 also requires the presence of the device clock. Data from the RAM is internally pipelined and the read strobe must be active for at least 5 clock periods (4 pipeline delays plus synchronization). The output bus will not go low impedance before this pipeline delay.

The matrix coefficients and the Control Register are directly addressable, and use the locations given in Table 1. Four addresses are used to access the three RAM's, and the scheme used is descibed in the section on the look up tables.

DEVICE CONFIGURATION

The device is configured by means of bits in a Control Register. A reset pulse must be applied, whilst the device clock is active, before loading the Control Register. The reset pulse will actually clear all the control bits to zero, and ensure that neither output bus is low impedance, even if OEN is low.

The significance of the bits is given below. For a fuller description of individual bits see the releant sections.

BIT	NAME	FUNCTION
0	OEI	This bit must be set and the OEN pin must be low for either the 24 or 16 bit output bus to be low impedance. The status of the MODE bit determines which bus is actually enabled as an output. With this arrangement either bus can be controlled by software or by driving a pin.
1	SEL	This bit controls the range of the luminance and chrominance data. When high the I/O range is 1-254. When low the luminance is 16-235 and the chrominance is 16-240.
2	MODE	This bit selects the direction of operation. When low the 24 bit bus represents RGB inputs and the 16 bit bus represents luminance and chrominance outputs. The filters then deci mate. When high the data flow reverses and the filters interpolate.

3 BYPASS This bit should be reset when Cr Cb data is to being processed. NORM must then be set. It should be set when the ranging and offset circuit is to be bypassed.

4 NORM When this bit is reset the chrominance outputs are not normalized, and the 8 bit outputs rep resent a range of ±1. When NORM is set the outputs will represents a range of ±0.5, still using 8 bits.

7:5 Reserved. Must all be reset.

CONVERSION BETWEEN RGB AND YUV

If incoming, gamma corrected, analog RGB is normalized to a range of 0 to 1, then the following coefficients will produce YUV outputs. Y will have a range of 0 to 1, U will have a range of ±0.436, and V will have a range of ±0.615. The NORM bit must be reset, and the BYPASS bit set. The 8 bit chrominance outputs then represent a possible range of ±1.

$$\begin{vmatrix} Y \\ U \\ V \end{vmatrix} = \begin{vmatrix} 0.299 & 0.587 & 0.114 \\ -0.147 & -0.289 & 0.436 \\ 0.615 & -0.51 & -0.100 \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

The coefficients given below will produce gamma corrected RGB normalized to a range of ±1, when YUV have the ranges given above.

$$\begin{vmatrix} R \\ G \\ B \end{vmatrix} = \begin{vmatrix} 1 & 0 & 1.140 \\ 1 & -0.395 & -0.581 \\ 1 & 2.032 & 0 \end{vmatrix} \begin{vmatrix} Y \\ U \\ V \end{vmatrix}$$

These coefficients translate to the following HEX values, which define the 12 bit number to be loaded. Note that these are given as simple three digit HEX values, without a separate 3 bit integer and 9 bit fractional part.

$$\begin{vmatrix} Y \\ U \\ V \end{vmatrix} = \begin{vmatrix} 099 & 12C & 03A \\ F64 & F6C & 0DF \\ 13A & EF8 & FCC \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

$$\begin{vmatrix} R \\ G \\ B \end{vmatrix} = \begin{vmatrix} 200 & 000 & 247 \\ 200 & F35 & ED6 \\ 200 & 410 & 000 \end{vmatrix} \begin{vmatrix} Y \\ U \\ V \end{vmatrix}$$

If normalized digital UV components are required, the coefficients must be modified as given below. The NORM and BYPASS bits should then be set. The U I/O range is expanded to ±0.5, and the V I/O range is compressed to the same values. Y has an I/O range of 0 to 255. The 8 bit chrominance outputs now represent a range of ±0.5.

$$\begin{vmatrix} Y \\ U \\ V \end{vmatrix} = \begin{vmatrix} 0.299 & 0.587 & 0.114 \\ -0.169 & -0.331 & 0.500 \\ 0.5 & -0.419 & -0.081 \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

$$\begin{vmatrix} R \\ G \\ B \end{vmatrix} = \begin{vmatrix} 1 & 0 & 1.42 \\ 1 & -0.344 & -0.714 \\ 1 & 1.772 & 0 \end{vmatrix} \begin{vmatrix} Y \\ U \\ V \end{vmatrix}$$

The equivalent HEX values which be loaded into the device are given below;

$$\begin{vmatrix} Y \\ U \\ V \end{vmatrix} = \begin{vmatrix} 099 & 12C & 03A \\ FA9 & F56 & 100 \\ 100 & F29 & FD6 \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

$$\begin{vmatrix} R \\ G \\ B \end{vmatrix} = \begin{vmatrix} 200 & 000 & 2CD \\ 200 & F4F & E92 \\ 200 & 38B & 000 \end{vmatrix} \begin{vmatrix} Y \\ U \\ V \end{vmatrix}$$

CONVERSION BETWEEN RGB AND YIQ

The coefficients for converting analog RGB to YIQ are given below. The gamma corrected RGB inputs have a range of 0 to 1. Analog I and Q have ranges of ±0.596 and ±0.525 respectively, and the NORM bit must be reset to produce 8 bit outputs representing a range of ±1. The BYPASS bit must be set.

$$\begin{vmatrix} Y \\ I \\ Q \end{vmatrix} = \begin{vmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.275 & -0.321 \\ 0.212 & -0.523 & 0.311 \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

In the opposite direction the following coefficients produce gamma corrected RGB, when the YIQ inputs have the ranges given above.

$$\begin{vmatrix} R \\ G \\ B \end{vmatrix} = \begin{vmatrix} 1 & 0.956 & 0.620 \\ 1 & -0.272 & -0.647 \\ 1 & -1.108 & 1.705 \end{vmatrix} \begin{vmatrix} Y \\ I \\ Q \end{vmatrix}$$

In HEX these values become;

$$\begin{vmatrix} Y \\ I \\ Q \end{vmatrix} = \begin{vmatrix} 099 & 12C & 03A \\ 131 & F73 & F5B \\ 06C & EF4 & 09F \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

$$\begin{vmatrix} R \\ G \\ B \end{vmatrix} = \begin{vmatrix} 200 & 1E9 & 139 \\ 200 & F74 & EB4 \\ 200 & SDC8 & 368 \end{vmatrix} \begin{vmatrix} Y \\ I \\ Q \end{vmatrix}$$

The conversion between digital RGB and normalized digital YIQ requires the following coefficients. I and Q are then compressed to fall in the range of ±0.5, and the NORM bit must be set since the 8 bit chrominance outputs now represent ±0.5. The BYPASS bit must also be set.

$$\begin{vmatrix} Y \\ I \\ Q \end{vmatrix} = \begin{vmatrix} 0.299 & 0.587 & 0.114 \\ 0.500 & -0.231 & -0.269 \\ 0.203 & -0.500 & 0.297 \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

$$\begin{vmatrix} R \\ G \\ B \end{vmatrix} = \begin{vmatrix} 1 & 1.139 & 0.648 \\ 1 & -0.324 & -0.677 \\ 1 & -1.321 & 1.783 \end{vmatrix} \begin{vmatrix} Y \\ I \\ Q \end{vmatrix}$$

These correspond to the HEX coefficients gven below;

$$\begin{vmatrix} Y \\ I \\ Q \end{vmatrix} = \begin{vmatrix} 099 & 12C & 03A \\ 100 & F89 & F76 \\ 068 & F00 & 098 \end{vmatrix} \begin{vmatrix} R \\ G \\ B \end{vmatrix}$$

R		200	247	146	Y
G	=	200	F5A	EA5	I
B		200	D5B	391	Q

CONVERSION FROM Y Cr Cb TO RGB

The analog conversion matrix is given below;

R = Y + 1.402(Cr - 128)
G = Y - 0.714(Cr - 128) - 0.344(Cb - 128)
B = Y +1.772(Cb - 128)

If the Y Cr Cb ranges are all 1 to 254, then the RGB range produced will be 1 to 254. If the Y range is 16 to 235 and the Cr Cb ranges are 16 to 240, then the expected RGB range is 16 to 235. Incoming Y CR Cb data can be adjusted to either of these ranges by using the SEL bit in the Control Register. The input circuit also does the necessary subtraction of 128 from the Cr and Cb values (the BYPASS bit must be reset). The resulting HEX values which must be loaded into the coefficient store are given below;

C1	C2	C3		200	2CE	0
C4	C5	C6	=	200	E92	F50
C7	C8	C9		200	0	38B

The digital conversion matrix is given below;

R = Y + 1.37(Cr - 128)
G = Y - 0.698(Cr - 128) - 0.336(Cb - 128)
B Y + 1.73(Cb - 128)

The corresponding HEX values are given below;

C1	C2	C3		200	2BD	0
C4	C5	C6	=	200	E9B	F54
C7	C8	C9		200	0	376

The digital matrix only functions correctly when the Y range is 16 to 235 and the Cr Cb ranges are 16 to 240. The RGB range produced should then be 16 to 235. Both the SEL and BYPASS bits should thus be reset.

CONVERSION FROM RGB TO Y Cr Cb

The analog matrix is given below;

Y = 0.299R + 0.587G + 0.114B
Cr = 0.5R - 0.419G - 0.081B + 128
Cb = -0.169R - 0.331G + 0.5B + 128

This can handle RGB ranges of either 1 to 254 or 16 to 235. If necessary the RAM based look up tables can be used to limit the range of the incoming RGB. The BYPASS bit must always be reset, and the NORM bit set, when producing Cr and Cb data.

The SEL bit is used to limit the range of the YCr Cb values which are outputed. When SEL is set all three output ranges are 1 to 254. When it is reset the Y range is 16 to 235, and the Cr Cb ranges are 16 to 240. Values ouside the range limits will be forced to the correct maximum or minimum value. The offset of 128 is added to the Cr Cb values before the ranging is done.

The HEX values which correspond to the analog matrix are given below;

C1	C2	C3		99	12D	3A
C4	C5	C6	=	100	F29	FD7
C7	C8	C9		FA9	F57	100

In the CCIR601 specification the digital matrix is expressed as fractions of 256, and is given below;

Y = 77/256R + 150/256G + 29/256B
Cr = 131/256R - 110/256G - 21/256B + 128
Cb = -44/256R - 87/256G + 131/256B + 128

The HEX values which correspond to this digital matrix are given below;

C1	C2	C3		9A	12C	3A
C4	C5	C6	=	106	F24	FD6
C7	C8	C9		FA8	F52	106

This matrix expects the RGB inputs to be in the range of 16 to 235, and also the SEL bit to detemine the output range.

CONVERSION BETWEEN RGB AND Y,R-Y,AND B-Y

The analog matrices used to convert between RGB and Y Cr Cb can also be used with normalized colour difference information. The BYPASS bit must, however, be reset to avoid the 128 offset circuitry. RGB and Y inputs and outputs will have a range of 0 to 255. Colour difference inputs and outputs will have a range of -128 to +127 (±0.5). The NORM bit should always be set.

When working with analog colour difference values the following coefficients should be used, with the NORM bit reset. R - Y will have a range of ±0.701, and B - Y a range of ±0.886.

Y		0.299	0.587	0.114	R
R-Y	=	0.701	-0.587	-0.114	G
B-Y		-0.299	-0.587	0.886	B

R	=	1	1	0	Y
G		1	-0.509	-0.194	R-Y
B		1	0	1	B-Y

The corresponding HEX values are given below;

Y		099	12C	03A	R
R-Y	=	167	ED3	FC6	G
B-Y		F67	ED3	1C6	B

R		200	200	0	Y
G	=	200	EFB	F9D	R-Y
B		200	0	200	B-Y

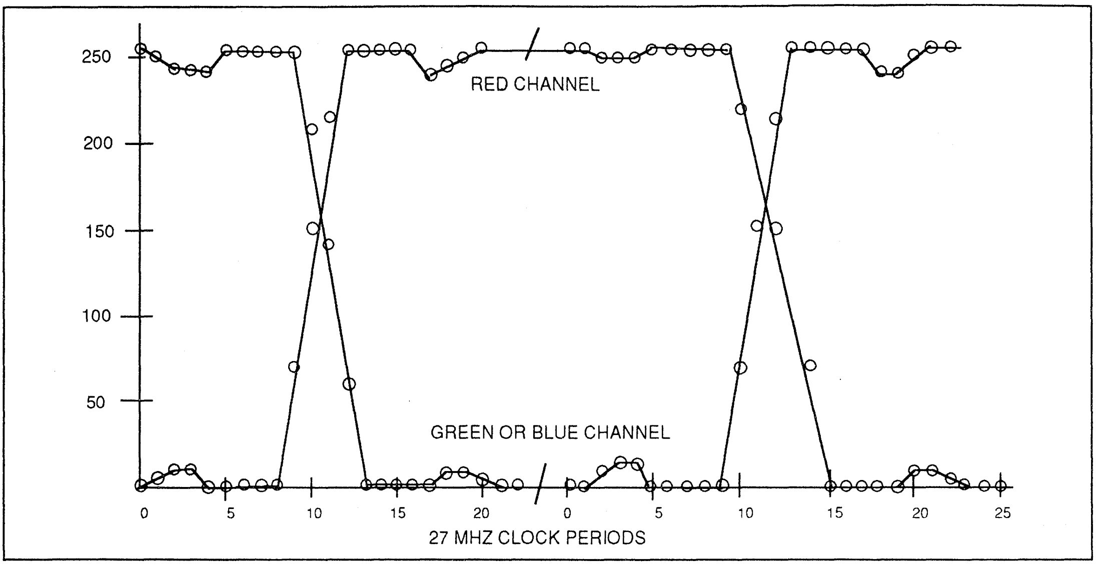

Figure 10. RGB Response to step changes in Y Cr Cb

RESPONSE TO INPUT STEP CHANGES

Figure 10 shows the actual response given by the RGB outputs to step changes in the Y Cr Cb inputs. Note that both negative undershoot and overshoot above 255 are prevented by the clipping circuit. The response of the Blue and Green filters will always be identical since they use identical circuits. The Red channel uses a different interpolating filter.

The Y Cr Cb changes were calculated to theoretically cause the RGB outputs to swing from maximum to minimum values, using the analog coefficients. Initial Y Cr Cb values were 151/20/43 with a step to 105/236/213. These should cause RGB to change from 0/255/0 to 255/0/255. The second transition was caused by changing the Y CR Cb values from 228/148/0 to 179/0/171. This should cause RGB to change from 255/255/0 to 0/255/255.

Figure 11 show the response of the Y Cr Cb outputs to maximum range step changes in RGB. The sequence used to cause the four transitions, and the theoretical results are given below.

	R	G	B	Y	Cr	Cb
Start	255	255	255	235	128	128
T1	255	255	0	226	149	16
T2	0	255	255	178	16	172
T3	0	255	0	148	21	44
T4	255	0	255	106	237	212

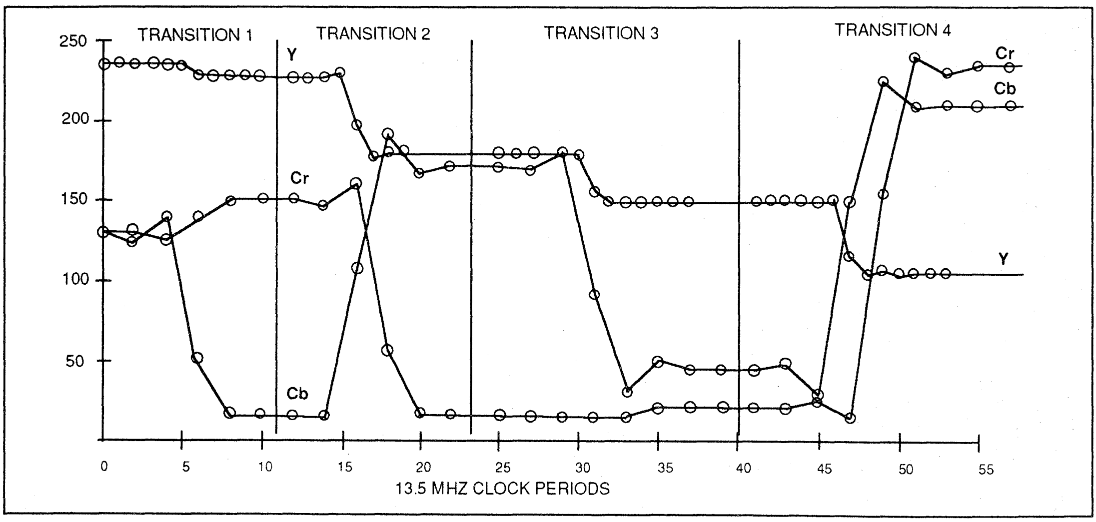

Figure 11. Y Cr Cb response to step changes in RGB

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage Vcc	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to Vcc + 0.5V
Output voltage V_{OUT}	-0.5V to Vcc + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	0°C to 70°C
Junction temperature	100°C
Package power dissipation	1000mW

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as negative into the device.

Test	Waveform - measurement level
Delay from output high to output high impedance	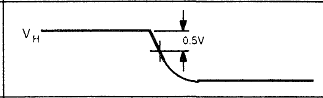
Delay from output low to output high impedance	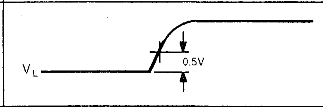
Delay from output high impedance to output low	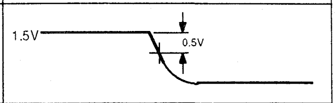
Delay from output high impedance to output high	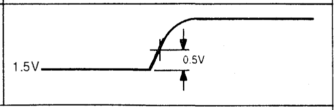

V_H - Voltage reached when output driven high
V_L - Voltage reached when output driven low

STATIC ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)
Tamb = 0 C to +70°C Vcc = 5.0v ± 10%

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	3.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.0		-	V	3V for CLK
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	μA	GND < V_{IN} < V_{CC}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	μA	GND < V_{OUT} < V_{CC}
Output S/C current	I_{SC}	10		300	mA	V_{CC} = Max

FUNCTION	PIN	FUNCTION	PIN	FUNCTION	PIN	FUNCTION	PIN
NC	1	R0	26	NC	51	C0	76
NC	2	GND	27	NC	52	NC	77
NC	3	NC	28	NC	53	GND	78
VDD	4	NC	29	VDD	54	NC	79
CLK	5	NC	30	Y7	55	NC	80
RES	6	VDD	31	Y6	56	VDD	81
GND	7	G7	32,	NC	57	D7	82
OEN	8	G6	33	Y5	58	D6	83
GND	9	G5	34	NC	59	D5	84
FI	10	G4	35	Y4	60	D4	85
NC	11	G3	36	Y3	61	D3	86
HREF	12	G2	37	Y2	62	D2	87
FO	13	G1	38	Y1	63	D1	88
HDLY	14	G0	39	Y0	64	D0	89
CRI	15	GND	40	GND	65	GND	90
CRO	16	VDD	41	VDD	66	VDD	91
GND	17	B7	42	C7	67	A4	92
VDD	18	B6	43	C6	68	A3	93
R7	19	B5	44	C5	69	A2	94
R6	20	B4	45	C4	70	A1	95
R5	21	B3	46	C3	71	A0	96
R4	22	B2	47	C2	72	CS	97
R3	23	B1	48	NC	73	RD	98
R2	24	B0	49	C1	74	WR	99
R1	25	GND	50	NC	75	GND	100

Pin Out Diagram

VP520

PAL/NTSC TO CIF/QCIF CONVERTER

IMPORTANT

At the time of publication of this handbook the VP520 datasheet was being extensively revised.
Please check with your nearest Customer Service Centre for the latest information.

FEATURES

- Converts CCIR601 luminance and chrominance to CIF or QCIF resolution, and vice versa, using a 27MHz system clock.
- Luminance and chrominance channels have their own sets of horizontal and vertical filters with on chip line stores
- Each filter set may be configured to either decimate or interpolate.
- NTSC line insertion or removal mode
- Produces / expects CIF/QCIF data in macroblock format.
- 120 Pin Package

ASSOCIATED PRODUCTS

- VP510 Colour Space Converter
- VP530 PAL/ NTSC Encoder
- VP2611 H261 Encoder
- VP2615 H261 Decoder
- VP2612 Video Multiplexer
- VP2614 Video Demultiplexer

DESCRIPTION

The VP520 is designed to convert 16 bit multiplexed luminance and chrominance data between CCIR601 and CIF/ QCIF resolutions. Vertical and horizontal FIR filters are provided, with the vertical filters supported by on chip line stores. The coefficients used by the filters are user definable, and are down loaded from an independent host data bus. An internal address generator supports an external DRAM frame store, and also provides line to macroblock conversion.

When producing CIF or QCIF video the horizontal filters precede the vertical filters, and are provided with between 8 and 16 taps. The vertical filters are provided with four CIF line delays which allow a 5 tap filter to be implemented. When producing QCIF the available RAM is used to provide six line delays, which thus allows 7 tap filters to be used.

When the device is producing CCIR601 video, the incoming data must be in macroblock format, and the vertical filters precede the horizontal filters The inputs are firstly written to a external CIF sized frame store, and are read out in line format. The VP520 will support two complete frame stores, and allows the CIF/QCIF data to be read out twice in order to produce two interlaced fields of video.

The VP520 supports the conversion between CIF/QCIF and NTSC video. An extra line is produced for every five lines when producing CIF data, and one line in six is removed when producing NTSC video. Poly phase filters are used to provide the correct decimation and interpolation ratios.

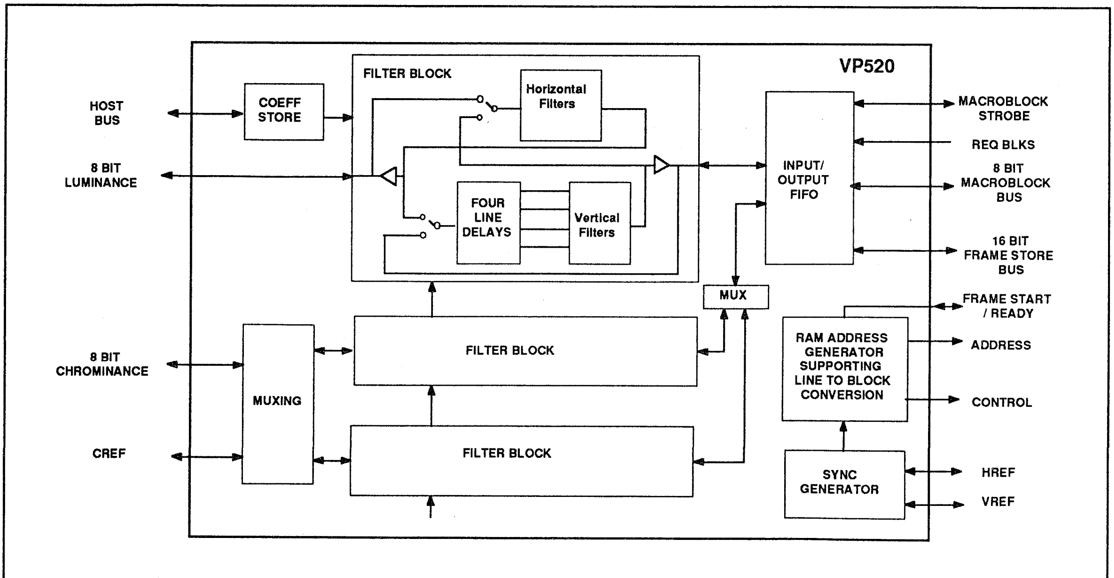

Fig 1 : Simplified Block Diagram

PIN DESCRIPTION

NAME	TYPE	FUNCTION
Y7:0	I/O	Luminance input or output bus
C7:0	I/O	Chrominance input or output bus
M7:0	I/O	Macroblock input or output bus
D15:0	I/O	16 bit data bus for DRAM frame store
A7:0	O	Multiplexed address bus to the DRAM
A8	O	Most sig address bit or second CAS
$\overline{\text{RAS}}$	O	Row strobe for the DRAM's
$\overline{\text{CAS}}$	O	Column strobe for the DRAM's
R/$\overline{\text{W}}$	O	Read/ write signal to the DRAM's
HREF	I/O	Horiz. reference in or horiz. sync out
VREF	I/O	Vertical reference in or vertical sync out
CREF	I/O	CREF in or CREF out
FREF	I/O	Field Indicator in or out
$\overline{\text{HBLNK}}$	O	Horizontal Blanking output
$\overline{\text{CSYNC}}$	O	Composite sync output in free run mode
CLMP	O	Defines a black level clamping period for A/D converters
REQYUV	I	Request macroblocks from encoder
MCLK	I/O	Macroblock I/O strobe
FSIG	I/O	Frame start/ ready signal
SCLK	I	System Clock. 27MHz in PAL/NTSC systems
HD7:0	I/O	Host data bus
HA3:0	I	Host controller address bits
$\overline{\text{RD}}$	I	An active low host read strobe
$\overline{\text{WR}}$	I	An active low host write strobe
$\overline{\text{CEN}}$	I	An active low enable for the strobes
$\overline{\text{RST}}$	I	Power on reset
TDI	I	JTAG I/P data
TDO	O	JTAG O/P data
TMS	I	Test mode select
TCK	I	JTAG clock
$\overline{\text{TRST}}$	I	JTAG reset
TOE	I	When high all O/P's are high impedance

NOTE:
"Barred" active low signals do not appear with a bar in the main body of the text.

VIDEO COMPRESS MODE (DECIMATE)

This mode is used when CCIR601 video is to be converted to CIF or QCIF spatial resolution prior to compression. Incoming luminance and chrominance data does not need any prior buffering, but must meet the timing requirements given in Figure 2. A bit in Control register 1 allows the Cb component to precede the Cr component if necessary. This data is passed through vertical and horizontal decimating filters before it is stored in an external frame store. When a complete field has been decimated it is read out in macroblock format and transferred to the next system component.

In this mode HREF, VREF, CREF, and FREF are normally inputs which are used to reference active video with respect to video synchronization pulses. The active going edges are used internally, and these must meet the set up time with respect to the system clock as given in Figure 2. The reference inputs need only stay active for one system clock period. Note that the active going edges for HREF and VREF can individually be defined to be high going or low going, through two bits in Control Register 0. Also note that CREF is used as a qualifier for SCLK and the actual edges of CREF are not used.

The internal sync generator can still be used in this mode, if there is a need to supply sync to the video source. The HREF and VREF pins are then used to output HSYNC and VSYNC. Composite sync is supplied on the CSYNC pin.

In addition the CLMP pin provides a pulse [13 SCLK's wide] which can be used to DC restore the black level in an A/D converter. It is active high during the back porch.

The horizontal blanking output (HBLANK) defines when the device expects the first pixel in a line to be supplied, and is derived from the user supplied HREF input. The delay between HREF and HBLANK is user definable in multiples of CREF periods. If the defined value is zero then the HREF input must be horizontal blanking with the minimum set time specified. The HBLANK output is then not defined.

All data changes are referenced to the system clock. The edge actually used is indicated by the CREF input signal, which has a period of double the clock period. The VP520 will strobe in data on the rising edge of the system clock which occurs whilst CREF is high.

The first video line to be filtered and stored will be derived from the vertical reference input (VREF). The user can choose

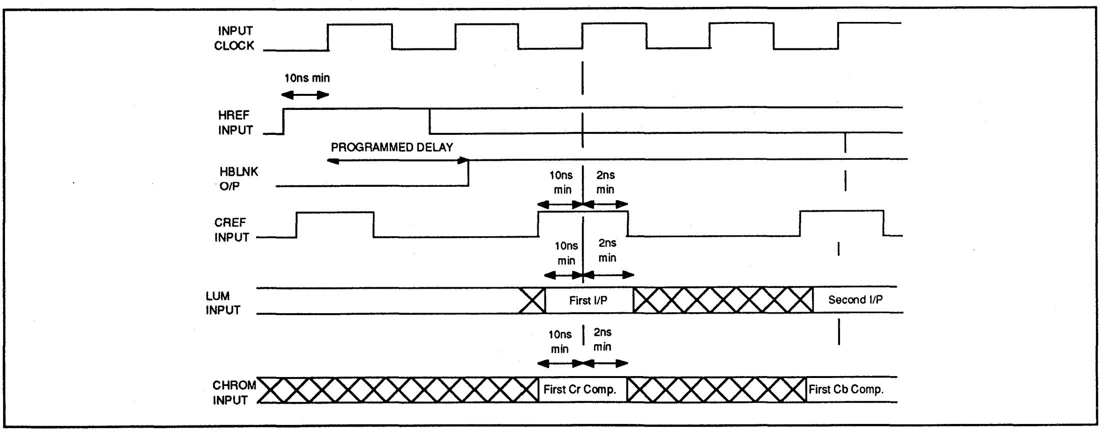

Fig 2 : Luminance and chrominance inputs in the decimate mode.

the number of transitions of the HREF input which must occur, after VREF has gone active, before starting the filter operation. Data is then not written to the DRAM until after the pipeline delay through the filters.

The VP520 only expects to use one field of CCIR601 video, which can be selected by the FREF input or internal logic. A bit in Control Register 1 (Internal / External Field Detect) determines which option is to be used. An additional Field Select Bit determines whether the field selected should correspond to FREF being high or low. When the Field Select Bit and the input are at the same logical level then that field is used. Note that FREF transitions must be coincident with active going VREF transitions.

Internal logic is provided which determines the field (Field 1) in which VREF goes active in less than half a line period after the HREF input last went active.The half line period is determined by VREF going active between 1 and 360 CREF qualified SCLK edges after HREF went active. Note that coincident VREF and HREF edges will indicate this field on the first CREF qualified SCLK edge.

This logic is used, rather than the FREF input, when the Internal / External Field Detect Bit is low. Field 1 is selected when the Field Select in Control register 1 is low, and Field 2 is used when the bit is high.

In the Split Screen mode this logic is overridden, and both fields are actually used. External logic is assumed to switch between two sources of video, one for each field. The internal DRAM address generator is modified such that half area pictures from the centre of each source are actually stored as CIF/QCIF data. The first line used in each field will be 72 line delays in addition to the number which has been defined by the user.

The VP520 will insert zero's into the line delays during vertical blanking. This ensures that all the filter accumulators are cleared and the edges of the picture are correctly processed. The horizontal filters always give the required results since four decimated values are ignored at either side of the picture.

Incoming luminance data could have a black level of 16, which will be shifted if the filter coefficients are not chosen to exactly give a gain of unity. A Control Bit is thus provided, which when set causes 16 to be subtracted from incoming luminance. A black level of zero will then stay as zero throughout the filter operation. At the output of the filters 16 is always added to the results, regardless of the state of the Control Bit. Saturation logic ensures that these addition / subtraction operations do not produce negative results or values greater than 254.

A Control Bit is also provided which selects between colour difference inputs and true Cr Cb chrominance values. Cr Cb values are 8 bit positive only numbers, with black levels of 128. These must be converted to two's complement signed numbers by subtracting 128, thus giving a black level of zero through the filters. The outputs of the filters are always converted to positive only Cr Cb values by adding 128 to the results, regardless of the state of the Control Bit.

CIF/QCIF MACROBLOCK OUTPUTS

When producing decimated CIF/QCIF data in macroblock format, the device raises a flag when a frame of data is ready for reading from the frame store (FSIG). The FSIG pin is automatically configured as an output in the decimate mode, but will only stay active (high) for the time given in Figure 2. If a Request Macroblock response (REQYUV) is not obtained during this period, then FSIG will be taken low and the frame of data presently available will be ignored. It will go high again when a new frame of data is available.

When it receives a REQYUV response from the next system component, it starts to output a macroblock by using an output strobe derived by dividing down the clock input. Detailed timing is given in Figure 2. This strobe only occurs when data is available at the output pins, and occurs in bursts at a rate of SYSCLK/2, but with an average rate of SYSCLK/4. The ' Request Macroblock ' flag must go inactive and then active again before a further macroblock is made available.

The Frame Ready flag is only available on the output pin if the Frame Enable Bit is set in Control Register 1. Through this control bit a host controller is able to determine whether a new frame is to be compressed and transmitted. In an alternative arrangement the control bit can be permanently set, and the Frame Ready Flag is then used as an interrupt to the host controller. It then generates a signal which is used as the Frame Ready signal for the next device.

The following sections describe this interface as it applies to the VP2611 H261 Video Encoder.

TRANSFERING MACROBLOCKS TO THE VP2611

When the VP520 has stored a complete field of decimated video in the DRAM, it raises a Frame Ready Flag (FSIG). If the bit in Control Register 1 does not inhibit the output, this flag becomes the FRMIN input on the VP2611. This responds to

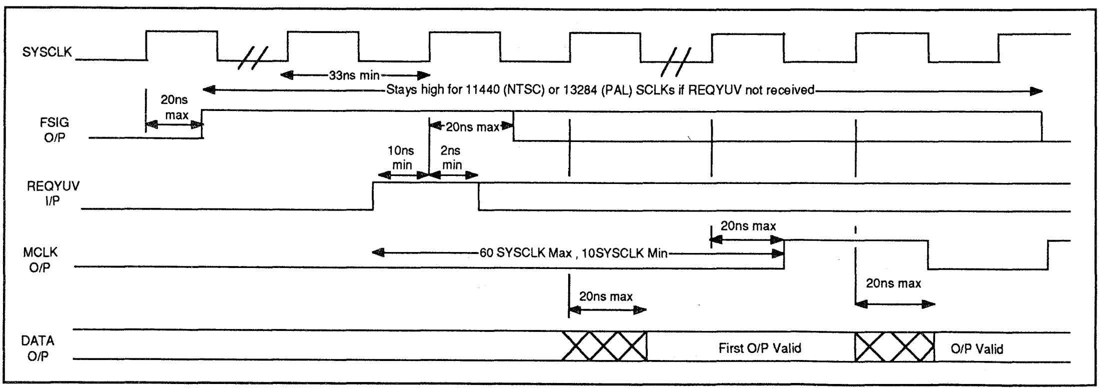

Fig 3 : Macroblock Output Timing

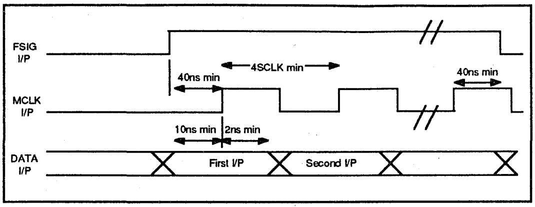

Fig 4 : Macroblock Input Timing

the FRMIN input by generating a Request for Macroblock Data (REQYUV). The VP2611 MUST then receive a complete macroblock (384 bytes) within 1870 cycles of the system clock. When the VP520 is producing decimated CIF/QCIF data, writing line data to the DRAM has priority, and only four macroblock read operations are possible in every 32 clock cycles i.e. one read takes eight cycles. These, however, are 16 bit word operations and it thus requires 384 x 8/2 = 1536 cycles to output the data. In addition there is a maximum delay of 60 clock periods from receiving REQYUV to producing the first output strobe (MCLK). This is still well within the time available.

The four 16 bit words are stored in the VP520 and transmitted to the VP2611 as eight bytes using a strobe (MCLK) derived from the system clock. This is only present when valid data is available, and it drives the PCLK input on the VP2611.

It takes the VP2611 almost exactly all the available time at 30 Hz frame rates to process all the macroblocks. After a field time (half an interlaced frame) the VP520 will start to write new data to the DRAM, and data could be overwritten during the last macroblocks. Since there is available space in the DRAM, a small address offset is used between video fields to avoid this problem.

INTERPOLATE MODE

In this mode the VP520 expects to receive CIF/QCIF data in macroblock format, which it then writes to an external frame store. This is then read back in line format and passed through vertical and horizontal interpolating filters to produce two fields of CCIR601 video. Detailed input timing is given in Figure 3.

FSIG automatically becomes an input which is used to identify the start of a frame and to reset the internal address counter. FSIG must stay high until a complete CIF/QCIF frame has been received (internal logic counts macroblocks). If FSIG goes low early then the complete frame will be ignored, and the previously received frame will continue to be displayed.

An input strobe, derived by dividing the system clock by four, must also be provided in order to input data. This must only be present when valid data is available on the input pins. Incoming macroblocks are byte wide, and these are internally buffered to allow four 16 bit words to be written to the DRAM every 32 system clock cycles. This is equivalent to a byte input rate of SCLK/4 which must not be exceeded.

The CIF frame store is double buffered such that a new frame can be received whilst the previous one is being displayed. In fact the use of 256K x 16 DRAM's gives sufficient capacity for more than three complete CIF frames, and the internal address generator will simply roll around to make full use of the available space.

Once a complete CIF/QCIF frame has been received, it will normally be used to generate two interlaced PAL or NTSC fields. These fields continue to be re-generated until a complete new CIF frame has been received. The rate of receiving frames depends on the transmission bandwidth, but the maximum rate is 30 Hz. The changeover to the newly received frame will occur when the VP520 has finished generating any one of the pair of interlaced fields for display, it does not necessarily have to have generated two interlaced fields from the received frame. If the VP520 is receiving frames at the full CIF 30 Hz frame rate but only displaying PAL frames at 25 Hz, then periodically one of the PAL frames (comprising two interlaced fields at 50 Hz) will be generated from two received CIF/QCIF frames. An incoming CIF/QCIF frame will always be used since the interlaced field rate is always greater than 30 Hz in either PAL or NTSC.

The data is read from the frame store such that interpolated data becomes available after programmed delays referenced to the VREF and HREF signals. The actual delays are greater than the programmed values because of the internal pipeline delays, which are also mode dependent.There are at least four internal line delays in CIF mode, and twelve in QCIF mode. Since the internal horizontal delay is greater than the time between HSYNC and the end of blanking, then the reference point is moved to the next HREF edge by adding a large offset of 704 CREF periods to the programmed value. (This is a convenient binary value which only modifies the 4 MSB's in the 10 bit number which defines the length of a complete line). The total delay to the first pixel is thus 704 + programmed value + mode dependent pipeline delay, the latter being of the order of 100 CREF periods but yet to be determined for each mode. There are 858 CREF periods in NTSC line, and 864 in a PAL line.

HREF and VREF can either be user supplied inputs, or are generated internally from a PAL/NTSC timing generator. A bit in Control Register 0 determines this option, and when the internal generator is specified the HREF pin becomes an

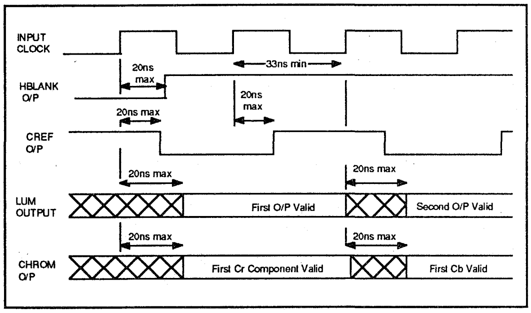

Fig 5 : Luminance and Chrominance Output Timing

SYMBOL	PARAMETER	MINIMUM	MAXIMUM
t RAC	Access time from RAS	-	105ns or under
t CAC	Access time from CAS	-	25ns or under
t RP	RAS precharge time	50ns or under	-
t CP	CAS precharge time	12ns or under	-
t RAS	RAS pulse width	80ns or under	-
t CAS	CAS pulse width	50ns or under	-
t REF	Time between complete refreshes	-	4 ms or over (8 ms with 256k x n)

N.B. All times are quoted assuming 27MHz operation. For lower clock frequencies increase the above values proportionately.

Table 1. External DRAM Timing Requirements

output which supplies horizontal sync and the VREF pin supplies vertical sync. A composite sync output is also provided for system level use. In this mode the VREF and HREF signals used internally are effectively vertical and horizontal sync, and the programmed delays should be chosen to reflect this condition.

The signals provided from the internal timing generator allow the VP520 to drive the VP510 Colour Space Converter and an RGB monitor. Detailed output timing is given in Figure 4. When a composite video monitor is required the VP530 Composite Video Encoder should be used, which will then generate HREF and VREF signals for the VP520 and also produce baseband composite video.

External chrominance data can have a zero colour difference value of either 0 or 128. This is defined using the Chrominance Control Bit. Where 128 is the zero colour difference value, 128 will be subtracted from incoming chrominance data and 128 will be added to output chrominace data. Output values will be limited to lie in the range 16 to 240.

External luminace data can have a black luminace level of either 0 or 16. This is defined using the Luminace Control Bit. Where 16 is the black value, 16 will be subtracted from incoming luminace data and 16 will be added to output luminace data. Output values will be limited to lie in the range 16 to 235.

The data stored in the CIF frame store will not contain the black levels normally present during horizontal and vertical flyback. This is inserted by the VP520 at the appropriate times in order to ensure that the correct filter operation occurs at the edges of the picture. In addition to these black levels during flyback, a bit in Control Register 1 allows all active video to be replaced by a fixed colour. This colour is user definable through YUV values in three registers.

FRAME STORE INTERFACE

All read and write operations to the external DRAM frame stores are based on the use of fast page mode with 13.5 MHz CAS cycles. Internally a 54 MHz clock is produced from the 27MHz System clock, and this determines the minimum time interval which can be used in the generation of pulses and defining precharge times. Any DRAM used must meet the timing constraints given in Table 1.

Reading and writing rates dictate the need for a 16 bit data interface, and line data is re-organized to allow a 16 bit word to consist of either two luminance values or two chrominance values. This gives compatibility with the macroblock requirements since a sub block is either all chrominance or all luminance data. Reading or writing macroblock data requires jumps between pages, but four words can always be read or written using fast page mode.

Read and write operations must be timeshared to meet the requirements of the system. This time-sharing is based on the use of 16 cycles of the 13.5 MHz clock. When reading or writing line data to the store, 10 cycles are used for eight words, and six cycles are left free for four exchanges with the encoder or decoder. The additional cycles are needed when using fast page mode in order to guarantee RAS precharge times and RAS to CAS delays.

The above time partitioning gives a line rate of 6.75 MHz, which meets real time CIF requirements. The exchange rate with the encoder or decoder is only half of this, but is adequate for CIF data at 30 Hz frame rates. In the decode mode the VP520 produces two fields at 60 Hz rates from every 30 Hz received frame, thus writing need only be half the rate of reading. In the decimate mode the VP520 produces a CIF frame using line rates which could have supported two 60Hz fields, but only one is used. Thus reading rates need only be half writing rates since the spare field time is available.

In the interpolate mode two complete CIF frame stores are required, which dictates the use of 256K word DRAM's. The A8 pin then provides the ninth address bit needed for such devices. In the decimate mode only one CIF frame store is required, and a Control Register Bit allows the user to select either 256K word DRAM's, or 64K x 16 devices. In the latter case two such devices are needed, and the A8 pin now supplies a second CAS strobe to enable the second device. Refresh cycles generate CAS before RAS sequences.

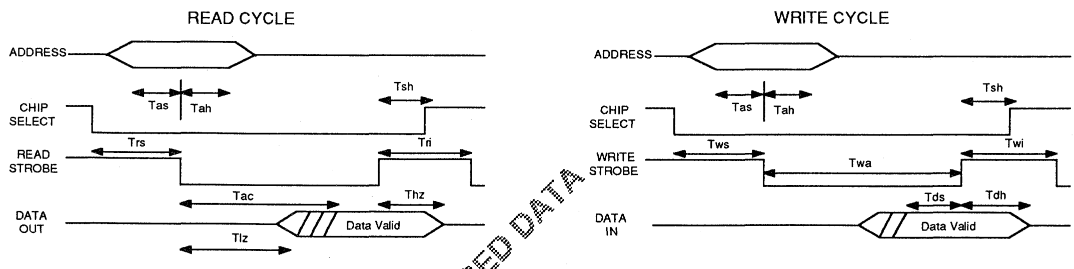

CHARACTERISTIC	SYMBOL	MIN	MAX
Addresss Set Up Time	Tas	10ns	
Address Hold Time	Tah	10ns	
Cip Select Set Up Time	Trs	10ns	
Chip Select Hold Time	Tsh	2ns	
Strobe Inactive Time	Tri	Øns	
Data Access Time	Tac		20 +3Øns
Delay to O/P's low Z	Tlz	2Øns	
Delay to O/P's high Z	Thz		25ns

NOTE

Ø is the period of the input clock

CHARACTERISTIC	SYMBOL	MIN	MAX
Addresss Set Up Time	Tas	10ns	
Address Hold Time	Tah	10ns	
Chip Select Set Up Time	Tws	10ns	
Chip Select Hold Time	Tsh	2ns	
Strobe Inactive Time	Twi	1Øns	
Strobe Active Time	Twa	3Øns	
Data Set Up Time	Tds	10ns	
Data Hold Time	Tdh	10ns	

Fig 6 : Host Interface Timing

HOST INTERFACE

The VP520 employs a conventional memory mapped host interface using a data bus and an address bus. To minimize on pin count the VP520 only uses four address lines, and all internal RAM is addressed through counters. All data is validated with a read or write strobe, and an active low enabling signal. These strobes can be asynchronous to the 27 MHz clock, but the latter must be present to move the data through several pipeline delays. Strobes must thus be valid for several clock periods. Timing is shown in Figure 5.

In the worst case mode (QCIF to NTSC video), the device must store 40 horizontal coefficients and 210 vertical coefficients. Internal storage must thus be provided for a total of 250 eight bit coefficients, and this is split into four blocks. These consist of storage for 24 horizontal luminance coefficients; storage for 16 horizontal chrominance coefficients; storage for 70 vertical luminance coefficients; and finally 140 vertical chrominance coefficients. Each block of RAM has its own internal address counter, and all counters are simultaneously reset with a write to address F hex. Each RAM area has an associated address as listed below, and a read or write using that address will increment the relevant counter. Attempts to use more addresses than are applicable to a particular area will cause undefined behaviour.

Address allocations are given below;

Addr	Function	
0	Reserved	
1	R/W horizontal luminance coefficients. Max 24	
2	R/W horizontal chrominance coefficients. Max 16.	
3, 4	Reserved for internal use	
5	R/W vertical luminance coefficients. Max 70.	
6	R/W vertical chrominance coefficients. Max 140.	
7	Reserved	
8	Control Register 0. See below.	
9	Control Register 1. See below	
A	Line delay from VREF to first active line.	See Page 4
B	Pixel delay from HREF to first active pixel	See Page 4
C	Blanked screen Y value	
D	Blanked screen U value	
E	Blanked screen V value	
F	Clear all address counters	

The bits in registers 0 and 1 are used individually, and are defined below. Where necessary the action caused when changing a control bit is delayed until the start of a new field.

REGISTER 0

BIT	FUNCTION
0	Interpolate if high, decimate if low
1	PAL if low, NTSC if high
2	QCIF if high, CIF if low
3	If low subtract 16 from Y, add 16 back after filtering
4	If low subtract 128 from chrominance, then add back
5	If low generate sync, if high lock to HREF and VREF
6	If low then active edge of VREF is low going.
7	If low then active edge of HREF is low going.

REGISTER 1

BIT	FUNCTION
0	If low then U inputs precede V inputs
1	If low use the internal field detect logic
2	Field Select. See text.
3	If low use 64Kx16 DRAM (encoder only)
4	When high specifies Split Screen mode (encoder only)
5	When low the Frame Ready Flag is enabled
6	When high the screen is blanked
7	When high DRAM writes are disabled

LOADING COEFFICIENTS

The following tables show the coefficient storage locations for different modes. The filter sections below describe the use of coefficient sets. Within a set, coefficients are stored in ascending order, ie. C0, C1, C2 etc. Note that some locations are shown as not used. However, since each store is loaded sequentially, the data stream used to load the coefficient stores must contain padding values corresponding to the unused addresses. Note also that only the address range shown in the tables have to be loaded with data.

A: Horizontal Luminance Store

This is a 24 byte RAM and coefficients will be stored as follows:

Mode	Addresses	Coefficient Set
CCIR -> CIF	0-7	1
CCIR -> QCIF	0-15	1
CIF -> CCIR	0-5	1
	6-11	2
QCIF -> CCIR	0-5	1
	6-11	2
	12-17	3
	18-23	4

B: Horizontal Chrominance Store

This is a 16 byte RAM and coefficients will be stored as follows:

Mode	Addresses	Coefficient Set
CCIR -> CIF	0-7	1
CCIR -> QCIF	0-15	1
CIF -> CCIR	0-3	1
	4-7	2
QCIF -> CCIR	0-3	1
	4-7	2
	8-11	3
	12-15	4

C: Vertical Luminance Store

This is a 70 byte RAM and coefficients will be stored as follows:

Mode	Addresses	Coefficient Set
625 line -> CIF	0-4	1
525 line -> CIF	0-4	1
	5-9	2
	10-14	3
	15-19	4

	20-24	5
	25-29	6
625 -> QCIF	0-6	1
525 -> QCIF	0-6	1
	7-13	2
	14-20	3
	21-27	4
	28-34	5
	35-41	6
CIF -> 625 line	0-4	1, even field
	5-9	1, odd field
CIF -> 525 line	0-4	1, even field
	5-9	2, even field
	10-14	3, even field
	15-19	4, even field
	20-24	5, even field
	25-29	not used
	30-34	1, odd field
	35-39	not used
	40-44	2, odd field
	45-49	3, odd field
	50-54	4, odd field
	55-59	5, odd field
QCIF -> 625 line	0-6	1, even field
	7-13	2, even field
	14-20	1, odd field
	21-27	2, odd field
OCIF -> 525 line	0-6	1, even field
	7-13	2, even field
	14-20	3, even field
	21-27	4, even field
	28-34	5, even field
	35-41	1, odd field
	42-48	2, odd field
	49-55	3, odd field
	56-62	4, odd field
	63-69	5, odd field

D: Vertical Chrominance Store

This is a 140 byte RAM and coefficients will be stored as follows:

Mode	Addresses	Coefficient Set
625 line -> CIF	0-4	1
525 line -> CIF	0-4	1
	5-9	2
	10-14	3
625 line -> QCIF	0-6	1
525 line -> QCIF	0-6	1
	7-13	2
	14-20	3
CIF -> 625 line	0-4	1, even field
	5-9	2, even field
	10-14	1, odd field
	15-19	2, odd field
CIF -> 525 line	0-4	1, even field
	5-9	2, even field
	10-14	3, even field
	15-19	4, even field
	20-24	5, even field
	25-29	not used
	30-34	1, odd used
	35-39	not used
	40-44	2, odd field
	45-49	3, odd field
	50-54	4, odd field
	55-59	5, odd field
QCIF -> 625 line	0-6	1, even field
	7-13	2, even field
	14-20	3, even field
	21-27	4, even field
	28-34	5, even field
	35-41	6, even field
	42-48	7, even field
	49-55	8, even field
	56-62	9, even field
	63-69	10, even field
	70-76	1, odd field
	77-83	2, odd field
	84-90	3, odd field
	91-97	4, odd field
	98-104	5, odd field
	105-111	6, odd field
	112-118	7, odd field
	119-125	8, odd field
	126-132	9, odd field
	133-139	10, odd field

HORIZONTAL FILTERS

Chrominance data is assumed to have already been decimated down to half the horizontal sampling rate of the luminance data, before it is applied to the VP520. When producing CIF data both luminance and chrominance are then both decimated by two, when producing QCIF data they are both decimated by four.

Simulations with actual video have shown that 8 tap CIF filters and 16 tap QCIF filters give more than adequate performance in the decimation mode. In the interpolation mode these same simulations have shown the need for longer filters in the luminance channel. The hardware thus supports a 12 tap filter when interpolating luminance from CIF inputs, but only 8 taps are provided for each chrominance channel. Even longer filters are needed when QCIF data must be interpolated, and the luminance channel is provided with 24 taps, and each chrominance channel with 16 taps.

Note that when interpolating by two the output rate is double the input rate, but every other input will be conceptually zero. Similarly when interpolating by four there are three zero's between every data point, even though the output rate is four times the input rate. Thus during any clock period only one half or one quarter of the coefficients are actually in use, and the computational burden is no greater than when doing the equivalent decimation.

Since all the coefficients are not in use during any clock cycle, it is convenient to refer to two smaller sets of coefficients. Thus the 12 tap CIF luminance filter, for example, can be considered to have two sets of 6 coefficients, and the 24 tap QCIF luminance filter to have four sets of 6 coefficients.

VERTICAL FILTERS

The vertical filters are designed to produce CIF with the spatial relationship shown in Figure 6, and QCIF with the spatial relationship shown in Figure 7. Original PAL or NTSC video contains lines of coincident luminance and chrominance, but the CIF specification requires that the decimated chrominance information is shifted such that it lies mid way between two luminance lines. This is achieved by choosing the centre outputs from the filter which best fit the requirements. The filter outputs actually used by the device are shown by the arrows in Figures 6 and 7, and are optimal when the even field provides the original video.

It is assumed that one of the interlaced fields has been discarded prior to the VP520, and thus no further decimation occurs when producing CIF luminance from PAL (NTSC in fact needs some interpolation - see the relevant section). Chrominance, however, is decimated by two. When producing QCIF data the luminance channel is decimated by two, and the chrominance by four.

When the VP520 is used to derive interlaced CCIR601 video, the internal address generator will read the CIF/QCIF frame store twice in order to produce the two fields. Each field has its own set of coefficients.

Internal RAM is provided which will support four CIF line delays for both chrominance and luminance. Five tap filters are thus possible for CIF conversions. With a QCIF system the internal RAM could theoretically be used to provide eight QCIF line delays. In practice, however, little benefit is obtained by using vertical filters with more than seven taps, and thus only six line delays are used.

Polyphase filters are used to support the spatial conversions. PAL conversion is relatively simple and only requires a set of coefficients for each mode. NTSC conversion requires several sets of coefficients since the 240 lines in a field must be converted to 288 lines of CIF. One line is repeated in every five to produce six lines which are then filtered with their own coefficients.

The generation of interpolated outputs requires CIF / QCIF data to be repeatedly read from the frame store at various line intervals. This is all handled by the internal address generator, and is transparent to the user. The device then produces coincident luminance and chrominance data which has been interpolated from data in the frame store. The first line will be produced to match the delay from the VREF input which has been pre-defined. This delay must be greater than the internal pipeline delay, which itself is mode dependent (delay yet to be determined).

The device introduces black lines at the top and bottom of the fields. Thus the first and last lines in the interpolated field will be filtered with varying amounts of black information.

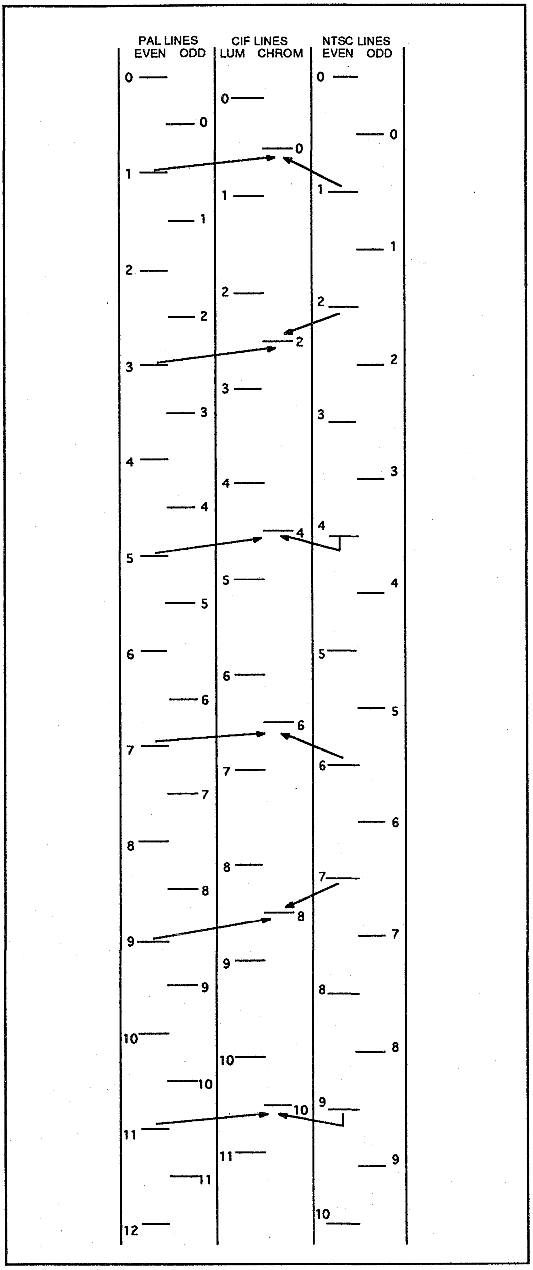

Fig 7 : CIF Spatial Relationships

PAL VERTICAL FILTERING

When producing CIF data the five tap filters provide outputs for every line at the 6.75 MHz decimated line rate. Every filtered luminance line is used but every other filtered chrominance line will be discarded. Filter outputs corresponding to odd numbered PAL chrominance lines being at the centre are used to provide the CIF chrominance lines. This is shown by the arrows in Figure 6.

When decimating down to QCIF seven tap filters are used, which provide outputs for every line at the 3.375 MHz line rate. Only every other filtered luminance line, and every fourth chrominance line are actually stored in the frame store. Different PAL lines are used to produce the offset luminance and chrominance lines as indicated by the arrows in Figure 7.

When interpolating from CIF the luminance channel conceptually uses a 10 tap filter, with every other input line containing only zero's. Thus only five coefficients are actually used when producing interpolated lines for the even field, and five different coefficients are used when producing the odd field. The device thus stores two sets of five coefficients; one set for each field produced by reading the CIF frame store twice.

The chrominance filter is conceptually a 20 tap filter with three lines of zero's for every actual input. Thus each chrominance channel needs four sets of five coefficients; two sets are needed to produce one field, and two sets are needed for the other field. The same chrominance data is read twice for a given pair of luminance lines, in order to provide inputs for the filter. Thus the internal line delays contain the same set of chrominance data on two consecutive lines supplying data to the filters.

When interpolating from QCIF, seven coefficients can be used in each set since six line delays are provided. The luminance filter conceptually contains 28 taps (four sets of seven coefficients with two sets used to produce each field). Similarly the chrominance filter consists of 56 taps arranged as eight sets of seven coefficients with four sets needed for each field. In order to provide data for the filters each luminance line is read twice, and each chrominance line is read four times to produce each field.

NTSC VERTICAL FILTERING

One field of NTSC video consists of 240 chrominance and luminance lines, which must be converted to 288 lines of CIF luminance and 144 line of CIF chrominance. The luminance increase is mechanized by repeating the first line in every five to produce six lines, which are then applied to the vertical filters. A different set of coefficients is used for each line, requiring a total of 30 to be stored within the device. The line repeat causes one set of line data to be used twice, but each time different coefficients are used by the filter. This technique is equivalent to interpolating the data by six, and then decimating by five. The required coefficients for each of the six sets can be derived by conceptually using this approach.

The line repeat requires an additional FIFO line delay before the four delays used by the filters. By reducing the horizontal blanking time it is possible to read six lines (one is repeated) from the FIFO in the time taken to acquire five lines of video with blanking.

Chrominance data also passes through the input FIFO and one line in every five is repeated. This is done in order to avoid differential delays with the luminance data. Three chrominance lines are only needed, however, for every five original lines. They are produced by using three sets of five coefficients and discarding two filtered lines in every five. The three selected filter outputs are chosen such that the centre line of

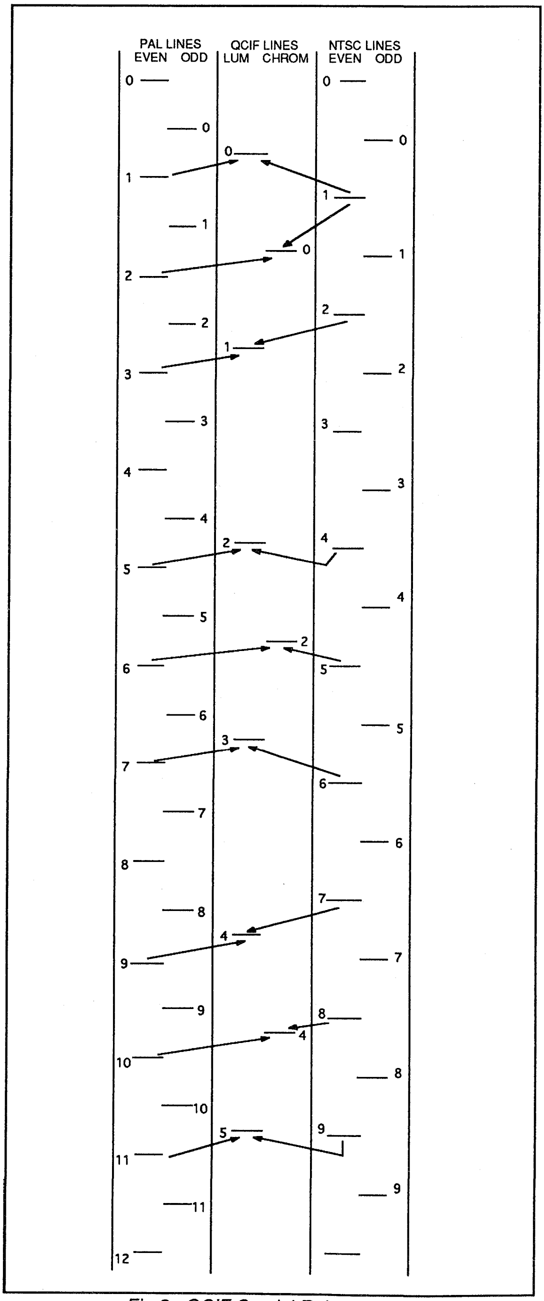

Fig 8 : QCIF Spatial Relationships

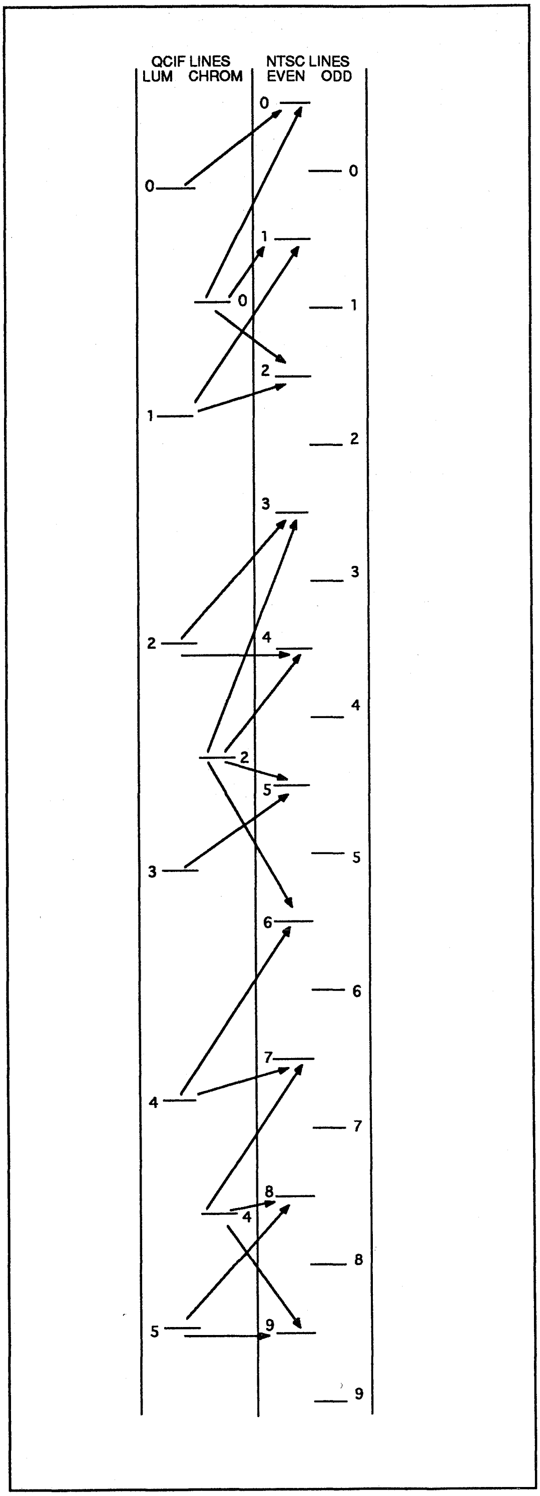

Fig 9 : Interpolating from QCIF to NTSC

the filter is closest to the CIF line number needed. The centre lines which are actually used are shown in Figure 6, and result in a sequence of two chosen outputs then a gap followed by one output then a gap. Simply using every other output would not give the best fit.

A simplified approach is used when decimating down to QCIF resolution, and the input FIFO is not used. Six luminance lines are derived from ten NTSC lines by choosing the six outputs produced when the centre line in the filter is closest to the QCIF line that is needed. Overall this results in a luminance sequence consisting of two outputs then a gap, followed by one output then a gap and is shown in Figure 7. Since the six lines are produced from two sub sequences of three lines, only three sets of coefficients are actually needed rather than six.

Three chrominance lines are derived from the same inputs by using three sets of seven coefficients. The chrominance sequence is also shown in Figure 7, and consists of an output then three gaps, followed by an output and two gaps.

When interpolating from CIF up to NTSC resolutions, it is necessary to read lines of data from the CIF frame store with reduced blanking periods. The timing is calculated such that six lines are read in the time that five lines would have been read if they had the correct blanking period. These fast lines are continuously filtered using all the available information, and the results are written to an output FIFO. This FIFO is then read with the correct blanking period inserted in order to provide NTSC data at the output pins. Thus five lines are read out in the time taken to load six lines (one of which need not actually be written since it is never used)

Five sets of coefficients are used to produce the five lines which are actually stored, but the coefficients are different for the even and odd field generation. Thus a total of ten sets of five coefficients are internally stored. In effect we have interpolated by five and then decimated by three in order to produce the complete NTSC frame.

Each CIF chrominance line is used to produce two filtered NTSC chrominance lines, and one filtered line in every six is then ignored. This is mechanized by reading each CIF chrominance line twice for every pair of luminance lines. The same filtering and discard technique as used in the luminance channel is then applied, using five sets of coefficients for each field. Ten sets are thus needed to produce two NTSC fields. We have effectively interpolated by ten and then decimated by three to produce 480 chrominance lines for the complete frame.

When interpolating from QCIF to NTSC the additional output buffering is not used. Instead a sequence is used which will generate 10 NTSC lines in any field from six QCIF luminance lines and three chrominance lines. Figure 8 illustrates how the first and fourth lines are used once and the second, third, fifth, and sixth used twice to produce QCIF luminance. Since this 1 - 2 - 2 sequence is used twice in every ten lines, only five rather than ten sets of coefficients are actually needed for each field (ten sets in total).

The first and third chrominance lines are used three times, and the second line is used four times. Thus ten sets of coefficients are needed for each field (twenty sets in total). Each luminance and chrominance set consists of seven coefficients, since six line delays are provided for the filters.

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage V_{DD}	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to V_{DD} + 0.5V
Output voltage V_{OUT}	-0.5V to V_{DD} + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	0°C to 70°C
Junction temperature	100°C
Package power dissipation	3000mW

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as negative into the device.

Test	Waveform - measurement level
Delay from output high to output high impedance	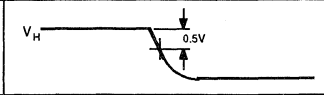
Delay from output low to output high impedance	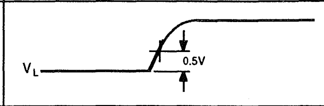
Delay from output high impedance to output low	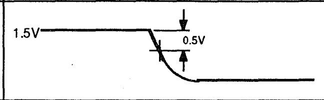
Delay from output high impedance to output high	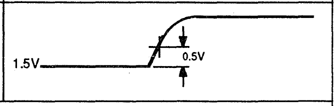

V_H - Voltage reached when output driven high
V_L - Voltage reached when output driven low

STATIC ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)
Tamb = 0 C to +70°C V_{DD} = 5.0v ± 5%

Characteristic	Symbol	Value Min.	Value Typ.	Value Max.	Units	Conditions
Output high voltage	V_{OH}	3.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.0		-	V	3.0V for SYSCLK and MCLK
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{DD}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{DD}
Output S/C current	I_{SC}	10		300	mA	

PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC
1	GND	25	M3	49	CREF	73	VDD	97	D14
2	A8	26	M2	50	GND	74	SCLK	98	D13
3	A7	27	M1	51	CSYNC	75	GND	99	D12
4	A6	28	M0	52	Y0	76	VDD	100	GND
5	A5	29	MCLK	53	Y1	77	HA0	101	VDD
6	VDD	30	VDD	54	Y2	78	HA1	102	D11
7	GND	31	GND	55	Y3	79	HA2	103	D10
8	A4	32	REQYUV	56	Y4	80	HA3	104	D9
9	A3	33	GND	57	Y5	81	WR	105	D8
10	A2	34	VDD	58	Y6	82	RD	106	GND
11	A1	35	FSIG	59	Y7	83	CEN	107	VDD
12	A0	36	GND	60	VDD	84	HD0	108	D0
13	VDD	37	VDD	61	GND	85	HD1	109	D1
14	GND	38	RST	62	HBLNK	86	HD2	110	D2
15	RW	39	TCK	63	C0	87	HD3	111	D3
16	VDD	40	TMS	64	C1	88	HD4	112	GND
17	GND	41	TRST	65	C2	89	HD5	113	VDD
18	RAS	42	TDI	66	C3	90	VDD	114	D4
19	VDD	43	TDO	67	C4	91	GND	115	D5
20	GND	44	TOE	68	C5	92	HD6	116	D6
21	M7	45	VDD	69	C6	93	HD7	117	D7
22	M6	46	VREF	70	C7	94	CLMP	118	CAS
23	M5	47	FREF	71	N/C	95	VDD	119	GND
24	M4	48	HREF	72	GND	96	D15	120	VDD

Pin out Table

Signal	Direction	JTAG Bit Number	Signal	Direction	JTAG Bit Number	Signal	Direction	JTAG Bit Number
A8	OUT	142	Y2	OUT	91	HD5	OUT	41
A7	OUT	141	Y2	IN	90	HD5	IN	40
A6	OUT	140	Y3	OUT	89	HD6	OUT	39
A5	OUT	139	Y3	IN	88	HD6	IN	38
A4	OUT	138	Y4	OUT	87	HD7	OUT	37
A3	OUT	137	Y4	IN	86	HD7	IN	36
A2	OUT	136	Y5	OUT	85	cdata_oeb*	OUT	35
A1	OUT	135	Y5	IN	84	CLMP	OUT	34
A0	OUT	134	Y6	OUT	83	D15	OUT	33
RW	OUT	133	Y6	IN	82	D15	IN	32
RAS	OUT	132	Y7	OUT	81	D14	OUT	31
DMHIZ*	OUT	131	Y7	IN	80	D14	IN	30
M7	OUT	130	HBLNK	OUT	79	D13	OUT	29
M7	IN	129	C0	OUT	78	D13	IN	28
M6	OUT	128	C0	IN	77	D12	OUT	27
M6	IN	127	C1	OUT	76	D12	IN	26
M5	OUT	126	C1	IN	75	D11	OUT	25
M5	IN	125	C2	OUT	74	D11	IN	24
M4	OUT	124	C2	IN	73	D10	OUT	23
M4	IN	123	C3	OUT	72	D10	IN	22
M3	OUT	122	C3	IN	71	D9	OUT	21
M3	IN	121	C4	OUT	70	D9	IN	20
M2	OUT	120	C4	IN	69	D8	OUT	19
M2	IN	119	C5	OUT	68	D8	IN	18
M1	OUT	118	C5	IN	67	D0	OUT	17
M1	IN	117	C6	OUT	66	D0	IN	16
M0	OUT	116	C6	IN	65	D1	OUT	15
M0	IN	115	C7	OUT	64	D1	IN	14
MCLK	OUT	114	C7	IN	63	D2	OUT	13
MCLK	IN	113	YCHIZ*	OUT	62	D2	IN	12
MHIZ*	OUT	112	CGTOUT (N/C)	OUT	61	D3	OUT	11
REQYUV	IN	111	CGHIZ*	OUT	60	D3	IN	10
FSIG	OUT	110	SCLK	IN	59	D4	OUT	9
FSIG	IN	109	HA0	IN	58	D4	IN	8
RST	IN	108	HA1	IN	57	D5	OUT	7
TOE	IN	107	HA2	IN	56	D5	IN	6
VREF	OUT	106	HA3	IN	55	D6	OUT	5
VREF	IN	105	WR	IN	54	D6	IN	4
FREF	OUT	104	RD	IN	53	D7	OUT	3
FREF	IN	103	CEN	IN	52	D7	IN	2
HREF	OUT	102	HD0	OUT	51	d-oeb*	OUT	1
HREF	IN	101	HD0	IN	50	CAS	OUT	0
CREF	OUT	100	HD1	OUT	49			
CREF	IN	99	HD1	IN	48			
RFHIZ*	OUT	98	HD2	OUT	47			
CSYNC	OUT	97	HD2	IN	46			
CNHIZ*	OUT	96	HD3	OUT	45			
Y0	OUT	95	HD3	IN	44			
Y0	IN	94	HD4	OUT	43			
Y1	OUT	93	HD4	IN	42			
Y1	IN	92						

JTAG Register Allocation

CGTOUT (N/C) This pin is only used for GPS test purposes and should not be used for system purposes.

VP8708

30MHz 8–BIT ANALOG VIDEO INPUT INTERFACE

The VP8708 is an analog input interface designed for video signal conditioning and digitisation.

Operating from a single +5V supply, the VP8708 includes an input multiplexer, video amplifier with clamp and gain control, an onboard reference and a buffered 8 bit ADC capable of digitising signals upto the Nyquist limit.

Video signals may be supplied to the VP8708 either directly to the ADC, or via one of three multiplexed inputs for signal conditioning. The analog signal (with appropriate gain and clamping) is available as output prior to digitisation. The coded data, updated at the sample rate, is available in either binary or two's complement format via the 3–state TTL output buffers.

FEATURES

- Direct replacement for the PhilipsTDA8708
- 30MHz Sampling rate
- Selectable data format:
- Internal ADC reference
- Clamp and AGC functions
- 3 state TTL outputs

ABSOLUTE MAXIMUM RATINGS

Supply voltages, AVCC, DVCC, OVC	7V
Supply differential	±1V
Ground differnetial	±1V
Video input voltage	AV_{CC}
Output current	10mA

THERMAL CHARACTERISTICS

Storage temperature range	−65°C to 150°C
Lead temperature (soldering 11 seconds)	+265°C

THERMAL RESISTANCES, DP PACKAGE

Junction to ambient (θ_{ja})	55°C/W
Junction to case (θ_{jc})	14°C/W

THERMAL RESISTANCES, MP PACKAGE

Junction to ambient (θ_{ja})	84°C/W
Junction to case (θ_{jc})	32°C/W

RECOMMENDED OPERATING CONDITIONS

Supply voltage AVCC, DVCC, OVCC	+5V
Supply differential	0V
Ground differential	0V
Video input voltage	1V

OPERATING TEMPERATURE RANGE

Commercial 0°C to+ 70°C(still air ambient)

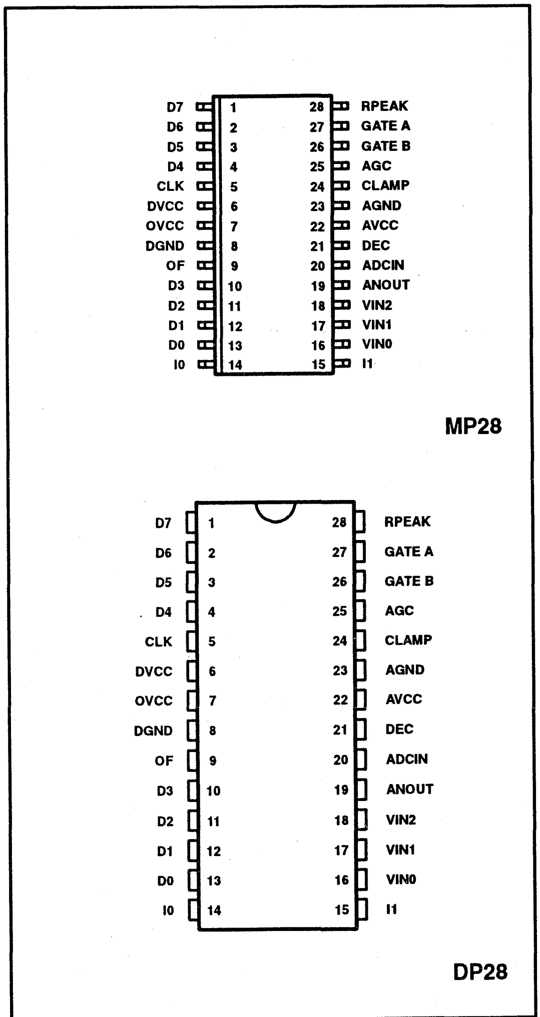

Fig. 1 Pin connections top view

ORDERING INFORMATION

VP8708G CG DPAS (Commercial – Plastic DIL)
VP8708G CG MPES (Commercial – Miniature Plastic DIL)

PIN	NAME	DESCRIPTION
1	D7	Data output Bit 7 (MSB)
2	D6	Data output Bit 6
3	D5	Data output Bit 5
4	D4	Data output Bit 4
5	CLK	Clock input
6	DVCC	Digital supply voltage
7	OVCC	Output buffer supply voltage
8	DGND	Digital gound
9	OF	Output format/Chip enable
10	D3	Data output Bit 3
11	D2	Data output Bit 2
12	D1	Data output Bit 1
13	D0	Data output Bit 0 (LSB)
14	I0	Input selection Bit 0
15	I1	Input selection bit 1
16	VIN0	Video input 0
17	VIN1	Video input 1
18	VIN2	Video input 2
19	ANOUT	Analog output
20	ADCIN	ADC input
21	DEC	Not connected VP8708 is internally stable
22	AVCC	Analog supply voltage
23	AGND	Analog ground
24	CLAMP	Clamp capacitor
25	AGC	AGC capacitor
26	GATE B	Black level control pulse
27	GATE A	Sync level control pulse
28	RPEAK	Peak current resistor

Table 1 Pin descriptions

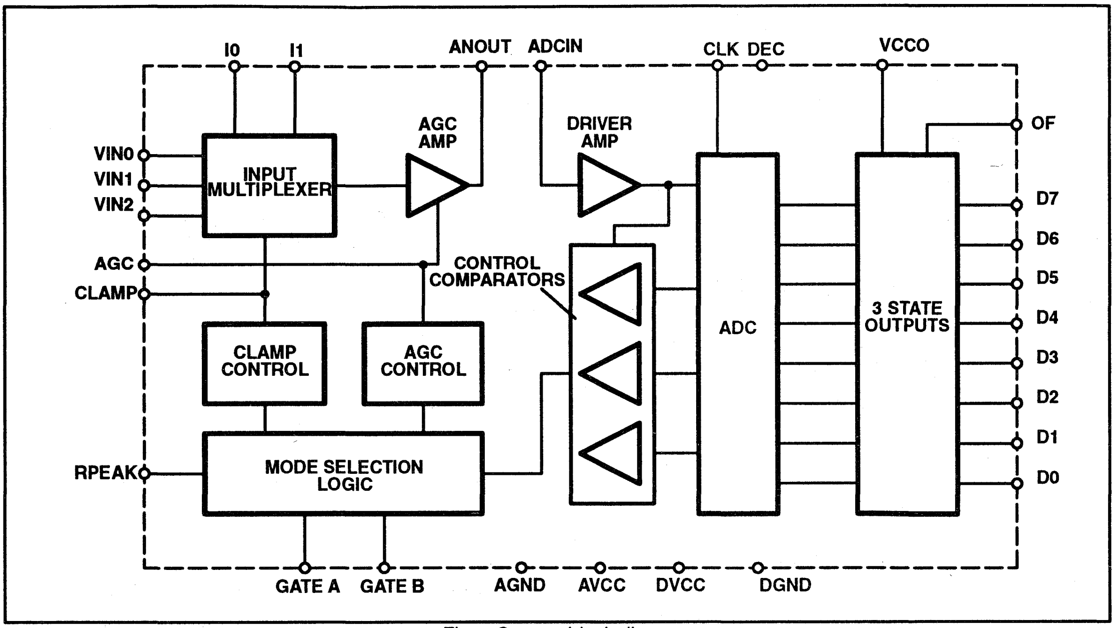

Fig. 2 System block diagram

GENERAL CIRCUIT DESCRIPTION

The VP8708 is an analog video input interface capable of digitising signals at sample rates upto 30MHz.

The multiplexer uses the logic conditions on the selection pins (IO, I1) to select one of upto three signals applied to the video inputs (VINO, VIN1, VIN2). This signal is then clamped to the required dc level by the action of the clamp control logic. The output of the multiplexer passes through the AGC amplifier, where the gain of the signal is adjusted such that after going through the driver amplifier, the signal fills the desired portion of the ADC range. This input to the ADC also drives three reference comparators which supply the signals necessary to control the clamp and AGC circuitry.

Two modes of operation are available; these being determined by the relative occurrences, during the sync and rear porch periods, of logic pulses at the GATE A and GATE B inputs.

Mode 1 (see Fig. 3) is employed initially to allow the device to reach its optimum operation point. The gain and dc level of the signal are roughly adjusted to set the sync level to ADC code 0 and the peak level to ADC code 255. This allows rapid recovery of the video synchronisation pulses.

If the GATE A and GATE B pulses become distinct (see Fig. 4) then the VP8708 will switch into mode 2. In this configuration a more sophisticated control scheme is used in order to produce a 'fixed'digitised output from the device: Whilst the GATE A pulse (which must be within the sync period) is high, the sync level is adjusted to code 0,.Similarly, the black level is adjusted to code 64 if the GATE B pulse occurs during the rear porch periods. Peak level control is always active such that maximum digital output will tend to lie below code 240. Nominal input signals, 1Vp–p, should have a peak output level equal to code 213.

For the device to operate, two external capacitors must be connected to the AGC and CLAMP pins. An optional external resistor may be attached to the RPEAK input to alter the maximum charge/discharge current into these capacitors. This varies the loop response time.

The format of the output data, updated at the sample rate, is determined by the logic level on the OF pin. The OF pin may also be used to force the outputs to a high impedance state.

The DEC ('decouple') pin is not connected on the VP8708. On similar devices (e.g. TDA8708) an external capacitor may be attached to this pin to stabilise the ADC reference voltage. The VP8708 has an internally stable reference and thus requires one less external component.

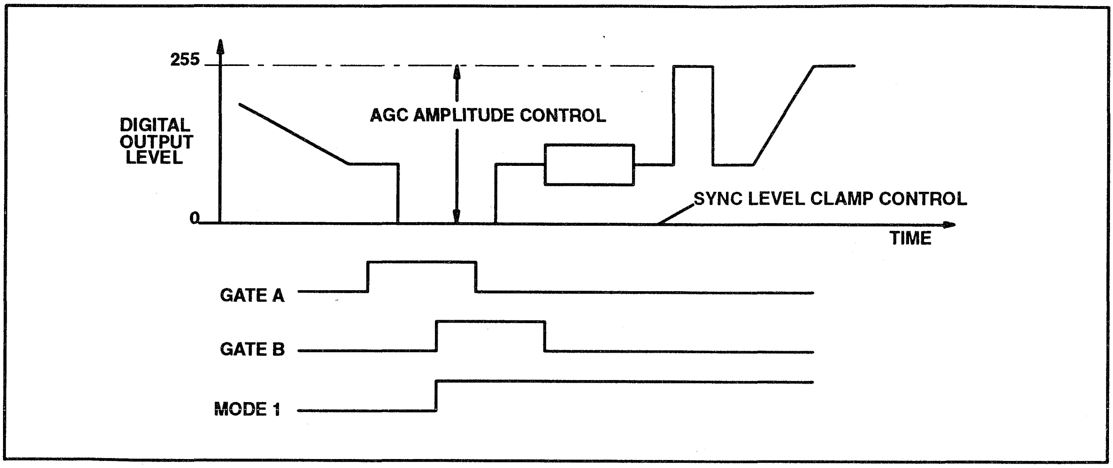

Fig. 3 Control mode 1

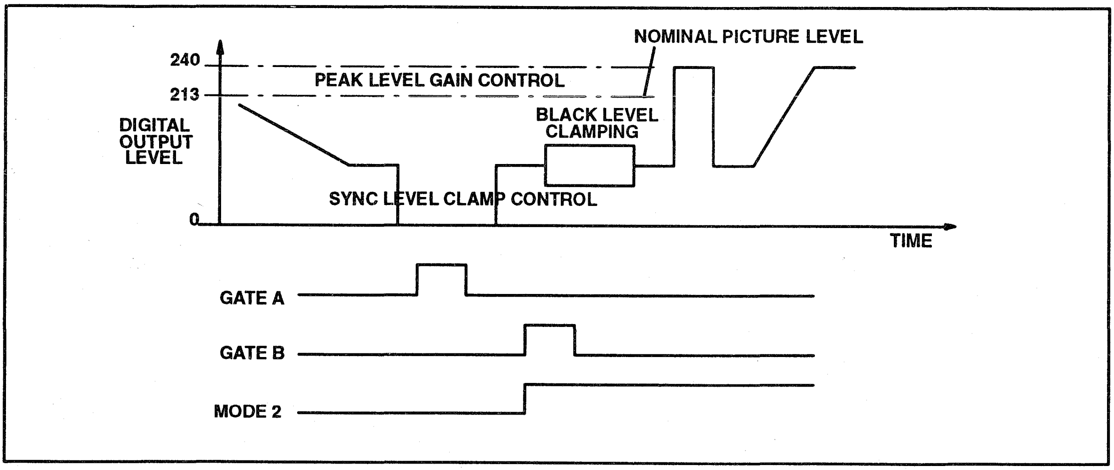

Fig. 4 Control mode 2

ELECTRICAL CHARACTERISTICS: DC CHARACTERISTICS

Test conditions (unless otherwise stated): AV_{CC}, DV_{CC}, OV_{CC}=5V±0.5V, AGND/DGND shorted together, Temp=$T_{full.}$ =0 to + 70°C.

Parameter	Symbol	Temp (°C)	Test level	Value			Units	Conditions
				Min	Typ	Max		
POWER SUPPLY								
Analog supply voltage	AV_{CC}	25 FULL	1 4	4.5 4.5	5.0 5.0	5.5 5.5	V V	
Digital supply voltage	DV_{CC}	25 FULL	1 4	4.5 4.5	5.0 5.0	5.5 5.5	V V	
Output supply voltage	OV_{CC}	25 FULL	1 4	4.5 4.5	5.0 5.0	5.5 5.5	V V	
Analog supply current	AI_{CC}	25 FULL	1 4		60	 72	mA mA	
Digital supply current	DI_{CC}	25 FULL	1 4		13	 19	mA mA	

ELECTRICAL CHARACTERISTICS: DC CHARACTERISTICS (cont.)

Test conditions (unless otherwise stated): AV_{CC}, DV_{CC}, OV_{CC}=5V±0.5V, AGND/DGND shorted together, Temp=$T_{full.}$ =0 to + 70°C.

Parameter	Symbol	Temp (°C)	Test level	Value Min	Value Typ	Value Max	Units	Conditions
Output supply current	OI_{CC}	25 FULL	1 4		12	 18	mA mA	
Power dissipation	P	25 FULL	1 4		425	 600	mW mW	
VIDEO INPUTS								
Input range	V_{IN} (p–p)	FULL	4	0.5	1.0	1.6	V	
Input impedance	$\lvert Z_{IN}\rvert$	FULL	4	18	20	22	kΩ	f_{in}=6MHZ
Input capacitance	C_{in}	25	5		2		pF	f_{in}=6MHZ
I1, IO GATE B,TTL INPUTS								
Input Voltage LOW	V_{il}	25 FULL	1 4			0.8 0.8	V V	
Input current HIGH	V_{ih}	25 FULL	1 4	2.0 2.0			V V	
Input current LOW	I_{il}	25 FULL	1 4	–150 –200			µA µA	V_i=0.4V V_i=0.4V
Input current HIGH	I_{ih}	25 FULL	1 4			10 20	µA µA	V_i=3.6V V_i=3.6V
RPEAK Input								
Peak capacitor charge/discharge current	I_{PEAK}	FULL	4		80		µA	R_{peak} = 0Ω
AGC Control Input								
AGC voltage for minimum gain	V_{agc}	25	4		2.7		V	
AGC voltage for maximum gain	V_{agc}	25	4		3.9		V	
AGC output current	I_{agc}	25	1		*			* See table 4
CLAMP Control Input								
Clamp voltage	V_{clp}	25	4		3.5		V	ADC output = 128
Clamp output current	I_{clp}	25	1		*			* See table 4
VIDEO Amplifier Outputs								
Internal current source	I_{19}	25	4			2	mA	
DC Output voltage for black level	V_{blk}	FULL	4		AV_{CC} 3.25		V	
AC Output voltage (peak–peak)	V_{19}	FULL	4		1		V	
Output impedance	Z_{19}	25	4		45		Ω	
CLK INPUT								
Input voltage LOW	V_{il}	25 FULL	1 4			0.8 0.8	V V	
Input voltage HIGH	V_{ih}	25 FULL	1 4	2.0 2.0			V V	

ELECTRICAL CHARACTERISTICS: DC CHARACTERISTICS (cont.)

Test conditions (unless otherwise stated): AV_{CC}, DV_{CC}, OV_{CC}=5V±0.5V, AGND/DGND shorted together, Temp=$T_{full.}$ =0 to + 70°C.

Parameter	Symbol	Temp (°C)	Test level	Value			Units	Conditions
				Min	Typ	Max		
Input current LOW	I_{il}	25 FULL	1 4	–150 –200			µA µA	V_i = 0.4V V_i = 0.4V
Input current HIGH	I_{ih}	25 FULL	1 4			10 20	µA µA	V_i= 3.6V V_i= 3.6V
Input impedance	$\lvert Z_{IN}\rvert$	25	4		3.5		kΩ	
Input capacitance	C_{clk}	25			5		pF	
Maximum frequency	f_{max}	FULL	4	30			MHz	
OF INPUT (3–state control)								
Input voltage LOW	V_{il}	25 FULL	1 4			0.8 0.8	V V	
Input voltage HIGH	V_{ih}	25 FULL	1 4	2.0 2.0			V V	
Input voltage 3–STATE	V_z	25	1		1.4		V	
Input current LOW	I_{il}	25 FULL	1 4	–175 –200			µA µA	V_i=0.4V V_i=0.4V
Input current HIGH	I_{ih}	25 FULL	1 4			500 700	µA µA	V_i=3.6V V_i=3.6V
ADCIN INPUT								
Input voltage	V_{adc}	FULL	4		AV_{CC}–1.75		V	For Code 0
Input voltage	V_{adc}	FULL	4		AV_{CC}–1.25		V	For Code 255
Input voltage amplitude (p–p)	V_{20}	FULL	4		0.5		V	
Input current	I_{adc}	25	1		1		µA	
Input impedance	Z_{adc}	25	1		14		MΩ	
Input capacitance	C_{adc}	25	4		5		pF	
DIGITAL OUTPUTS								
Output voltage LOW	V_{ol}	25 FULL	4 4			0.4 0.4	V V	I_{ol}= 2mA I_{ol}= 2mA
Output voltage HIGH	V_{oh}	25 FULL	1 4	2.4 2.4			V V	I_{oh}=–0.4mA I_{oh}=–0.4mA
3–STATE Output current	I_{OZ}	25	4		2		µA	
ADC PERFORMANCE								
Static differential non – linearity	DNL	25 FULL	1 4		± 0.5 ± 0.5		lsb lsb	
Static integral non – linearity	INL	25 FULL	1 4		± 1 ± 1		lsb lsb	
Dynamic integral non – linearity	INL	25 FULL	1 4		± 2 ± 2		lsb lsb	

VIDEO AMPLIFIER DYNAMIC PERFORMANCE								
–3dB Bandwidth	f3dB	25	4		20		MHz	
Differential gain	G_d	25	4		2		%	
Differential phase	$\oslash_d$	25	4		2		degrees	
Gain range	ΔG	25	4	–3		7	dB	
Crosstalk between VIN inputs		25	4		–60		dB	
Signal–to–noise ratio	SNR	25	4		55		dB	
ANALOG SIGNAL PROCESSING								f_{clk} = 30MHz
–3dB Bandwidth	f3dB	25	1		15		MHz	
Differential gain	G_d	25	4		2		%	
Differential phase	$\oslash_d$	25	4		2		degrees	
Total harmonic distortion	THD	25	4		–55		dB	
Supply voltage ripple rejection	SVRR	25	4		5		%/V	
TIMING*								* f_{clk}=30MHz C_1=15pF I_{ol}=2mA
Sampling delay	t_{ds}	FULL	4		3		ns	
Output hold time	t_{ho}	FULL	4	5			ns	
Output delay time	t_d	FULL	4			20	ns	
3 State delay time for enble	t_{ez}	FULL	4			25	ns	
3 State delay time for disable	t_{dz}	FULL	4			25	ns	

ELECTRICAL CHARACTERISTIC DEFINITIONS

Analog –3dB Bandwidth ADC

The analog input frequency at which the spectral power of the fundamental frequency, as determined by Fast Fourier Transform analysis, is 3dB down on the DC level.

Differential Non–Linearity (DNL)

The deviation of any code width from an ideal 1LSB step size.

Integral Non–Linearity (INL)

The deviation of the centre of each code from a reference line which has been determined by a least–square curve–fit.

Differential Gain (G_d) and Phase ($\oslash_d$)

The difference in gain/phase at the ADC output when a 17.5 IRE units peak to peak, 3.58MHz input signal is superimposed on a dc level at 1/16 and 15/16 full scale input.

Supply Voltage Rejection Ratio (SVRR)

The variation in the amplitude of the given signal when the supply voltage is changed by 1V.

Signal–to–Noise Ratio

The ratio of the RMS signal amplitude to the RMS value of "noise" which is defined as the sum of all other spectral components including harmonics, but excluding dc with a full–scale analog input signal.

Total Harmonic Distortion (THD)

The RMS addition of all peaks in a Fast Fourier Transform measurement, which occur at integer multiples of the fundamental frequency of the input signal.

Test Levels

Level 1	100% production tested
Level 2	100% production tested at 25°C and sample tested at specified temperatures
Level 3	Sample tested only
Level 4	Parameter is guaranteed by design and characteristics testing
Level 5	Parameter is a typical value only

VIDEO INPUTS

Each of the three video inputs may be selected with the selection pins I0 and I1. Table 2 shows the action of these pins.

I1	I0	SELECTED INPUTS
0	0	VIN0
0	1	VIN1
1	0	VIN2
1	1	VIN2

Table 2 Video input selection

OUTPUT FORMAT (OF PIN)

The format of the output data is selectable via the OF pin as shown in Table 3. To improve noise immunity, a small (10nF) capacitor may be connected between this pin and ground if binary operation is required.

OF	D0 TO D7
0	Active, two's complement
Open	Active, binary
1	High impedance

Table 3 Output format control

ANOUT AND ADCIN

The analog output (ANOUT) and ADC input (ADCIN) pins should have an external anti–aliasing filter connected betwen them. Care must be taken to ensure that the filter input impedance and filter output levels are in accordance with the specifications.

As an evaluatuion tool, two 1.5kΩ resistors may be used to form a potential divider between ANOUT and AV_{CC}. The centre tap of this divider can then be used to connect the signal to ADCIN. It should be noted however, that dynamic device performance will not be maximum.

TIMING INFORMATION

Fig. 4 depicts the system relationship between sampling edge offset and output data.

The analog input signal is sampled t_{ds} seconds after the rising edge of the clock signal. Old data will remain valid for at least t_{ho} and new data will become valid after at most t_d. Data may be latched on either the falling or rising edges of the clock signal.

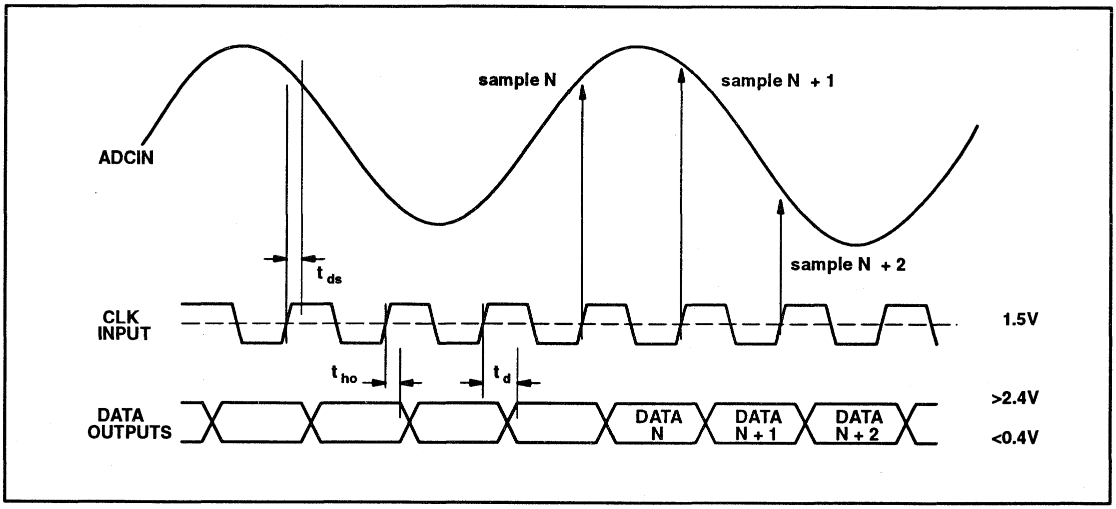

Fig. 4 System timing

AGC AND CLAMP CONTOL

In both mode 1 and 2, AGC and clamp control is acheived by the charging and discharging of two external capacitors attached to the AGC and CLAMP pins. Suggested values for these capacitors are C_{agc}=220nF and C_{clp}=18nF.

For correct device operation, the relative occurrences of the control pulses at the GATE A and GATE B pins must be as given below:

MODE 1 – GATE A and GATE B must overlap, GATE B being delayed with respect to GATE A. This will ensure that the signal fills the complete ADC range and thus allows quick recovery of the video sync pulses.

MODE 2 – GATE A must occur wholly within the sync period GATE B must occur wholly within the rear porch. This will tend to hold the output signal amplitude below code 240, the black level at code 64, and the sync tip at code 0.

Table 4 shows the control action of the device with reference to the logic states of the GATE A and GATE B inputs.

GATE A	GATE B	MODE	DIGITAL OUTPUT CODE	I_{AGC}	I_{CLP}
1	↑	Device will enter mode 1			
		1	Output >255	$-I_{peak}$	+5μA
		1	Output <255	+5μA	–
		1	Output >0	–	+5μA
		1	Output <0	+5μA	$-I_{peak}$
↑ ↓ 0	0 0 ↑	This control sequence will switch the device into mode 2.			
0	1	2	Output >240	$-I_{peak}$	+50μA
0	1	2	Output >64	–	+50μA
0	1	2	Output <64	–	–50μA
0	0	2	Output >240	$-I_{peak}$	–
1	0	2	Output >240	$-I_{peak}$	–
1	0	2	Output >0	–5μA	–
1	0	2	Output <0	+5μA	–

Table 4: Mode 1/Mode 2 control

PCB CONSTRUCTION

As with all high speed analog to digital converters, careful consideration must be given to the PCB layout.

In general, the best results will be obtained by tying all grounds to a 'solid' low impedance ground plane. Separate analog and digital ground planes will also help. Device connections to the ground plane should be as short as possible.

Supply decoupling is important when dealing with mixed analog and digital signals, it can provide a feedback path for the digital output currents. The VP8708 should therefore be decoupled as close to the supply pins as possible. Good quality, high frequency, low inductance capacitors. Isolation may be further improved by adding series inductors to the supplies.

Jitter and noise on the clock pin and its reference to ground must be minimised. Long clock lines should be avoided and all lines correctly terminated.

A typical application circuit is shown below.

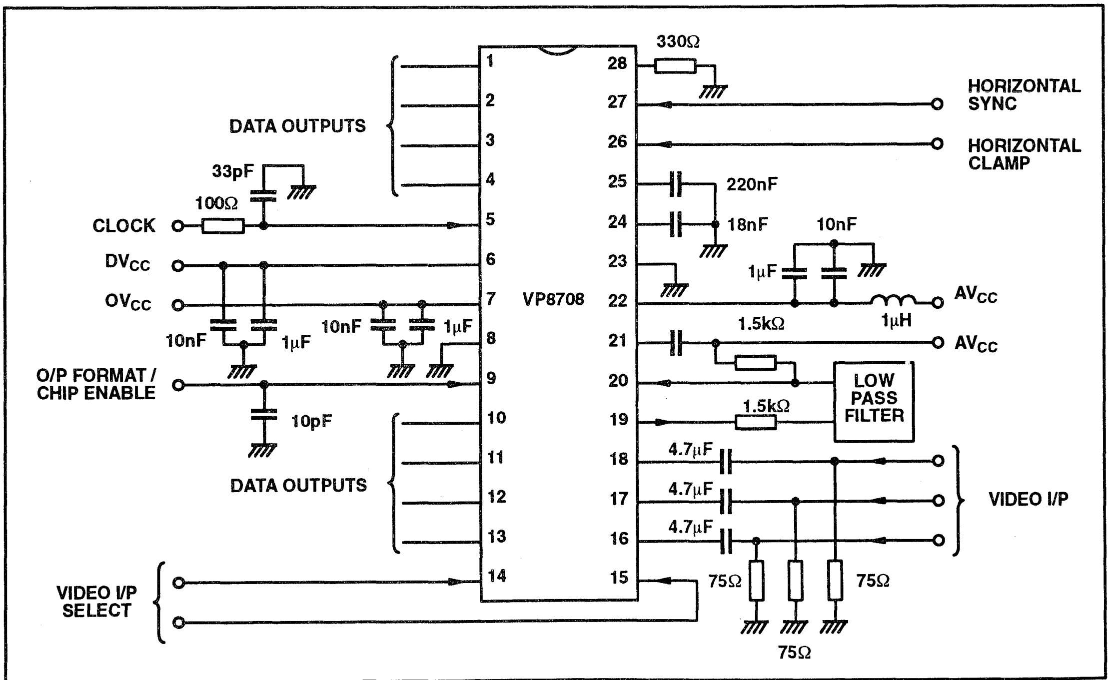

Fig. 5 Typical application circuit

Digital Signal Processing

GEC PLESSEY
SEMICONDUCTORS

PDSP1601/PDSP1601A

ALU AND BARREL SHIFTER

The PDSP1601 is a high performance 16-bit arithmetic logic unit with an independent on-chip 16-bit barrel shifter. The PDSP1601A has two operating modes giving 20MHz or 10MHz register-to-register transfer rates.

The PDSP1601 supports Multicycle multiprecision operation. This allows a single device to operate at 20MHz for 16-bit fields, 10MHz for 32-bit fields and 5MHz for 64-bit fields. The PDSP1601 can also be cascaded to produce wider words at the 20MHz rate using the Carry Out and Carry In pins. The Barrel Shifter is also capable of extension, for example the PDSP1601 can used to select a 16-bit field from a 32-bit input in 100ns.

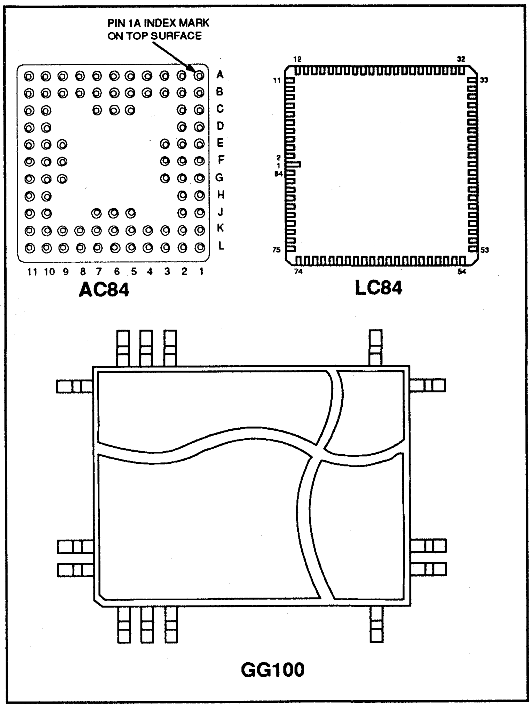

Fig. 1 Pin connections - bottom view

FEATURES

- 16-bit, 32 instruction 20MHz ALU
- 16-bit, 20MHz Logical, Arithmetic or Barrel Shifter
- Independent ALU and Shifter Operation
- 4 x 16-bit On Chip Scratchpad Registers
- Multiprecision Operation; e.g. 200ns 64-bit Accumulate
- Three Port Structure with Three Internal Feedback Paths Eliminates I/O Bottlenecks
- Block Floating Point Support
- 300mW Maximum Power Dissipation
- 84-pin Pin Grid Array or 84 Contact LCC Packages or 100 pin Ceramic Quad Flat Pack

APPLICATIONS

- Digital Signal Processing
- Array Processing
- Graphics
- Database Addressing
- High Speed Arithmetic Processors

ASSOCIATED PRODUCTS

PDSP16112 Complex Multiplier
PDSP16116 16 x 16 Complex Multiplier
PDSP16318 Complex Accumulator
PDSP16330 Pythagoras Processor

PIN DESCRIPTION

LC pin	AC pin	Function	LC pin	AC pin	Function	LC pin	AC pin	Function	LC pin	AC pin	Function
1	C6	IA4	22	F3	GND	43	J6	IS0	64	F9	GND
2	A6	MSB	23	G3	MSA0	44	J7	IS1	65	F11	C8
3	A5	MSS	24	G1	MSA1	45	L7	IS2	66	E11	C9
4	B5	B15	25	G2	A15	46	K7	IS3	67	E10	C10
5	C5	B14	26	F1	A14	47	L6	SV0	68	E9	C11
6	A4	B13	27	H1	A13	48	L8	SV1	69	D11	C12
7	B4	B12	28	H2	A12	49	K8	SV2	70	D10	C13
8	A3	B11	29	J1	A11	50	L9	SV3	71	C11	C14
9	A2	B10	30	K1	A10	51	L10	$\overline{SVOE}$	72	B11	C15
10	B3	B9	31	J2	A9	52	K9	RS0	73	C10	$\overline{OE}$
11	A1	B8	32	L1	A8	53	L11	RS1	74	A11	BFP
12	B2	B7	33	K2	A7	54	K10	VCC	75	B10	VCC
13	C2	B6	34	K3	A6	55	J10	RS2	76	B9	CO
14	B1	B5	35	L2	A5	56	K11	C0	77	A10	RA0
15	C1	B4	36	L3	A4	57	J11	C1	78	A9	RA1
16	D2	B3	37	K4	A3	58	H10	C2	79	B8	RA2
17	D1	B2	38	L4	A2	59	H11	C3	80	A8	CI
18	E3	B1	39	J5	A1	60	F10	C4	81	B6	IA0
19	E2	B0	40	K5	A0	61	G10	C5	82	B7	IA1
20	E1	$\overline{CEB}$	41	L5	$\overline{CEA}$	62	G11	C6	83	A7	IA2
21	F1	CLK	42	K6	MSC	63	G9	C7	84	C7	IA3

GG	SIG	GG	SIG	GG	SIG	GG	SIG
1	N/C	26	N/C	51	N/C	76	N/C
2	N/C	27	N/C	52	N/C	77	N/C
3	N/C	28	N/C	53	N/C	78	N/C
4	N/C	29	N/C	54	N/C	79	N/C
5	VCC	30	B7	55	A7	80	VCC
6	C0	31	B6	56	A6	81	RS2
7	RA0	32	B5	57	A5	82	C0
8	RA1	33	B4	58	A4	83	C1
9	RA2	34	B3	59	A3	84	C2
10	CI	35	B2	60	A2	85	C3
11	IA0	36	B1	61	A1	86	C4
12	IA1	37	B0	62	A0	87	C5
13	IA2	38	CEB	63	CEA	88	C6
14	IA3	39	CLK	64	MSC	89	C7
15	IA4	40	GND	65	IS0	90	GND
16	MSB	41	MSA0	66	IS1	91	C8
17	MSS	42	MSA1	67	IS2	92	C9
18	B15	43	A15	68	IS3	93	C10
19	B14	44	A14	69	SV0	94	C11
20	B13	45	A13	70	SV1	95	C12
21	B12	46	A12	71	SV2	96	C13
22	B11	47	A11	72	SV3	97	C14
23	B10	48	A10	73	SVOE	98	C15
24	B9	49	A9	74	RS0	99	OE
25	B8	50	A8	75	RS1	100	BFP

N/C = not connected - leave open circuit

All GND and VDD pin must be used

PIN DESCRIPTIONS

Symbol	Pin No. (LC84 Package)	Description
MSB	2	**ALU B-input multiplexer select control.**[1] This input is latched internally on the rising edge of CLK.
MSS	3	**Shifter Input multiplexer select control.**[1] This input is latched internally on the rising edge of CLK.
B15 - B0	4 - 19	**B Port data input.** Data presented to this port is latched into the input register on the rising edge of CLK. B15 is the MSB.
$\overline{\text{CEB}}$	20	**Clock enable, B Port input register.** When low the clock to this register is enabled.
CLK	21	**Common clock to all internal registered elements.** All registers are loaded, and outputs change on the rising edge of CLK.
MSA0 - MSA1	23 - 24	**ALU A-input multiplexer select control.**[1] These inputs are latched internally on the rising edge of CLK.
A15 - A0	25 - 40	**A Port data input.** Data presented to this port is latched into the input register on the rising edge of CLK. A15 is the MSB.
$\overline{\text{CEA}}$	41	**Clock enable, A Port input register.** When low the clock to this register is enabled.
MSC	42	**C-Port multiplexer select control.**[1] This input is latched internally on the rising edge of CLK.
IS0 - IS3	43 - 46	**Instruction inputs to Barrel Shifter, IS3 = MSB.**[1] These inputs are latched internally on the rising edge of CLK.
SV0 - SV3	47 - 50	**Shift Value I/O Port.** This port is used as an input when shift values are supplied from external sources, and as an output when Normalise operations are invoked. The I/O functions are determined by the IS0 - IS3 instruction inputs, and by the $\overline{\text{SVOE}}$ control. The shift value is latched internally on the rising edge of CLK.
$\overline{\text{SVOE}}$	51	**SV Output enable.** When high the SV port can only operate as an input. When low the SV port can act as an input or as an output, according to the IS0 - IS3 instruction. This pin should be tied hihg or low, depending upon the application.
RS0, RS1 RS2	52 - 53 55	**Instruction inputs to Barrel Shifter registers.**[1] These inputs are latched internally on the rising edge of CLK.
C0 - C15	56 - 63 65 - 72	**C Port data output.** Data output on this port is selected by the C output multiplexer. C15 is the MSB.
$\overline{\text{OE}}$	73	**Output enable.** The C Port outputs are in high impedance condition when this control is high.
BFP	74	**Block Floating Point Flag** from ALU, active high.
CO	76	**Carry out** from MSB of ALU.
RA0 - RA2	77 - 79	**Instruction inputs to ALU registers.**[1] These inputs are latched internally on the rising edge of CLK.
CI	80	**Carry in** to LSB of ALU.
IA0 - IA3 IA4	81 - 84 1	**Instruction inputs to ALU.**[1] IA4 = MSB. These inputs are latched internally on the rising edge of CLK.
Vcc	54 & 75	**+5V supply:** Both Vcc pins must be connected.
GND	22 & 64	**0V supply:** Both GND pins must be connected.

NOTES

1. All instructions are executed in the cycle commencing with the rising edge of the CLK which latches the inputs.

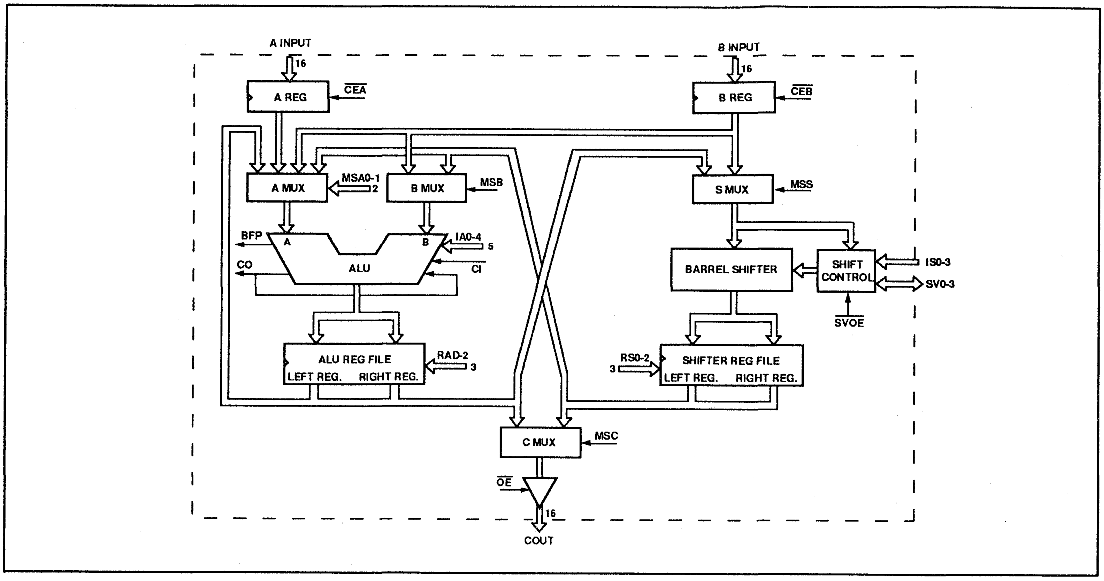

Fig.2 PDSP1601 block diagram

FUNCTIONAL DESCRIPTION

The PDSP1601 contains four main blocks: the ALU, the Barrel Shifter and the two Register Files.

The ALU

The ALU supports 32 instructions as detailed in Table 1.

The inputs to the ALU are selected by the A and B MUXs. Data will fall through from the selected register through the A or B input MUXs and the ALU to the ALU output register file in 50ns for the PDSP1601A (100ns for the PDSP1601).

The ALU instructions are latched, such that the instruction will not start executing until the rising edge of CLK latches the instruction into the device.

The ALU accepts a carry in from the CI input and supplies a carry out to the CO output. Additionally, at the end of each cycle, the carry out from the ALU is loaded into an internal 1 bit register, so that it is available as an input to the ALU on the next cycle. In the manner, multicycle, multiprecision operations are supported. (See MULTICYCLE CASCADE OPERATIONS).

BFP Flag

The ALU has a user programmable BFP flag. This flag may be programmed to become active at any one of four conditions. Two of these conditions are intended to support Block Floating Point operations, in that they provide flags indicating that the ALU result is within a factor of two or four of overflowing the 16 bit number range. For multiprecision operations the flag is only valid whilst the most significant 16 bit byte is being processed. In this manner the BFP flag may be used over any extended word width.

The remaining two conditions detect either an overflow condition or a zero result. For the overflow condition to be active the ALU result must have overflowed into the 16th (sign) bit, (this flag is only valid whilst the most significant 16 bit byte is being processed). The zero condition is active if the result from the ALU is equal to zero. For multiprecision operations the zero flag must be active for all of the 16 bit bytes of an extended word.

The BFP flag is programmed by executing on of the four SBFXX instructions (see Table 1). During the execution of any of these four instructions, the output of the ALU is forced to zero.

Multicycle/Cascade Operation

The ALU arithmetic instructions contain two or three options for each arithemtic operation.

The ALU is designed to operate with two's complement arithmetic, requiring a one to be added to the LSB for all subtract operations. The instructions set includes instructions that will force a one into the LSB, e.g. MIAX1, AMBX1, BMAX1 (see Table 1).

These instructions are used for the least significant 16 bit byte of any subtract operation.

The user has an option of cascading multiple devices, or multicycling a single device to extend the arithmetic precision. Should the user cascade multiple devices, then the cascade arithmetic instructions using the external CI input should be employed for all but the least significant 16 bit byte, e.g. MIACI, APBCI, BMACI (see Table 1).

Should the user multicycle a single device, then the Multicycle Arithmetic instructions, using the internally registered CO bit should be employed for all but the least significant 16 bit byte, e.g. MIACO, APBCO, AMBCO, BMACO (see Table 1).

Table 1 ALU instructions

1a. ARITHMETIC INSTRUCTIONS

Inst	IA4-AI0	Mnemonic	Operation	Function	Mode
00	00000	CLRXX	RESET	CLEAR ALL REGISTERS	---------
01	00001	MIAX1	MINUS A	NA Plus 1	LSBYTE
02	00010	MIACI	MINUS A	NA Plus CI	CASCADE
03	00011	MIACO	MINUS A	NA Plus CO	MULTICYCLE
04	00100	A2SGN	A/2	A/2 Sign Extend	MSBYTE
05	00101	A2RAL	A/2	A/2 with RAL LSB	MULTICYCLE
06	00110	A2RAR	A/2	A/2 with RAR LSB	MULTICYCLE
07	00111	A2RSX	A/2	A/2 with RSX LSB	MULTICYCLE
08	01000	APBCI	A PLUS B	A Plus B Plus CI	CASCADE
09	01001	APBCO	A PLUS B	A Plus B Plus CO	MULTICYCLE
0A	01010	AMBX1	A MINUS B	A Plus NB Plus 1	LSBYTE
0B	01011	AMBCI	A MINUS B	A Plus NB Plus CI	CASCADE
0C	01100	AMBCO	A MINUS B	A Plus NB Plus CO	MULTICYCLE
0D	01101	BMAX1	B MINUS A	NA Plus B Plus 1	LSBYTE
0E	01110	BMACI	B MINUS A	NA Plus B Plus CI	CASCADE
0F	01111	BMACO	B MINUS A	NA Plus B Plus CO	MULTICYCLE

1b. LOGICAL INSTRUCTIONS

Inst	IA4-AI0	Mnemonic	Operation	Function
10	10000	ANXAB	A AND B	A. B
11	10001	ANANB	A AND NB	A. NB
12	10010	ANNAB	NA AND B	NA. B
13	10011	ORXAB	A OR B	A + B
14	10100	ORNAB	NA OR B	NA + B
15	10101	XORAB	A XOR B	A XOR B
16	10110	PASXA	PASS A	A
17	10111	PASNA	INVERT A	NA

1c. CONTROL INSTRUCTIONS

Inst	IA4-AI0	Mnemonic	Operation
18	11000	SBFOV	Set BFP Flag to OVR, Force ALU output to zero
19	11001	SBFU1	Set BFP Flag to UND 1 Force ALU output to zero
1A	11010	SBFU2	Set BFP Flag to UND 2 Force ALU output to zero
1B	11011	SBFZE	Set BFP Flag to ZERO Force ALU output to zero
1C	11100	OPONE	Output 0001 Hex
1D	11101	OPBYT	Output 00FF Hex
1E	11110	OPNIB	Output 000F Hex
1F	11111	OPALT	Output 5555 Hex

KEY

A	= A input to ALU
B	= B input to ALU
CI	= External Carry in to ALU
CO	= Internally Registered Carry out from ALU
RAL	= ALU Register (Left)
RAR	= ALU Register (Right)
RSX	= Shifter Register (Left or Right)

MNEMONICS

CLRXX	Clear All Registers to zero		
MIAXX	Minus A,	XX	= Carry in to LSB
A2XXX	A Divided by 2,	XXX	= Source of MSB
APBXX	A Plus B,	XX	= Carry in to LSB
AMBXX	A Minus B,	XX	= Carry in to LSB
BMAXX	B Minus A,	XX	= Carry in to LSB
ANX-Y	AND	X	= Operand 1, Y = Operand 2
ORX-Y	OR	X	= Operand 1, Y = Operand 2
XORXY	Exclusive OR	X	= Operand 1, Y = Operand 2
PASXX	Pass	XX	= Operand
SBFXX	Set BFP Flag	XX	= Function
OPXXX	Output Constant	XXX	

Divide by Two

The ALU has four (A2SGN, A2RAL, A2RAR, A2RSX) instructions used for right shifting (dividing by two) extended precision words. These words, (up to 64 bits) may be stored in the two on-chip register files. When the least significant 16 bit word is shifted, the vacant MSB must be filled with the LSB from the next most significant 16 bit byte. This is achieved via the A2RAL, A2RAR or A2RSX instructions which indicate the source of the new MSB (see ALU INSTRUCTION SET).

When the most significant 16 bit byte is right shifted, the MSB must be filled with a duplicate of the original MSB so as to maintain the correct sign (Sign Extension). This operation is achieved via the A2SGN instruction (see Table 1).

Constants

The ALU has four instructions (OPONE, OPBYT, OPNIB, OPALT) that force a constant value onto the ALU output. These values are primarily intended to be used as masks, or the seeds for mask generation, for example, the OPONE instruction will set a single bit in the least significant position. This bit may be rotated any where in the 16 bit field by the Barrel Shifter, allowing the AND function of the ALU to perform bit-pick operations on input data.

CLR

The ALU instruction CLRXX is used as a Master Reset for the entire device. This instruction has the effect of:

1. Clearing ALU and Barrel Shifter register files to zero.
2. Clearing A and B port input registers to zero.
3. Clearing the R1 and R2 shift control registers to zero.
4. Clearing the internally registered CO bit to zero.
5. Programming the BFP flag to detect *overflow* conditions.

The Barrel Shifter

The Barrel Shifter supports 16 instructions as detailed in Table 2. The input to the Barrel Shifter is selected by the S MUX. Data will fall through from the selected register, through the S MUX and the Barrel Shifter to the shifter output register file in 50ns for the PDSP1601A (100ns for the PDSP1601).

The Barrel Shifter instructions are latched, such that the instructions will not start executing until the rising edge of CLK latches the instruction into the device.

The Barrel Shifter is capable of Logical Arithmetic or Barrel Shifts in either direction.

A. Logical shifts discard bits that exit the 16 bit field and fill spaces with zeros.
B. Arithmetic shifts discard bits that exit the 16 bit field and fill spaces with duplicates of the original MSB.
C. Barrel Shifts rotate the 16 bit fields such that bits tha exit the 16 bit field to the left or right reappear in the vacant spaces on the right or left.

The amount of shift applied is encoded onto the 4 bit Barrel Shifter input as illustrated in Table 3. The type of shift and the amount are determined by the shift control block. The shift control block (see Fig.3) accepts and decodes the four bit ISO-3 instruction. The shift control block contains a priority encoder and two user programmable 4 bit registers R1 and R2.

There are four possible sources of shift value that can be passed onto the Barrel Shifter, there are:

1. The Priority Encoder
2. The SV input
3. The R1 register
4. The R2 register

Inst	IS3-IS0	Mnemonic	Operation	I/O
0	0000	LSRSV	Logical Shift Right by SV	I
1	0001	LSLSV	Logical Shift Left by SV	I
2	0010	BSRSV	Barrel Shift Right by SV	I
3	0011	BSLSV	Barrel Shift Left by SV	I
4	0100	LSRR1	Logical Shift Right by R1	X
5	0101	LSLR1	Logical Shift Left by R1	X
6	0110	LSRR2	Logical Shift Right by R2	X
7	0111	LSLR2	Logical Shift Left by R2	X
8	1000	LR1SV	Load Register 1 From SV	I
9	1001	LR2SV	Load Register 2 From SV	I
A	1010	ASRSV	Arithmetic Shift Right by SV	I
B	1011	ASRR1	Arithmetic Shift Right by R1	X
C	1100	ASRR2	Arithmetic Shift Right by R2	X
D	1101	NRMXX	Normalise Output PE	O
E	1110	NRMR1	Normalise Output PE, Load R1	O
F	1111	NRMR2	Normalise Output PE, Load R2	O

Table 2 Barrel shifter instructions

KEY

SV	= Shift Value
R1	= Register 1
R2	= Register 2
PE	= Priority Encoder Output
I	=> SV Port operates as an Input
O	=> SV Port operates as an Output
X	=> SV Port in a High Impedance State

MNEMONICS

LSXYY	Logical Shift,	X	= Direction	YY = Source of Shift Value
BSXYY	Barrel Shift,	X	= Direction	YY = Source of Shift Value
ASXYY	Arithmetic Shift,	X	= Direction	YY = Source of Shift Value
LXXYY	Load	XX	= Target	YY = Source
NRMYY	Normalise by PE, Output PE value on SV Port, Load YY Reg			

SV3	SV2	SV1	SV0	Shift
0	0	0	0	No shift
0	0	0	1	1 place
0	0	1	0	2 places
0	0	1	1	3 places
0	1	0	0	4 places
0	1	0	1	5 places
0	1	1	0	6 places
0	1	1	1	7 places
1	0	0	0	8 places
1	0	0	1	9 places
1	0	1	0	10 places
1	0	1	1	11 places
1	1	0	0	12 places
1	1	0	1	13 places
1	1	1	0	14 places
1	1	1	1	15 places

Table 3 Barrel shifter codes

Priority Encoder

If the priority encoder is selected as the source of the shift value (instructions:- NRMXX, NRMR1, MRMRZ), then within one 100ns cycle or two 50ns cycles for the PDSP1601A (one 200ns or two 100ns cycles for the PDSP1601), the shift circuitry will:

(1) Priority encode the 16 bit input to the Barrel Shifter and place the 4 bit value in either of the R1 or R2 registers and output the value on the SV port (if enabled by $\overline{\text{SVOE}}$).

(2) Shift the 16 bit input by the amount indicated by the Priority Encoder such that the output from the Barrel Shifter is a normalised value.

SV Input

If the SV port is selected as the source of the shift value, then the input to the Barrel Shifter is shifted by the value stored in the internal SV register.

$\overline{\text{SVOE}}$

The SV port acts as an input or an output depending upon the IS0-3 instruction. If the user does not wish to use the normalise instructions, then the SV port mat be forced to be input only by typing $\overline{\text{SVOE}}$ control high. In this mode the SV port may be considered an extension of the instruction inputs.

R1 and R2 Registers

The R1 and R2 registers may be loaded from the Priority Encoder (NRMR1 and NRMR2) or from the SV input (LR1SV, LR2SV).

Whilst the latter two instructions are executing, the Barrel Shifter will pass its input to the output unshifted.

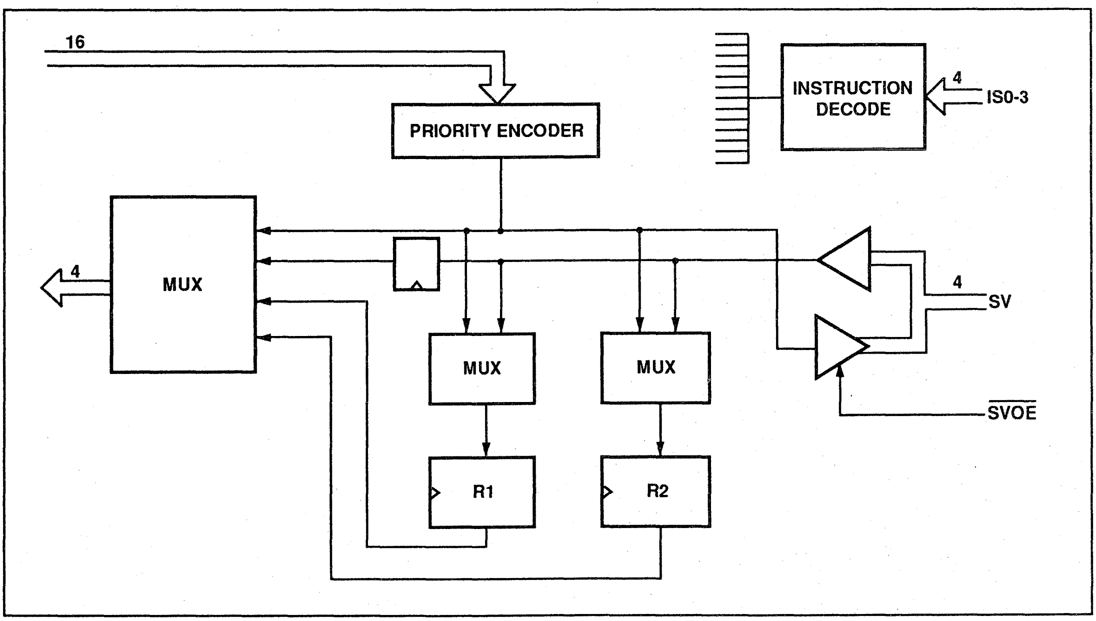

Fig.3 Shift control block

The Register Files

There are two on-chip register files (ALU and Shifter), each containing two 16 bit registers and each supporting 8 instructions (see Table 4). The instructions for the ALU register file and the Barrel Shifter Register file are the same.

The Inputs to the register files come from either the ALU or the Barrel Shifter, and are loaded into the Register files on the rising edge of CLK.

The register file instructions are latched such that the instruction will not start executing until the rising edge of the CLK latches the instruction into the device.

The register file instructions (see Table 4) allow input data to be loaded into either, neither or both of the registers. Data is loaded at the end of the cycle in which the instruction is executing.

The register file instructions allow the output to be sourced from either of the two registers, the selected output will be valid during the cycle in which the instruction is executing.

ALU REGISTER INSTRUCTIONS			
Inst	**RA2-RA0**	**Mnemonic**	**Operation**
0	000	LLRRR	Load Left Reg Output Right Reg
1	001	LRRLR	Load Right Reg Output Left Reg
2	010	LLRLR	Load Left Register, Output Left Reg
3	011	LRRRR	Load Right Register, Output Right Reg
4	100	LBRLR	Load Both Registers, Output Left Reg
5	101	NOPRR	No Load Operation, Output Right Reg
6	110	NOPLR	No Load Operation, Output Left Reg
7	111	NOPPS	No Load Operation, Pass ALU Result
SHIFTER REGISTER INSTRUCTIONS			
Inst	**RA2-RA0**	**Mnemonic**	**Operation**
0	000	LLRRR	Load Left Reg Output Right Reg
1	001	LRRLR	Load Right Reg Output Left Reg
2	010	LLRLR	Load Left Register, Output Left Reg
3	011	LRRRR	Load Right Register, Output Right Reg
4	100	LBRLR	Load Both Registers, Output Left Reg
5	101	NOPRR	No Load Operation, Output Right Reg
6	110	NOPLR	No Load Operation, Output Left Reg
7	111	NOPPS	No Load Operation, Pass Barrel Shifter Result

Table 4 ALU and shift register instructions mnemonics

MNEMONICS

LXXYY	Load XX = Target,	YY	= Source of Output
LBOXX	Load Both Registers,	XX	= Source of Output
NOPXX	No Load Operation,	XX	= Source of Output

Multiplexers

There are four user selectable on-chip multiplexers (A-MUX, B-MUX, S-MUX and C-MUX).

These four multiplexers support instructions as tabulated in Table 5.

The MUX instructions are latched such that the instruction will not start executing until the rising edge of CLK latches the instruction onto the device.

		MSA1	**MSA0**	**Output**
A-MUX	MARAX	0	0	ALU REGISTER FILE OUPUT
	MAAPR	0	1	A-PORT INPUT
	MABPR	1	0	B-PORT INPUT
	MARSX	1	1	SHIFTER REGISTER FILE OUTPUT
		MSB		**Output**
B-MUX		0		B-PORT INPUT
		1		SHIFTER REGISTER FILE OUTPUT
		MSS		**Output**
S-MUX		0		B-PORT INPUT
		1		SHIFTER REGISTER FILE OUTPUT
		MSC		**Output**
C-MUX		0		ALU REGISTER FILE OUTPUT
		1		SHIFTER REGISTER FILE OUTPUT

Table 5

INSTRUCTION SET

ALU Arithmetic Instructions

Mnemonic	Op Code	Function
CLRXX	<00>	On the rising edge of CLK at the end of the cycle in which this instruction is executing, the A Port, B Port, ALU, Barrel Shifter, and Shift Control Registers will be loaded with zeros. The internal registered CO will also be set to zero, and the BFP flag will be set to activate on overflow conditions.
MIAX1	<01>	The A input to the ALU is inverted and a one is added to the LSB.
MIAC1	<02>	The A input to the ALU is inverted and the CI input is added to the LSB.
MIACO	<03>	The A input to the ALU is inverted and the CO output from the ALU on the previous cycle is added to the LSB.
A2SGN	<04>	The A input to the ALU is right shifted one bit position. The LSB is discarded, and the vacant MSB is filled by duplicating the original MSB (Sign Extension).
A2RAL	<05>	The A input to the ALU is right shifted one bit position. The LSB is discarded, and the vacant MSB is filled with the LSB from the ALU register.
A2RAR	<06>	The A input to the ALU is right shifted one bit position. The LSB is discarded, and the vacant MSB is filled with the LSB from the ALU register.
A2RSX	<07>	The A input to the ALU is right shifted one bit position. The LSB is discarded, and the vacant MSB is filled with the LSB from the B input to the ALU.
APBCI	<08>	The A input to the ALU is added to the B input, and the CI input is added to the LSB.
APBCO	<09>	The A input to the ALU is added to the B input, and the CO out from the ALU on the previous cycle is added to the LSB.
AMBX1	<0A>	The A input to the ALU is added to the inverted B input, and a one is added to the LSB.
AMBCI	<0B>	The A input to the ALU is added to the inverted B input, and the CI input is added to the LSB.
AMBCO	<0C>	The A input to the ALU is added to the inverted B input, and the CO out from the ALU on the previous cycle is added to the LSB.
BMAX1	<0D>	The inverted A input to the ALU is added to the B input, and a one is added to the LSB.
BMAC1	<0E>	The inverted A input to the ALU is added to the B input, and the CI input is added to the LSB.
BMACO	<0F>	The inverted A input to the ALU is added to the B input, and the CO out from the ALU on the previous cycle is added to the LSB.

ALU Logical Instructions

Mnemonic	Op Code	Function
ANXAB	<10>	The A input to the ALU is logically 'ANDed' with the B input.
ANANB	<11>	The A input to the ALU is logically 'ANDed' with the inverse of the B input.
ANNAB	<12>	The inverse of the A input to the ALU is logically 'ANDed' with the B input.
ORXAB	<13>	The A input to the ALU is logically 'ORed' with the B input.
ORNAB	<14>	The inverse A input to the ALU is logically 'ORed' with the B input.
XORAB	<15>	The A input to the ALU is logically Exclusive-ORed with the B input.
PASXA	<16>	The A input to the ALU is passed to the output.
PASNA	<17>	The inverse of the A input to the ALU is passed to the output.

ALU Control Instructions

Mnemonic	Op Code	Function
SBFOV	<18>	The BFP flag is programmed to activate when an ALU operation causes an overflow of the 16 bit number range. This flag is logically the exclusive-or of the carry into and out of the MSB of the ALU. For the most significant Byte this flag indicates that the result of an arithmetic two's complement operation has overflowed into the sign bit. The output of the ALU is forced to zero for the duration of this instruction.
SBFU1	<19>	The BFP flag is programmed to activate when an ALU operation comes within a factor of two of causing an overflow of the 16 bit number range. For the most significant Byte this flag indicates that the result of an arithmetic two's complement operation is within a factor of two of overflowing into the sign bit. The output of the ALU is forced to zero for the duration of this instruction.
SBFU2	<1A>	The BFP flag is programmed to activate when an ALU operation comes within a factor of four of causing an overflow of the 16 bit number range. For the most significant Byte this flag indicates that the result of an arithmetic two's complement operation is within a factor of four of overflowing into the sign bit. The output of the ALU is forced to zero for the duration of this instruction.
SBFZE	<1B>	The BFP flag is programmed to activate when an ALU operation causes a result of zero. The output of the ALU is forced to zero for the duration of this instruction. During the execution of this instruction the BFP flag will become active.
OPONE	<1C>	The ALU will output the binary value 0000000000000001, the MSB on the left.
OPBYT	<1D>	The ALU will output the binary value 0000000011111111, the MSB on the left.
OPNIB	<1E>	The ALU will output the binary value 0000000000001111, the MSB on the left.
OPALT	<1F>	The ALU will output the binary value 0101010101010101, the MSB on the left.

Barrel Shifter Instructions

Mnemonic	Op Code	Function
LSRSV	<0>	The 16 bit input to the Barrel Shifter is right shifted by the number of places indicated by the magnitude of the four bit number present in the SV register. The LSBs are dicarded, and the vacant MSBs are filled with zeros.
LSLSV	<1>	The 16 bit input to the Barrel Shifter is left shifted by the number of places indicated by the magnitude of the four bit number present in the SV register. The LSBs are dicarded, and the vacant MSBs are filled with zeros.
BSRSV	<2>	The 16 bit input to the Barrel Shifter is rotated to the right by the number of places indicated by the magnitude of the four bit number present in the SV register. The LSBs that exit the 16 bit field to the right, reappear in the vacant MSBs on the left.
BSLSV	<3>	The 16 bit input to the Barrel Shifter is rotated to the left by the number of places indicated by the magnitude of the four bit number present in the SV register. The LSBs that exit the 16 bit field to the right, reappear in the vacant MSBs on the right.
LSRR1	<4>	The 16 bit input to the Barrel Shifter is right shifted by the number of places indicated by the magnitude of the four bit number resident within the R1 register. The LSBs are discarded, and the vacant MSBs are filled with zeros.
LSLR1	<5>	The 16 bit input to the Barrel Shifter is left shifted by the number of places indicated by the magnitude of the four bit number resident within the R1 register. The LSBs are discarded, and the vacant LSBs are filled with zeros.
LSRR2	<6>	The 16 bit input to the Barrel Shifter is right shifted by the number of places indicated by the magnitude of the four bit number resident within the R2 register. The LSBs are discarded, and the vacant MSBs are filled with zeros.
LSLR2	<7>	The 16 bit input to the Barrel Shifter is left shifted by the number of places indicated by the magnitude of the four bit number resident within the R2 register. The LSBs are discarded, and the vacant LSBs are filled with zeros.

Mnemonic	Op Code	Function
LR1SV	<8>	On the rising edge of CLK at the end of the cycle in which this instruction is executing, the R1 register will be loaded with the data present on the SV port. The input to the Barrel Shifter will be passed onto the output unshifted.
LR2SV	<9>	On the rising edge of CLK at the end of the cycle in which this instruction is executing, the R2 register will be loaded with the data present on the SV port. The input to the Barrel Shifter will be passed onto the output unshifted.
ASRSV	<A>	The 16 bit input to the Barrel Shifter is right shifted by the number of places indicated by the magnitude of the four bit number present in the SV register. The LSBs are discarded, and the vacant MSBs are filled with duplicates of the original MSB. (Sign Extension).
ASRR1	<B>	The 16 bit input to the Barrel Shifter is right shifted by the number of places indicated by the magnitude of the four bit number resident within the R1 register. The LSBs are discarded, and the vacant MSBs are filled with duplicates of the original MSB. (Sign Extension).
ASRR2	<C>	The 16 bit input to the Barrel Shifter is right shifted by the number of places indicated by the magnitude of the four bit number resident within the R2 register. The LSBs are discarded, and the vacant MSBs are filled with duplicates of the original MSB. (Sign Extension).
NRMXX	<D>	The 16 bit input to the Barrel Shifter is left shifted by the number of places indicated by the magnitude of the four bit number output from the Priority Encoder. This value is also output on the SV port (provided $\overline{SVOE}$ is low). The effect of this operation is to left shift the input by the necessary amount (max 15 places) to result in the MSB and the next most significant bit being different. This has the effect of eliminating unnecessary Sign Bits, and hence Normalising the input data. The MSBs shifted out to the left are discarded, and the vacant LSBs on the right are filled with zeros.
NRMR1	<E>	The 16 bit input to the Barrel Shifter is left shifted by the number of places indicated by the magnitude of the four bit number output from the Priority Encoder. This value is also loaded into the R1 register at the end of the cycle, and is output on the SV port (provided $\overline{SVOE}$ is low). The effect of this operation is to left shift the input by the necessary amount (max 15 places) to result in the MSB and the next most significant bit being different. This has the effect of eliminating unnecessary Sign Bits, and hence Normalising the input data. The MSBs shifted out to the left are discarded, and the vacant LSBs on the right are filled with zeros.
NRMR2	<F>	The 16 bit input to the Barrel Shifter is left shifted by the number of places indicated by the magnitude of the four bit number output from the Priority Encoder. This value is also loaded into the R2 register at the end of the cycle, and is output on the SV port (provided $\overline{SVOE}$ is low). The effect of this operation is to left shift the input by the necessary amount (max 15 places) to result in the MSB and the next most significant bit being different. This has the effect of eliminating unnecessary Sign Bits, and hence Normalising the input data. The MSBs shifted out to the left are discarded, and the vacant LSBs on the right are filled with zeros.

Barrel Shifter or ALU Register Instructions

Mnemonic	Op Code	Function
LLRRR	<0>	After the rising edge of CLK at the beginning of the cycle in which this instruction is executed, the contents of the Right register will appear on the output. On the rising edge of CLK at the end of the cycle, and the data on the register inputs will be loaded into the Left Register.
LRRLR	<1>	After the rising edge of CLK at the beginning of the cycle in which this instruction is executed, the contents of the Left register will appear on the output. On the rising edge of CLK at the end of the cycle, the data on the register inputs will be loaded into the Right Register.
LLRLR	<2>	After the rising edge of CLK at the beginning of the cycle in which this instruction is executed, the contents of the Left register will appear on the output. On the rising edge of CLK at the end of the cycle, the data on the register inputs will be loaded into the Left Register.
LRRRR	<3>	After the rising edge of CLK at the beginning of the cycle in which this instruction is executed, the contents of the Right register will appear on the output. On the rising edge of CLK at the end of the cycle, the data on the register inputs will be loaded into the Right Register.
LBRLR	<4>	After the rising edge of CLK at the beginning of the cycle in which this instruction is executed, the contents of the Left register will appear on the output. On the rising edge of CLK at the end of the cycle, and the data on the register inputs will be loaded into both Left and Right Register.
NOPRR	<5>	After the rising edge of CLK at the beginning of the cycle in which this instruction is executed, the contents of the Right register will appear on the output. On the rising edge of CLK at the end of the cycle no load operation will occur, the register contents will remain unchanged.
NOPLR	<6>	After the rising edge of CLK at the beginning of the cycle in which this instruction is executed, the contents of the Left register will appear on the output. On the rising edge of CLK at the end of the cycle no load operation will occur, the register contents will remain unchanged.
NOPPS	<7>	After the rising edge of CLK at the beginning of the cycle in which this instruction is executed, the input to the registers will appear on the output. On the rising edge of CLK at the end of the cycle no load operation will occur, the register contents will remain unchanged.

TYPICAL APPLICATION

Select a 16 bit field from each word in a block of 32 bit words with a 10MHz throughput.

The 16 bit field indicated is to be selected from each 32 bit word.

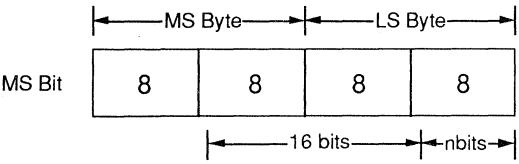

The 32 bit words are fed into the B port of the PDSP1601 in two cycles, MS byte first.

The PDSP1601 shift control is initiated by programming the R1 and R2 registers with n and 16-n respectively.

The shift operation is implemented in three steps:-

(1) The MS byte is logically left shifted (16-n) places, the MSBs being discarded and the LSB spaces being filled with zeros. This shifted data is loaded into the shifter register file left register.

(2) The LS byte is logically right shifted, n-places, the LSBs being discarded and the MSBs being filled with zeros. This shifted data is loaded into the shifter register file left register.

During this cycle the previous contents of this register are passed through the ALU to the ALU register file left register.

(3) While the MS byte of the next 32 bit word is shifted in the Barrel Shifter, the two previous results, resident within the left registers of the ALU and Shifter Register files are 'ORed' by the ALU, the result being the desired 16 bit field is loaded into the ALU register file right register ready to be output on the next cycle.

The instructions from initialisation are given in Table 6.

CLK	CEB	MSA	MSB	MSS	MSC	IA	IS	SV	RA	RS	Comment
1/	1	MARSX	1	0	0	CLRXX	X	X	NOPLR	NOPLR	Clear
2/	1	MARSX	1	0	0	PASXA	LR1SV	n	NOPLR	NOPLR	Load R1 with n
3/	0	MARSX	1	0	0	PASXA	LR2SV	(16-n)	NOPLR	NOPLR	Load R2 with (16-n)
4/	0	MARSX	1	0	0	PASXA	LSLR2	X	NOPLR	LLRLR	Shift 1st MS byte
5/	0	MARSX	1	0	0	PASXA	LSRR1	X	LLRRR	LLRLR	Shift 1st LS byte
6/	0	MARAX	1	0	0	ORXAB	LSLR2	X	LRRLR	LLRLR	OR 1st bytes and shift 2nd MS byte
7/	0	MARSX	1	0	0	PASXA	LSRR1	X	LLRRR	LLRLR	Shift 2nd LS byte and output first result
8/	0	MARAX	1	0	0	ORXAB	LSLR2	X	LRRLR	LLRLR	Shift 3rd LS byte

Repeat instruction pair 5/ and 6/ until all 16 bit fields have been selected.

Table 6

ABSOLUTE MAXIMUM RATINGS (Note 1)

Supply voltage Vcc	-0.5V to 7.0V
Input voltage V_{IN}	-0.9 to Vcc + 0.9V
Output voltage V_{OUT}	-0.9 to Vcc + 0.9V
Clamp diode current per pin Ik (see note 2)	±18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to +150°C
Ambient temperature with power applied T_{amb}	
Military	-40°C to +125°C
Industrial	-40°C to +85°C
Package power dissipation P_{TOT}	
AC	1000mw
LC	1000mw

NOTES

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.

THERMAL CHARACTERISTICS

Package type	Θ_{JC} °C/W	Θ_{JA} °C/W
AC	12	36
LC	12	35

ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)

T_{AMB} (Commercial) = 0°C to +70°C, V_{CC} = 5.0V±5%, Ground = 0V
T_{AMB} (Industrial) = -40°C to +85°C, V_{CC} = 5.0V±10%, Ground = 0V
T_{AMB} (Military) = -55°C to +125°C, V_{CC} = 5.0V±10%, Ground = 0V

Static Characteristics

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4			V	I_{OH} = 8mA
Output low voltage	V_{OL}			0.4	V	I_{OL} = -8mA
Input high voltage	V_{IH}	3.5			V	
Input low voltage	V_{IL}			0.5	V	
Input leakage current	I_{IL}	-10		+10	µA	GND < V_{IN} < V_{CC}
Vcc current	I_{CC}			60	mA	T_{amb} = -40°C to +85°C
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{CC}
Output S/C current	I_{SC}	12		80	mA	V_{CC} = Max
Input capacitance	C_{IN}		5		pF	

Switching Characteristics

Characteristic	Value				Units	Conditions
	PDSP1601		PDSP1601/A			
	Min.	Max.	Min.	Max.		
CLK rising edge to C-PORT	5	40	5	25	ns	2 x LSTTL + 20pF
CLK rising edge to CO	5	100	5	50	ns	1 x LSTTL + 5pF
CLK rising edge to BFP	5	100	5	50	ns	1 x LSTTL + 5pF
Setup $\overline{CEA}$ or $\overline{CEB}$ to CLK rising edge	30		15		ns	
Hold $\overline{CEA}$ or $\overline{CEB}$ after CLK rising edge		0		0	ns	
Setup A or B port inputs to CLK rising edge	40		20		ns	
Hold A or B port inputs after CLK rising edge		0		0	ns	
Setup MSA0-1, MSB, MSS, MSC, RA2-0, RS0-2, IA0-4, IS0-3, to CLK rising edge	40		20		ns	
Hold RS0-2, IA0-4 after CLK rising edge		0		0	ns	
Hold IS0-3 after CLK rising edge		3		3	ns	
Hold MSA0-1, MSB, MSS, MSC, RA0-2 after CLK rising edge		0		0	ns	
Setup SV to CLK rising edge	40		20		ns	Input mode
Hold SV after CLK rising edge		3		3	ns	Input mode
CLK rising edge to SV	5	100	5	50	ns	20pF load, SV O P mode
$\overline{OE}$ _/‾ C-PORT ‾_ Z		40		25	ns	2 x LSTTL + 20pF
$\overline{OE}$ _/‾ C-PORT _/‾ Z		40		25	ns	2 x LSTTL + 20pF
$\overline{OE}$ ‾_ C-PORT Z _/‾		40		25	ns	2 x LSTTL + 20pF
$\overline{OE}$ ‾_ C-PORT Z ‾_		40		25	ns	2 x LSTTL + 20pF
Clock period (ALU & Barrel Shifter, serial mode)	200		100		ns	
Clock period (ALU & Barrel Shifter, parallel mode)	100		50		ns	
Clock high time	40		20		ns	
Clock low time	40		20		ns	

NOTES
1. LSTTL is equivalent to I_{OH} at 20µA I_{OL} of -0.4mA
2. Current is defined as negative into the device.

ORDERING INFORMATION

PDSP1601 BO AC	10MHz Industrial - PGA package
PDSP1601 BO LC	10MHz Industrial - LCC package
PDSP1601 AO AC	10MHz Military - PGA package
PDSP1601 AO LC	10MHz Military - LCC package
PDSP1601 MC GGCR	10MHz MIL883 Screened - QFP package
PDSP1601A CO AC	20MHz Commerical - PGA package
PDSP1601A BO AC	20MHz Industrial - PGA package
PDSP1601A BO LC	20MHz Industrial - LCC package

PDSP16112/PDSP16112A

16 x 12 BIT COMPLEX MULTIPLIER

The PDSP16112/PDSP16112A will multiply a complex (16 + 16) bit word by a complex (12 + 12) bit coefficient word and produce a complex (17 + 17) bit rounded product. The input data format is two's complement. The device consists of four 16 x 12 multiplier sections based on Booth's '2 bits at a time' algorithm and is pipelined to achieve a 20MHz (PDSP16112A) or 10MHz (PDSP16112) throughput.

FEATURES

- 20MHz Complex Number (16 + 16) x (12 + 12) Multiplication
- Pipeline Architecture
- Power Dissipation only 500mW
- TTL Compatible Inputs
- 120 pin PGA or QFP packages

APPLICATIONS

- Digital Filtering
- Fast Fourier Transforms
- Radar and Sonar Processing
- Instrumentation
- Automation
- Image Processing

ASSOCIATED PRODUCTS

PDSP1601 Arithmetic Logic Unit
PDSP16318 40MHz Address Generator
PDSP16330 Pythagoras Processor

PIN 1A INDEX MARK ON TOP SURFACE

	1	2	3	4	5	6	7	8	9	10	11	12	13
A	VCC	GND	PI07	PI10	PI13	PI16	IC	GND	PR15	PR12	PR09	PR06	VCC
B	PI03	NC	PI06	PI09	PI12	PI15	CLK	IC	PR14	PR11	PR08	GND	PR03
C	PI02	PI05	NC	PI08	PI11	PI14	GND	PR16	PR13	PR10	PR07	PR05	PR02
D	PI00	PI01	PI04	LOCATION PIN (NC)							PR04	PR01	PR00
E	XI00	VCC	GND								NC	GND	VCC
F	XI03	XI02	XI01								GND	XR00	XR01
G	XI05	XI04	VCC								XR03	XR04	XR02
H	XI06	XI07	XI08								XR07	XR06	XR05
J	XI09	XI10	XI11								XR10	XR09	XR08
K	XI12	XI13	YI11								XR15	XR12	XR11
L	XI14	NC	YI10	YI07	YI04	YI01	CLK	YR01	YR04	YR07	YR10	YR11	XR13
M	XI15	NC	YI09	YI06	YI03	YI00	GND	YR00	YR03	YR06	YR09	NC	XR14
N	VCC	GND	YI08	YI05	YI02	GND	GND	GND	YR02	YR05	YR08	GND	VCC

PIN GRID ARRAY
PACKAGE
AC120

XRxx : X REAL INPUTS
XIxx : X IMAGINARY INPUTS
YRxx : Y REAL INPUTS
YIxx : Y IMAGINARY INPUTS
PRxx : PRODUCT REAL OUTPUTS
PIxx : PRODUCT IMAGINARY OUTPUTS

AC120

Fig. 1 Pin connections - top view (AC120 - PGA)

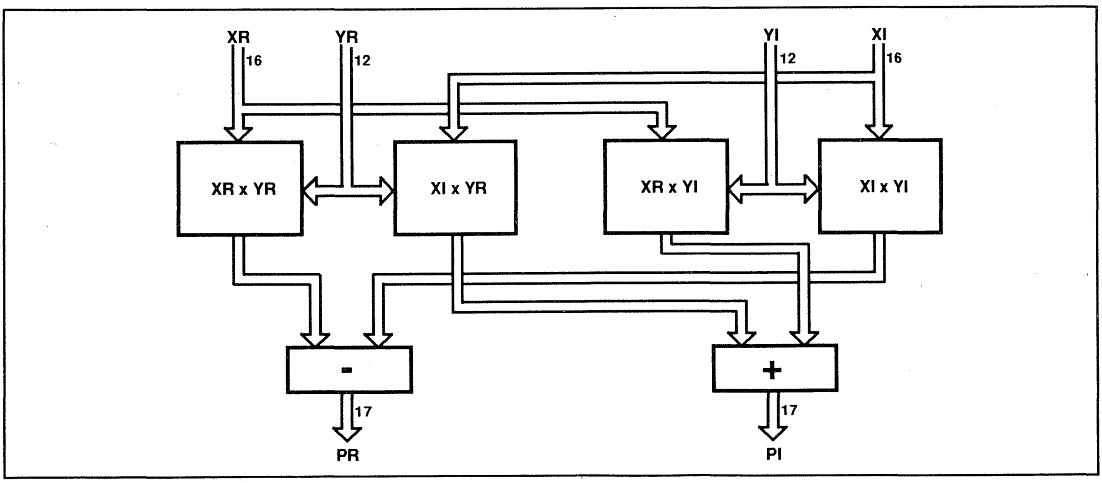

Fig. 2 Multiplier block diagram

PIN OUT - FUNCTION TO PIN (PGA Package - AC120)

Symbol	Pin No.	Symbol	Pin No.	Symbol	Pin No.	Symbol	Pin No.
PR00	D13	PR09	A11	PI00	D1	PI09	B4
PR01	D12	PR10	C10	PI01	D2	PI10	A4
PR02	C13	PR11	B10	PI02	C1	PI11	C5
PR03	B13	PR12	A10	PI03	B1	PI12	B5
PR04	D11	PR13	C9	PI04	D3	PI13	A5
PR05	C12	PR14	B9	PI05	C2	PI14	C6
PR06	A12	PR15	A9	PI06	B3	PI15	B6
PR07	C11	PR16	C8	PI07	A3	PI16	A6
PR08	B11	CLK	L7	PI08	C4	CLK	B7
XR00	F12	XI00	E1	YR00	M8	YI00	M6
XR01	F13	XI01	F3	YR01	L8	YI01	L6
XR02	G13	XI02	F2	YR02	N9	YI02	N5
XR03	G11	XI03	F1	YR03	M9	YI03	M5
XR04	G12	XI04	G2	YR04	L9	YI04	L5
XR05	H13	XI05	G1	YR05	N10	YI05	N4
XR06	H12	XI06	H1	YR06	M10	YI06	M4
XR07	H11	XI07	H2	YR07	L10	YI07	L4
XR08	J13	XI08	H3	YR08	N11	YI08	N3
XR09	J12	XI09	J1	YR09	M11	YI09	M3
XR10	J11	XI10	J2	YR10	L11	YI10	L3
XR11	K13	XI11	J3	YR11	L12	YI11	K3
XR12	K12	XI12	K1	NC	B2	NC	M12
XR13	L13	XI13	K2	NC	L2	NC	M2
XR14	M13	XI14	L1	VCC	A1	NC	E11
XR15	K11	XI15	M1	VCC	G3	NC	C3
GND	N12	GND	C7	VCC	E2	GND	N8
GND	N7	GND	A2	VCC	A13	GND	N6
GND	M7	GND	E12	VCC	E13	GND	F11
GND	N2	GND	E3	VCC	N1	IC	B8
GND	A8	GND	B12	VCC	N13	IC	A7

NOTE
IC = Internally connected - do not connect to these pins.
All inputs are internally connected to Vcc by 10k (nominal) resistors.

PIN OUT - PIN TO FUNCTION (PGA Package - AC120)

	1	2	3	4	5	6	7	8	9	10	11	12	13
A	VCC	GND	PI07	PI10	PI13	PI16	IC	GND	PR15	PR12	PR09	PR06	VCC
B	PI03	NC	PI06	PI09	PI12	PI15	CLK	IC	PR14	PR11	PR08	GND	PR03
C	PI02	PI05	NC	PI08	PI11	PI14	GND	PR16	PR13	PR10	PR07	PR05	PR02
D	PI00	PI01	PI04								PR04	PR01	PR00
E	XI00	VCC	GND								NC	GND	VCC
F	XI03	XI02	XI01								GND	XR00	XR01
G	XI05	XI04	VCC								XR03	XR04	XR02
H	XI06	XI07	XI08								XR07	XR06	XR05
J	XI09	XI10	XI11								XR10	XR09	XR08
K	XI12	XI13	YI11								XR15	XR12	XR11
L	XI14	NC	YI10	YI07	YI04	YI01	CLK	YR01	YR04	YR07	YR10	YR11	XR13
M	XI15	NC	YI09	YI06	YI03	YI00	GND	YR00	YR03	YR06	YR09	NC	XR14
N	VCC	GND	YI08	YI05	YI02	GND	GND	GND	YR02	YR05	YR08	GND	VCC

PIN OUT - PIN TO FUNCTION (PGA Package - AC120)

GG	SIG	GG	SIG	GG	SIG	GG	SIG
84	PR00	95	PR09	8	PI00	115	PI09
85	PR01	96	PR10	7	PI01	114	PI10
86	PR02	97	PR11	6	PI02	113	PI11
87	PR03	98	PR12	5	PI03	112	PI12
88	PR04	99	PR13	4	PI04	111	PI13
89	PR05	100	PR14	3	PI05	110	PI14
92	PR06	101	PR15	118	PI06	109	PI15
93	PR07	102	PR16	117	PI07	108	PI16
94	PR08	46	CLK	116	PI08	105	CLK
79	XR00	11	XI00	49	YR00	43	YI00
78	XR01	12	XI01	50	YR01	42	YI01
77	XR02	13	XI02	51	YR02	41	YI02
76	XR03	14	XI03	52	YR03	40	YI03
75	XR04	15	XI04	53	YR04	39	YI04
74	XR05	17	XI05	54	YR05	38	YI05
73	XR06	18	XI06	55	YR06	37	YI06
72	XR07	19	XI07	56	YR07	36	YI07
71	XR08	20	XI08	57	YR08	35	YI08
70	XR09	21	XI09	58	YR09	34	YI09
69	XR10	22	XI10	59	YR10	33	YI10
68	XR11	23	XI11	63	YR11	28	YI11
67	XR12	24	XI12	1	N/C	29	N/C
66	XR13	25	XI13	16	N/C	31	N/C
65	XR14	26	XI14	2	VCC	61	N/C
64	XR15	27	XI15	10	VCC	83	N/C
9	GND	45	GND	30	VCC	44	GND
32	GND	47	GND	62	VCC	48	GND
60	GND	104	GND	81	VCC	80	GND
82	GND	106	GND	90	VCC	103	I/C
91	GND	120	GND	119	N/C	107	I/C

N/C = Not connected - leave open circuit
I/C = Internally connected - leave open circuit

All GND and VDD pins must be used

PIN DESCRIPTION

XR00 - XR15	X Real Inputs : Two's Complement Format XR15 = MSB (Sign) XR00 = LSB For Fractional Arithmetic the Weighting of XR15 = 1 i.e. $-1 \leq XR < 1$	PR00 - PR16	P Real Inputs : Two's Complement Format PR16 = MSB (Sign) PR00 = LSB For Fractional Arithmetic the Weighting of PR16 = 2 i.e. $-2 \leq PR < 2$
XI00 - XI15	X Imag Inputs : Two's Complement Format XI15 = MSB (Sign) XI00 = LSB For Fractional Arithmetic the Weighting of XI15 = 1 i.e. $-1 \leq XI < 1$	PI00 - PI16	P Imag Outputs : Two's Complement Format PI16 = MSB (Sign) PI00 = LSB For Fractional Arithmetic the Weighting of PI16 = 2 i.e. $-2 \leq PI < 2$
YR00 - YR11	Y Real Inputs : Two's Complement Format YR11 = MSB (Sign) YR00 = LSB For Fractional Arithmetic the Weighting of YR11 = 1 i.e. $-1 \leq YR < 1$	CLK pin B7 and Pin L7 VCC GND	Common Clock to all on chip registers, both pins must be connected All VCC and GND pins must be connected
YI00 - YI11	Y Imag Inputs : Two's Complement Format YI11 = MSB (Sign) YI00 = LSB For Fractional Arithmetic the Weighting of YI11 = 1 i.e. $-1 \leq YI < 1$	IC	Internally connected - do not use

FUNCTIONAL DESCRIPTION

The PDSP16112 Complex Multiplier contains four pipeline 16 x 12 Array Multipliers, a 17-bit adder and a 17-bit subtractor.

The multipliers accept data from the XR, XI, YR, and YI inputs and perform the four multiplies necessary to implement a Complex Multiply Operation.

$$(XR \times YR,\ XR \times YI,\ XI \times YR,\ XI \times YI).$$

The 28-bit results from these operations are rounded to the most significant 16-bits before being passed to the adder and subtractor. The subtractor calculates

$$(XR \times YR) - (XI \times YI)$$

to form a 17-bit result representing the real result of the complex multiplication. The adder calculates

$$(XR \times YI) + (XI \times YR)$$

to form a 17-bit result that represents the imaginary result of the complex multiplication. These real and imaginary results are passed to the PR and PI outputs respectively.

The add and subtract operations may (depending upon the data) cause the multiplier results to grow by one bit hence requiring 17-bit outputs to represent the results. The PDSP16112 is designed to operate with two's complement arithmetic, hence if the Fractional two's complement format is used the outputs will lie in the range

$$-2 \leq P < 2$$

for inputs in the range

$$-1 \leq X \text{ or } Y < 1$$

If the output magnitude lies in the range

$$-1 \leq P < 1,$$

then the 17th (MSB) bit of the outputs will duplicate the 16th (Sign) bit of the output.

In common with other Array multipliers, the operation

$$-1 \times -1$$

will yield an incorrect result for fractional two's complement formats, and hence should be avoided.

Both X and Y inputs are registered as are the PR and PI outputs. On the rising edge of CLK data present on the XR, XI, YR and YI inputs is clocked into the input registers. At the same time a new result is clocked into the output registers and made available on the PR and PI output ports.

Pipelined Operation

The internal Multiply and Add operations are divided into stages by six internal pipeline registers giving a total latency through the device of eight clock cycles. This means that the result from data loaded into the device on the first clock cycle appears at the outputs during the seventh clock cycle, and may be loaded into another device on the eigth clock cycle.

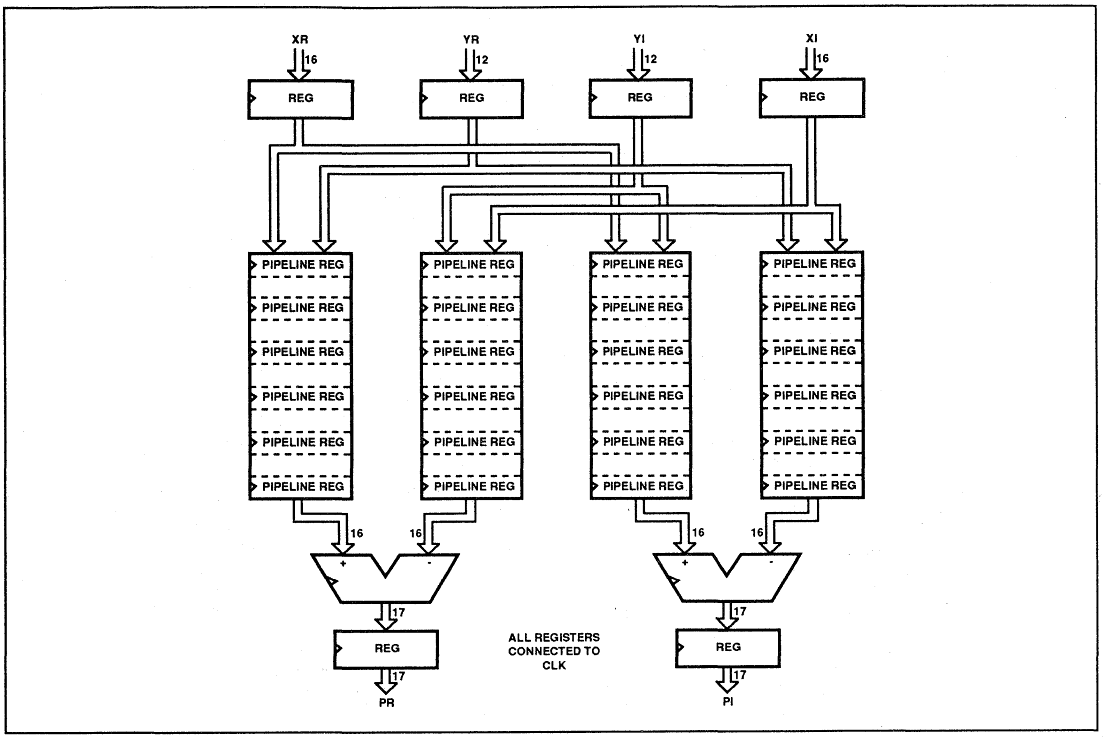

Fig.3 Pipeline multiplier structure

TYPICAL APPLICATION

The PDSP16112A may be configured as the main arithmetic element in the FFT Butterfly calculation. A single PDSP16112A together with two PDSP16318As will produce an arithmetic processor capable of executing a new Radix 2 DIT Butterfly every 50ns using 16-bit data and 12-bit coefficients. The PDSP16318A provides flags that monitor the magnitude of the output data, together with on chip shift circuits.

A single Butterfly processor of this type will allow the following FFT benchmarks.

1024 point complex radix 2 transform in 256µsecs
512 point complex radix 2 transform in 115µsecs
256 point complex radix 2 transform in 51µsecs

The arithmetic operation required to realise a radix 2 decimation in time algorithm is as follows.

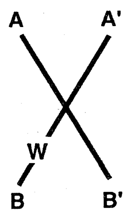

$$A' = A + (B \times W)$$

$$B' = A - (B \times W)$$

Where A and B are the data inputs, A' and B' are the data outputs, and W is the coefficient. A,B,A',B' and W are all complex numbers i.e. they all have real and imaginary components. The Butterfly therefore requires one complex multiply and two complex adds to execute, which is equivalent to four real multiplies and six real adds.

Fig.4 illustrates the interconnection of the PDSP16112A with the two PDSP16318A Complex Accumulators. The PDSP16112A performs the complex multiply operation at the full 20MHz rate to provide the real and imaginary components of the (B x W) to the two ALUs. The PDSP16318A is capable of 16-bit operations at 20MHz and has on chip register storage and Shifter. In every 20MHz cycle each PDSP16318A performs two arithmetic operations to calculate the real or imaginary parts of A + (B x W) and A - (B x W). One of the PDSP16318As calculates the real parts and the other calculates the imaginary parts.

For greater throughput one chip-set mat be allocated to each column of the FFT. For example, a 1K complex FFT could be calculated by 10 chip-sets every 26µs.

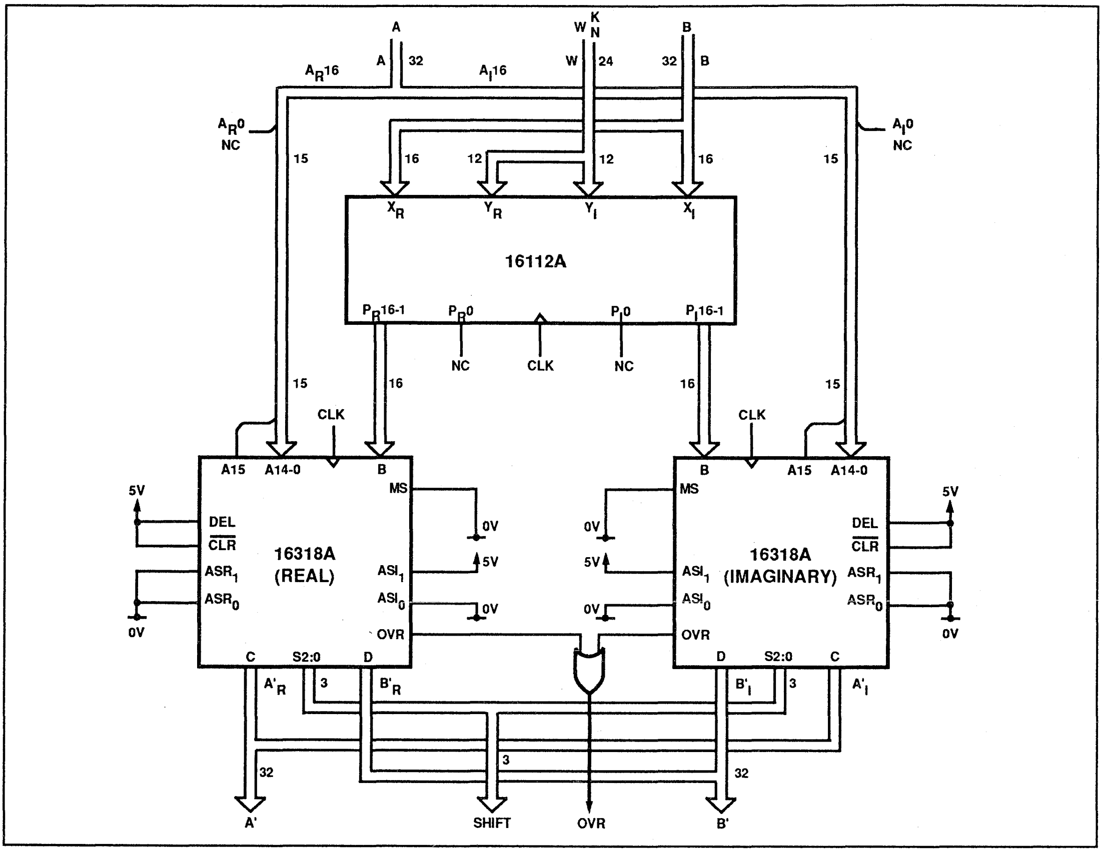

Fig.4 Radix 2 DIT butterfly processor

ELECTRICAL CHARACTERISTICS

Test conditions (unless otherwise stated):

T_{amb} (Industrial) =-40°C to +85°C, Vcc = 5.0V ± 10%, GND = 0V

T_{amb} (Military) =-55°C to +125°C, Vcc = 5.0V ± 10%, GND = 0V

T_{amb} (Commercial) = 0°C to +70°C, Vcc = 5.0V ± 5%, GND = 0V

Static Characteristics

Characteristics	Symbol	Value						Units	Conditions
		PDSP16112			PDSP16112A				
		Min.	Typ.	Max.	Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4			2.4			V	I_{OH} = 4mA
Output low voltage	V_{OL}			0.6			0.6	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.8			2.8			V	
Input low voltage	V_{IL}			0.8			0.8	V	
Input leakage current *	I_{IL}	-1.2		+0.01	-1.2		+0.01	mA	GND $\leq V_{IN} \leq V_{CC}$
Output short circuit current	I_{OS}	30		200	40		200	mA	V_{CC} = max
Input capacitance	C_I		10			10		pF	

* All inputs have a nominal 10K pull resistor to Vcc.

AC Characteristics

Characteristic	Symbol	Value Industrial						Value Military		Units	Conditions
		PDSP16112			PDSP16112A						
		Min.	Typ.	Max.	Min.	Typ.	Max.	Min.	Typ.		
Vcc current	I_{CC}			90			170		90	mA	Vcc = max Outputs unloaded f_{CLK} = max
Max. CLK frequency	f_{CLK}	10			20			10		MHz	
Min. CLK frequency				DC			DC		DC		
Input setup time	t_{SU}			30			20		30	ns	
Input hold time	t_{IH}			5			5		5	ns	
CLK to output delay	t_d	5		50	5		30	5	50	ns	
CLK Mark/Space ratio		40		60	40		60	40	60	%	
Drive capability		2 x LSTTL +20pF									

ABSOLUTE MAXIMUM RATINGS (Note 1)

Supply voltage V_{CC}	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to Vcc +0.5V
Output voltage V_{OUT}	-0.5V to Vcc +0.5V
Clamp diode current per I_k (see Note 2)	±18mA
Static discharge voltage	500V
Storage temperature range T_S	-65°C to +150°C
Junction temperature	150°C
Ambient temperature with power applied T_{amb}	
Commercial	0°C to +70°C
Industrial	-44°C to +85°C
Military	-55°C to +125°C
Package power dissipation P_{TOT}	1000mW

NOTES
1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliability.

THERMAL CHARACTERISTICS

Package Type	θ_{JC} °C/W	θ_{JA} °C/W
AC	12	35

ORDERING INFORMATION

Commercial (0°C to +70°C)

PDSP16112 C0 AC (10MHz - PGA)
PDSP16112A C0 AC (20MHz - PGA)
PDSP16112A C0 GG (20MHz - QFP)

Industrial (-40°C to +85°C)

PDSP16112 B0 AC (10MHz - PGA)
PDSP16112A B0 AC (20MHz - PGA)
PDSP16112A B0 GG (20MHz - QFP)

Military (-55°C to +125°C)

PDSP16112 A0 AC (10MHz - PGA)
PDSP16112A A0 AC (20MHz - PGA)
PDSP16112A A0 GG (20MHz - QFP)

Call for availability on High Reliability parts and MIL-883C screening.

GEC PLESSEY
SEMICONDUCTORS

PRELIMINARY INFORMATION

DS3707 - 2.0

PDSP16116/A

16 BY 16 BIT COMPLEX MULTIPLIER

The PDSP16116A will multiply two complex (16 + 16) bit words every 50ns and can be configured to output the complete complex (32+32) bit result within a single cycle. The data format is fractional two's complement.

The PDSP16116/A contains four 16 x 16 Array Multipliers, two 32 bit Adder/Subtractors and all the control logic required to support Block Floating Point Arithmetic as used in FFT applications. In combination with a PDSP16318, the PDSP16116A forms a two chip 10MHz Complex Multiplier Accumulator with 20 bit accumulator registers and output shifters. The PDSP16116 in combination with two PDSP16318s and two PDSP1601s forms a complete 10MHz Radix 2 DIT FFT Butterfly solution which fully supports Block Floating Point Arithmetic. The PDSP16116/A has an extremely high throughput that is suited to recursive algorithms as all calculations are performed with a single pipeline delay (two cycle fall-through).

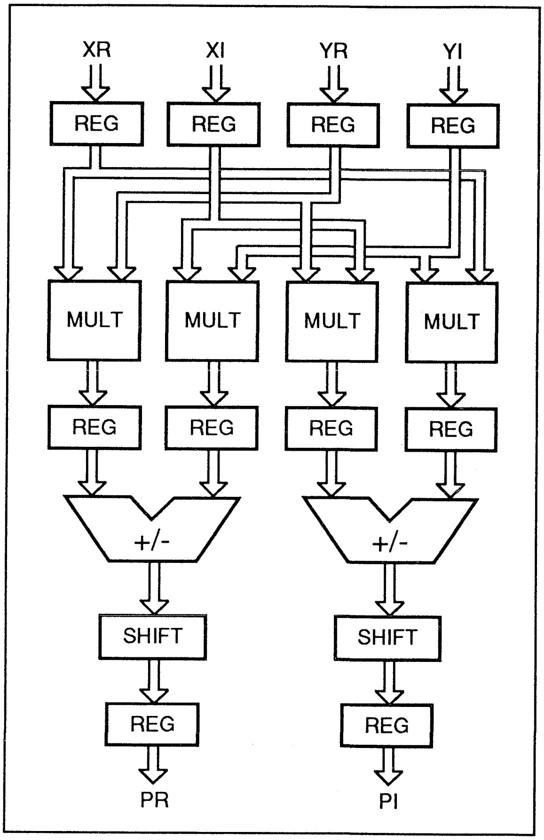

Fig. 1 Simplified Block Diagram

FEATURES

- Complex Number (16 + 16) X (16 + 16) Multiplication
- Full 32 bit Result
- 20MHz Clock Rate
- Block Floating Point FFT Butterfly Support
- -1 times -1 Trap
- Two's Complement Fractional Arithmetic
- TTL Compatible I/O
- Complex Conjugation
- 2 Cycle Fall Through
- 144 pin PGA or QFP packages

APPLICATION

- Fast Fourier Transforms
- Digital Filtering
- Radar and Sonar Processing
- Instrumentation
- Image Processing

ASSOCIATED PRODUCTS

PDSP16318/A	Complex Accumulator
PDSP16112/A	(16 + 16) X (12 + 12) Complex Multiplier
PDSP16330/A	Pythagoras Processor
PDSP1601/A	ALU and Barrel Shifter
PDSP16350	Precision Digital Modulator
PDSP16256	Programmable FIR Filter
PDSP16510	Single Chip FFT Processor

The PDSP16116 has a number of features tailored for System applications.

-1 x -1 Trap

In multiply operations utilising Twos Complement Fractional notation, the -1 x -1 operation forms an invalid result as +1 is not representable in the fractional number range. The PDSP16116/A eliminates this problem by trapping the -1 x -1 operation and forcing the Multiplier result to become the most positive representable number.

Complex Conjugation

Many algorithms utilising complex arithmetic require conjugation of complex data stream. This operation has traditionally required an adiditional ALU to multiply the imaginary component by -1. The PDSP16116 eliminates the requirement for the extra ALU by offering on chip complex conjugation of either of the two incoming complex data words with no loss in throughput.

Easy Interfacing

As with all PDSP family members the PDSP16116 has registered I/O for data and control. Data inputs have independent clock enables and data outputs have independent three state output enables.

Signal	Type	Description	Normal mode Configuration
XR15:0	INPUT	16 bit input for real x data	
XI15:0	INPUT	16 bit input for imag x data	
YR15:0	INPUT	16 bit input for reaal y data	
YI15:0	INPUT	16 bit input for imag y data	
PR15:0	OUTPUT	16 bit output for real p data	
PI15:0	OUTPUT	16 bit output for img p data	
CLK	INPUT	Clock, new data is loaded on rising edge of CLK	
$\overline{\text{CEX}}$	INPUT	Clock, enable X-port input register	
$\overline{\text{CEY}}$	INPUT	Clock, enable Y-port input register	
CONX	INPUT	Conjugate X data	
CONY	INPUT	Conjugate Y data	
ROUND	INPUT	Rounds the real & imag results	
MBFP	INPUT	Mode select (BFP/Normal)	Tie Low
$\overline{\text{SOBFP}}$	INPUT	Start of BFP operations **	Tie Low
$\overline{\text{EOPSS}}$	INPUT	End of pass **	Tie Low
AR15:13	INPUT	3 MSB's from real part of A-word **	Tie Low
AI15:13	INPUT	3 MSB's from imag part of A-word **	Tie Low
WTA1:0	INPUT	Word tag from A-word	Tie Low
WTB1:0	INPUT	Word tag from B-word / shift control *	
WTOUT1:0	OUTPUT	Word tag output **	
SFTA1:0	OUTPUT	Shift control for A-word / overflow flag *	
SFTR2:0	OUTPUT	Shift control for accumulator resul **	
GWR4:0	OUTPUT	Global weighting register contents **	
OSEL1:0	INPUT	Selects the desired output configuration	
$\overline{\text{OER}}$, $\overline{\text{OEI}}$	INPUT	Output enables	
VDD	POWER	+5V Supply All supply pins	
GND	POWER	0V Supply must be connected	

* Indicates pin performs different functions in BFP / Normal modes.
** Indicates pin is used only in BFP mode

Table. 1 Signal Descriptions

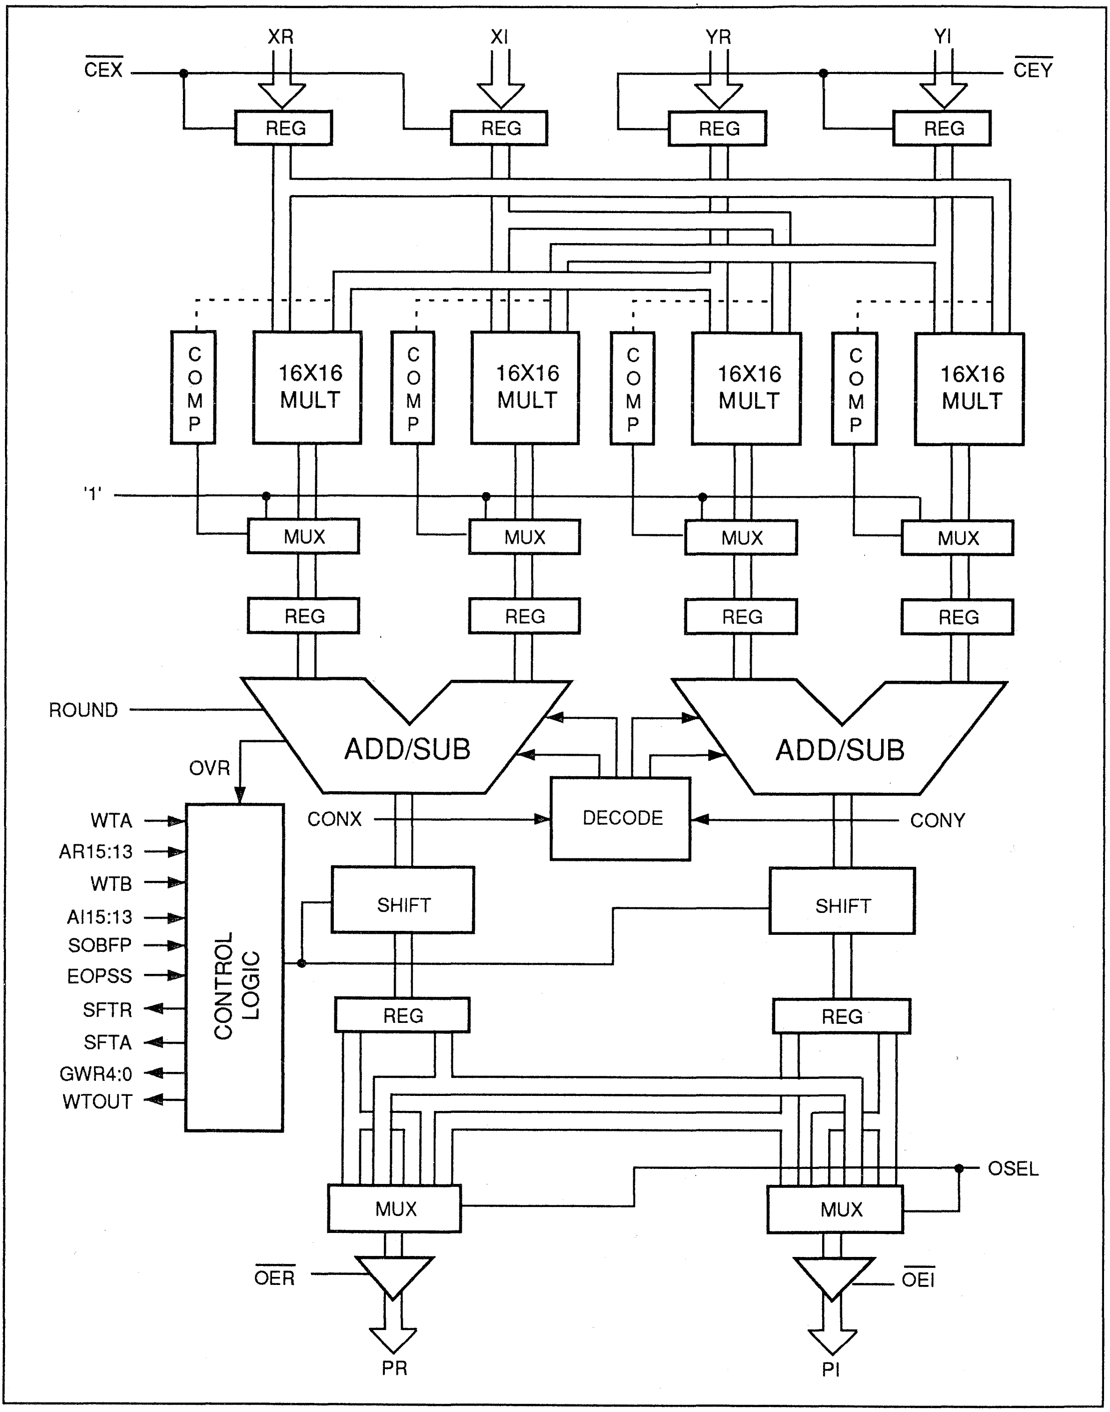

Fig.2 Block Diagram

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	VDD	XI3	XI6	XI8	XI11	XI14	XI15	$\overline{CEX}$	XR13	XR12	XR8	XR6	XR5	XR1	YR13
P	YI1	GND	XI2	XI4	XI7	XI10	XI12	$\overline{CEY}$	XR11	XR10	XR7	XR4	XR2	YR14	YR12
N	YI5	YI2	XI0	XI1	XI5	XI9	XI13	XR15	XR14	XR9	XR3	XR0	YR15	VDD	YR9
M	YI6	YI4	YI0										GND	YR11	YR7
L	YI8	YI7	YI3										YR10	YR8	YR4
K	YI12	YI10	YI9										YR6	YR5	YR1
J	YI13	YI11	YI14										YR2	YR3	YR0
H	AR15	AR14	YI15										$\overline{SOBFP}$	$\overline{EOPSS}$	VDD
G	AR13	AI13	AI14										WTB1	WTA0	WTB0
F	AI15	CONX	CONY										CLK	MBFP	WTA1
E	ROUND	SFTR2	SFTR0										GWR0	OSEL0	OSEL1
D	$\overline{OEI}$	WTOUT0	PI14	Location pin									GWR3	SFTA1	$\overline{OER}$
C	SFTR1	PI15	PI13	PI12	PI9	VDD	PI0	PR0	VDD	PR7	PR11	VDD	GWR4	GWR1	SFTA0
B	WTOUT1	GND	PI10	PI8	GND	PI4	PI3	PR2	GND	PR6	PR9	PR12	GND	PR15	GWR2
A	VDD	PI11	PI7	PI6	PI5	PI2	PI1	PR1	PR3	PR4	PR5	PR8	PR10	PR13	PR14

Fig.3 Pin Allocation Diagram (Bottom View)
144 pin PGA - AC144

GG	SIG	GG	SIG	GG	SIG	GG	SIG
1	PI14	37	XI1	73	GND	109	GND
2	PI15	38	XI2	74	VDD	110	VDD
3	WTOUT1	39	XI3	75	YR12	111	PR13
4	WTOUT0	40	XI4	76	YR11	112	PR12
5	SFTR0	41	XI5	77	YR10	113	PR11
6	SFTR1	42	XI6	78	YR9	114	PR10
7	SFTR2	43	XI7	79	YR8	115	PR9
8	$\overline{\text{OEI}}$	44	XI8	80	YR7	116	PR8
9	CONX	45	XI9	81	YR6	117	PR7
10	CONY	46	XI10	82	YR5	118	PR6
11	ROUND	47	XI11	83	YR4	119	PR5
12	AI13	48	XI12	84	YR3	120	GND
13	AI14	49	XI13	85	YR2	121	VDD
14	AI15	50	XI14	86	YR1	122	PR4
15	AR13	51	XI15	87	YR0	123	PR3
16	AR14	52	$\overline{\text{CEY}}$	88	$\overline{\text{EOPSS}}$	124	PR2
17	AR15	53	$\overline{\text{CEX}}$	89	VDD	125	PR1
18	YI15	54	XR15	90	$\overline{\text{SOBFP}}$	126	PR0
19	YI14	55	XR14	91	WTB1	127	PI0
20	YI13	56	XR13	92	WTB0	128	PI1
21	YI12	57	XR12	93	WTA1	129	PI2
22	YI11	58	XR11	94	WTA0	130	PI3
23	YI10	59	XR10	95	MBFP	131	PI4
24	YI9	60	XR9	96	CLK	132	VDD
25	YI8	61	XR8	97	OSEL1	133	PI5
26	YI7	62	XR7	98	OSEL0	134	GND
27	YI6	63	XR6	99	$\overline{\text{OER}}$	135	PI6
28	YI5	64	XR5	100	SFTA0	136	PI7
29	YI4	65	XR4	101	SFTA1	137	PI8
30	YI3	66	XR3	102	GWR0	138	PI9
31	YI2	67	XR2	103	GWR1	139	PI10
32	YI1	68	XR1	104	GWR2	140	PI11
33	YI0	69	XR0	105	GWR3	141	PI12
34	XI0	70	YR15	106	GWR4	142	PI13
35	GND	71	YR14	107	PR15	143	GND
36	VDD	72	YR13	108	PR14	144	VDD

All GND and VDD pins must be used.

Fig.3A Pin Allocation Diagram - 144 pin ceramic QFP - GG144

NORMAL MODE OPERATION

When the MBFP mode select input is held low the 'Normal' mode of operation is selected. This mode supports all Complex Multiply operations that do not require Block Floating Point arithmetic.

Multiplier Satge

Complex two's complement fractional data is loaded into the X and Y input registers via the X and Y Ports on the rising edge of CLK. The Real and Imaginary components of the fractional data are each assumed to have the following format

BIT NUMBER	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
WEIGHTING	S	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-8}	2^{-9}	2^{-10}	2^{-11}	2^{-12}	2^{-13}	2^{-14}	2^{-15}

Where S = sign bit which has an effective weighting -2^0

The value of the 16 bit two's complement word is

Value = $(-1xS)+(bit14x2^{-1})+(bit13x2^{-2})+(bit12x2^{-3})$. . .

The X & Y port registers are individually enabled by the $\overline{CEX}$ & $\overline{CEY}$ signals respectvely. If the registers are required to be permanently enabled, then these signals may be tied to ground. On each clock cycle the contents of the input registers are passed to the four multipliers to start a new Complex Multiply operation. Each Complex Multiply operation requires four partial products (Xr x Yr), (Xr x Yi), (Xi x Yr), (Xi x Yi), all of which are calculated in parallel by the four 16 x 16 Multipliers. Only one clock cycle is required to complete the multiply stage before the Mutliplier results are loaded into the Multiplier output registers for passing on to the Adder/Subtractors in the next cycle. Each multiplier produces a 31 bit result with the duplicate sign bit eliminated. The format of the output data from the Multipliers is

BIT NUMBER	30	29	28	27	26	25	24	...	7	6	5	4	3	2	1	0
WEIGHTING	S	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	...	2^{-23}	2^{-24}	2^{-25}	2^{-26}	2^{-27}	2^{-28}	2^{-29}	2^{-30}

The effective weighting of the sign bit is -2^0

Result Correction

Due to the nature of the fraction twos complement representation it is possible to represent -1 exactly but not 1. With conventional multipliers this causes a problem when -1 is multiplied by -1 as the multiplier produces an incorrect result. The PDSP16116 includes a trap to ensure that the most positive number (value = 1.2^{-30}), (hex = 7FFFFFFF) is subsituted for the incorrect result. The multiplier result is therefore always a (correct) fractional value.

Complex Conjugation

Either the X or Y input data may be complex conjugated by asserting the CONX or CONY signals respectively. Asserting either of these signals has the effect of inverting (multiplying by -1) the imaginary component of the respective input. Table 3 shows the effect of CONX and CONY on the X and Y inputs.

FUNCTION	OPERATION	CONX	CONY
X x Y	(XR+XI)x(YR+YI)	low	low
X x Conj Y	(XR+XI)x(YR-YI)	low	high
Conj X x Y	(XR-XI)x(YR+YI)	high	low
Invalid	Invalid	high	high

Table 3 Conjugate Functions

Adder / Subtractor Stage

The 31 bit Real and Imaginary results from the Multipliers are passed to two 32 bit Adder/Subtractors. The Adder calculates the imaginary result ((Xr x Yi) + (Xi x Yr)) and the Subtractor calculates the Real result ((Xr x Yr) = (Xi x Yi)). Each Adder/Subtractor produces a 32 bit result with the following format.

BIT NUMBER	31	30	29	28	27	26	...	8	7	6	5	4	3	2	1	0
WEIGHTING	S	2^{0}	2^{-1}	2^{-2}	2^{-3}	2^{-4}	...	2^{-22}	2^{-23}	2^{-24}	2^{-25}	2^{-26}	2^{-27}	2^{-28}	2^{-29}	2^{-30}

The effective weighting of the sign bit is -2^1

Rounding

The ROUND control when asserted rounds the most significant 16 bits of the full 32 bit result from the Adder/Subtractor. If the ROUND signal is active (High), then bit 16 is set to a one, rounding the most significant 16 bits of the Adder/Subractor result. (The least siginificant 16 bits are unaffected). Inserting a one ensures that the rounding error is never greater than 1LSB, and that no DC bias is introduced as a result of the rounding processes.

The format of the Rounded result is;

BIT NUMBER	31	30	29	28	27	...	18	17	16	15	14	13	...	2	1	0
WEIGHTING	S	2^{0}	2^{-1}	2^{-2}	2^{-3}	...	2^{-12}	2^{-13}	2^{-14}	2^{-15}	2^{-16}	2^{-17}	...	2^{-28}	2^{-29}	2^{-30}

← ROUNDED VALUE →← LBS's →

The effective weighting of the sign is -2^1

Shifter

Each of the two Adder/Subtractors are followed by Shifters controlled via the WTB control input. These shifters can each apply four different shifts, however the same shift is applied to both real and imaginary components. The four shift options are:

i) WTB1:0 = 11 Shift complex product one place to the left giving a shifter output format:

BIT NUMBER	31	30	29	28	27	26	25	...	7	6	5	4	3	2	1	0
WEIGHTING	S	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	...	2^{-24}	2^{-25}	2^{-26}	2^{-27}	2^{-28}	2^{-29}	2^{-30}	2^{-31}

The effective weighting of the sign bit is -2^0

ii) WTB1:0 = 00 No shift applied giving a shifter output format:

BIT NUMBER	31	30	29	28	27	26	...	8	7	6	5	4	3	2	1	0
WEIGHTING	S	2^{0}	2^{-1}	2^{-2}	2^{-3}	2^{-4}	...	2^{-22}	2^{-23}	2^{-24}	2^{-25}	2^{-26}	2^{-27}	2^{-28}	2^{-29}	2^{-30}

The effective weighting of the shift bit is -2^{1}.

iii) WTB1:0 = 01 Shift complex product one place to the right giving a shifter output format:

BIT NUMBER	31	30	29	28	27	26	25	24	...	6	5	4	3	2	1	0
WEIGHTING	S	2^{1}	2^{0}	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	...	2^{-23}	2^{-24}	2^{-25}	2^{-26}	2^{-27}	2^{-28}	2^{-29}

The effective weighting of the sign bit is -2^{2}.

iv) WTB1:0 = 10 Shift complex product two places to the right giving a shifter output format:

BIT NUMBER	31	30	29	28	27	26	26	24	...	6	5	4	3	2	1	0
WEIGHTING	S	2^{2}	2^{1}	2^{0}	2^{-1}	2^{-2}	2^{-3}	2^{-4}	...	2^{-22}	2^{-23}	2^{-24}	2^{-25}	2^{-26}	2^{-27}	2^{-28}

The effective weighting of the sign bit is -2^{3}.

Overflow

If the left shift option is selected and the Adder/Subtractor contain a 32 bit word, then an invalid result will be passed to the output. An invalid output arising from this combination of events will be flagged by the SFTA0 flag output. The SFTA0 Flag will go high if either the real or imaginary reslut is invalid.

Output Select

The output from the Shifters is passed to the Output Select Mux, which is controlled via the OSEL inputs. These inputs are not registered and hence allow the output combination to be changed within each cycle. The full complex 64 bit result from the multiplier may therefore be output within a single cycle. The OSEL control selects four different output combinations as as summarised in Table 4.

OSEL1	OSEL0	PR	PI
0	0	MSR	MSI
0	1	LSR	LSI
1	0	MSR	LSR
1	1	MSI	LSI

Table 3 Output Selection

(Where MSR and LSR are the most and least siginificant 16 bit words of the Real Shifter output, MSI and LSI are the most and least significant 16 bit words of the imaginary Shifter output).

The output select options allow two different modes for extracting the full 32 bit result from the PDSP16116. The first mode treats the two 16 bit outputs as real and imaginary ports allowing the real and imaginary results to be output in two halves on the real and imaginary output ports. The second mode treats the two 16 bit outputs as one 32 bit output and allows the real and imaginary results to be output as 32 bit words.

PIN DESCRIPTIONS

XR, XI, YR, YI

Data inputs 16 bits: Data is loaded into the input registers from these ports on the rising edge of CLK. The data format is Twos Complement Fractional, where the MSB (sign bit) is bit 15. In normal mode the weighting of the MSB is -20 ie -1.

PR, PI

Data outputs 16 bits: Data is clocked into the output registers and passed to the PR and PI outputs on the rising edge of CLK. The data format is Twos Complement Fractional. The field of the internal result selected for output via PR and PI is controlled by signals OSEL1:0 (see Table 4).

CLK

Common Clock to all internal register.

$\overline{\text{CEX}}$, $\overline{\text{CEY}}$

Clock enables for X and Y input ports: When low these inputs enable the CLK signal to the X or Y input registers allowing new data to be clocked into the Multiplier.

CONX, CONY

If either of these inputs are high on the rising edge of CLK, then the data in the associated input has its imaginary component inverted (multiplied by -1), see Table 3. CONX and CONY affect data input on the same clock rising edge.

ROUND

The ROUND control is used to round the most siginficant 16 bits of the Adder/Subtractor result prior to being passed to the output register. The rounding operation takes place one cycle after the ROUND input is taken high. The ROUND input is not latched and is intended to be tied high or low depending upon the application.

MBFP

Mode select: When high, Block Floating Point (BFP) mode is selected. This allows the device to maintain the dynamic range of the data using a series of word tags. This is especially useful in FFT appllications. When low, the chip operates in normal mode for more general applications. This pin is intended to be tied high or low, depending on application.

$\overline{SOBFP}$ (BFP MODE ONLY)

Start of BFP: This input should be held low for the first cycle of the first pass of the BFP calculations (see Fig.7). It serves to reset the internal registers associated with BFP control. When operating in normal mode this input should be tied low.

$\overline{EOPSS}$ (BFP MODE ONLY)

End of pass: This input should be held low for the last cycle of each pass and for the lay time between passes. It instructs the control logic to update the value of the global weighting register and prepare the BFP circuitry for the next pass. When operating in normal mode this input should be tied low.

AR15:13 (BFP MODE ONLY)

Three Msbs of the real part of the A-word : These are used in the FFT butterfly application to deteremine the magnitude of the real part of the A-word and, hence, to determine if there will be any chage of word growth in the PDSP16318 Complex Accumulator. When operating in normal mode, these inputs are not used and may be tied low.

AI15:13 (BFP MODE ONLY)

Three Msbs of the imaginary part of the A-word : used in the same fashion as AR.

SFTR2:0 (BFP MODE ONLY)

Accumulator result shift control. These pins should be linked directly to the S2:0 pins on the PDSP16318 Complex Accumulator. They control the accumulator's barrel shifter (see Table 5). The purpose of this shift is to minimise sign extension in the multiplier or accumulator ALU's. When operating in normal mode, these output are superfluous.

SFTR2:0	FUNCTION
0 0 0	Reserved
0 0 1	Reserved
0 1 0	Reserved
0 1 1	Shift right by one
1 0 0	No shift
1 0 1	Shift left by one
1 1 0	Shift left by two
1 1 1	Reserved

Table 5 Auccumulator Shifts (BFP mode)

GWR4:0 (BFP MODE ONLY)

Contents of the global weighting register: This stores the weighting of the largest word present with respect to the weighting of the original input words. Hence, if the contents of the GWR are 00010, this indicates that the largest word currently being processed has its binary point two bits to the right of the original data at the start of the BFP calculations. The contents of this register are updated at the end of each pass, according to the largest value of WTOUT occuring during that pass. (i.e. If WTOUT = 11, then GWR will be increased by 2). The GWR is presented in two's complement format. These outputs are superfluous in normal mode.

WTOUT1:0 (BFP MODE ONLY)

Word tag output. This tag records the weighting of the output words from the current cycle relative to the current global weighting register (see Table 6). It should be stored along with the A' and B' words as it will form the input word tags, WTA and WTB, for each complex word during the next pass. These outputs are superfluous in normal mode.

WTOUT1:0	Weighting of the output relative to the current global weighting register
0 0	One less
0 1	The same
1 0	One more
1 1	Two more

Table 6 Word Tag Weightings

WTA1:0 (BFP MODE ONLY)

Word tag from the A-word. This word records the weighting of the A-word relative to the global weighting register on the previous pass. Although the A-word inself is not processed in the PDSP16116, this information is required by the control logic for the radix-2 butterfly FFT application. These inputs should be tied low in normal mode.

WTB1:0 (BFP & NORMAL MODES)

In BFP mode, this is the word tag from the B-word. This is operated in the same manner as WTA but for the B-word. The value of the word tags are used to ensure that the binary weighting of the A word and the product of the complex multiplier are the same at the inputs to the complex accumulator. Depending on which word is the larger, the weighting adjustment is performed using either the internal shifter or an external shifter controlled by SFTA. The word tags are also used to maintain the weighting of the final result to within plus two and minus one binary points relative to the new GWR. (On the first pass all word tags will be ignored).

In normal mode, these inputs perform a different function. They directly control the internal shifter at the output port as shown in Table 7.

WTB1:0	FUNCTION
11	shift complex product one place to the left
00	no shift applied
01	shift complex product one place to the right
10	shift complex product two places to the right

Table 7 Normal Mode Shift Control

SFTA1:0 (BFP & NORMAL MODES)

In BFP mode, these signals act as as the A-word shift control. They allow shifting from one to four places to the right, see Table 8. Depending on the relative weightings of the A-words and the complex product, the A-word may have to be shifted to the right to ensure compatible weightings at the inputs to the PDSP16318 complex accumulator. (The two words must have the same weighting if they are to be added).

In normal mode, SFTA0 performs a different a different function. If WTB1:0 is set to implement a left shift, then overflow will occur if the data is fully 32 bits wide. This pin is used to flag such an overflow. SFTA1 is not used in normal mode.

WTB1:0	FUNCTION
0 0	Shift A-word 1 places to the right
0 1	Shift A-word 2 places to the right
1 0	Shift A-word 3 places to the right
1 1	Shift A-word 4 places to the right

Table 8 External A-word shift control

OSEL1:0

The outputs from the device are selected by the OSEL0 & OSEL1 instruction bits. These controls allow selection of the output combination during the current cycle. (They are not registered). These are four possible output configurations that allow either complex outputs of the most or least significant bytes, or real or imaginary outputs of the full 32 bit word (see Table 4). OSEL0 and OSEL1 should both be tied low when in BFP mode.

BFP MODE FFT APPLICATION

The PDSP16116 may be used as the main arithmetic unit of the butterfly processor which will allow the following FFT benchmarks:

1024 point complex radix-2 transform in 517us
512 point complex radix-2 transform in 235us
256 point complex radix-2 transform in 106us

In addition, with pin MBFP tied high, the BFP circuitry within the PDSP16116 can be used to adaptively rescale data throughout the course of the FFT so as to give high-resolution results.

The BFP system on the PDSP16116 can be used with any variation of the Radix-2 Decimation-In-Time FFT - e.g. the Constant Geometry algorithm, the In-Place algorithm etc. An N-point Radix-2 DIT FFT is split into log (N) passes. Each pass consists of N/2 'butterflies', each performing the operation:

A' = A + B.W
B' = A - B.W

Where W is the complex coefficient and A & B are the complex data.

Fig.4 illustrates how a single PDSP16116 may be combined with two PDSP1601's and two PDSP16318's to form a complete BFP butterfly processor. The PDSP16318's are used to perform the complex addition and subtraction of the butterfly operation, while the PDSP1601's are used to match the data path of the A-word to the pipelining and shifting operations within the PDSP16116.

For more information on the theory and construction of this butterfly processor, refer to application note AN59.

BFP MODE OPERATION

The BFP mode on the PDSP16116 is intended for use in the FFT application described above. i.e. it is intended to prevent data degredation during the course of an FFT calculation. The operation of the PDSP16116 based BFP butterfly processor (see Fig.4) is described below.

The Block Floating Point System

A block floating point system is essentially an ordinary integer arithmetic system with some clever logic bolted on. The object of the extra logic is to lend the system some of the enormous dynamic range afforded by a true floating point system without suffering the corresponding loss in performance.

The initial data used by the FFT should all have the same binary arithmetic weighting. i.e. the binary point should occupy the same position in every data word, as is normal in integer arithmetic. However, during the course of the FFT, a variety of weightings are used in the data words to increase the dynamic range available. This situation is similar to that within a true floating point system, though the range of numbers representable is more limited. In the BFP system used in the PDSP16116, there are, within any one pass of the FFT, four possible positions of the binary point wihin the integer words. To record the position of its binary point, each word has a 2-bit word tag associated with it. By way of example, in a particular pass we may have the following four positions of binary point avaiable, each denoted by a certain value of word tag:

XX.XXXXXXXXXXXX	word tag = 00
XXX.XXXXXXXXXXX	word tag = 01
XXXX.XXXXXXXXXX	word tag = 10
XXXXX.XXXXXXXXX	word tag = 11

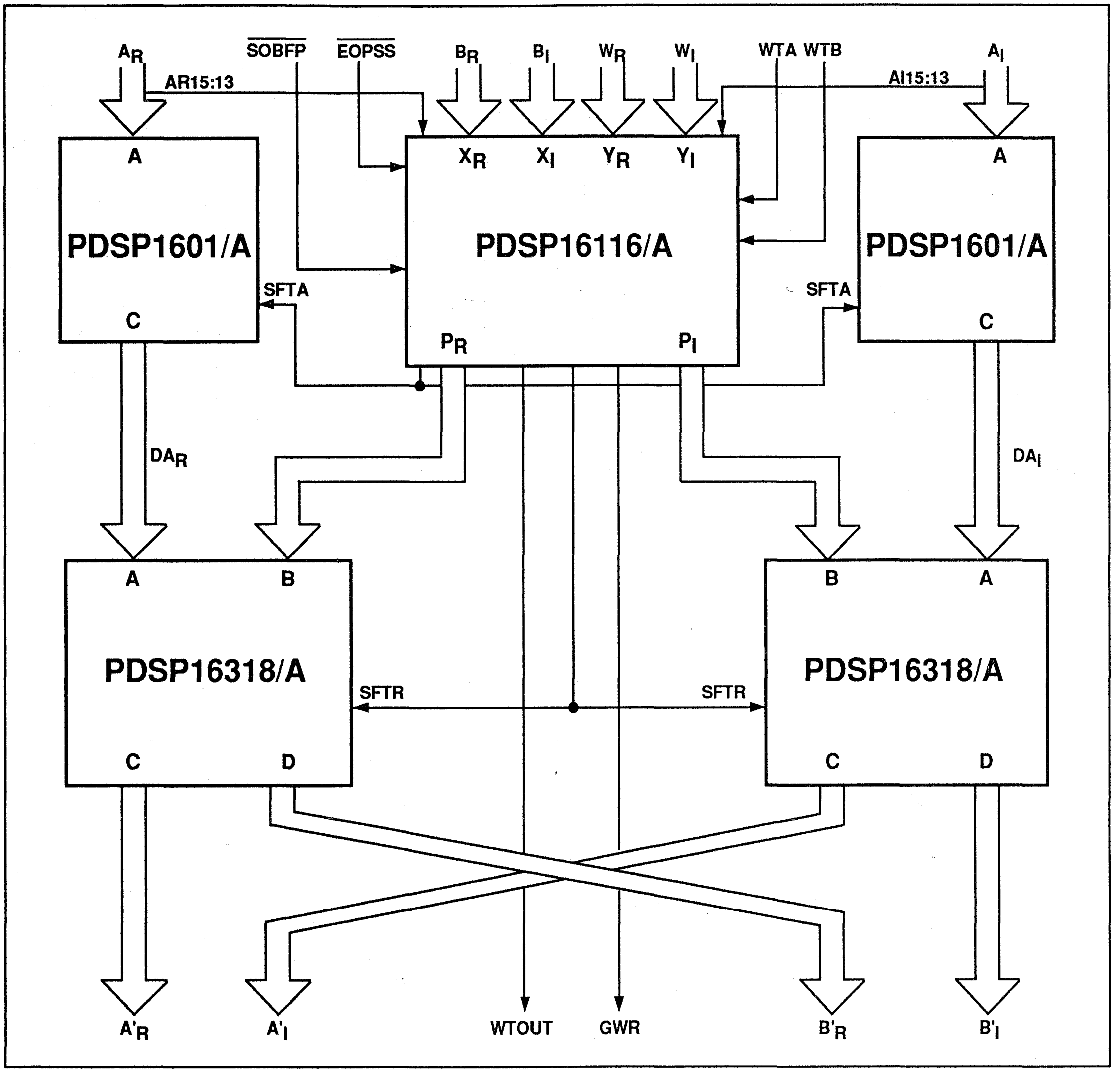

Fig.4 FFT Butterfly Processor

At the end of each constituent pass of the FFT, the positions of the binary point supported may change to reflect the trend of data increase or decreases in magnitude. Hence, in the pass following that of the above example, the four positions of binary point supported may be change to:

XX.XXXXXXXXXXXX	word tag = 00
XXX.XXXXXXXXXXX	word tag = 01
XXXX.XXXXXXXXXX	word tag = 10
XXXXX.XXXXXXXXX	word tag = 11

This variation in the range of binary points supported from pass to pass (i.e. the movement of the binary point relative to its position in the original data) is recorded in the GWR.

Thus we can determine the position of the binary point relative to its initial position by modifying the value of GWR by WTOUT for a given word as shown in Table 6.

As an example, if GWR=01001 and WTOUT=10 then the binary point has moved 10 places to the right of its original position.

The butterfly operation

The butterfly operation is the arithmetic operation which is repeated many times to produce an FFT. The PDSP16116A based butterfly processor performs this operation in a low power high accuracy chip set.

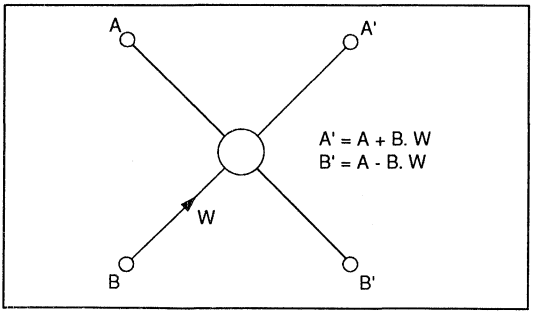

Fig.5 Butterfly Operation

A new butterfly operation is commenced each cycle, requiring a new set of data for , B, W, WTA and WTB. Five cycles later, the corresponding results A' and B' are produced along with their associated WTOUT. In between, the signals SFTA and SFTR are produced and acted upon by the shifters in the PDSP1601/A and PDSP16318/A. The timing of the data and control signals is shown in Fig.6.

The results (A' and B') of each butterfly calculation in a pass must be stored away to be used later as the input data (A and B) in the next pass. Each result must be stored together with its associated word tag, WTOUT. Although WTOUT is common to both A' and B', it must be stored separately with each word as the words are used on different cycles during the next pass. At the inputs, the word tag associated with the A word is known as WTA and the word tag associated with the B word is known as WTB. Hence, the WTOUTs from one pass will become the WTAs and WTBs for the following pass. It should be noted that the first pass is unique in that word tags need not be input into the butterfly as all data initially has the same weighting. Hence, during the first pass alone, the inputs WTA and WTB are ignored.

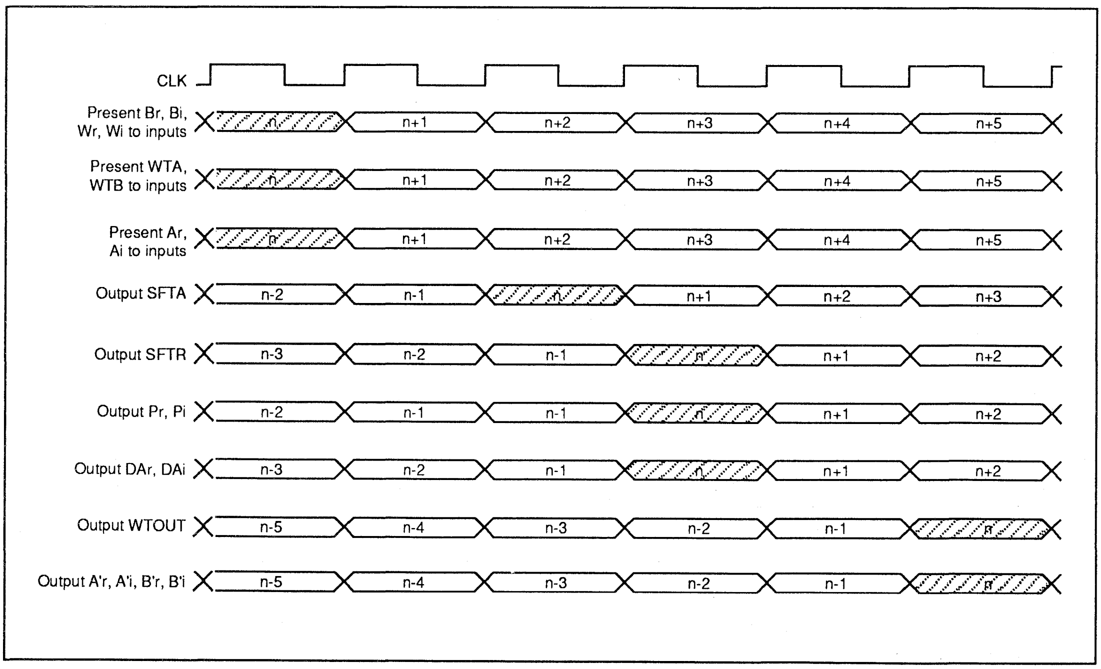

Fig.6 Butterfly Data and Control Signals

Control of the FFT

To enable the block floating point hardware to keep track of the data, the following signals are provided :
$\overline{\text{SOBFP}}$ - start of the FFT
$\overline{\text{EOPSS}}$ - end of current pass

These inform the PDSP16116/A when an FFT is starting and when each pass is complete. Fig.7 shows how these signals should be used and a commentary is provided below.

To commence the FFT, the signal $\overline{\text{EOPSS}}$ should be set high (where it will remain for the duration of the pass). $\overline{\text{SOBFP}}$ should be pulled low during the initial cycle when the first data words A and B are presented to the inputs of the butterfly processor. The following cycle $\overline{\text{SOBFP}}$ must be pulled high where it should remain for the duration of the FFT. New data is presented to the processor each successive cycle until the end of the first pass of the FFT. On the last cycle of the pass, the signal $\overline{\text{EOPSS}}$ should be pulled low and remain low for a minimum of five cycles *, the time required to clear the pipeline of the butterfly processor so that all the results from one pass are obtained before commencing the following pass. On the initial cycle of each new pass, the signal $\overline{\text{EOPSS}}$ should be pulled high and it should remain high until the final cycle of that pass, when it is pulled low again.

* Should a longer pause be required between passes - to arrange the data for the next pass, for example, then $\overline{\text{EOPSS}}$ may be kept low as long as necessary - the next pass cannot commence until it is brought high again.

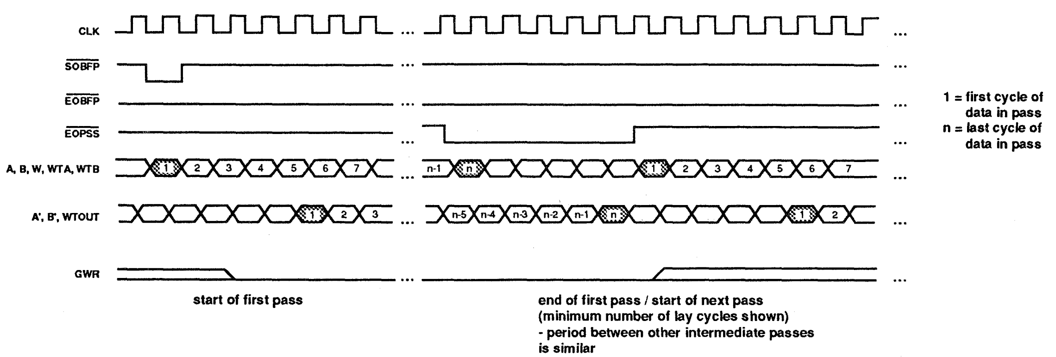

Fig.7 Use of the BFP Control Signals

FFT Output Normalisation

When an FFT system outputs a series of FFT results for display, storage or transmission, it is essential that all results are compatible, i.e. with the binary point in the same position. However, in order to preserve the dynamic range of the data in the FFT calculation, the PDSP1601/A employs a range of different weightings. Therefore, data must be re-formatted at the end of the FFT to be pre-determined common weighting. This can be done by comparing the exponent of given data word with the pre-determined unversial exponent and then shifting the data word by the difference. The PDSP1601/A, with its multifunction 16 bit barrel shifter, is ideally suited to this task.

What value should the Unversal Exponent take? Well, according to theory, the largest possible data result from an FFT is N times the largest input data. This means that the binary point can move a maximum of log2(N) places to the right. Hence, if we choose the Unverisal Exponent to be log2(N) this should give us sufficient range to represent all data points faithfully.

In practice, data output may never approach the theoretical maximum. Hence, it may be worthwhile to try various Unverisal Exponents and choose the one best suited to the particular application.

Data is output from the butterfly processor with a two-part exponent: the 5-bit GWR applicable to all data words from a given FFT and a 2-bit WTOUT associated with each individual data word. To find the complete exponent for a given word, the GWR for that FFT must be modified by its WTOUT as shown in Table 6. The result is the number of places the binary point has shifted to the right during the course of the FFT.

This value must be compared with the Unversial Exponent to determine the shift required. This is done by subtracting it from the Unversial Exponent. The number of places to be shifted is equal to the difference between the two exponents. The shift can be implemented in a PDSP1601/A. The shift value is fed into the SV port.

As FFT data consists of real and imaginary parts, either two PDSP1601As must be used (controlled by the same logic) or a single PDSP1601/A could be used handling real and imaginary data on alternate cycles (using the same instructions for both cycles).

An example of an output normalisation circuit is shown in Fig.8. Only 4 bit data paths are used in calculating the shift. This means that we must be able to trap very small values negative of GWR and force a 15-bit right shift in such cases.

N.B.

It is easier to simply add the word tag to the exponent for the purpose of determing the shift required, instead of modifying it according to Table.6. To compensate for this, the Universal Exponent may be increased by one.

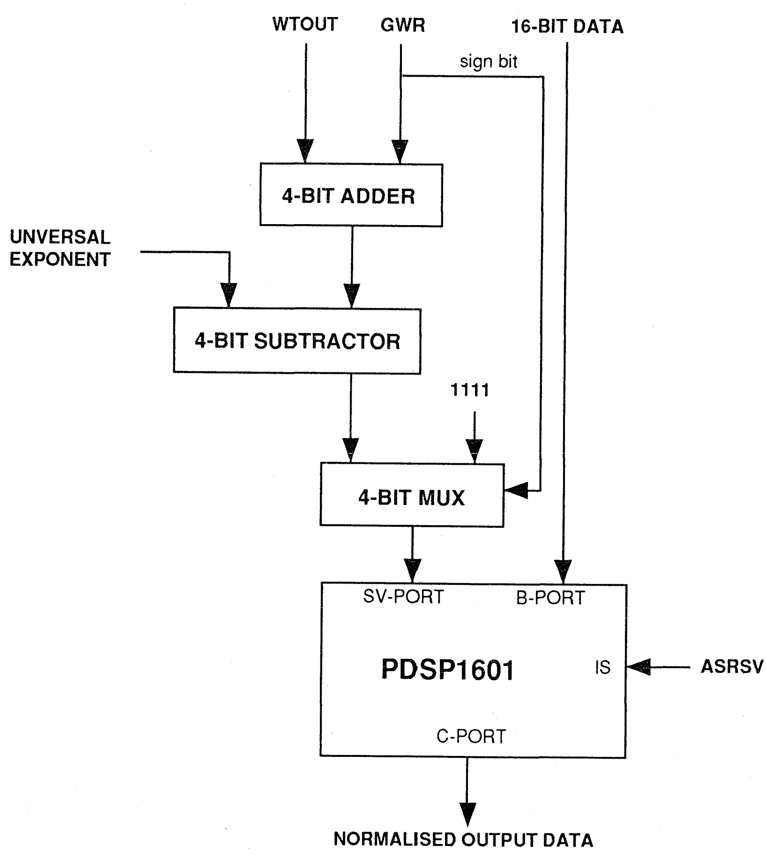

Fig.8 Output Normalisation Circuitry

ABSOLUTE MAXIMUM RATINGS (Note 1)

Supply voltage V_{CC}	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to V_{CC} +0.5V
Output voltage V_{OUT}	-0.5V to V_{CC} +0.5V
Clamp diode current per I_k (see note 2)	18mA
Static discharge voltage (HBM)	500V
Storage temperature range T_S	-65°C to +150°C
Ambient temperature with power applied T_{AMB}	
Military	-55°C to +125°C
Industrial	-40°C to +85°C
Junction temperature	150°C
Package power dissipation	1000mW
Thermal resistances	
Junction to case θ_{JC}	12°C/W
Junction to case θ_{JA}	29°C/W

NOTES
1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceedeed, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliability.

ELECTRICAL CHARACTERISTICS

Operating conditions (unless otherwise stated):

Industrial: T_{AMB} = -40°C to +85°C, V_{CC} = 5.0V ± 10%, GND = 0V

Military: T_{AMB} = -55°C to +125°C, V_{CC} = 5.0V ± 10%, GND = 0V

Static Characteristics

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Min.		
Output high voltage	V_{OH}	2.4		-	V	I_{OH} = 8mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -8mA
Input high voltage	V_{IH}	3.0		-	V	CLK input only
Input high voltage	V_{IH}	2.2		-	V	All other inputs
Input low voltage	V_{IL}	-		0.8	V	GND <V_{IN}<V_{CC}
Input leakage current	I_{IN}	-10		+10	µA	
Input capacitance	C_{IN}		10		pF	GND <V_{IN}<V_{CC}
Output leakage current	I_{OZ}	-50		+50	µA	V_{CC} = Max
Output S/C current	I_{OS}	10		300	mA	

Switching Characteristics

Characteristic	PDSP16116		PDSP16116A		Units	Conditions
	Min.	Max.	Min.	Max.		
CLK rising edge to P-PORTS	5	45	5	23	ns	2 x LSTTL + 20pF
CLK rising edge to WTOUT1:0	5	30	5	20	ns	2 x LSTTL + 20pF
CLK rising edge to GWR4:0	5	30	5	20	ns	2 x LSTTL + 20pF
CLK rising edge to SFTA1:0	5	60	5	30	ns	2 x LSTTL + 20pF
CLK rising edge to SFTR2:0	5	50	5	28	ns	2 x LSTTL + 20pF
Setup $\overline{\text{CEX}}$ or $\overline{\text{CEY}}$ to CLK rising edge	11	-	8	-	ns	
Hold $\overline{\text{CEX}}$ or $\overline{\text{CEY}}$ to CLK rising edge	-	0	-	0	ns	
Setup X or Y port inputs to CLK rising edge	11	-	8	-	ns	
Hold X or Y port inputs to CLK rising edge	-	2	-	0	ns	
Setup WTA1:0, WTB1:0, $\overline{\text{SOBFP}}$ or $\overline{\text{EOPSS}}$ inputs to CLK rising edge	14	-	8	-	ns	
Hold WTA1:0, WTB1:0, $\overline{\text{SOBFP}}$ or $\overline{\text{EOPSS}}$ inputs to CLK rising edge	-	0	-	0	ns	
Setup CONX or CONY inputs to CLK rising edge	14	-	8	-	ns	
Hold CONX or CONY inputs to CLK rising edge	-	0	-	0	ns	
Setup AR15:13 or AI15:13 to CLK rising edge	14	-	-	-	ns	
Hold AR15:13 or AI15:13 to CLK rising edge	-	0	-	0	ns	
OPSEL to valid P-PORTS	-	35	-	20	ns	2 x LSTTL + 20pF
$\overline{\text{OER}}$ or $\overline{\text{OEI}}$ rising PR-PORT or PI-PORT high to Z	-	35	-	25	ns	see Fig.9
$\overline{\text{OER}}$ or $\overline{\text{OEI}}$ rising PR-PORT or PI-PORT low to Z	-	45	-	25	ns	see Fig.9
$\overline{\text{OER}}$ or $\overline{\text{OEI}}$ falling PR-PORT or PI-PORT Z to high	-	22	-	18	ns	see Fig.9
$\overline{\text{OER}}$ or $\overline{\text{OEI}}$ falling PR-PORT or PI-PORT Z to low	-	24	-	18	ns	see Fig.9
Clock period	100	-	50	-	ns	
Clock high time	30	-	12	-	ns	
Clock low time	20	-	12	-	ns	
Vcc Current (CMOS input levels)	-	60	-	80	mA	see Note 4
Vcc Current (TTL input levels)	-	100	-	130	mA	see Note 4

NOTE 4 :- V_{CC} = Max Outputs unloaded, clock freq = Max

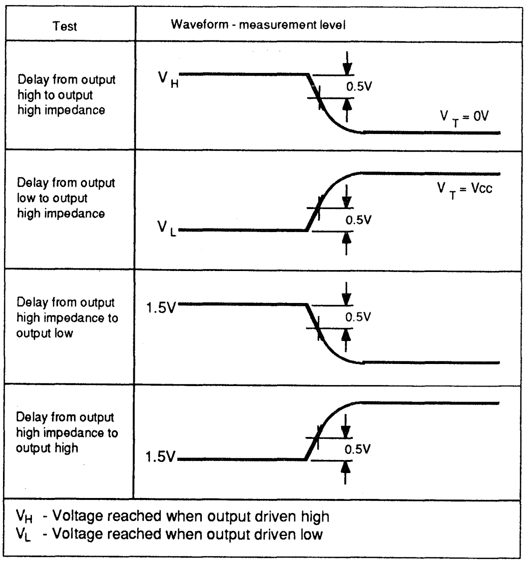

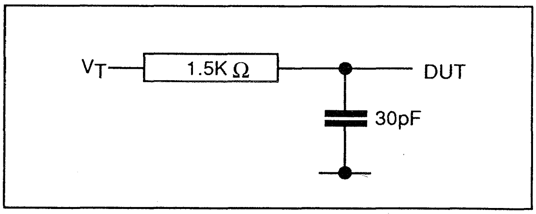

Fig.9 Three state delay measurement load

ORDERING INFORMATION

PDSP16116 C0 AC	**10MHz**	**Commercial**
PDSP16116 B0 AC	**10MHz**	**Industrial**
PDSP16116 A0 AC	**10MHz**	**Military**
PDSP16116 MC GGDR	**10MHz**	**MIL-883 screened**
PDSP16116A C0 AC	**20MHz**	**Commerical**
PDSP16116A B0 AC	**20MHz**	**Industrial**
PDSP16116A A0 AC	**20MHz**	**Military**
PDSP16116A C0 GG	**20MHz**	**Commerical**
PDSP16116A B0 GG	**20MHz**	**Industrial**
PDSP16116A A0 GG	**20MHz**	**Military**
PDSP16116A MC GGDR	**20MHz**	**MIL-883 screened**

PDSP16256 / A

PROGRAMMABLE FIR FILTER

(Supersedes January 1991 (or later) Edition)

The PDSP16256 contains sixteen multiplier - accumulators, which can be multi cycled to provide from 16 to 128 stages of digital filtering. It accepts 16 bit data and coefficients, and accumulates results upto 32 bits.

In 16 tap mode the device samples data at the 25MHz system clock rate. If a lower sample rate is acceptable then the number of stages can be increased in powers of two upto a maximum of 128. Each time the number of stages is doubled, the sample clock rate must be halved with respect to the system clock. With 128 stages the sample clock is therefore one eighth of the system clock.

In all speed modes devices can be cascaded to provide filters of any length, only limited by the possibility of accumulator overflow. The 32 bit results are passed between cascaded devices without any intermediate scaling and subsequent loss of precision.

The device can be configured as either, one long filter, or two separate filters with half the number of taps in each. Both networks can have independent inputs and outputs.

Both single and cascaded devices can be operated in decimate by two mode. The output rate is then half the input rate, but twice the number of stages are possible at a given sample rate. A single device with a 20MHz clock would then, for example, provide a 128 stage low pass filter, with a 5MHz input rate and 2.5MHz output rate.

Coefficients are stored internally and can be down loaded from a host system or an EPROM. The latter requires no additional support, and is used in stand alone applications. A full set of coefficients is then automatically loaded at power on, or at the request of the system. A single EPROM can be used to provide coefficients for upto 16 devices.

FEATURES

- Sixteen MACs in a single device
- Basic mode is 16 tap filter with 25MHz sample rates
- 16 bit data and 32 bit accumulators
- Programmable to give up to 128 taps with sampling rates proportionally reducing to 3.13MHz
- Can be configured as one long filter or two half length filters
- Decimate by two option will double the filter length
- Coefficients supplied from a host system or a local EPROM
- Advanced 144 PGA package with integral ground and supply Splanes

APPLICATIONS

- High Performance Digital Filters
- Pulse Compression for Radar & Sonar
- Matrix Multiplication
- Correlation

ASSOCIATED PRODUCTS

PDSP16350 I/Q Splitter / NCO

PDSP16510 FFT Processor

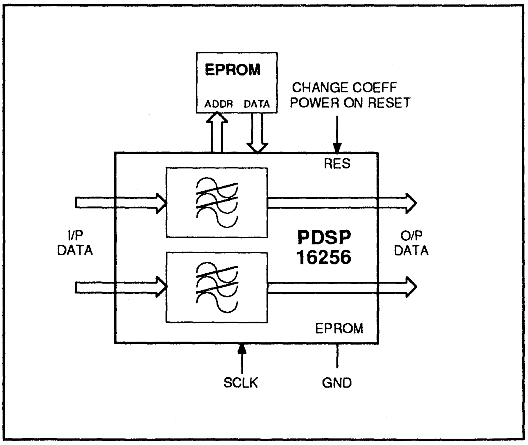

Fig. 1 Dual Filter

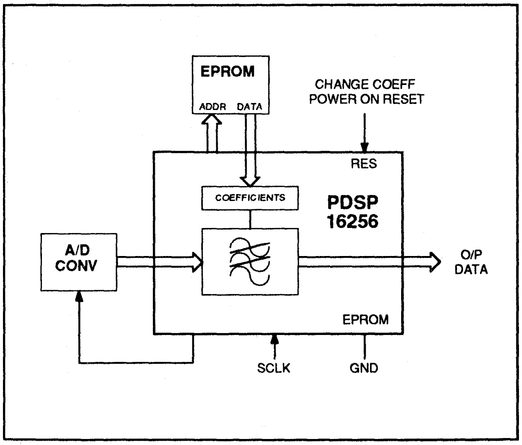

Fig. 2 Typical System Application

SIGNAL	DESCRIPTION
DA15:0	16 bit data input bus to Network A.
DB15:0	Delayed data output bus in the single filter mode. Connected to the data input bus of the next device in a cascaded chain. Input to Network B in the dual filter modes.
X31:0	Expansion input bus in the single filter mode. Connected to the previous filter output in a cascaded chain. The inputs are not used on a single device system or on the Termination device in a cascaded chain. The output from Network B in the dual modes.
F31:0	In single filter mode this bus holds the main device output. In dual mode it holds the output from Network A.
FEN	Filter enable. The first high present on an SCLK rising edge defines the first data sample. The signal must stay active whilst valid data is being received.
DFEN	Delayed filter enable. This output is connected to the Filter Enable input of the next device in a cascaded chain, when moving towards the termination device. It is used to coordinate the control logic within each device.
SWAP	Selects either the upper or lower set of coefficients for Bank Swap. A low selects the lower bank, a high the upper bank.
FRUN	When high this signal allows continuous filter operations to occur without the need for the initial FEN edge. If the device is not a single or interface device then this pin must be tied low.
$\overline{\text{DCLR}}$	A low on this signal on the SCLK rising edge will clear all the internal accumulators. DCLR need only remain low for a single cycle, signal BUSY will indicate when the internal clearing is complete. After a clear the device must be re-synchronised to the data stream using FEN. It is recommended the FEN is taken low at the same time as clear. FEN may then be taken high to synchronise the data stream once BUSY has returned low.
C15:0	16 bit coefficient input bus. In the Byte mode of operation, C15:8 have alternative uses as explained in the text.
A7:0	Coefficient address bus. In the EPROM mode A7:0 are address outputs for an EPROM. In the remote host mode they are inputs from the host. A7 is not used when coefficients are loaded as 16 bit words.
CCS	This pin is similar in operation to A7:0 and provides a higher order address bit. When low the coefficients are loaded, when high the control register is loaded.
$\overline{\text{WEN}}$	In the remote mode this pin is an input which when low enables the load operation. In the EPROM mode it is an output which provides the write enable for other slave devices.
$\overline{\text{CS}}$	This pin is always an input and must also be low for the internal write operation to occur.
$\overline{\text{BYTE}}$	When this pin is tied low, coefficients are loaded as two bytes. When the pin is high they are loaded as 16 bit words. In the EPROM mode this pin is ignored.
$\overline{\text{EPROM}}$	When this pin is tied low coefficients are loaded as bytes from an external EPROM. The device outputs an address on A7:0. When the pin is high coefficients must be loaded from a remote master. They can then be transferred individually rather than as a complete set.
SCLK	The main system clock, all operations are synchronous with this clock. The clock rate must be either 1, 2, 4, or 8 times the required data sampling rate. The factor used depends on the required filter length.
CLKOP	This output when used to enable SCLK can provide a data sampling clock. It has the effect of dividing the SCLK rate by 1, 2, 4 or 8 depending on the filter mode selected.
$\overline{\text{OEN}}$	Tri-state enable for the F bus. When high the outputs will be high impedance. OEN is registered onto the device and does not therefore take effect until the first SCLK rising edge

SIGNAL	DESCRIPTION
BUSY	A high on this signal indicates that the device is completing internal operations and is not yet able to accept new data. The signal is used during automatic EPROM loading, reset and accumulator clearing.
$\overline{\text{RES}}$	When this pin is low the control logic and accumulators are reset. In the EPROM mode it will initiate a load sequence when it goes high.

NOTE unused busses (e.g. X31:0 when the device is configured in single or termination mode) can be set to any value. They should however be maintained at a valid logic level to avoid an increase in power consumption.

To ensure correct input voltage thresholds are maintained all the VDD and GND pins must be connected to adequate power and ground planes.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R	C14	C11	C8	C5	C2	DA15	VDD	DA12	GND	DA8	DA5	DA3	DA0	SWAP	$\overline{\text{DCLR}}$
P	C15	C12	C9	C7	C4	C1	DA14	DA11	DA9	DA7	DA4	DA1	CLKOP	FEN	F31
N	CCS	$\overline{\text{WEN}}$	C13	C10	C6	C3	C0	DA13	DA10	DA6	DA2	$\overline{\text{OEN}}$	DFEN	F29	F28
M	SCLK	CS	GND										F30	F27	F25
L	$\overline{\text{EPROM}}$	$\overline{\text{BYTE}}$	$\overline{\text{RES}}$										F26	F24	F22
K	A2	A1	A0										F23	F21	F19
J	GND	A3	A4										F20	F18	GND
H	A6	A5	A7										F17	F15	F16
G	VDD	DB0	DB2										F14	F13	VDD
F	DB1	DB3	DB5										F10	F11	F12
E	DB4	DB6	DB8										F6	F8	F9
D	DB7	DB9	DB12										F2	F5	F7
C	DB10	DB11	DB15	X1	X5	X9	X13	X16	X19	X22	X25	X29	FRUN	F3	F4
B	DB13	DB14	X2	X4	X7	X10	X12	X14	X17	X20	X23	X26	X28	X31	F1
A	BUSY	X0	X3	X6	X8	X11	VDD	X15	GND	X18	X21	X24	X27	X30	F0

Fig. 3A Device Pinout - Bottom view (144 pin PGA - AC144)

GC	SIG	GC	SIG	GC	SIG	GC	SIG
1	F0	44	SWAP	87	C15	130	GND
2	F1	45	GND	88	GND	131	BUSY
3	F2	46	$\overline{\text{OEN}}$	89	GND	132	X0
4	F3	47	CLKOP	90	$\overline{\text{WEN}}$	133	VDD
5	VDD	48	VDD	91	CCS	134	X1
6	F4	49	DA0	92	$\overline{\text{CS}}$	135	X2
7	F5	50	DA1	93	VDD	136	X3
8	GND	51	DA2	94	$\overline{\text{RES}}$	137	X4
9	F6	52	DA3	95	SCLK	138	X5
10	F7	53	DA4	96	GND	139	X6
11	F8	54	DA5	97	VDD	140	GND
12	F9	55	GND	98	$\overline{\text{BYTE}}$	141	X7
13	F10	56	DA6	99	$\overline{\text{EPROM}}$	142	X8
14	F11	57	DA7	100	A0	143	VDD
15	F12	58	DA8	101	A1	144	X9
16	GND	59	DA9	102	A2	145	X10
17	F13	60	VDD	103	A3	146	X11
18	F14	61	DA10	104	A4	147	X12
19	F15	62	DA11	105	VDD	148	X13
20	VDD	63	DA12	106	A5	149	X14
21	F16	64	DA13	107	A6	150	GND
22	F17	65	DA14	108	GND	151	X15
23	F18	66	DA15	109	A7	152	X16
24	F19	67	GND	110	DB0	153	X17
25	VDD	68	C0	111	DB1	154	X18
26	F20	69	C1	112	DB2	155	X19
27	F21	70	C2	113	GND	156	X20
28	GND	71	C3	114	DB3	157	X21
29	F22	72	C4	115	DB4	158	X22
30	F23	73	C5	116	DB5	159	GND
31	F24	74	VDD	117	DB6	160	X23
32	F25	75	C6	118	DB7	161	X24
33	F26	76	C7	119	VDD	162	X25
34	F27	77	C8	120	DB8	163	VDD
35	F28	78	C9	121	DB9	164	X26
36	GND	79	C10	122	DB10	165	X27
37	F29	80	GND	123	DB11	166	X28
38	F30	81	C11	124	DB12	167	X29
39	F31	82	C12	125	DB13	168	X30
40	VDD	83	C13	126	DB14	169	GND
41	FEN	84	VDD	127	GND	170	X31
42	DFEN	85	GND	128	DB15	171	VDD
43	$\overline{\text{DCLR}}$	86	C14	129	VDD	172	FRUN

Fig. 3B Device Pinout (172 pin QFP - GC172)

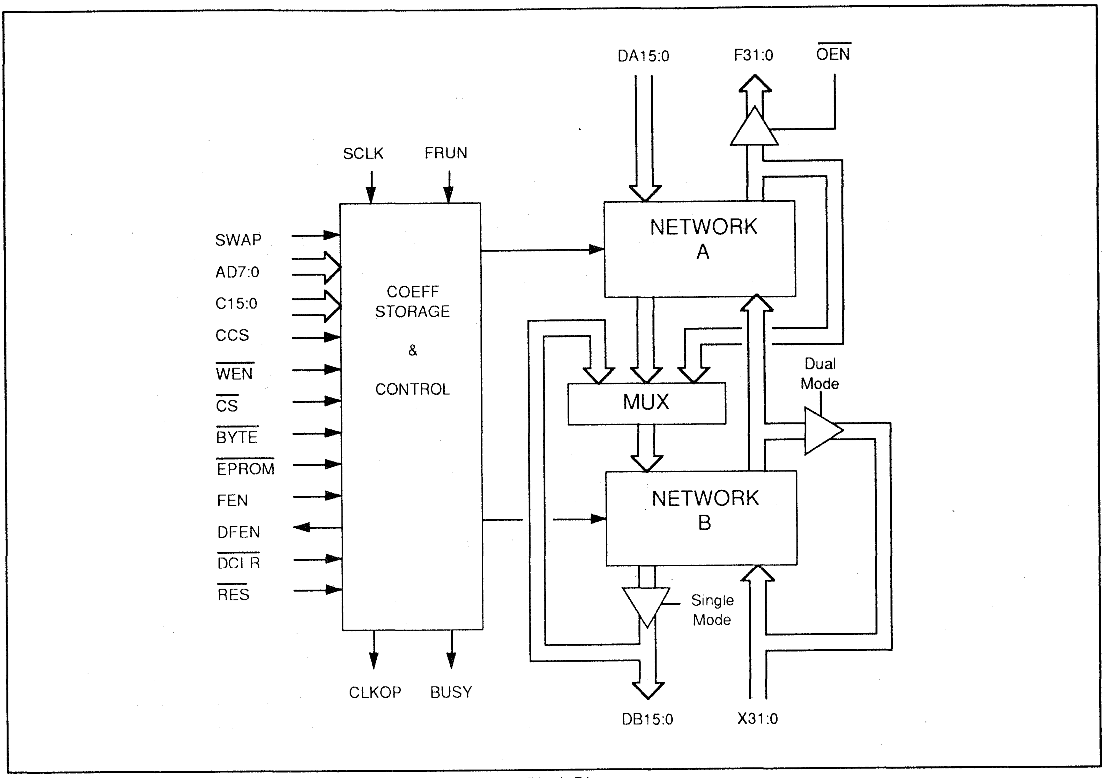

Fig. 4 Block Diagram

OPERATIONAL OVERVIEW

The PDSP16256 is an application specific FIR filter for use in high performance digital signal processing systems. Sampling rates can be upto 25MHz The device provides the filter function without any software development, and the options are simply selected by loading a control register. The device can be user configured as either a single filter, or as two separate filters. The latter can provide two independent filters for the in-phase and quadrature channels after IQ splitting, or can provide two filters in cascade for greater stop band rejection.

The device operates from a system clock, with rates up to 25MHz. This clock must be 1, 2, 4, or 8 times the required sampling frequency, with the higher multiplication rates producing longer filter networks at the expense of lower sampling rates. Devices can be connected in cascade to produce longer filter lengths. This can be accomplished without the need for any additional external data delays, and all the single device options remain available.

Continuous inputs are accepted, and continuous results produced after the internal pipeline delay. Connection can be made directly to an A/D converter. The filter operation can be synchronised to a Filter Enable signal whose active going edge marks the first data sample. The internal multiplier - accumulator array can be cleared with a dedicated input. This is necessary if erroneous results obtained during the normal data 'flush through' are not permissible.

Coefficients can be loaded from a host system using a conventional peripheral interface and separate data bus. Alternatively, they can be loaded as a complete set from a byte wide EPROM. The device produces addresses for the EPROM and a BUSY output indicates that the transfer is occurring. Up to sixteen devices can have their coefficients supplied from a single EPROM. These devices need not necessarily be part of the same filter network.

Each of the filter networks shown in Fig. 4 contains eight systolic multiplier accumulator stages, an example with four stages is shown in Fig. 5. Input data flows through the delay lines and is presented for multiplication with the required coefficient. This is added to either the last result from this accumulator or the result from the previous accumulator. The filter results progress along the adders at the data sample rate. If the sample rate equals SCLK divided by four, for example, then the accumulated result is passed onto the next stage every forth cycle. The structure described is highly efficient when used to calculate filtered results from continuous input data.

A comprehensive digital filter design program is available for PC compatible machines. This will optimise the filter coefficients for the filter type required and number of taps available at the selected sample rate within the PDSP16256 device. An EPROM file can be automatically generated in Motorola S-record format.

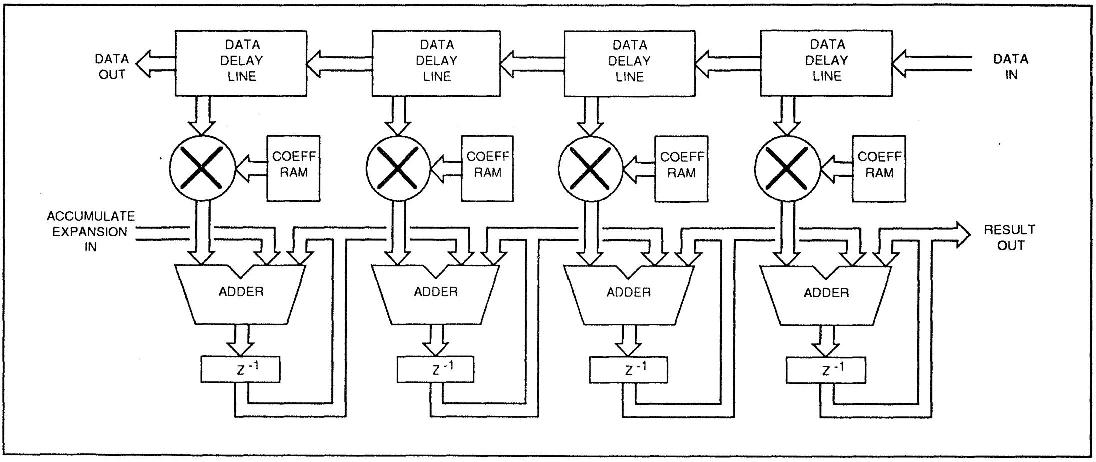

Fig. 5 Filter Network Diagram

SINGLE FILTER OPTIONS

When operating as a single filter the device accepts data on the 16 bit DA bus at the selected sample rate, see Figs 6 and 7. Results are presented on the 32 bit F bus, which may be tristated using the OEN input. Signal OEN is registered onto the device and does not therefore take effect until the first SCLK rising edge. Devices may be cascaded this allows filters with more taps than available from a single device. To accomplish this two further busses are utilised. The DB bus presents the input data to the next device in cascade after the appropriate delay, while, partial results are accepted on the X bus.

Single filter mode is selected by setting control register bit 15 to a one. The required filter length is then selected using control register bits 14 and 13 as summarised in Table 3. The options define the number of times each multiplier - accumulator is used per sample clock period. This can be once, twice, four times, or eight times.

In addition a normal/decimate bit (CR12) allows the filter length to be doubled at any sample rate. This is possible when the filter coefficients are selected to produce a low pass filter, since the filtered output would then not contain the higher frequency components present in the input. The Nyquist criterion, specifying that the sampling rate must be at least double the highest frequency component, can still then be satisfied even though the sampling rate has been halved.

CR 14 13 12	Input Rate	Output Rate	Filter Length	Setup Latency
0 0 0	SCLK	SCLK	16 Taps	16
0 0 1	SCLK	SCLK/2	32 Taps	17
0 1 0	SCLK/2	SCLK/2	32 Taps	16
0 1 1	SCLK/2	SCLK/4	64 Taps	18
1 0 0	SCLK/4	SCLK/4	64 Taps	20
1 0 1	SCLK/4	SCLK/8	128 Taps	24
1 1 0	SCLK/8	SCLK/8	128 Taps	24

Table 3. Single Filter Options

The system clock latency for a single device is shown in Table 3. This is defined as the delay from a particular data sample being available on the input pins to the first result including that input appearing on the output pins. It does not include the delay needed to gather N samples, for an N tap filter, before a mathematically correct result is obtained.

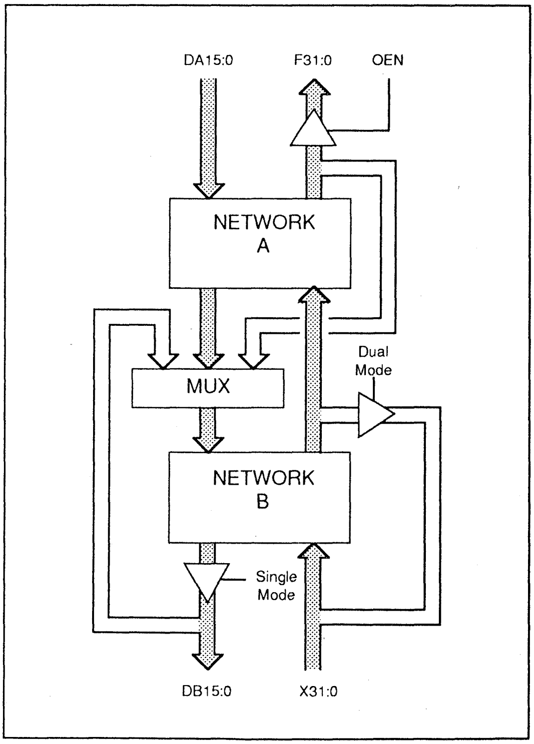

Fig. 6 Single Filter Bus Utilisation

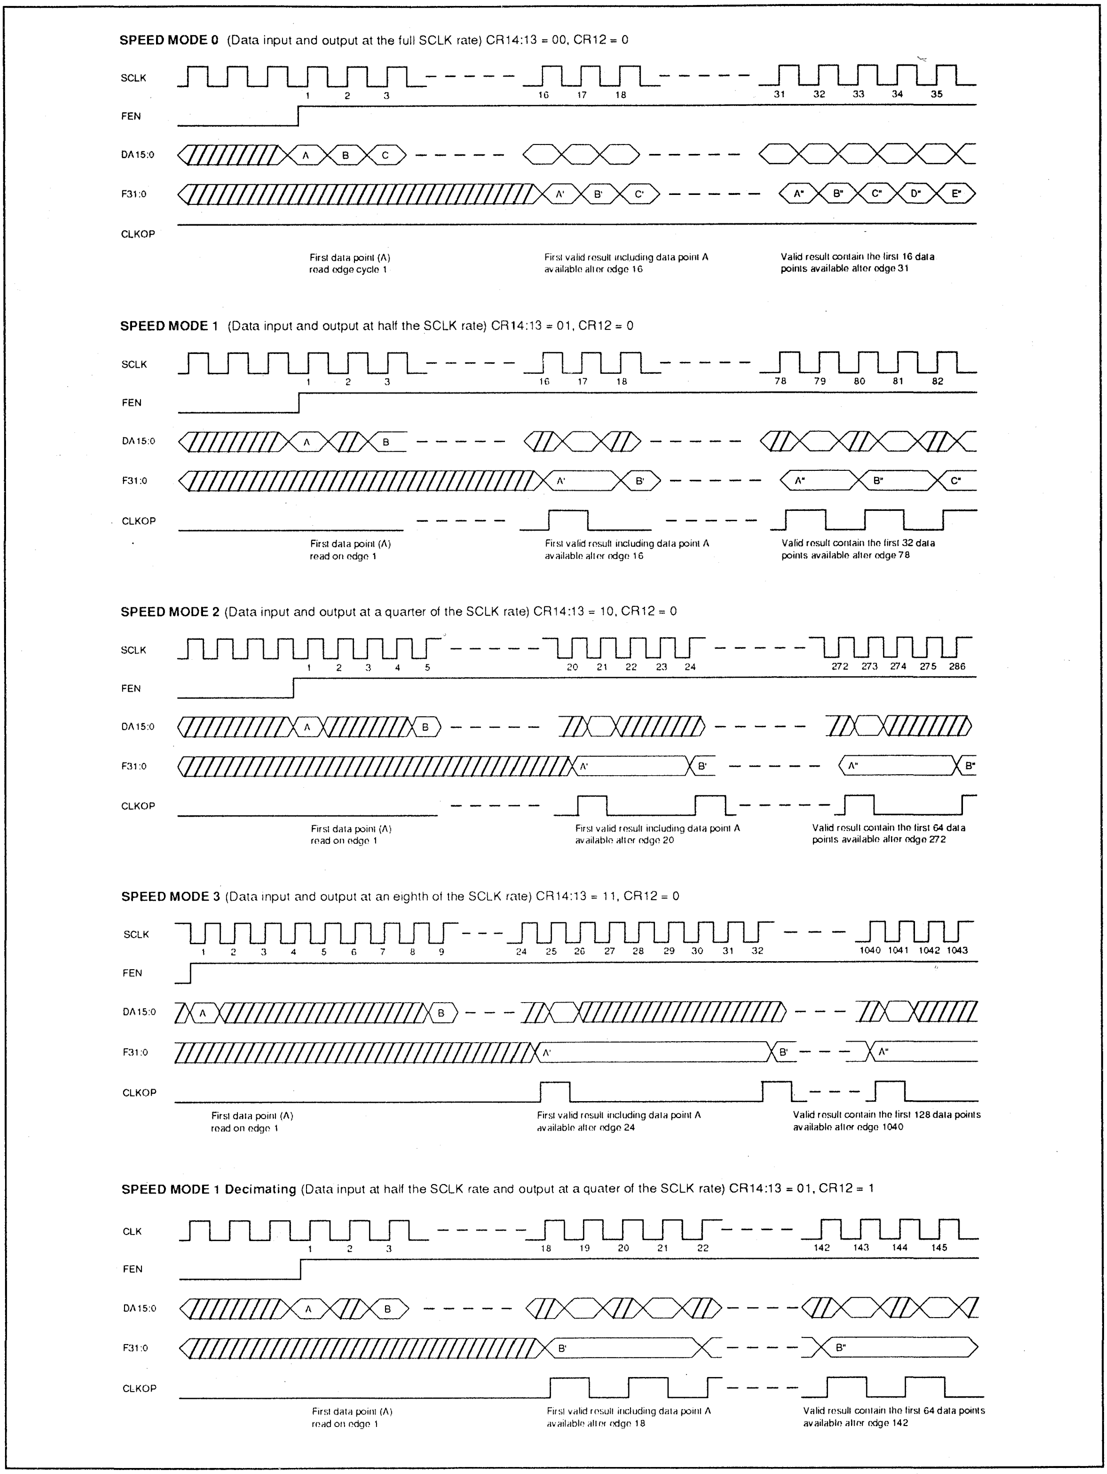

Fig. 7 Single Filter Timing Diagrams

DUAL INDEPENDENT FILTER OPTIONS

When operating as two independent filters the device accepts 16 bit data on both the DA and DB buses at the selected sample rate, see Fig. 8. Results are available from both the F and X buses. The F bus may be tristated using the OEN input. Signal OEN is registered onto the device and does not therefore take effect until the first SCLK rising edge

Each filter must be configured in the same manner, and multiple device expansion is not possible due to the pin reorganization. The latter requirement can, of course, still be satisfied by several devices configured as single filters.

Dual independent filter mode is selected by setting control register bits 15 and 4 to a zero. The required filter length is selected using control register bits 14 and 13 as summarised in Table 4, which also shows the resulting latency. As in single filter mode normal or decimate by two operation can be selected using control register bit 12.

CR 14 13 12	Input Rate	Output Rate	Filter Length	Setup Latency	
				Ind	Cas
0 0 0	SCLK	SCLK	8 Taps	16	27
0 0 1	SCLK	SCLK/2	16 Taps	17	-
0 1 0	SCLK/2	SCLK/2	16 Taps	16	28
0 1 1	SCLK/2	SCLK/4	32 Taps	18	-
1 0 0	SCLK/4	SCLK/4	32 Taps	20	36
1 0 1	SCLK/4	SCLK/8	64 Taps	24	-
1 1 0	SCLK/8	SCLK/8	64 Taps	24	40

Table 4. Dual Filter Options

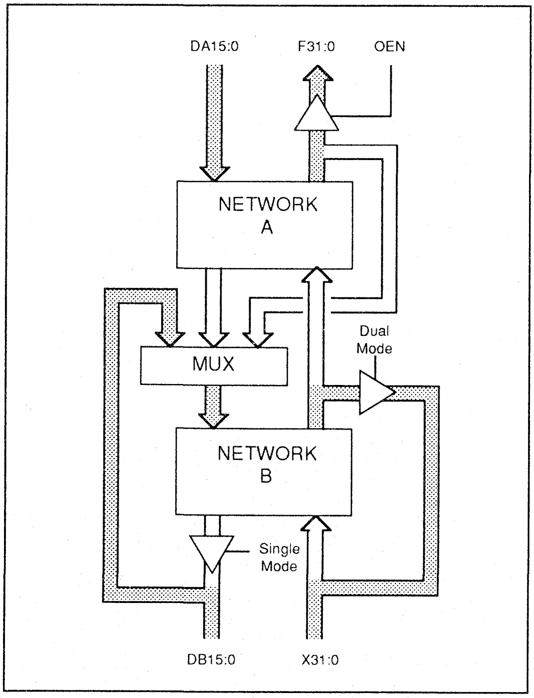

Fig. 8 Dual Independent Filter Bus Utilisation

DUAL CASCADED FILTER OPTIONS

When operating as two cascaded filters the device accepts 16 bit data on the DA bus at the selected sample rate. Results are presented on the 32 bit X bus, see Fig. 9. Each filter must be configured in the same manner. Multiple device expansion is not possible in this mode.

Dual cascaded filter mode is selected by setting control register bit 15 to a zero and bit 4 to a one. The required filter length is selected using control register bits 14 and 13 as summarised in Table 4, which also shows the resulting latency. The decimate by two option is not available in this mode.

The data for the second filter network is extracted as the middle 16 bits from the first networks accumulated result. For successful operation the first filter network must have unity gain. See the section on filter accuracy for more details.

The cascade option is used to increase the stop band rejection in a practical filter application. Theoretically, increasing the number of taps in an FIR filter will increase the stop band rejection, but this assumes floating point calculations with no accuracy limitations. In practice, with fixed point arithmetic, better performance is achieved with two smaller filters in series.

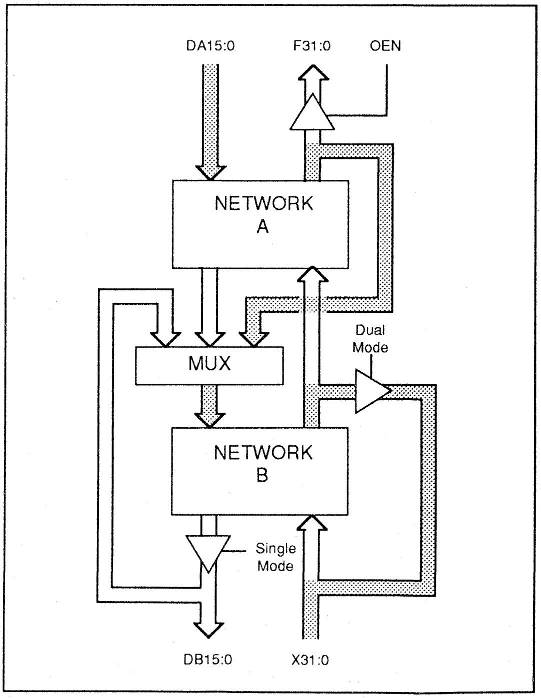

Fig. 9 Dual Cascaded Filter Bus Utilisation

FILTER ACCURACY

Input data and coefficients are both represented by 16bit two's complement numbers. The coefficients are converted to twelve bits by rounding towards zero. This is achieved as follows. If the coefficient is positive then the least significant 4 bits are discarded. If the coefficient is negative then the logical 'OR' of the least significant 4 bits are added to the remainder of the word. Twelve bit coefficients can be used directly provided the least significant four bits are set to zero.

The FIR filter results are calculated using a multiplier accumulator structure as shown in Fig. 10. The truncation and word growth allowed for in the data path are explained in Fig. 11. The 16 bit data and 12 bit coefficient inputs, (each with one sign bit before the binary point), are presented to the multiplier. This produces a 28 bit result with two bits before the binary point. Producing the full 28 bit result ensures that if both the data and coefficients are set to -1 a valid result is generated. Prior to entering the accumulator the least significant 4 bits of the multiplier result are truncated and the resulting 24 bits sign extended to 32 bits. The final accumulator result is 32 bits with 10 bits before the binary point. Thus 9 bits of word growth are allowed within the accumulator. All accumulator bits are made available on the output pins.

In cascade mode the middle 16 bits from the network A accumulator are fed round to the network B data inputs, see Fig. 11.

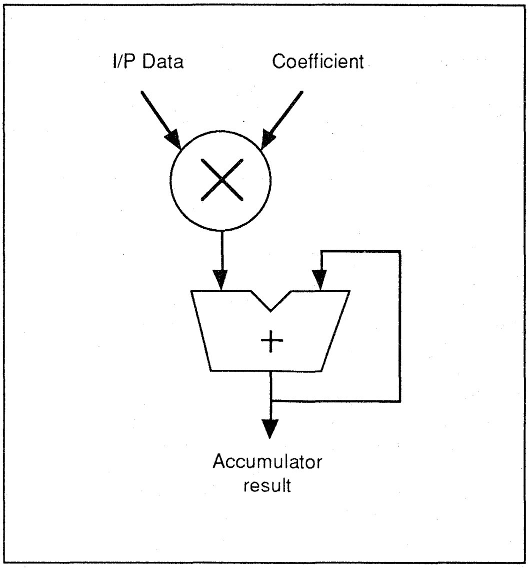

Fig. 10 Multiplier Accumulator

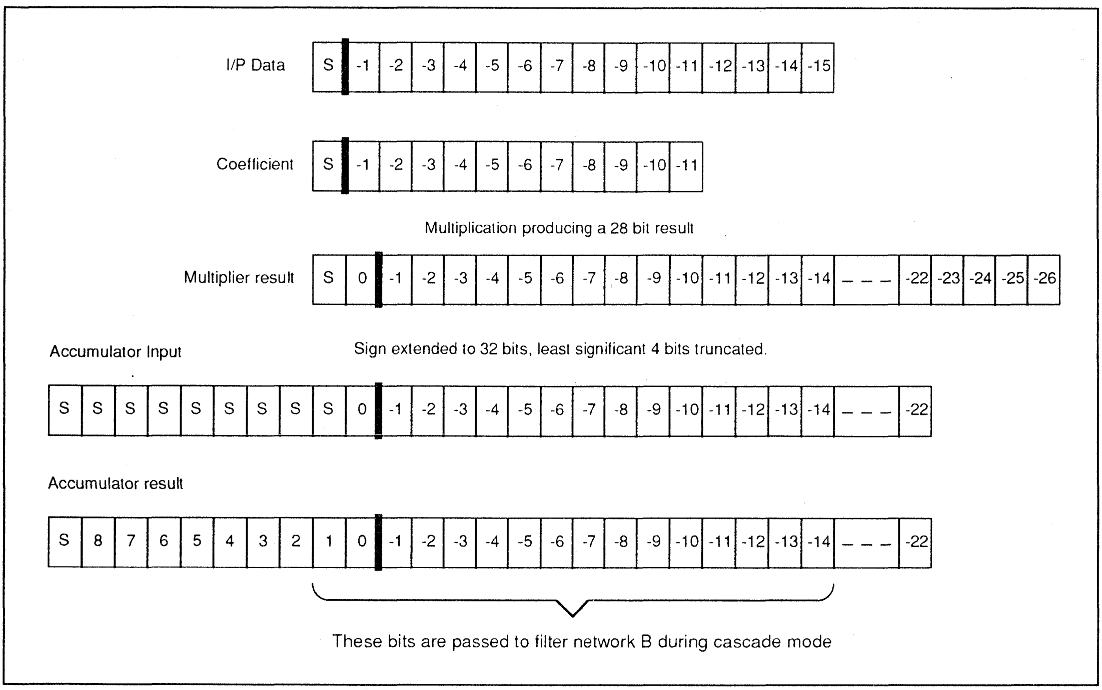

Fig. 11 Filter Accuracy

CASCADING DEVICES

When the filter requirements are beyond the capabilities of a single device, it is possible to connect several devices in cascade increasing the number of taps available at the required sample rate. Within each device all filter length, decimate, and bank swap options are still possible, but each device in the chain must be similarly programmed and configured as a single filter.

The number of devices which can be cascaded is only limited by the possibility of overflow in the 32 bit intermediate accumulations. If more than sixteen devices are cascaded in auto EPROM load mode, then an additional EPROM will be needed.

In modes where the data sample rate does not equal the clock rate. Then the cascade arrangement shown in Fig. 12 is utilised. Delayed data is passed from device to device in one direction, while intermediate results flow in the opposite direction. The interface device both accepts the input data and produces the final result. It is not necessary for each device to know its exact position in the chain, but the device which receives the input data and produces the final result must be identified, as must the device which terminates the chain. The former is known as the Interface device and the latter as the Termination device, all others are Intermediate devices. Control Register bits CR11:10 are used to define these positions as shown in Table 6.

The control logic in each of the devices must be synchronised with respect to the Interface device. This is achieved by connecting the Delayed Filter Enable output (DFEN) to the Filter Enable input (FEN) of the next device in the chain. The Interface device, itself, needs a Filter Enable signal produced by the system, unless the Free Run pin is pulled high. Even when the latter is true, the Filter Enable connection must be made between the remaining devices in the chain.

When devices are cascaded such that the data sample rate equals the clock rate, (Control register bits 14:13 = 00), then a different cascade configuration must be used. This is shown in Fig. 13. The number of devices which can be cascaded is, again, only limited by the 32 bit accumulators.

In this mode the delayed data is passed from device to device in the same direction as the intermediate results. The device which accepts the input data is now at the opposite end of the chain to the device which produces the final result. The control logic in each of the devices must be synchronised this is achieved by connecting all the device FEN inputs to the global Filter enable.

AVAILABLE OPTIONS

No more than 128 coefficients can be stored internally. This limits the filter length / decimate / bank swap options to those which do not require more than that number of coefficients. Thus when a filter with 128 taps is to be implemented in a single device, it is not possible to decimate or bank swap. When a filter with 64 taps is implemented, decimate or bank swap are possible, but not both. With all other filter lengths, all decimate and bank swap configurations are possible.

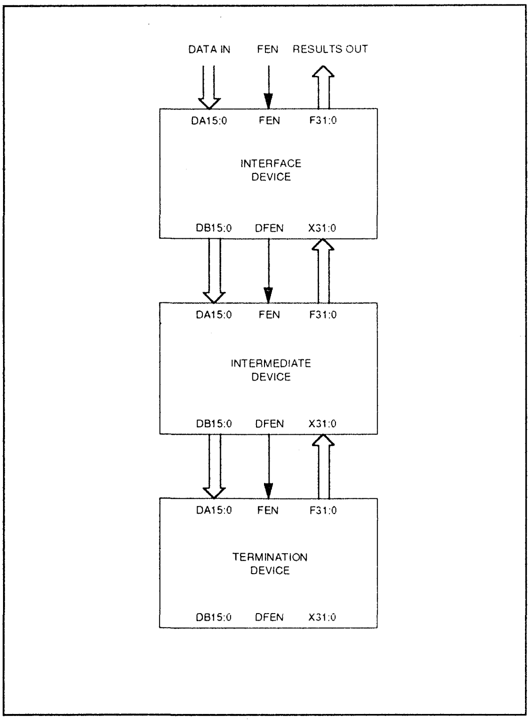

Fig. 12 Three Device Cascaded System

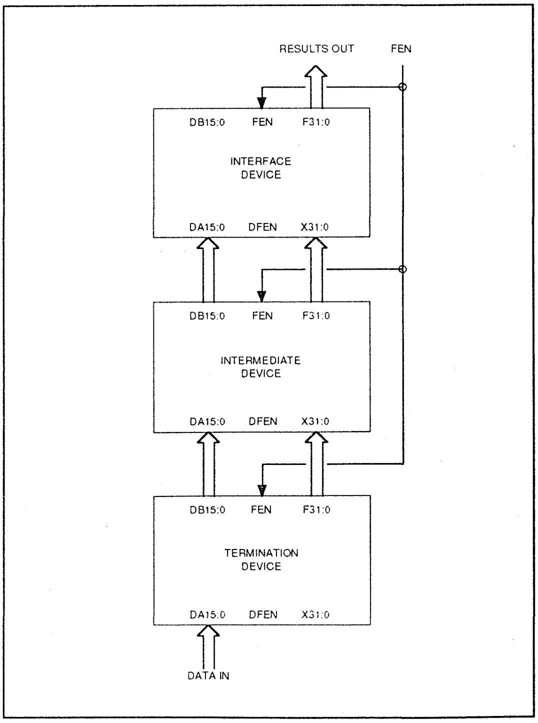

Fig. 13 Full Speed Cascaded System

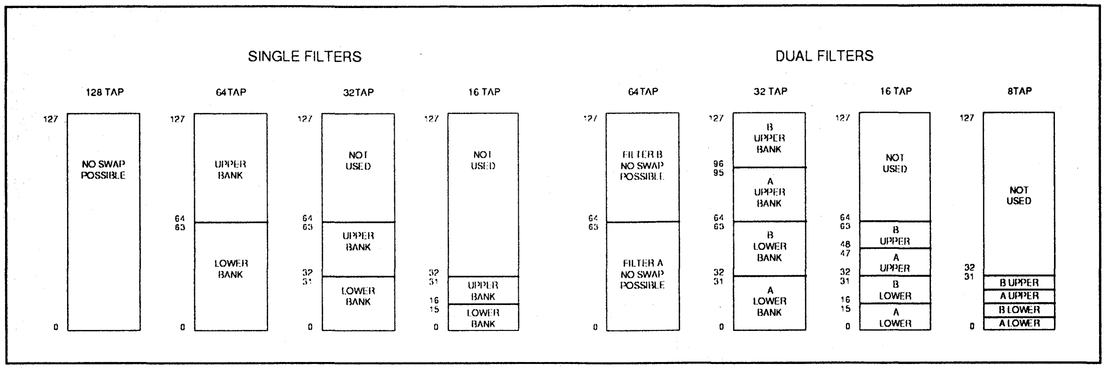

Fig. 14 Coefficient Memory Map

FILTER CONTROL

Two control modes are available selected by input signal FRUN. When FRUN is tied high the device will commence operation once the coefficients have been loaded. The CLKOP signal indicating when new input data is required and that new results are available, see Fig. 7. When FRUN is tied low filter operation will not commence until a high has been detected on signal FEN. This mode allows synchronisation to an existing data stream. Signal FEN should be taken high when the first valid data sample is available so that both are read into the device on the next SCLK rising edge.

During device reset the RES signal must be held low for a minimum of 16 SCLK cycles. After a reset the control register returns to it's default state of 8C80 Hex. This places the device into the following mode :-

- Single filter
- Sample rate equal to the clock rate
- Non-decimating
- A single device (Not in a cascade chain)
- Bank swap selected by bit in the control register

COEFFICIENT BANK SWAP

A Bank Swap feature is provided which allows ALL coefficients to be simultaneously replaced with a different set. A bit in the Control Register (CR7) allows the swap to be controlled by either input signal SWAP or Control Register bit (CR6). The latter is useful if the device is controlled by a microprocessor, when driving a separate pin would entail additional address decoding logic and an external latch.

If the pin or control register bit is low, the coefficients used will be those loaded into the lower banks illustrated in Fig. 14. When the pin or bit is high, the upper banks are used.

The actual swap will occur when the next sampling clock active going transition occurs. This can be up to seven system clocks later than the swap transition, and is filter length dependent. The first valid filtered output will then occur after the pipeline latancies given in Tables 3 and 4.

By setting a bit in the Control Register it is possible to bank swap on every data sampling clock. This function does not depend on the status of the SWAP pin or bit, and the lower bank will be initially selected after FEN goes active. The option can be used to implement filters with complex coefficients.

LOADING COEFFICIENTS

When the device is to operate in a stand alone application then the coefficients can be down loaded as a complete set from a previously programmed EPROM. Alternatively if the system contains a microprocessor they can be individually transferred from a remote master under software control. In any mode the system clock must be present and stable during the transfer, and the addressing scheme is such that the least significant address specifies the coefficient applied to the first multiplier seen by incoming data.

The addresses used during the load operation are those illustrated in Fig. 14. The Control Register is loaded when CCS is high. In BYTE mode address A0 is used to select the portion of control register loaded, otherwise the address bits are redundant. When an EPROM is used to provide coefficients, this redundancy causes the number of locations needed for any device to be double that for the coefficients alone.

AUTO EPROM LOAD

When the EPROM pin is tied low, the PDSP16256 assumes the role of a master device in the system and controls the loading of coefficients from an external EPROM, see Fig.15. A load sequence commences when the RESET input goes inactive, and will continue until every coefficient has been loaded. The BUSY pin goes high to indicate that a load sequence is occurring and the filter output is invalid. The device will not commence a filter operation until the Filter Enable edge is received (FEN) after BUSY has gone low. This requirement can be avoided if the Free Run pin (FRUN) is tied high.

The address bus pins become outputs on the Master device, and produce a new address every four system clock periods. This four clock interval, minus output delays and the data set up time, defines the available EPROM access time.

The coefficients are always loaded as bytes. The state of the BYTE pin on the master device is ignored. This arrangement also allows the eight, most significant, coefficient bus pins (C15:8) to be used for other purposes as described later. Since the 16 bit coefficients are loaded in two bytes the A0 pin specifies the required byte. The maximum number of stored coefficients is 128, eight address outputs are therefore provided for the EPROM. These eight outputs from the Master

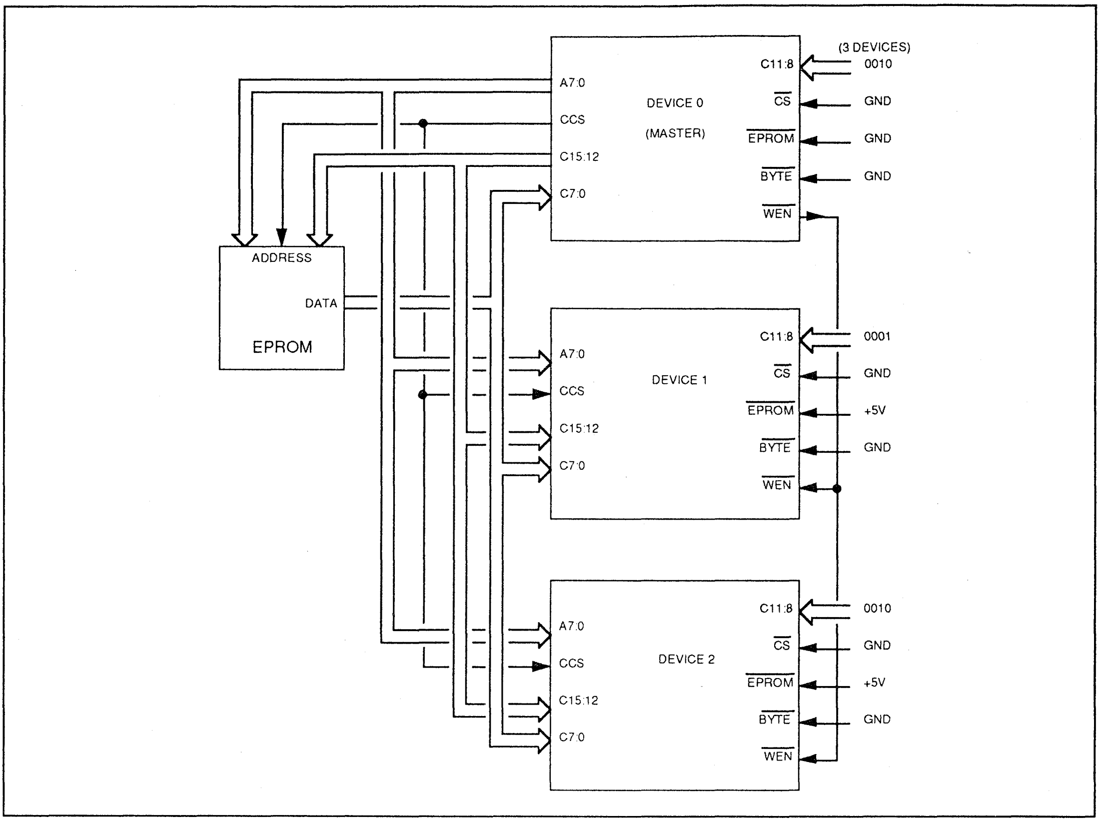

Fig. 15 Three device auto EPROM load

must also drive the address inputs on the slave devices.

When the filter length is less than the maximum, the PDSP16256 will only transfer the correct number of coefficients, and one or more significant address bits will remain low. Sufficient coefficients are always loaded to allow for a possible Bank Swap to occur, and the EPROM allocation must allow for this even if the feature is not to be used. Table 5 shows the number of coefficients loaded for each of the modes.

If several devices are cascaded, only one device assumes the role of the Master by having its EPROM pin grounded. It produces a Write Enable signal for the other devices, plus four higher order address outputs on C15:12. The extra address bits on C15:12 define separate areas of EPROM, containing coefficients for up to fifteen additional devices. The least significant block of memory must always be allocated to the Master device. The additional devices need not in practice be all part of the same cascaded chain, but can consist of several independent filters. They must, however, all have their BYTE pins tied low.

When one EPROM is supplying information for several devices, some means of selectively enabling each additional device must be provided. This is achieved by using the C11:8 pins on the slave devices as binary coded inputs to define one to fifteen extra devices. These coded inputs always correspond to the block address used for the segment of EPROM allocated to that device. Code 'all zeros' must not be used since the Master device has implied use of the bottom segment. This is necessary since the C11:8 pins are alternatively used on the Master device to define the number of devices supported by the EPROM.

In addition to providing the most significant addresses to the EPROM, the C15:12 address outputs from the master device must also drive the C15:12 inputs on the slave devices. These C15:12 inputs are internally compared to the C11:8 inputs to decide if that device is currently to be loaded. This approach avoids the need for external decoders and makes the Chip Enable input redundant. This input, however, must be tied low on every device in an EPROM supported system.

The Control Coefficient pin (CCS) is used to define when the control register is to be loaded. It becomes an output on the Master device which provides an EPROM address bit next in significance above A7:0, and also drives the CCS inputs on the slave devices. This output is high for the first two EPROM transfers in order to access the control information, and then remains low whilst the coefficients are loaded. This control information is thus not stored adjacent to the coefficients within the EPROM, and in fact the EPROM must provide twice the storage necessary to contain the coefficients alone. All but two of the bytes in the additional half are redundant. See Fig.16 for the EPROM memory map.

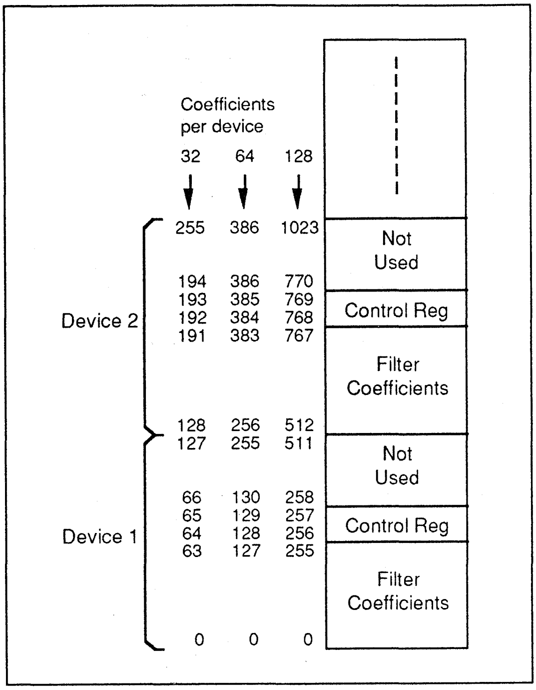

Fig. 16 EPROM Memory Map

Control Register 14 13 12	Number of Coefficients Loaded
0 0 0	32
0 0 1	64
0 1 0	64
0 1 1	128
1 0 0	128
1 0 1	128
1 1 0	128
1 1 1	Invalid Mode

Table 5. Number of Coefficients loaded

NOTE the EPROM memory map Fig. 16 assumes that, for the 32 and 64 coefficient per device options, that the unused address pins are unconnected. If all address pins are connected as shown in Fig. 15 then the 128 coefficients per device memory map column should be used. Only those coefficients required will be read, hence the upper portions of the coefficient address space will be ignored.

USING A REMOTE MASTER

When a remote master is used to load coefficients, the EPROM pin must be tied high and a conventional peripheral interface is then provided. It is not possible, however, to read coefficients already stored. The master supplies an address and data bus, and writes to the PDSP16256 occur under the control of synchronous Chip Enable and Write Strobe inputs. The Coefficient Control Register pin (CCS) must be driven by a master address line higher in significance than A7:0. Both the WEN and CS signals must be low for the load operation to occur. When loading the control register the CS signal must be held low for a further 2 cycles see Fig. 17. Since the internal write operation is actually performed with the system clock, it is necessary for the clock to be present during the transfer.

The BYTE input defines whether coefficients are loaded as a single 16 bit word or two 8 bytes. The latter saves on connections to the remote master. Address bits A7:0 are used in BYTE mode. 16 bit word mode uses bits A6:0, A7 being redundant. When writing in byte mode the least significant byte (A0 = 0) must be written first followed by the most significant byte (A0 = 1).

In the byte mode of working the internal comparison between C15:12 and C11:8 is made, regardless of the state of the EPROM pin. For this reason pins C15:8 should all be tied low when a remote master is used with byte transfers. This ensures that the internal comparison gives equality and allows the load operation to occur.

The address and coefficient busses plus the Write Enable and CS signals must all meet the specified set up and hold times with respect to the system clock, see Fig 17. This synchronous interface is optimum for the majority of high end applications, when individual coefficients must be updated at sample clock rates. If, for convenience reasons, the coefficients are loaded under software control from a general purpose microprocessor, the Write Enable will probably be asynchronous to the system clock used by the PDSP16256. In this case external synchronising logic is needed, see Fig.18.

Fig. 19 shows the recommended loading sequence and filter operation initiation. The simplest technique is to reset the device prior to loading a set of coefficients. Coefficients may be loaded once BUSY returns low or 22 cycles after RESET is taken high.

When loading a device from a remote master the control register must be loaded first followed by the filter coefficients. Fig. 19 shows the required loading sequence, two examples are given one for byte mode the other for word mode. A gap of at least one cycle must be left after loading the control register before loading the first coefficient.

Filter operations are started by presenting the first data word at the same time as raising signal FEN.

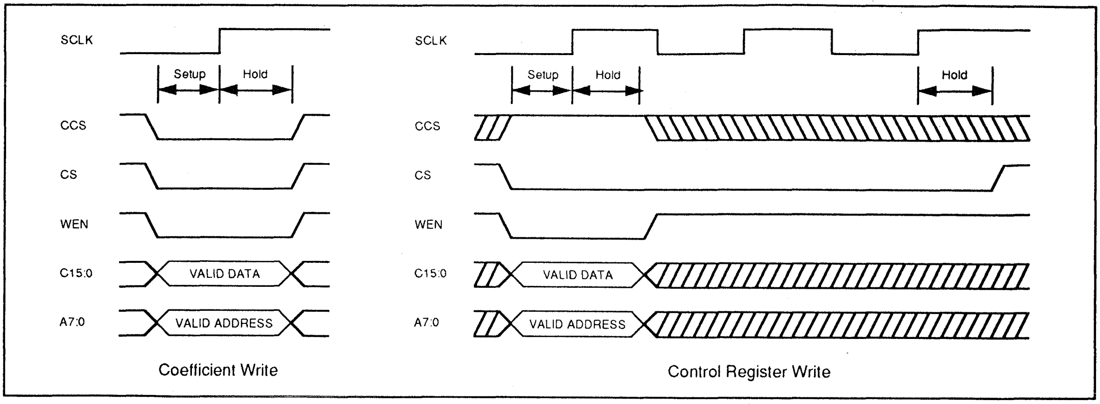

Fig. 17 Remote Master Setup & Hold Timings

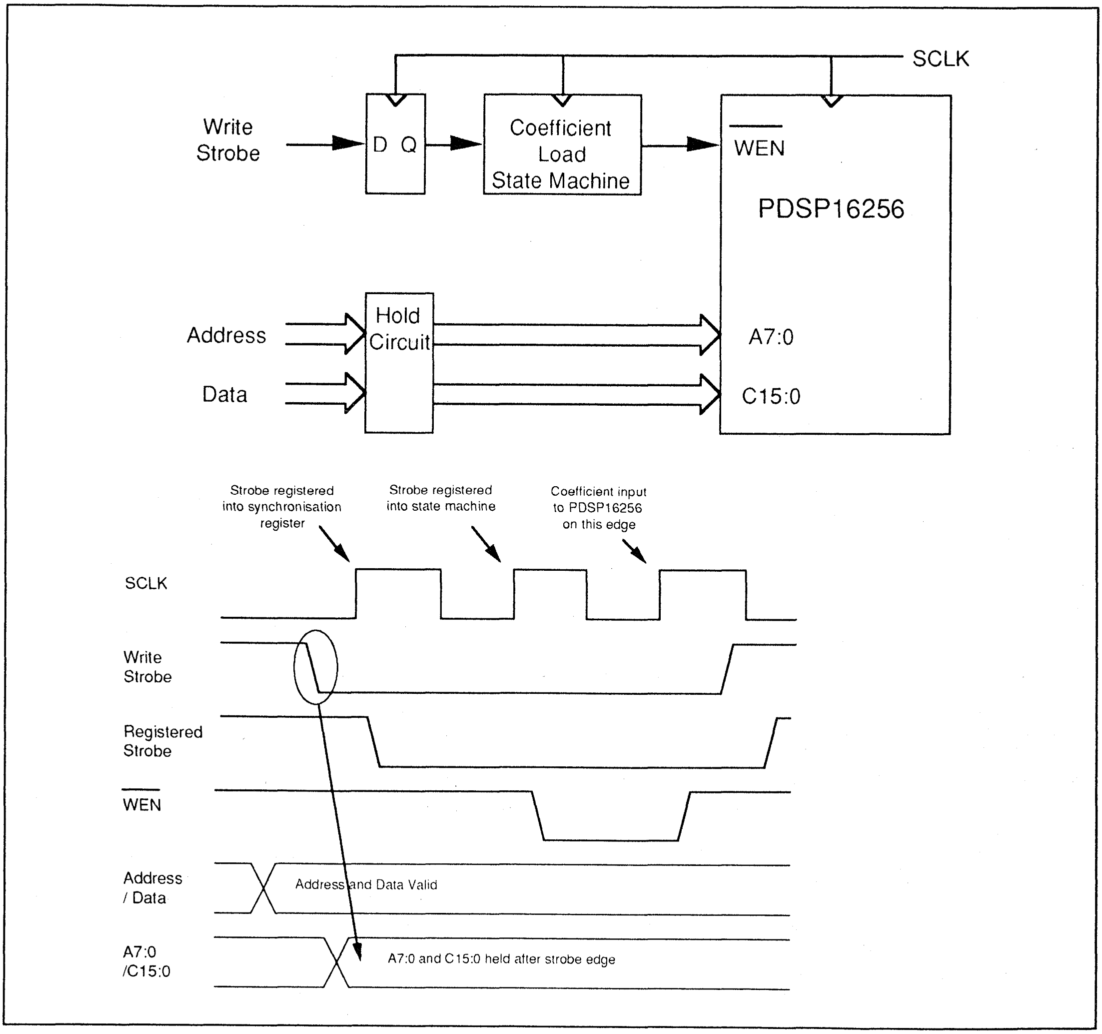

Fig. 18 Remote Master Synchronisation

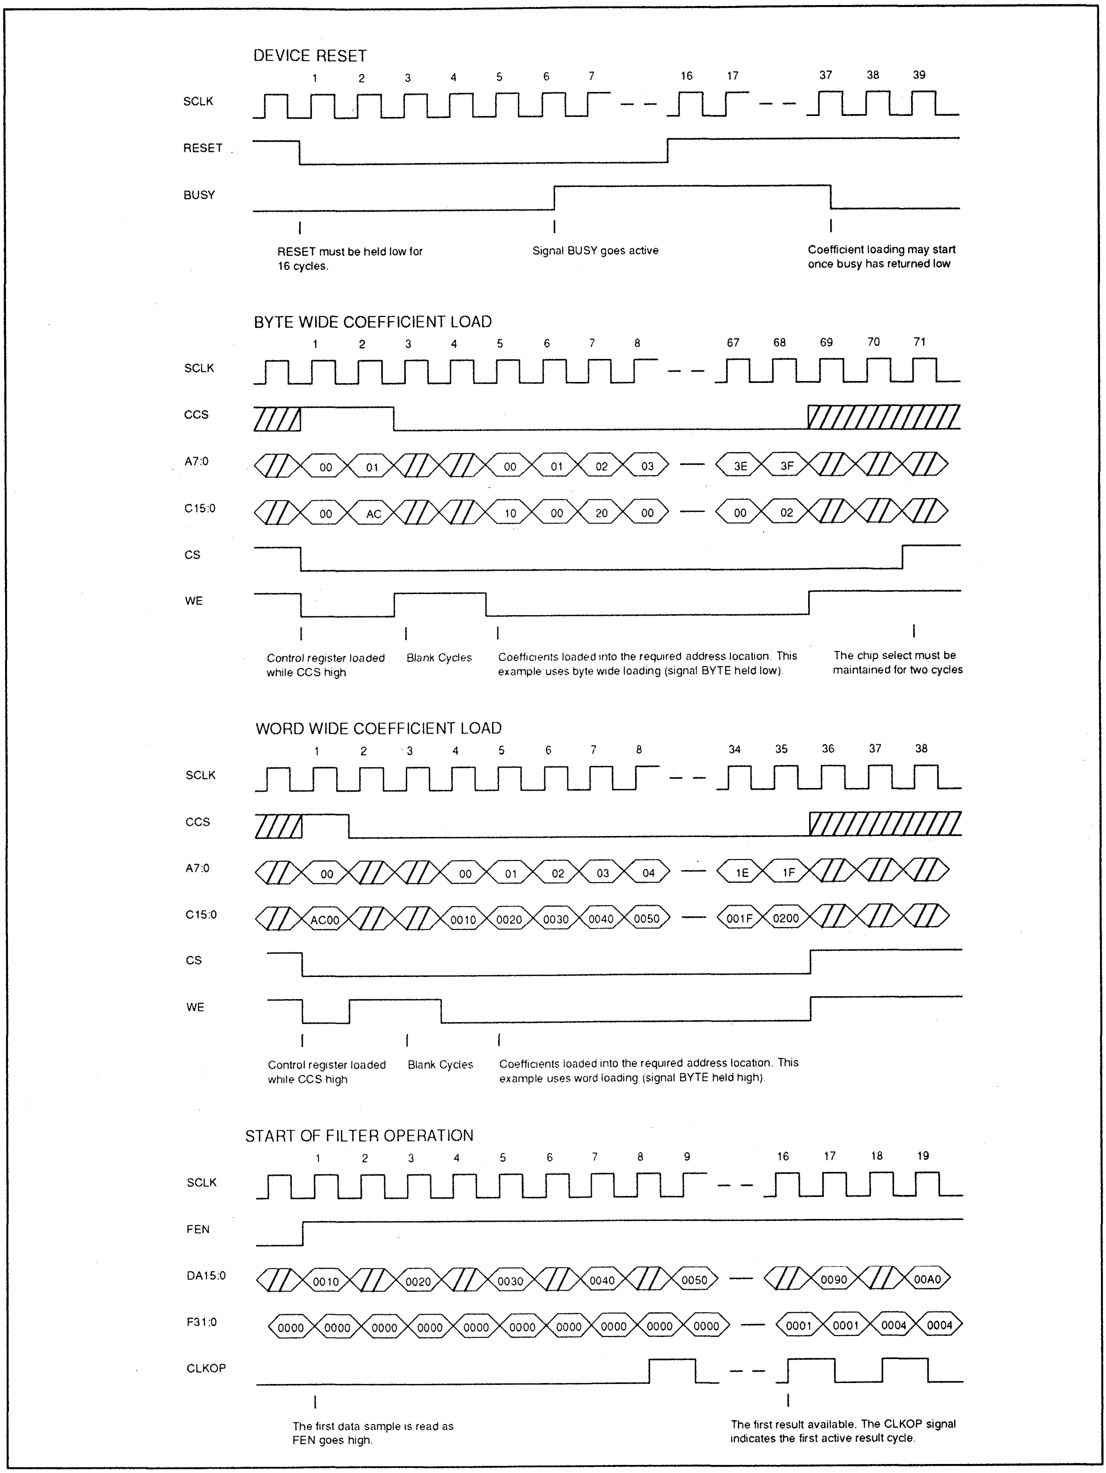

Fig. 19 Device Startup

CONTROL REGISTER

The internal operation of the PDSP16256 is controlled by the status of a 16 bit control register. In the dual filter modes both networks are controlled by the same register. The significance of the various bits are shown in Table 6. Tables 7 and 8 define the control register bit interdependence for the filter and bank swapping modes.

The control register is double buffered. This allows the writing of a new control word without affecting the current operation of the device. To activate the new control register after it has been written to the device the bank swap signal must be toggled. After a reset the active control register is loaded directly and bank swap need not be used.

Control Register Bits		Function
15	4	
0	0	Two independent filters
0	1	Two filters in cascade
1	X	Single Filter

Table 7 Control Register Filter Mode Bits

Control Register Bits			Function
7	6	5	
0	X	0	Control by input pin
1	0	0	Lower bank selected
1	1	0	Upper bank selected
X	X	1	Swap on every sample clock

Table 8 Control Register Bank Swap bits

Bits	Decode	Function
15	0	Dual filter mode
15	1	Single filter mode
14:13	00	Sample rate is the system clock
14:13	01	Sample rate is half the system clock
14:13	10	Sample rate is quarter the system clock
14:13	11	Sample rate is eighth the system clock
12	0	Output rate equals the input rate
12	1	Decimate bt two
11:10	00	Intermediate device
11:10	01	Interface device
11:10	10	Termination device
11:10	11	Single device
9:8	00	These bits MUST be at logical zero
7	0	Bank swap is controlled by input pin
7	1	Bank swap is controlled by Bit 6
6	0	Lower bank if Bit 7 is set
6	1	Upper bank if Bit 7 is set
5	0	Normal Bank Swap
5	1	Bank swap on every sample clock
4	0	Two independent filters
4	1	Two filters in cascade
3:0		These bits MUST be at logical zero

Table 6. Control Register Bit Allocation

ABSOLUTE MAXIMUM RATINGS (Note 1)

Supply voltage Vcc	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to Vcc + 0.5V
Output voltage V_{OUT}	-0.5V to Vcc + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HBM)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	-55°C to +125°C
Junction temperature with power applied T_J	150°C
Package power dissipation	3000mW
Thermal resistances	
Junction to Case $\varnothing_{JC}$	5°C/W

NOTES
1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as negative into the device
5. Vcc = Max, Outputs Unloaded, Clock Freq = Max
6. The $\varnothing_{JC}$ data assumes that heat is extracted from the top face of the package.

ELECTRICAL CHARATERISTICS

Operating Conditions (unless otherwise stated)

Commercial:	T_{AMB} = 0°C to +70°C	$T_{J\,(MAX)}$ = 100°C	Vcc = 5.0V±5%	Ground = 0V
Industrial:	T_{AMB} = -40°C to +85°C	$T_{J\,(MAX)}$ = 110°C	Vcc = 5.0V±10%	Ground = 0V
Military:	T_{AMB} = -55°C to +125°C	$T_{J\,(MAX)}$ = 150°C	Vcc = 5.0V±10%	Ground = 0V

Static Charateristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage (CMOS)	V_{IH}	3.5		-	V	SCLK input only
Input low voltage (CMOS)	V_{IL}	-		1.0	V	SCLK input only
Input high voltage (TTL)	V_{IH}	2.0		-	V	All other inputs
Input low voltage (TTL)	V_{IL}	-		0.8	V	All other inputs
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{CC}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{CC}
Output S/C current	I_{OS}	10		300	mA	V_{CC} = Max

Switching Characteristic	Commercial		Industrial		Military		Units	Conditions
	Min.	Max.	Min.	Max.	Min.	Max.		
Input signal setup to clock rising edge	8	-	8	-	8	-	ns	
Input signal hold after clock rising edge	4	-	4	-	4	-	ns	
OEN setup to clock rising edge	20	-	20	-	20	-	ns	
OEN hold after clock rising edge	4	-	4	-	4	-	ns	
Clock rising edge to output signal valid	5	26	5	28	5	28	ns	30pF
Clock Frequency	-	25	-	20	-	20	MHz	
Clock High Time	18	-	20	-	20	-	ns	
Clock Low Time	11	-	12	-	12	-	ns	
Clock to data valid from high impedance	-	30	-	30	-	30	ns	see Fig. 20
Clock to data high impedance	-	30	-	30	-	30	ns	see Fig. 20
Vcc Current	-	320	-	250	-	250	mA	see Note 5

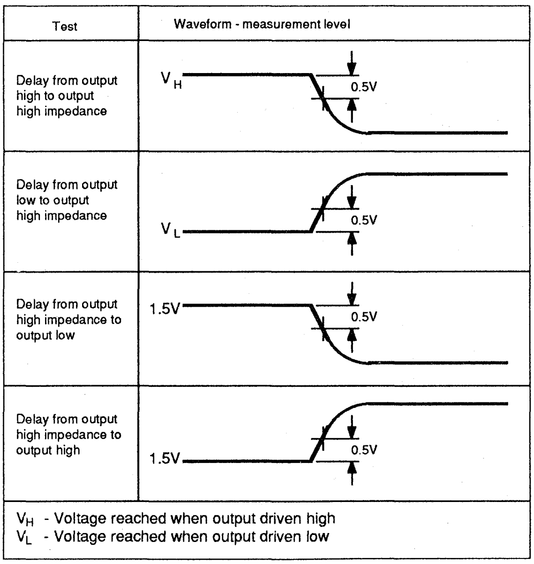

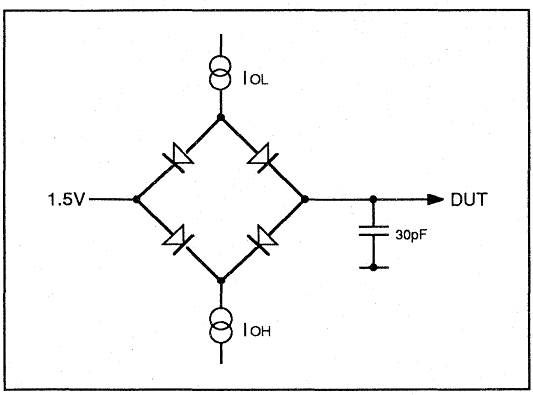

Fig. 20 Three state delay measurement load.

ORDERING INFORMATION

PDSP16256A	**C0 AC**	25MHz	Commercial
PDSP16256	**B0 AC**	20MHz	Industrial
PDSP16256	**A0 AC**	20MHz	Military

GC132 surface mount package under development

PDSP16318/PDSP16318A

COMPLEX ACCUMULATOR

The PDSP16318 contains two independent 20-bit Adder/Subtractors combined with accumulator registers and shift structures. The four port architecture permits full 20MHz throughout in FFT and filter applications.

Two PDSP16318As combined with a single PDSP16112A Complex Multiplier provide a complete arithmetic solution for a Radix 2 DIT FFT Butterfly. A new complex Butterfly result can be generated every 50ns allowing 1K complex FFTs to be executed in 256µs.

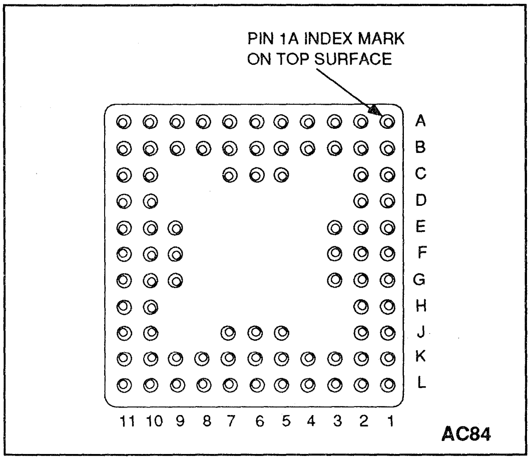

Fig. 1 Pin connections - bottom view (AC84 - PGA)

FEATURES

- Full 20MHz Throughout in FFT Applications
- Four Independent 16-bit I/O Ports
- 20-bit Addition or Accumulation
- Fully Compatible with PDSP16112 Complex Multiplier
- On Chip Shift Structures for Result Scaling
- Overflow Detection
- Independent Three-State Outputs and Clock Enables for 2 Port 20MHz Operation
- 1.4 micron CMOS
- 500mW Maximum Power Dissipation
- 84 Pin PGA or QFP packages

APPLICATIONS

- High speed Complex FFT or DFTs
- Complex Finite Impulse Response (FIR) Filtering
- Complex Conjugation
- Complex Correlation/Convolution

ASSOCIATED PRODUCTS

PDSP16112 16 x 12 Complex Multiplier
PDSP16116 16 x 16 Complex Multiplier
PDSP1601 ALU and Barrel Shifter
PDSP16330 Pythagoras Processor

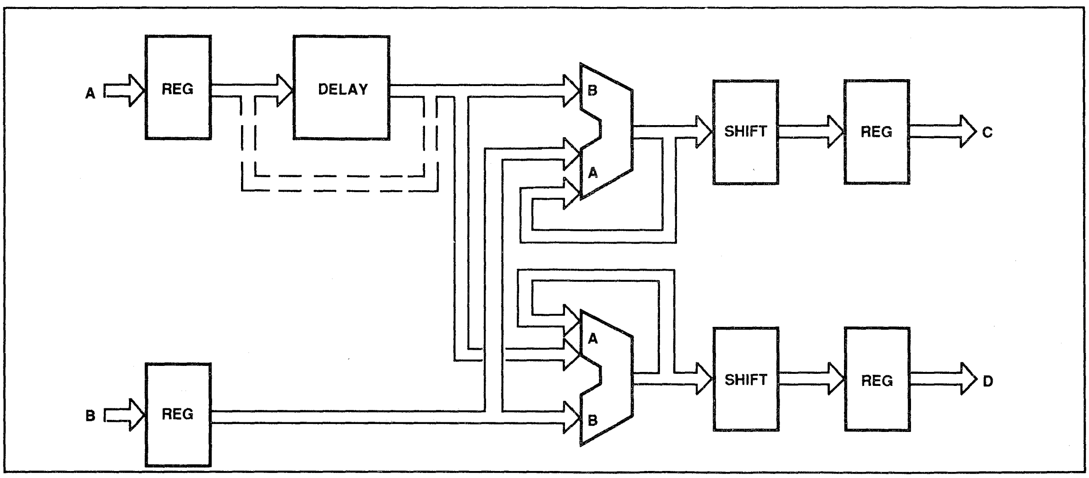

Fig. 2 PDSP16318 simplified block diagram

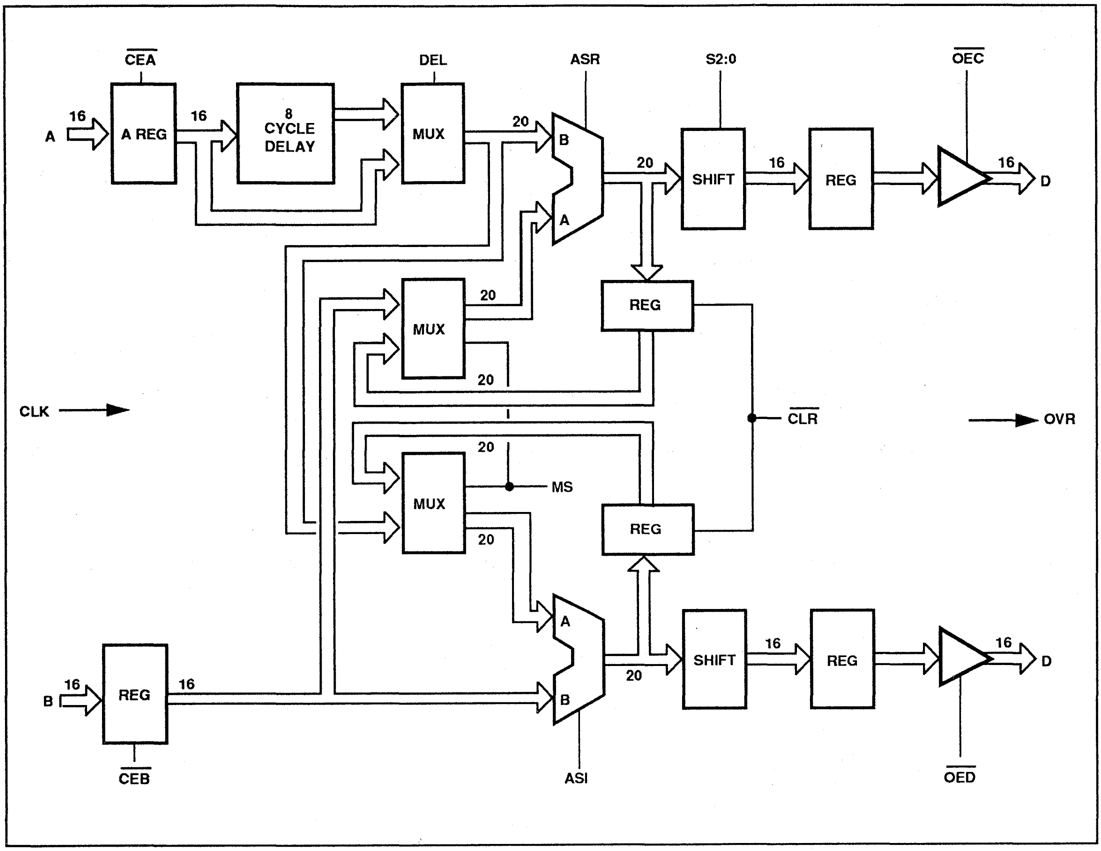

Fig. 3 Block diagram

FUNCTIONAL DESCRIPTION

The PDSP16318 is a Dual 20-bit Adder/Subtractor configured to support Complex Arithmetic. The device may be used with each of the adders allocated to real or imaginary data (e.g. Complex Conjugation), the entire device allocated to Real or Imaginary Data (e.g. Radix 2 Butterflys) or each of the adders configured as accumulators and allocated to real or imaginary data (Complex Filters). Each of these modes ensures that a full 20MHz throughput is maintained through both adders, the first and last mode illustrating true Complex operation, where both real and imaginary data is handled by the single device.

Both Adder/Subtractors may be controlled independently via the ASR and ASI inputs. These controls permit A + B, A - B, B - A or pass A operations, where the A input to the Adder is derived from the input multiplexer. The $\overline{\text{CLR}}$ control line allows the clearing of both accumulator registers. The two multiplexers may be controlled via the MS inputs, to select either new input data, or fed-back data from the accumulator registers. The PDSP16318 contains an 8-cycle deskew register selected via the DEL control. This deskew register is used in FFT applications to ensure correct phasing of data that has not passed through the PDSP16112 Complex Multiplier.

The 16-bit outputs from the PDSP16318 are derived from the 20-bit result generated by the Adders. The three bit S2:0 input selects eight different shifted output formats ranging from the most significant 16 bits of the 20-bit data, to the least significant 13 bits of the 20-bit data. In this mode the 14th, 15th and 16th bits of the output are set to zero. The shift selected is applied to both adder outputs, and determines the function of the OVR flag. The OVR flag becomes active when either of the two adders produces a result that has more significant digits than the MSB of the 16-bit output from the device. In this manner all cases when invalid data appears on the output are flagged.

Symbol	Type	Description
A15:0	Input	**Data** presented to this input is loaded into the input register on the rising edge of CLK. A15 is the MSB.
B15:0	Input	**Data** presented to this input is loaded into the input register on the rising edge of CLK. B15 is the MSB and has the same weighting as A15.
C15:0	Output	New **data** appears on this output after the rising edge of CLK. C15 is the MSB.
D15:0	Output	New **data** appears on this output after the rising edge of CLK. C15 is the MSB.
CLK	Input	**Common Clock** to all internal registers
$\overline{\text{CEA}}$	Input	**Clock enable:** when low the clock to the A input register is enabled.
$\overline{\text{CEB}}$	Input	**Clock enable:** when low the clock to the B input register is enabled.
$\overline{\text{OEC}}$	Input	**Output enable:** Asynchronous 3-state output control: The C outputs are in a high impedance state when this input is high.
$\overline{\text{OED}}$	Input	**Output enable:** Asynchronous 3-state output control: The D outputs are in a high impedance state when this input is high.
OVR	Output	**Overflow flag:** This flag will go high in any cycle during which either the output data overflows the number range selected or either of the adder results overflow. A new OVR appears after the rising edge of the CLK.
ASR1:0	Input	**Add/subtract Real:** Control input for the 'Real' adder. This input is latched by the rising edge of clock.
ASI1:0	Input	**Add/subtract Imag:** Control input for the 'Imag' adder. This input is latched by the rising edge of clock.
$\overline{\text{CLR}}$	Input	**Accumulator Clear:** Common accumulator clear for both Adder/Subtractor units. This input is latched by the rising edge of CLK.
MS	Input	**Mux select:** Control input for both adder multiplexers. This input is latched by the rising edge of CLK. When high the feedback path is selected.
S2:0	Input	**Scaling control:** This input selects the 16-bit field from the 20-bit adder result that is routed to the outputs. This input is latched by the rising edge of CLK.
DEL	Input	**Delay Control:** This input selects the delayed input to the real adder for operations involving the PDSP16112. This input is latched by the rising edge of CLK.
VCC	Power	**+5V supply:** Both Vcc pins must be connected.
GND	Ground	**0V supply:** Both GND pins must be connected.

GG pin	AC pin	Function	GG pin	AC pin	Function	GG pin	AC pin	Function	GG pin	AC pin	Function
77	B2	D7	6	K2	C7	31	K10	A1	56	B10	B10
82	C2	D8	7	K3	C6	32	J10	A2	57	B9	B9
83	B1	D9	8	L2	C5	33	K11	A3	58	A10	B8
84	C1	D10	9	L3	C4	34	J11	A4	59	A9	B7
85	D2	GND	10	K4	C3	35	H10	A5	60	B8	B6
86	D1	VCC	11	L4	C2	36	H11	A6	61	A8	B5
87	E3	D11	12	J5	C1	37	F10	A7	62	B6	B4
88	E2	D12	13	K5	C0	38	G10	A8	63	B7	B3
89	E1	D13	14	L5	$\overline{\text{OED}}$	39	G11	A9	64	A7	B2
90	F2	D14	15	K6	$\overline{\text{OEC}}$	40	G9	A10	65	C7	B1
91	F3	D15	16	J6	S2	41	F9	A11	66	C6	B0
92	G3	C15	17	J7	S1	42	F11	A12	67	A6	CLK
93	G1	C14	18	L7	S0	43	E11	A13	68	A5	$\overline{\text{CEB}}$
94	G2	C13	19	K7	MS	44	E10	A14	69	B5	OVR
95	F1	C12	20	L6	ASI1	45	E9	A15	70	C5	D0
96	H1	VCC	21	L8	ASI0	46	D11	$\overline{\text{CEA}}$	71	A4	D1
97	H2	GND	22	K8	DEL	47	D10	B15	72	B4	D2
98	J1	C11	23	L9	$\overline{\text{CLR}}$	48	C11	B14	73	A3	D3
99	K1	C10	24	L10	ASR1	49	B11	B13	74	A2	D4
100	J2	C9	25	K9	ASR0	50	C10	B12	75	B3	D5
5	L1	C8	26	L11	A0	51	A11	B11	76	A1	D6

Device Pinout for ceramic 84 - pin PGA (AC84) and ceramic QFP (GG100)

ASR or ASI ASX1	 ASX0	ALU Function
0	0	A + B
0	1	A
1	0	A - B
1	1	B - A

DEL	Delay Mux Control
0	A port input
1	Delayed A port input

MS	Real and Imag' Mux Control
0	B port input/Del mux output
1	C accumulator/D accumualtor

S2:0			Adder result																			
S2	S1	S0	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0				
0	0	1		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
0	1	0			15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
0	1	1				15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
1	0	0					15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	1						15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1	1	0							15	14	13	12	11	10	9	8	7	6	5	4	3	2
1	1	1								15	14	13	12	11	10	9	8	7	6	5	4	3

NOTE

This table shows the portion of the adder result passed to the D15:0 and C15:0 outputs. Where fewer than 16 adder bits are selected the output data is padded with zeros.

ABSOLUTE MAXIMUM RATINGS (Note 1)

Supply voltage V_{CC}	-0.5V to 7.0V
Input voltage V_{IN}	-0.9V to V_{CC} +0.9V
Output voltage V_{OUT}	-0.9V to V_{CC} +0.9V
Clamp diode current per pin I_K (see Note 2)	18mA
Static discharge voltage (HMB) V_{STAT}	500V
Storage temperature range T_S	-65°C to +150°C
Ambient temperature with power applied T_{amb}	
Industrial	-40°C to +85°C
Military	-55°C to +125°C
Junction temperature	150°C
Package power dissipation P_{TOT}	1000mW

NOTES

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliability.

THERMAL CHARACTERISTICS

Package Type	θ_{JC} °C/W	θ_{JA} °C/W
LC	12	35
AC	12	36

Test	Waveform - measurement level
Delay from output high to output high impedance	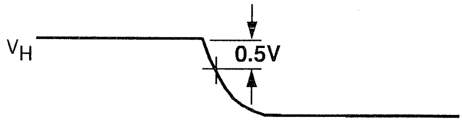
Delay from output low to output high impedance	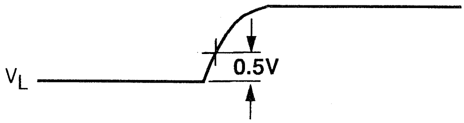
Delay from output high impedance to Output low	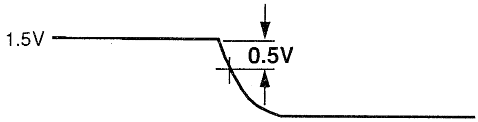
Delay from output high impedance to Output high	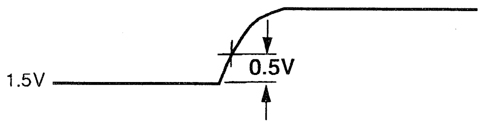

NOTES
1. V_H - Voltage reached when output driven high
2. V_L - Voltage reached when output driven low

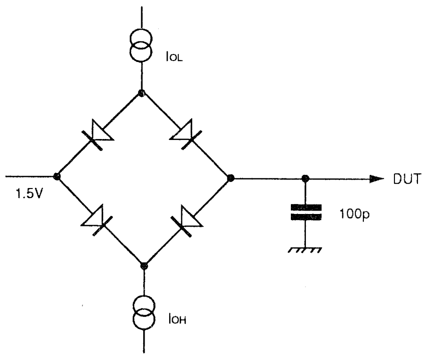

ELECTRICAL CHARACTERISTICS

Test conditions (unless otherwise stated):

T_{amb} (Commercial) = 0°C to +70°C, Vcc = 5.0V ± 5%, GND = 0V

T_{amb} (Industrial) =-40°C to +85°C, Vcc = 5.0V ± 10%, GND = 0V

T_{amb} (Military) =-55°C to +125°C, Vcc = 5.0V ± 10%, GND = 0V

STATIC CHARACTERISTICS

Characteristic	Symbol	Value Min.	Value Typ.	Value Max.	Units	Conditions
Output high voltage	V_{OH}	2.4		-	V	I_{OH} = 3.2mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL}=-3.2mA
Input high voltage	V_{IH}	3.5		-	V	
Input low voltage	V_{IL}	-		0.5	V	
Input leakage current	I_{IL}	-10		+10	µA	GND $\leq V_{IN} \leq V_{CC}$
Output leakage current	I_{oz}	-50	-	+50	µA	GND $\leq V_{OUT} \leq V_{CC}$
Output SC current	I_{OS}	20	-	200	mA	V_{cc} = Max
Input capacitance	C_{IN}	-	9	-	pF	

SWITCHING CHARACTERISTICS

Characteristic	Value Industrial + Commercial PDSP16318 Min.	PDSP16318 Max.	PDSP16318A Min.	PDSP16318A Max.	Value Military PDSP16318 Min.	Max.	Units	Conditions
Clock period	100	-	50	-	100	-	ns	
Clock High Time	20	-	15	-	20	-	ns	
Clock Low Time	20	-	15	-	20	-	ns	
A15:0, B15:0 setup to clock rising edge	8	-	5	-	8	-	ns	
A15:0, B15:0 hold after clock rising edge	2	-	2	-	2	-	ns	
MS, S2:0, ASI setup to clock rising edge	10	-	10	-	10	-	ns	
DEL, ASR, $\overline{CLR}$ setup to clock rising edge	8	-	5	-	8	-	ns	
DEL, ASR, $\overline{CLR}$, MS, S2:0, ASI hold after clock rising edge	2	-	2	-	2	-	ns	
$\overline{CEA}$, $\overline{CEB}$ setup to clock falling edge	2	-	2	-	2	-	ns	
$\overline{CEA}$, $\overline{CEB}$ hold after clock rising edge	8	-	8	-	8	-	ns	
Clock rising edge to OVR, C15:0, D15:0	5	40	5	30	5	40	ns	2 x LSTTL + 20pF
$\overline{OEC}/\overline{OED}$ low to C15:0/D15:0 high data valid	-	40	-	30	-	40	ns	2 x LSTTL + 20pF
$\overline{OEC}/\overline{OED}$ low to C15:0/D15:0 low data valid	-	40	-	30	-	40	ns	2 x LSTTL + 20pF
$\overline{OEC}/\overline{OED}$ high to C15:0/D15:0 high impedance	-	40	-	30	-	40	ns	2 x LSTTL + 20pF
Vcc current	-	70	-	110	-	70	mA	Vcc = max, TTL input levels Outputs unloaded, f_{CLK} = max
Vcc current	-	30	-	60	-	30	mA	Vcc = max, CMOS input levels Outputs unloaded, f_{CLK} = max

NOTES
1. LSTTL is equivalent to I_{OH} = 20 microamps, I_{OL} = -0.4mA
2. Current is defined as negative into the device
3. CMOS input levels are defined as:
   V_{IL} = 0.5
   $V_{IH} = V_{DD}$ - 0.5

ORDERING INFORMATION

Comercial (0°C to +70°C)

PDSP16318/C0/AC	(10MHz - PGA)
PDSP16318/C0/GG	(10MHz - QFP)
PDSP16318A/C0/AC	(20MHz - PGA)
PDSP16318A/C0/GG	(20MHz - QFP)

Industrial (-40°C to +85°C)

PDSP16318/B0/AC	(10MHz - PGA)
PDSP16318/B0/GG	(10MHz - QFP)
PDSP16318A/B0/AC	(20MHz - PGA)
PDSP16318A/B0/GG	(20MHz - QFP)

Military (-55°C to +125°C)

PDSP16318/A0/AC	(10MHz - PGA)
PDSP16318/A0/GG	(10MHz - QFP)
PDSP16318A/A0/AC	(20MHz - PGA)
PDSP16318A/A0/GG	(20MHz - QFP)

Call for availability on High Reliability parts and MIL 883C screening.

GEC PLESSEY
SEMICONDUCTORS

PRELIMINARY INFORMATION

DS3884 - 1.0

PDSP16330/A/B

PYTHAGORAS PROCESSOR

The PDSP16330 is a high speed digital CMOS IC that converts Cartesian data (Real and Imaginary) into Polar form (Magnitude and Phase), at rates up to 20MHz. Cartesian 16+16 bit 2's complement or Sign-Magnitude data is converted into 16 bit Phase format. The Magnitude output may be scaled in amplitude by powers of 2. The Phase output represents a full 2 x π field to eliminate phase ambiguities.

Polyimide is used as an inter-layer dielectric and as glassivation.

The PDSP16330 is offered in three speed grades: a basic 10MHz part (PDSP16330), a 20MHz version (PDSP16330A) and a 25MHz version (PDSP16330). A MIL-STD-883 version is also detailed in a separate datasheet.

FEATURES

- 25MHz Cartesian to Polar Conversion
- 16-Bit Cartesian Inputs
- 16-Bit Magnitude Output
- 12-Bit Phase Output
- 2's Complement or Sign-Magnitude Input Formats
- Three-state Outputs and Independent Data Enables Simplify System Interfacing
- Magnitude Scaling Facility with Overflow Flag
- Less than 400 mW Power Dissipation at 10MHz
- 84-pin PGA or 100 pin QFP Package or 84 LCC

APPLICATIONS

- Digital Signal Processing
- Digital Radio
- Radar Processing
- Sonar Processing
- Robotics

ASSOCIATED PPODUCTS

PDSP16112 16 X 12 Complex Multiplier
PDSP16116 16 X 16 Complex Multiplier
PDSP16318 Complex Accumulator
PDSP16340 Polar to Cartesian Converter
PDSP16350 I/Q Splitter and NCO
PDSP16510A Stand Alone FFT Processor

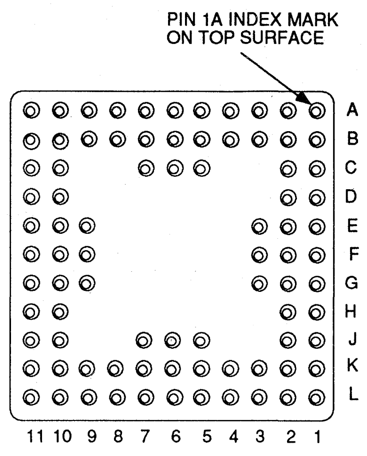

Fig.1 Pin connections - bottom view (PGA)

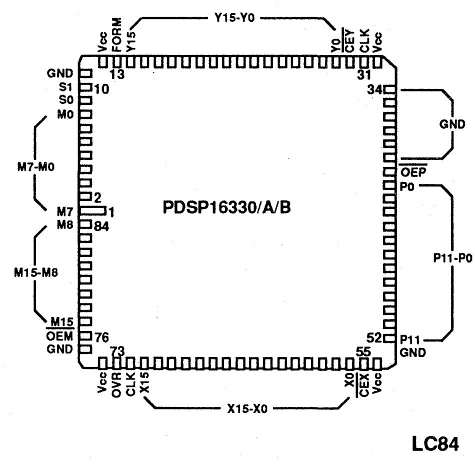

Fig.2 Pin connections - LCC Package

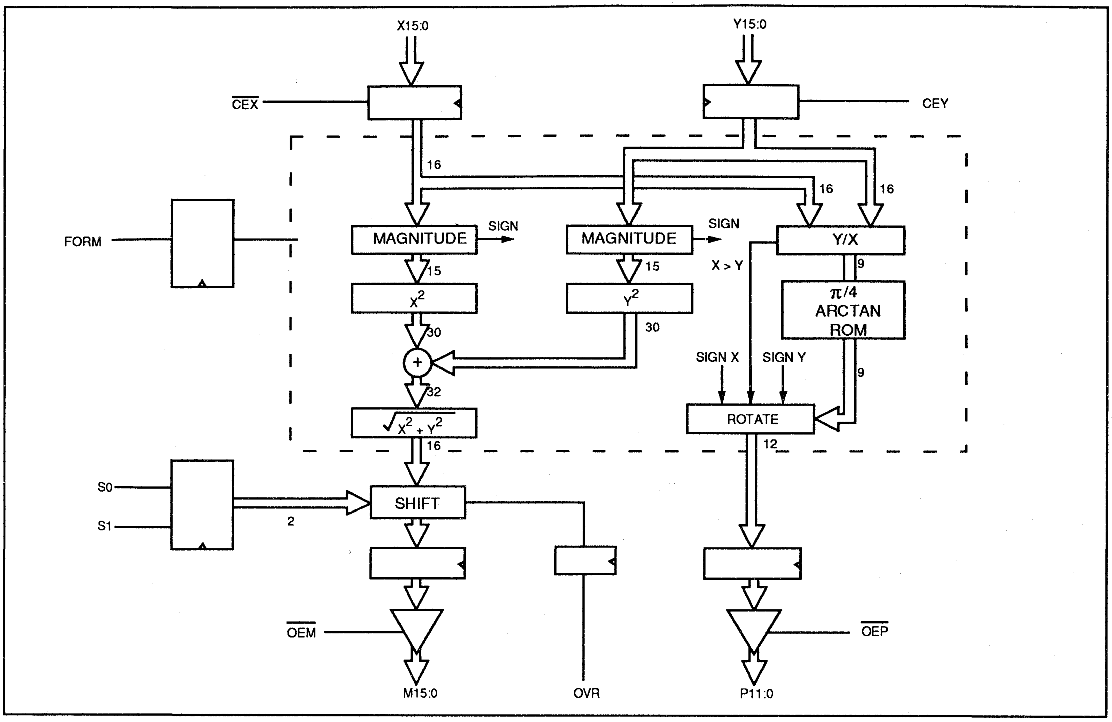

Fig.2 Block diagram

FUNCTIONAL DESCRIPTION

The PDSP16330 converts incoming Cartesian Data into the equivalent Polar Values. The device accepts new 16 + 16 bit complex data every cycle, and delivers a 16 bit + 12 bit Polar equivalent after 24 clock cycles. The input data can be in 2s' Complement or Sign Magnitude format selected via the FORM input. The output is in a magnitude format for both the Magnitude output and the Phase. Phase data is zero for data with a zero Y input and positive X, and is 400 hex for zero X data and positive Y, is 800 hex for zero Y data and negative X, and is C00 hex for zero X and negative Y. The LSB weighting (bit 0) is 2 x π/4096 radians. The 16 bit Magnitude result may be scaled by shifting one, two, or three places in the more significant direction, effectively multiplying the Magnitude result by 2,4 or 8 respectively. Any of these shifts can under certain conditions cause an invalid result to be output from the device. Under these circumstances the OVR output will become active. The PDSP16330 has independent clock enables and three state output controls for all ports.

FORM

This input selects the format of the X and Y input data. A low level on FORM indicates that the Input data is twos' complement format (Note: input data 8000 hex is not valid in 2s' complement mode). This input refers to the format of the current Input data and may be changed on a per cycle basis if desired. The level of FORM is latched at the same time as the data to which it refers.

S1-0

These inputs select the scaling factor to be applied to the Magnitude output. They are latched by the rising edge of CLK and determine the scaling of the output in the cycle after they are loaded into the device. The scale factor applied is determined by the table. Should the scaling factor applied cause an invalid Magnitude result to be output on the M Port, then the OVR Flag will become active for the period that the M Port output is invalid.

S1	S0	Scaling Factor
0	0	x1
0	1	x2
1	0	x4
1	1	x8

The output number range is from 0 to 2 when the scaling factor is set at x1.

PIN DESCRIPTIONS

Symbol	Pin Name and Description
CLK	**Clock:** Common Clock to device Registers. Register contents change on the rising edge of clock. Both pins must be connected.
$\overline{\text{CEX}}$	**Clock Enable:** Clock Enable for X Port. The clock to the X port is enabled by a low level.
$\overline{\text{CEY}}$	**Clock Enable:** Clock Enable for Y Port The clock to the Y port is enabled by a low level.
X15-X0	**X Data Input** Data presented to this input is loaded into the device by the rising edge of CLK. X15 is the MSB
Y15-Y0	**Y Data Input** Data presented to this input is loaded into the device by the rising edge of CLK. Y15 is the MSB
M15-M0	**M Data Output:** Magnitude data generated by the device is output on this port. Data changes on the rising edge of CLK, M15 is the MSB. The weighting of M15 is determined by the Scale factor selected .
P11-P0	**P Data Output:** Phase data generated by the device is output on this port. Data changes on the rising edge of CLK, P11 is the MSB. The weighting of P11 is π radians.
$\overline{\text{OEM}}$	**Output Enable:** Output Enable for M Port. The M Port is in a high impedance state when this input is high.
$\overline{\text{OEP}}$	**Output Enable:** Output Enable for P Port. The P Port is in a high impedance state when this input is high.
FORM	**Format Select** This input selects the format of the Cartesian Data input on the X and Y ports. This input is latched by the rising edge of CLK, and is applied at the same time as the data to which it refers. A low !evel indicates that two's complement data is applied, a high indicates Sign-Magnitude
S1-S0	**Scaling Control:** Control input for scaling of Magnitude Data. This input is latched by the rising edge of CLK, and determines the scaling to be applied to the Magnitude result. The Scaling is applied to the output data in the cycle following the cycle in which the control was latched.
OVR	**Overflow:** Overflow flag. This signal becomes active if the scaling currently selected causes an invalid value to be presented to the Magnitude output.
Vcc	**+5V supply.** All Vcc pins must be connected.
GND	**0V supply.** All GND pins must be connected.

INPUT DATA RANGE

2's Complement	Sign Magnitude
7FFF	7FFF
.	.
.	.
.	.
0001	0001
0000	0000
FFFF	8000
.	.
.	.
.	.
8001	FFF

PIN FUNCTION

Pin No. AC	GG	LC	Function	Pin No. AC	GG	LC	Function	Pin No. AC	GG	LC	Function
F3	91	1	M7	L9	23	29	YO	A9	59	57	X1
G3	92	2	M6	L10	24	30	$\overline{CEY}$	B8	60	58	X2
G1	93	3	M5	K9	25	31	CLK	A8	61	59	X3
G2	94	4	M4	L11	26	32	Vcc	B6	62	60	X4
F1	95	5	M3	K10	31	33	GND	B7	63	61	X5
H1	96	6	M2	J10	32	34	GND	A7	64	62	X6
H2	97	7	M1	K11	33	35	GND	C7	65	63	X7
J1	98	8	M0	J11	34	36	GND	C6	66	64	X8
K1	99	9	S0	H10	35	37	GND	A6	67	65	X9
J2	100	10	S1	H11	36	38	GND	A5	68	66	X10
L1	1	11	GND	F10	37	39	GND	B5	69	67	X11
K2	6	12	Vcc	G10	38	40	$\overline{OEP}$	C5	70	68	X12
K3	7	13	FORM	G11	39	41	P0	A4	71	69	X13
L2	8	14	Y15	G9	40	42	P1	B4	72	70	X14
L3	9	15	Y14	F9	41	43	P2	A3	73	71	X15
K4	10	16	Y13	F11	42	44	P3	A2	74	72	CLK
L4	11	17	Y12	E11	43	45	P4	B3	75	73	OVR
J5	12	18	Y11	E10	44	46	P5	A1	76	74	Vcc
K5	13	19	Y10	E9	45	47	P6	B2	81	75	GND
L5	14	20	Y9	D11	46	48	P7	C2	82	76	$\overline{OEM}$
K6	15	21	Y8	D10	47	49	P8	B1	83	77	M15
J6	16	22	Y7	C11	48	50	P9	C1	84	78	M14
J7	17	23	Y6	B11	49	51	P10	D2	85	79	M13
L7	18	24	Y5	C10	50	52	P11	D1	86	80	M12
K7	19	25	Y4	A11	51	53	GND	E3	87	81	M11
L6	20	26	Y3	B10	52	54	Vcc	E2	88	82	M10
L8	21	27	Y2	B9	57	55	$\overline{CEX}$	E1	89	83	M9
K8	22	28	Y1	A10	58	56	X0	F2	90	84	M8

ELECTRICAL CHARACTERISTICS

Test conditions (unless otherwise stated): T_{amb} (Commercial) = 0°C to + 70°C, T_{amb} (Industrial) = -40°C to + 85°C
V_{cc} (Commercial) = 5.0V ± 5%, V_{cc} (Industrial and Military) = 5.0V ± 1%, GND = 0V

STATIC CHARACTERISTICS

Characteristic	Symbol	Value			Units	Sub-group	Conditions
		Min.	Typ.	Max.			
* Output high voltage	V_{OH}	2.4			V	1,2,3	IOH = 3.2mA
* Output low voltage	V_{OL}			0.6	V	1,2,3	IOL=-3.2mA
* Input high voltage (CMOS)	V_{IH}	3.0			V	1,2,3	Inputs $\overline{CEX}$, $\overline{CEY}$ and CLK only
* Input low voltage (CMOS)	V_{IL}			1.0	V	1,2,3	Inputs $\overline{CEX}$, $\overline{CEY}$ and CLK only
* Input high voltage (TTL)	V_{IH}	2.2			V	1,2,3	All other inputs
* Input low voltage (TTL)	V_{IL}			0.8	V	1,2,3	All other inputs
* Input leakage current (Note 1)	I_{IL}	-10		+ 120	µA	1,2,3	GND $\leq V_{IN} \leq V_{CC}$
† Input capacitance	C_{IN}		10		pF		
* Output leakage current	I_{oz}	-50		+ 50	µA	1,2,3	GND $\leq V_{IN} \leq V_{CC}$
† Output SC current	I_{OS}	-50		230	mA		V_{cc} = Max

NOTES
1. All inputs except clock inputs have high value pull-down resistors
2. All parameters marked * are tested during production. Parameters marked † are guaranteed by design and characterisation.

SWITCHING CHARACTERISTICS

Characteristic	Value						Units	Conditions
	PDSP16330		PDSP16330A		PDSP16330B			
	Min.	Max.	Min.	Max.	Min.	Max.		
† Input data setup to clock rising edge	15		12		12		ns	
† Input data Hold after clock rising edge	2		2		2		ns	
† $\overline{CEX}$, $\overline{CEY}$ Setup to clock rising edge	30		12		12		ns	
† $\overline{CEX}$, $\overline{CEY}$ Hold aher clock rising edge	0		0		0		ns	
† FORM, S1:0 Setup to clock rising edge	15		12		12		ns	
† FORM, S1:0 Hold after clock rising edge	7		2		2		ns	
† Clock rising edge to valid data	5	40	5	25	5	25	ns	2 x LSTTL + 20pF
* Clock period	100		50		40		ns	
† Clock high time	25		15		15		ns	
† Clock low time	25		15		15		ns	
† Latency	24	24	24	24	24	24	cycles	
† $\overline{OEM}$, $\overline{OEP}$ low to data high data valid		30		25	25		ns	2 x LSTTL + 20pF
† $\overline{OEM}$, $\overline{OEP}$ low to data low data valid		30		25	25		ns	2 x LSTTL + 20pF
† $\overline{OEM}$, $\overline{OEP}$ high to data high impedance		30		25	25		ns	2 x LSTTL + 20pF
† $\overline{OEM}$, $\overline{OEP}$ low to data high impedance		30		25	25		ns	2 x LSTTL + 20pF
† Vcc current (TTL input levels)		110		180	225		mA	V_{CC} = Max Outputs unloaded Clock freq. = Max
† Vcc current (CMOS input levels)		70		120	150		mA	V_{CC} = Max Outputs unloaded Clock freq. = Max

NOTES
1. LSTTL is equivalent to I_{OH} = 20µA, I_{OL} = -0.4mA
2. Current is defined as negative into the device
3. CMOS input levels are defined as: $V_{IH} = V_{DD}$ - 0.5V, V_{IL} = +0.5V
4. All parameters marked * are tested during production.
   Parameters marked † are guaranteed by design and characterisation.
5. All timings are dependent on silicon speed. This speed is tested by measuring clock period.
   This guarantees all other timings by characterisation and design.

ABSOLUTE MAXIMUM RATINGS

Supply voltage, V_{cc}	-0.5V to + 7.0V
Input voltage, V_{IN}	-0.5V to VCC + 0.5V
Output voltage, V_{our}	-0.5V to VCC + 0.5V
Clamp diode current per pin, I_K (see Note 2)	±18mA
Static discharge voltage (HMB), V_{STAT}	500V
Storage temperature. T_{stg}	-65°C to + 150°C
Ambient temperature with power applied T_{amb}:	
Commercial	0°C to + 70°C
Industrial	-40°C to + 85°C
Military	-55 °C to + 125°C
Package power dissipation P_{TOT}	1200mW
Junction temperature	150°C

THERMAL CHARACTERISTICS

Package Type	θ_{JC}°C/W	θ_{JA}°C/W
AC	12	36
LC	12	35

NOTES
1. Exceeding these ratings may cause permanent damage. Functional operatlon under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded; only one output to be tested at any one time.
3. Exposure to Absoulte Maximum Ratings for extended periods may affect device reliability.

ORDERING INFORMATION

Commercial (0°C to +70°C)

PDSP16330	CO	LC	(10MHZ - LCC Package)
PDSP16330	CO	AC	(10MHZ - PGA Package)
PDSP16330	CO	GG	(10MHZ - GG Package)
PDSP16330A	CO	LC	(20MHZ - LCC Package)
PDSP16330A	CO	AC	(20MHZ - PGA Package)
PDSP16330A	CO	GG	(20MHZ - GG Package)
PDSP16330B	CO	AC	(25MHZ - PGA Package)

Industrial (-40°C to +85°C)

PDSP16330	BO	LC	10MHZ - LCC Package
PDSP16330	BO	AC	10MHZ - PGA Package
PDSP16330	BO	GG	10MHZ - GG Package
PDSP16330A	BO	LC	20MHZ - LCC Package
PDSP16330A	BO	AC	20MHZ - PGA Package
PDSP16330A	BO	GG	20MHZ - GG Package
PDSP16330B	BO	AC	25MHZ - PGA Package

Military (-55°C to +125°C)

PDSP16330	AO	LC	10MHZ - LCC Package
PDSP16330	AO	AC	10MHZ - PGA Package
PDSP16330	AO	GG	10MHZ - GG Package
PDSP16330A	AO	LC	20MHZ - LCC Package
PDSP16330A	AO	AC	20MHZ - PGA Package
PDSP16330A	AO	GG	20MHZ - GG Package

Test	Waveform - measurement level
Delay from ouput high to output high impedance	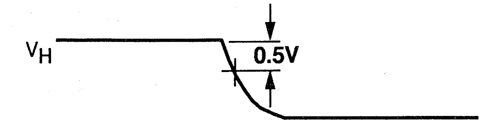
Delay from ouput low to output high impedance	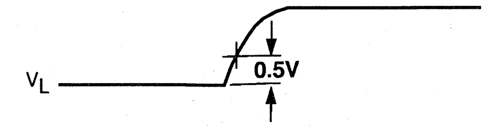
Delay from ouput high impedance to Output low	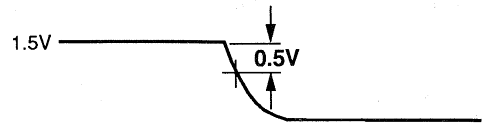
Delay from ouput high impedance to Output high	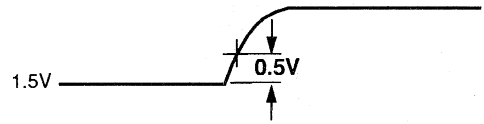

NOTES
1. V_H - Voltage reached when output driven high
2. V_L - Voltage reached when output driven low

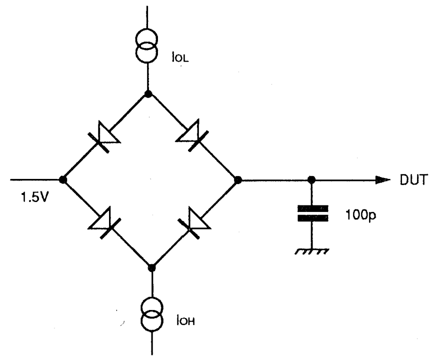

Fig.3 Three state delay measurement load

PDSP16340

POLAR TO CARTESIAN CONVERTER

The PDSP16340 can be configured to perform either a coordinate conversion function, or simply to provide a sine / cosine look-up table. When employed as a coordinate conversion processor, the device converts data from 16 bit polar coordinates (R,Ø) into 16 bit cartesian coordinates (Real, Imaginary). The translation is illustrated in Fig. 1, and uses the formula:-

$Xr = R \cos(Ø)$
$Xi = R \sin(Ø)$

In look-up table mode, the user enters 16 bit phase data, and the chip outputs the corresponding sine and cosine values. A typical application is shown in Fig. 5.

The PDSP16340 is pipelined to process a continuous stream of data at 20 MHz, and outputs a new (16 + 16) bit result every clock cycle. The RANGE control signal allows the user to select the input range most appropriate to the system. Data is produced in Two's Complement Fractional format.

APPLICATIONS

- Digital Signal Processing
- Radar Systems
- Sonar Systems
- Robotics
- Medical Imaging

FEATURES

- Provides R cos(Ø) and R sin(Ø) in 16 bit streams using a CORDIC processor
- Look-up table equivalent to 64k by 32 bit ROM
- 20MHz clock rate
- Tri-state outputs and independent data enables
- 84 Pin PGA or 132 pin QFP

ASSOCIATED PRODUCTS

PDSP16330	Pythagoras Processor
PDSP16256	Programmable FIR Filter
PDSP16510	FFT Processor
PDSP16350	I/Q Splitter and NCO
PDSP16116	16 Bit Complex Multiplier
PDSP16318	Complex Accumulator

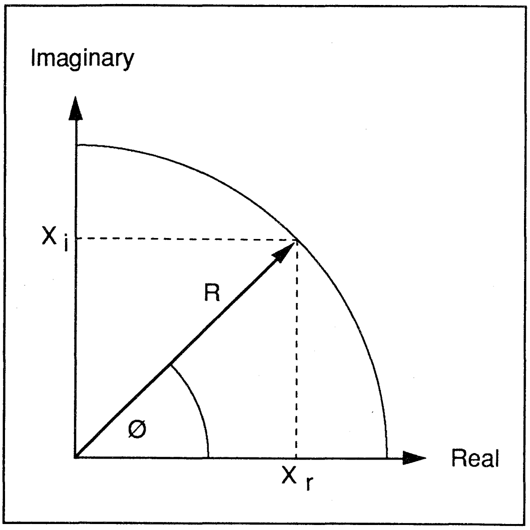

Fig. 1. Cartesian to Polar Coordinates

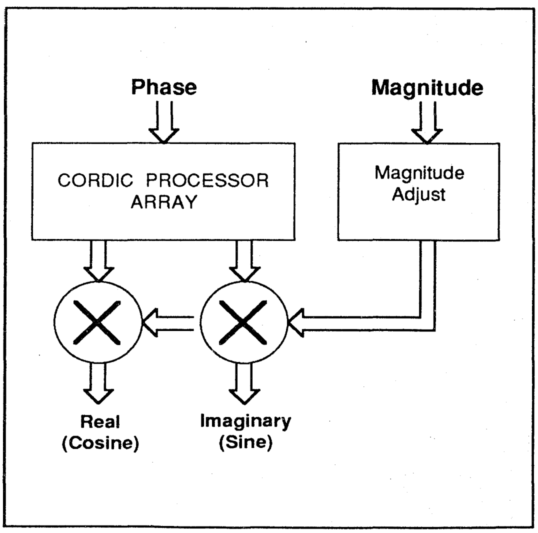

Fig. 2. Simplified Block Diagram

	1	2	3	4	5	6	7	8	9	10	11	12	13
N	PEN	MODE	M1	M3	M5	VDD	M8	GND	M10	M12	M14	VOUT	SAT
M	RANGE		M0	M2	M4	M6	M7	M9	M11	M13	M15		VIN
L	P15	O/C										XI15	XI14
K	P13	P14										XI13	XI12
J	P11	P12										XI11	XI10
H	GND	P10										XI9	GND
G	P9	P8										XI8	XI7
F	VDD	P7										XI6	VDD
E	P6	P5										XI4	XI5
D	P4	P3										XI2	XI3
C	P2	P1										XI0	XI1
B	P0		XR15	XR13	XR11	XR9	XR7	XR6	XR4	XR2	XR0		$\overline{\text{MEN}}$
A	CLOCK	GND	XR14	XR12	XR10	VDD	XR8	GND	XR5	XR3	XR1	$\overline{\text{OEI}}$	$\overline{\text{OER}}$

Fig. 3A Device Pinout - Bottom View (84 pin PGA - AC84)

GC	SIG	GC	SIG	GC	SIG	GC	SIG
1	N/C	34	N/C	67	GND	100	GND
2	$\overline{\text{MEN}}$	35	VOUT	68	RANGE	101	VDD
3	N/C	36	M15	69	N/C	102	GND
4	XI0	37	GND	70	N/C	103	N/C
5	XI1	38	VDD	71	P15	104	XR15
6	XI2	39	M14	72	GND	105	XR14
7	GND	40	N/C	73	VDD	106	N/C
8	VDD	41	M13	74	P14	107	XR13
9	XI3	42	M12	75	P13	108	XR12
10	XI4	43	N/C	76	P12	109	N/C
11	N/C	44	M11	77	N/C	110	XR11
12	XI5	45	M10	78	P11	111	N/C
13	XI6	46	N/C	79	P10	112	XR10
14	N/C	47	M9	80	N/C	113	XR9
15	XI7	48	GND	81	P9	114	VDD
16	XI8	49	VDD	82	GND	115	GND
17	VDD	50	M8	83	VDD	116	XR8
18	GND	51	M7	84	P8	117	XR7
19	XI9	52	M6	85	P7	118	N/C
20	N/C	53	M5	86	P6	119	XR6
21	XI10	54	VDD	87	N/C	120	XR5
22	XI11	55	M4	88	P5	121	N/C
23	N/C	56	VDD	89	N/C	122	XR4
24	XI12	57	M3	90	P4	123	N/C
25	XI13	58	GND	91	P3	124	XR3
26	XI14	59	M2	92	VDD	125	XR2
27	VDD	60	M1	93	GND	126	N/C
28	GND	61	GND	94	P2	127	XR1
29	XI15	62	VDD	95	P1	128	VDD
30	VIN	63	M0	96	N/C	129	GND
31	N/C	64	MODE	97	P0	130	XR0
32	N/C	65	$\overline{\text{PEN}}$	98	N/C	131	$\overline{\text{OE1}}$
33	SAT	66	VDD	99	CLK	132	$\overline{\text{OE2}}$

Fig.3B Pin out Table (132 pin ceramic QFP - GC132)

SIGNAL	DESCRIPTION
M15:0	16 bit 2's complement data representing the magnitude of the phase angle. Data is loaded into the input register on the rising edge of CLK. These inputs are not used in look-up table mode, however, they should be tied high or low for electrical, rather than logical, reasons. M15 is the MSB.
P15:0	16 bit data representing the phase angle. Data is loaded into the input register on the rising edge of CLK. P15 is the MSB.
XR15:0	16 bit 2's complement real data output, or cosine output in the table look-up mode. Data is passed to the XR outputs on the rising edge of CLK.
XI15:0	16 bit 2's complement imaginary data output, or sine output in the table look-up mode. Data is passed to the XI outputs on the rising edge of CLK.
RANGE	Magnitude range select. When this pin is high, the MSB of the M input bus (also the sign bit) will represent 2^1. When low, it will represent 2^0
SAT	Input data saturated flag. This output goes high to indicate that input data of magnitude greater than SQRT(2) has been saturated to SQRT(2). It is internally delayed such that it appears at the output at the same time as the data which resulted from the saturated input value.
$\overline{\text{MEN}}$	Clock enable for the magnitude input port. When low new data may be latched in the input register; when high the register remains in its previous state.
$\overline{\text{PEN}}$	Clock enable for the phase input port. When low new data may be latched in the input register: when high the register remains in its previous state.
$\overline{\text{OER}}$	Output enable for the XR output port. When high the XR output is forced into a high impedance state.
$\overline{\text{OEI}}$	Output enable for the XI output port. When high the XI output is forced into a high impedance state.
VIN	Valid data input flag. This input is connected to VOUT via a pipeline delay which matches the data path pipeline delay. Hence, if VIN is set high when valid data is input, then VOUT will go high when valid results are output. It performs no internal control function.
VOUT	Valid data output flag which is a delayed version of VIN as explained above.
MODE	When high, this input configures the chip into look-up table mode in which the M inputs are redundant and internally replaced by a unity magnitude. When low, the chip is configured in coordinate conversion mode.
CLK	Common clock to all internal registers.
VDD	Four +5V power pins. All power supply pins must be connected.
GND	Four ground pins. All pins must be connected.

Table 1. Signal description

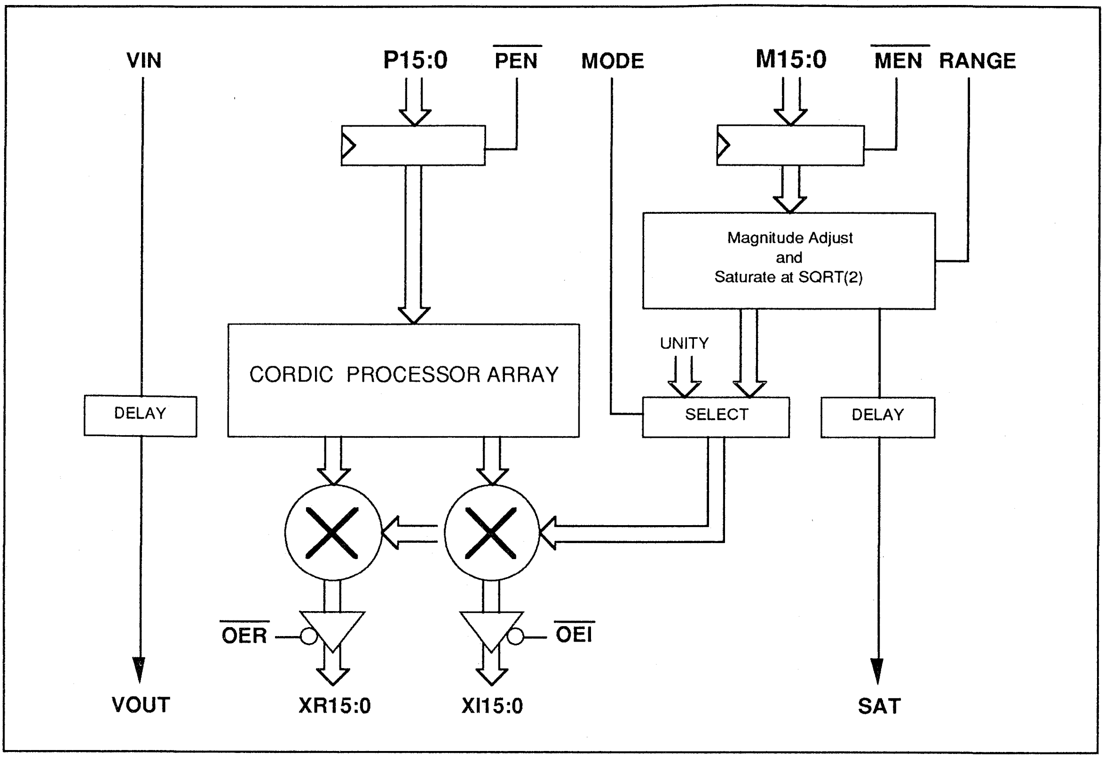

Fig. 4. Internal Block Diagram

OPERATION

The functional blocks used within the device are illustrated by Fig. 4. Both input data and output data are fully registered to allow the device to be easily incorporated into data flow DSP systems. The sine and cosine values are actually calculated in a 26 stage pipelined arithmetic processor, and are accurate to 16 bits. This technique allows high data throughputs, and requires less die area than the equivalent ROM.

The PDSP16340 has two modes of operation, which are selected by the logical state of the MODE input pin. This pin should be tied high or low to suite the particular application.

Look-up Mode

In the Table Look-up mode the MODE pin is tied high, and the device is used to provide simultaneous sine and cosine values at rates up to the maximum clock frequency. A new phase value is clocked into the Phase Port (P15:0) on each cycle, and the corresponding sine and cosine values appear at the XI and XR ports 29 clock cycles later. In this operating mode the MAGNITUDE inputs, the $\overline{\text{MEN}}$, and the RANGE inputs are logically redundant. They must, however, be tied either high or low for electrical reasons. If the Phase Port is disabled by pulling $\overline{\text{PEN}}$ high, then the look-up table will continue to provide the sine and cosine outputs corresponding to the value of P15:0 present during the active clock edge before the $\overline{\text{PEN}}$ level change.

Fig. 5. illustrates a typical FFT arrangement with the PDSP16340 providing sine and cosine 'twiddle' factors for use by the butterfly processor. Use of the PDSP16520 Quad Port RAM, and the PDSP16116 / 318 complex arithmetic element, allows butterfly calculations to be performed at rates up to 20 MHz.

Coordinate Conversion

In the Coordinate Conversion Processor mode the MODE pin is tied low, and the PDSP16340 converts data from polar format into the corresponding real and imaginary Cartesian co-ordinates. The coordinate conversion operation is equivalent to the inverse of the function performed by the PDSP16330 Pythagoras Processor. The device produces simultaneous sine and cosine values from the incoming phase angle, and then multiplies these results with the appropriate magnitude value. The $\overline{\text{MEN}}$ input allows the value in the input latch to be retained in a similar manner to the use of the $\overline{\text{PEN}}$ control.

The RANGE control allows the device to accept magnitude data in the range of, either, -1 to within one LSB of +1, or from -2 to within one LSB of +2. The smaller range option allows maximum accuracy to be preserved, if fractional inputs are expected.The latter option enables the theoretical maximum polar magnitude of SQRT(2) to be accommodated. A negative magnitude introduces a 180° phase shift.

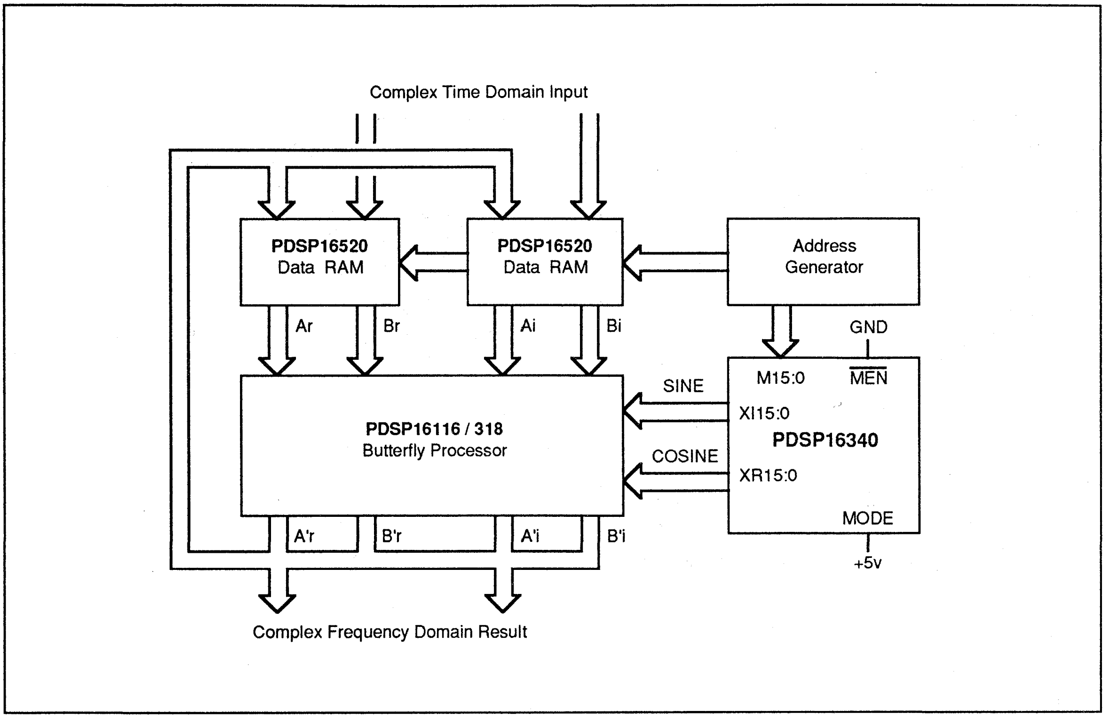

Fig. 5. Sin / Cos generator for 20 MHz FFT System

The device will replace all incoming values above the square root of two with the maximum value. The SAT output indicates when this replacement has internally occurred. The flag is delayed such that it is valid at the same time as the output data which was calculated from the saturated input.

DATA FORMATS

When the device is configured in the co-ordinate conversion mode (MODE pin is low), the magnitude (M) input bus can have one of the following data formats:

BIT NUMBER	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
WEIGHTING RANGE = 1	S	2^{0}	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-8}	2^{-9}	2^{-10}	2^{-11}	2^{-12}	2^{-13}	2^{-14}
RANGE = 0	S	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-8}	2^{-9}	2^{-10}	2^{-11}	2^{-12}	2^{-13}	2^{-14}	2^{-15}

The sign bit is provided to maintain compatibility with normal arithmetic procedures, but in most applications the value will always be positive. The sign bit could then be tied low, and the lower fifteen bits used to define the input. If a negative value is used this will introduce a 180° phase shift. When the MODE pin is high the state of the RANGE pin is irrelevant, and the magnitude is internally defined to be unity.

The PHASE port has the following data format:

BIT NUMBER	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
WEIGHTING in Π radians	2^{0}	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-8}	2^{-9}	2^{-10}	2^{-11}	2^{-12}	2^{-13}	2^{-14}	2^{-15}

Thus, for example :

+90° (= - 270°) = 0100000000000000
-180° (= +180°) = 1000000000000000
-90° (= +270°) = 1100000000000000

The 16 bit radius value is multiplied with the 16 bit internally generated sine and cosine values, to produce a 16 bit result. The RANGE input controls the format of the output data as given below:

BIT NUMBER	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
WEIGHTING RANGE = 1	S	2^{0}	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-8}	2^{-9}	2^{-10}	2^{-11}	2^{-12}	2^{-13}	2^{-14}
RANGE = 0 OR MODE = 1	S	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-8}	2^{-9}	2^{-10}	2^{-11}	2^{-12}	2^{-13}	2^{-14}	2^{-15}

ABSOLUTE MAXIMUM RATINGS (Note 1)

Supply voltage Vcc	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to Vcc + 0.5V
Output voltage V_{OUT}	-0.5V to Vcc + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	
Military	-55°C to +125°C
Industrial	-40°C to 85°C
Junction temperature	150°C
Package power dissipation	3500mW
Thermal resistances	
Junction to Case $Ø_{JC}$	5°C/W

NOTES
1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Vcc = Max. Outputs Unloaded, Clock Freq = Max.
5. CMOS levels are defined as
   V_{IH} = Vcc - 0.5v
   V_{IL} = +0.5v
6. Current is defined as negative into the device.
7. $Ø_{JC}$ data assumes that heat is extracted from the top face of the package.

ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)

Commercial: T_{AMB} = -0°C to +70°C $T_{J(MAX)}$ = 95°C V_{CC} = 5.0V±5% Ground = 0V
Industrial: T_{AMB} = -40°C to +85°C $T_{J(MAX)}$ = 110°C V_{CC} = 5.0V±10% Ground = 0V
Military: T_{AMB} = -55°C to +125°C $T_{J(MAX)}$ = 150°C V_{CC} = 5.0V±10% Ground = 0V

Static Characteristics

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	3.0		-	V	
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{CC}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{CC}
Output S/C current	I_{SC}	10		250	mA	V_{CC} = Max

Switching Characteristics

Characteristic	Industrial			Military			Units	Conditions
	Min.	Typ.	Max.	Min.	Typ.	Max.		
M15:0 or P15:0 setup to clock rising edge	15		-	15		-	ns	
M15:0 or P15:0 hold after clock rising edge	4		-	4		-	ns	
$\overline{\text{MEN}}$ or $\overline{\text{PEN}}$ setup to clock rising edge	20		-	20		-	ns	
$\overline{\text{MEN}}$ or $\overline{\text{PEN}}$ hold after clock rising edge	0		-	0		-	ns	
RANGE setup to clock rising edge	15		-	15		-	ns	
RANGE hold after clock rising edge	8		-	8		-	ns	
Clock rising edge to all outputs valid	5		30	5		30	ns	30pF
Clock freq	DC		20	DC		20	MHz	
Clock High Time	15		-	15		-	ns	
Clock Low Time	20		-	20		-	ns	
$\overline{\text{OER}}$,$\overline{\text{OEI}}$ low to data valid	-		20	-		20	ns	see Fig. 6
$\overline{\text{OER}}$,$\overline{\text{OEI}}$ high to data high impedance	-		20	-		20	ns	see Fig. 6
Pipeline delay VIN to VOUT	28		28	28		28	CLKs	
Vcc Current (CMOS inputs)	-		430	-		450	mA	see Note 4
Vcc Current (TTL inputs)	-		460	-		500	mA	see Note 4

Test	Waveform - measurement level
Delay from output high to output high impedance	V_H 0.5V
Delay from output low to output high impedance	V_L 0.5V
Delay from output high impedance to output low	1.5V 0.5V
Delay from output high impedance to output high	1.5V 0.5V
V_H - Voltage reached when output driven high V_L - Voltage reached when output driven low	

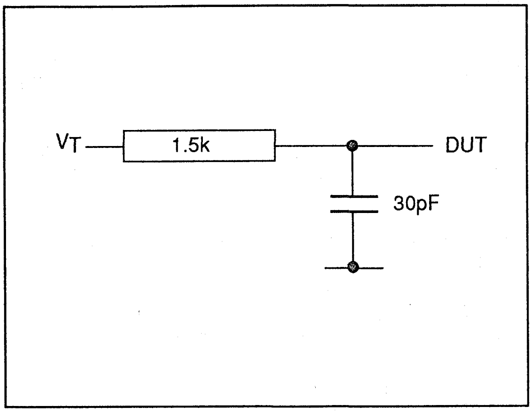

Fig. 6 Tri-state delay measurement load.

ORDERING INFORMATION

Commercial (0°C to +70°C)

PDSP16340 / C0 / AC (20MHz - PGA)
PDSP16340 / C0 / GC (20MHz - QFP)

Industrial (-40°C to +85°C)

PDSP16340 / B0 / AC (20MHz - PGA)
PDSP16340 / B0 / GC (20MHz - QFP)

Military (-55°C to +125°C)

PDSP16340 / A0 / AC (20MHz - PGA)
PDSP16340 / A0 / GC (20MHz - QFP)

Call for availability of High Rel parts and MIL 883C screening.

PDSP16350

I/Q SPLITTER / NCO

The PDSP16350 provides an integrated solution to the need for very accurate, digitised, sine and cosine waveforms. Both these waveforms are produced simultaneously, with 16 bit amplitude accuracy, and are synthesised using a 34 bit phase accumulator. The more significant bits of this provide 16 bits of phase accuracy for the sine and cosine look up tables.

With a 20 MHz system clock, waveforms up to 10 MHz can be produced, with 0.001 Hz resolution. If frequency modulation is required with no discontinuities, the phase increment value can be changed linearly on every clock cycle. Alternatively absolute phase jumps can be made to any phase value.

The provision of two output multipliers allows the sine and cosine waveforms to be amplitude modulated with a 16 bit value present on the input port. This option can also be used to generate the in-phase and quadrature components from an incoming signal. This I/Q split function is required by systems which employ complex signal processing.

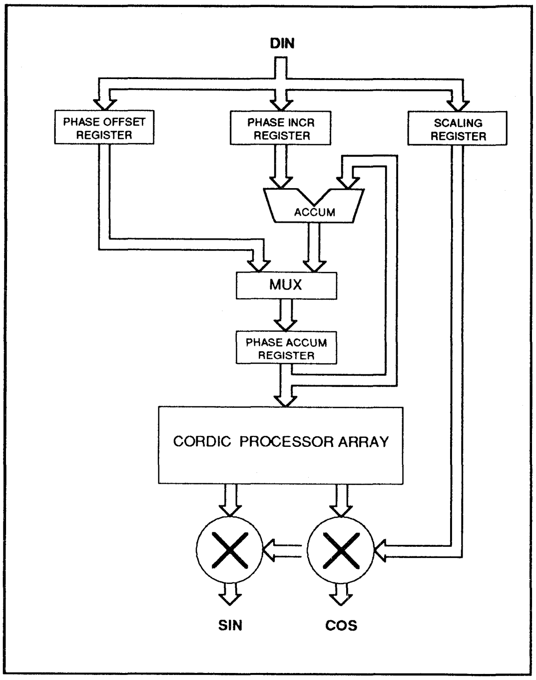

Fig. 1 Block Diagram

FEATURES

- Direct Digital Synthesiser producing simultaneous sine and cosine values
- 16 bit phase and amplitude accuracy, giving spur levels down to - 90 dB
- Synthesised outputs from DC to 10 MHz with accuracies better than 0.001 Hz
- Amplitude and Phase modulation modes
- 84 pin PGA or 132 pin QFP

APPLICATIONS

- Numerically controlled oscillator (NCO)
- Quadrature signal generator
- FM, PM, or AM signal modulator
- Sweep Oscillator
- High density signal constellation applications with simultaneous amplitude and phase modulation
- VHF reference for UHF generators
- Signal demodulator

ASSOCIATED PRODUCTS

PDSP16256 Programmable FIR Filter

PDSP16510 FFT Processor

PDSP16340 Polar to Cartesian Converter

PDSP16488 2D Convolver

N	JUMP	MODE	DIN19	DIN21	DIN23	VDD	DIN26	GND	DIN28	DIN30	DIN32	VOUT	RES
M	DIN17		DIN18	DIN20	DIN22	DIN24	DIN25	DIN27	DIN29	DIN31	DIN33		VIN
L	DIN15	DIN16										SIN15	SIN14
K	DIN13	DIN14										SIN13	SIN12
J	DIN11	DIN12										SIN11	SIN10
H	GND	DIN10										SIN9	GND
G	DIN9	DIN8										SIN8	SIN7
F	VDD	DIN7										SIN6	VDD
E	DIN6	DIN5										SIN4	SIN5
D	DIN4	DIN3										SIN2	SIN3
C	DIN2	DIN1										SIN0	SIN1
B	DIN0		COS15	COS13	COS11	COS9	COS7	COS6	COS4	COS2	COS0		CEN
A	CLOCK	GND	COS14	COS12	COS10	VDD	COS8	GND	COS5	COS3	COS1	$\overline{OES}$	$\overline{OEC}$
	1	2	3	4	5	6	7	8	9	10	11	12	13

Fig. 2 A. Pin out - bottom view (84 pin PGA - AC84)

GC	SIG	GC	SIG	GC	SIG	GC	SIG
1	N/C	34	N/C	67	GND	100	GND
2	$\overline{\text{CEN}}$	35	VOUT	68	DIN17	101	VDD
3	N/C	36	DIN33	69	N/C	102	GND
4	SIN0	37	GND	70	DIN16	103	N/C
5	SIN1	38	VDD	71	DIN15	104	COS15
6	SIN2	39	DIN32	72	GND	105	COS14
7	GND	40	N/C	73	VDD	106	N/C
8	VDD	41	DIN31	74	DIN14	107	COS13
9	SIN3	42	DIN30	75	DIN13	108	COS12
10	SIN4	43	N/C	76	DIN12	109	N/C
11	N/C	44	DIN29	77	N/C	110	COS11
12	SIN5	45	DIN28	78	DIN11	111	N/C
13	SIN6	46	N/C	79	DIN10	112	COS10
14	N/C	47	DIN27	80	N/C	113	COS9
15	SIN7	48	GND	81	DIN9	114	VDD
16	SIN8	49	VDD	82	GND	115	GND
17	VDD	50	DIN26	83	VDD	116	COS8
18	GND	51	DIN25	84	DIN8	117	COS7
19	SIN9	52	DIN24	85	DIN7	118	N/C
20	N/C	53	DIN23	86	DIN6	119	COS6
21	SIN10	54	VDD	87	N/C	120	COS5
22	SIN11	55	DIN22	88	DIN5	121	N/C
23	N/C	56	GND	89	N/C	122	COS4
24	SIN12	57	DIN21	90	DIN4	123	N/C
25	SIN13	58	VDD	91	DIN3	124	COS3
26	SIN14	59	DIN20	92	VDD	125	COS2
27	VDD	60	DIN19	93	GND	126	N/C
28	GND	61	GND	94	DIN2	127	COS1
29	SIN15	62	VDD	95	DIN1	128	VDD
30	VIN	63	DIN18	96	N/C	129	GND
31	N/C	64	MODE	97	DIN0	130	COS0
32	N/C	65	JUMP	98	N/C	131	$\overline{\text{OES}}$
33	RESET	66	VDD	99	CLK	132	$\overline{\text{OEC}}$

Fig.2B Pin out (132 pin ceramic QFP - GC132)

SIGNAL	DESCRIPTION
DIN33:0	Data bus for the input register. This input register provides a 34 bit, incremental or absolute, phase value, if the mode pin is low. Alternatively if the mode pin is high, it provides either an 18 bit phase increment value, via D17:0, and a 16 bit scale value via D33:18 or a 34 bit phase increment value depending on the JUMP input see below.
SIN15:0	16 bit sine output data in fractional two's complement format.
COS15:0	16 bit cosine output data in fractional two's complement format.
$\overline{\text{CEN}}$	Clock enable for the data input register. When low, data will be latched on the rising edge of the clock. When high data will be retained in the input register.
MODE	Mode control input. When low, data in the input register is interpreted as either a 34 bit phase increment value or a 34 bit absolute phase value. When high, the output multipliers are enabled and will scale the waveforms with the upper 16 bits in the input register. The phase increment is loaded from the the lower 18 bits. The full 34 bit phase increment register can also be loaded using JUMP see below.
JUMP	With MODE low (Frequency or Phase Modulation) When low JUMP will allow normal phase incrementing to occur. When high, the data on the input pins will be interpreted as a 34 bit absolute phase value to replace the present value in the accumulator. JUMP is internally latched to match the delay through the data input register, and to allow data in the internal pipeline to be correctly processed. $\overline{\text{CEN}}$ must also be low to latch the required data from DIN. When Mode is high (Amplitude Modulation) When low JUMP will allow normal phase incrementing to occur, with the phase increment value taken from the lower 18 data inputs. When high, the data on the input pins will replace the full 34 bits of the phase increment register. $\overline{\text{CEN}}$ must also be low to latch the required data.
RES	When high will clear the phase accumulator and phase increment registers, after data in the internal pipeline has been correctly processed.
CLK	Input clock.
$\overline{\text{OES}}$	Output enable for SIN 15:0. Outputs are high impedance when $\overline{\text{OES}}$ is high.
$\overline{\text{OEC}}$	Output enable for COS15:0. Outputs are high impedance when $\overline{\text{OEC}}$ is high.
VIN	Valid input flag. A delayed version of this input is available on the VOUT pin, with the delay matching the data processing pipeline delay. This input has no other internal function.
VOUT	Valid output flag. See above.
GND	Five ground pins. All must be connected.
VCC	Four +5V pins. All must be connected.

Table 1. Pin Description

DEVICE OPERATION

Sine and cosine are simultaneously produced by the Cordic processor, which is addressed by the upper 16 bits of the output from a 34 bit phase accumulator. The accumulator divides the digital phase circle into a number of steps, one step for each state of the accumulator. When the accumulator reaches its maximum value it overflows back to zero and the sequence is repeated.

The accumulator is incremented once per incoming clock cycle, by an amount which defines the frequency which is to be generated. The increment required is defined by :

$$\text{Increment} = \frac{\text{Desired O/P Frequency}}{\text{Incoming Clock Frequency}} \times 2^N$$

where N is the number of bits in the accumulator. Since the Nyquist criteria for proper waveform reconstruction must still be obeyed, the maximum output frequency is half the incoming frequency. In practice, when a return is made to the analog world, just meeting the minimum Nyquist requirement would require a 'brick wall' low pass filter to remove the alias signals. A more useful 'rule of thumb' is to limit the generated waveforms to less than 40% of the clock frequency.

The resolution, or tuning sensitivity, of the waveform generator is given by :

$$\text{Resolution} = \frac{\text{Incoming Clock Frequency}}{2^N} \text{ Hz}$$

These equations illustrate some very important features of direct digital synthesisers :-

1) Tuning sensitivity is defined by both the number of bits in the accumulator and the incoming time base frequency.
2) The oscillator tunes linearly over its entire range.
3) The frequency accuracy matches the accuracy of the incoming increment value.
4) DC can be generated since the increment value can be zero.
5) Frequency stability will match the stability of the incoming frequency when the increment is fixed.

The residual noise characteristics of an oscillator are very important in modern communication systems. This parameter defines how well the device maintains its set frequency for very short periods (nanoseconds to seconds) of time. Poor figures will significantly affect the system signal to noise ratio and limit the dynamic range.

The PDSP16350 will, of course, inherit the residual noise characteristics of the source of the incoming frequency. The output frequency is, however, always less than half the incoming frequency in order to satisfy the Nyquist criterion. This is in contrast to a phase locked loop synthesiser, when a small input frequency controls a high output frequency.

The commonly used 20 log N rule states that the phase noise at the output of a synthesiser will be no better than twenty times the log of the ratio of the output frequency to the input frequency. In a phase locked loop synthesiser N is large, in the PDSP16350 it is less than half. Log N is thus less than zero and phase noise improvement is obtained.

The output waveforms are produced after a pipeline delay with respect to the DIN inputs. The effects of the JUMP or RES commands are delayed such that all data in the internal pipe will be processed before the discontinuity occurs. New data may be presented to the device on the cycle following the JUMP or RES and a valid result will be obtained after 31 clock cycles.

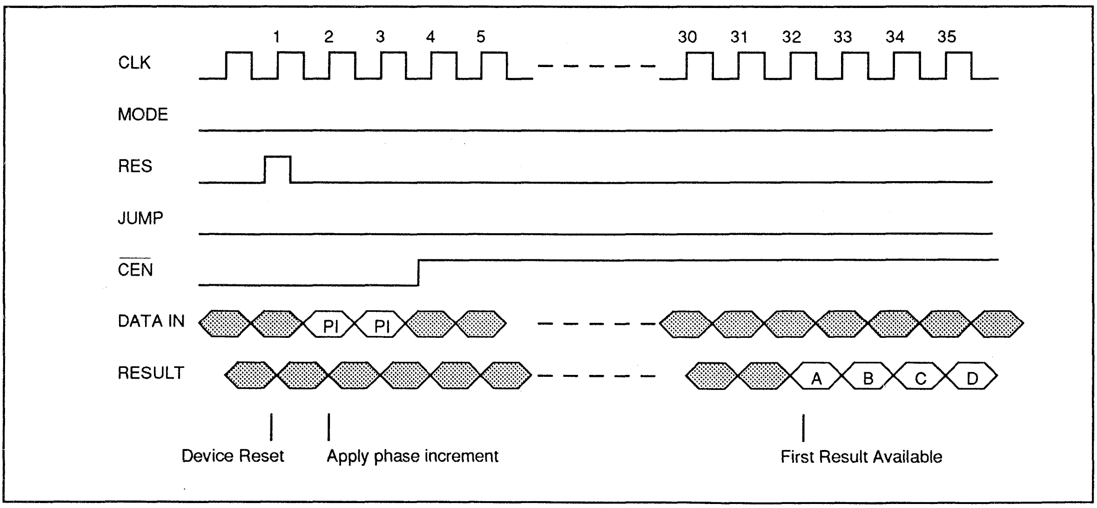

Fig. 3 Fixed Frequency Timing Diagram

USING THE PDSP16350

Frequency, phase, and amplitude modulation are all possible with the PDSP16350. The former two requirements are satisfied by the ability to change the phase increment value on every clock cycle. The latter needs the addition of two multipliers, which allow both sine and cosine to be modified by an incoming waveform.

Fixed Frequency, Constant Amplitude

To generate sine and cosine outputs at a fixed frequency, the MODE pin should be tied low, see Fig. 3. The phase increment value required to generate the desired frequency should be clocked into the internal phase increment register. This value is entered via the DIN port with $\overline{\text{CEN}}$ low. If $\overline{\text{CEN}}$ subsequently goes inactive (high), the value need not be maintained on the input pins.

The correct phase increment value can be calculated as follows :

$$\text{DIN value} = \frac{\text{Desired O/P Frequency}}{\text{Clock Frequency}} \times 2^{34}$$

This will give a decimal value which must be converted to a 34 bit binary number. The frequency resolution of the generated waveforms will be :

$$\text{Resolution} = \frac{\text{Clock Frequency}}{2^{34}} \ \text{Hz}$$

With a 20 MHz clock this results in a frequency resolution of 0.001 Hz. This can be improved by reducing the clock frequency, with the Nyquist restraint being the limiting factor. The latter states that the frequency of the generated waveform must be no more than 50% of the input clock. In practice 40% is a better limit to use, as previously discussed.

A practical example can be used to illustrate the calculation. With a clock frequency of 10.73864 MHz, and the need to generate an output frequency of 20 kHz, then the above equation tells us we need a DIN value of 31996359. This corresponds to a binary value of:

DIN33:0 = 00 0000 0001 1110 1000 0011 1001 1100 0111

The resolution would be 0.0006 Hz. It should be noted that the accuracy of the PDSP16350 cannot be any better than the accuracy of the incoming clock, and these resolutions are based on perfect incoming waveforms.

Fixed Frequency, Modulated Amplitude

The MODE pin should be high if modulation of the output waveforms is required. In this mode each of the output waveforms is multiplied by the 16 bit, two's complement, value, present on the most significant 16 bits of the DIN port. The phase increment register is normally loaded with the 18 bit value on the least significant portion of the DIN bus. It is also possible to load the full 34 bits of the phase increment register when greater accuracy is required, this is explained below. When using the full 34 bits it is possible to obtain the same frequency resolution as in the fixed amplitude mode described earlier. When using 18 bit accuracy directly from the DIN bus the correct phase increment value can be calculated as follows :

$$\text{DIN value} = \frac{\text{Desired O/P Frequency}}{\text{Clock Frequency}} \times 2^{18}$$

The frequency resolution is correspondingly reduced and given by :

$$\text{Resolution} = \frac{\text{Clock Frequency}}{2^{18}} \ \text{Hz}$$

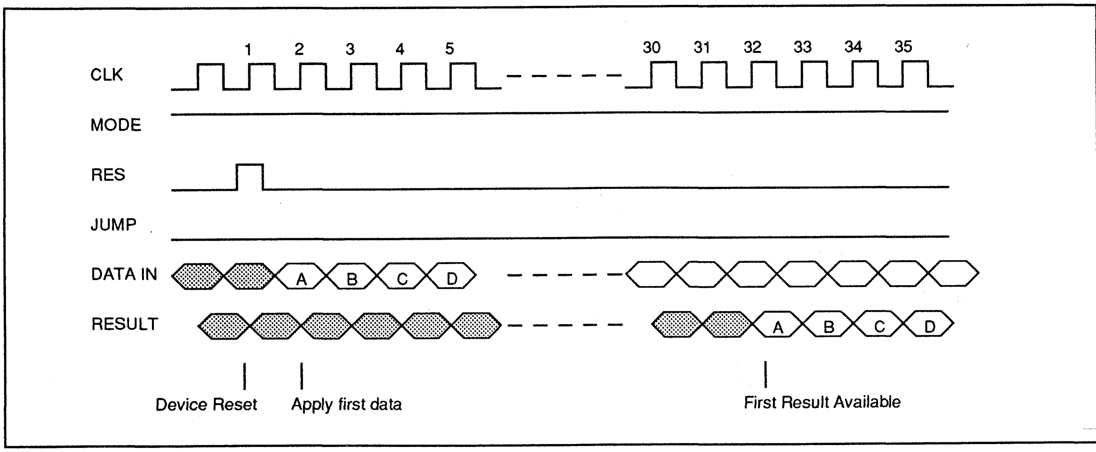

Fig. 4 Amplitude Modulation (18bit frequency accuracy)

Fig. 4 shows the operation of the device when loading the phase increment directly from the DIN bus. First the device must be reset then data is presented on each clock cycle. The amplitude modulation value is presented on the most significant 16 bits while the phase increment is presented on the least significant 18 bits. The first valid result is obtained after 31 cycles. (In this mode the least significant 16 bits of the phase increment register remain low).

Fig.6 shows the operation of the device when using the full 34 bits of the phase increment register. First the device must be reset, then the full 34 bits of the phase increment register are loaded from the DIN bus by taking signal JUMP high before the rising edge of the clock. Following this new data can be presented on each cycle of the clock. The amplitude modulation value is presented on the most significant 16 bits while the phase increment is presented on the least significant 18 bits. The least significant 16 bits of the phase increment register remain fixed at the value loaded using JUMP. The first valid result is obtained after 31 cycles. When using JUMP to load the phase increment register, normal operation cannot be maintained. This is because the amplitude modulation value normally presented on the most significant 16 bits of the DIN bus are replaced by part of the new phase increment value.

The AM mode is useful in systems requiring frequency sweeps. By varying the amplitudes at different frequencies, it is possible to compensate for the analog gain characteristics of amplifiers further along in the system.

It can also be used to generate the in-phase and quadrature components of an analog waveform, which has been digitized and which is to be processed using complex techniques. Such a quadrature heterodyning system, alternatively known as an IQ splitter, is shown in Fig. 5.

The output from an A/D converter drives the D33:18 inputs of the PDSP16350. If all sixteen inputs are not required, the unused least significant bits should be tied to ground, and the more significant inputs connected to the A/D converter. Multiplying an input signal with a local oscillator in this manner produces both sum and difference components. The former can be removed by using the PDSP16256 Programmable FIR Filter.

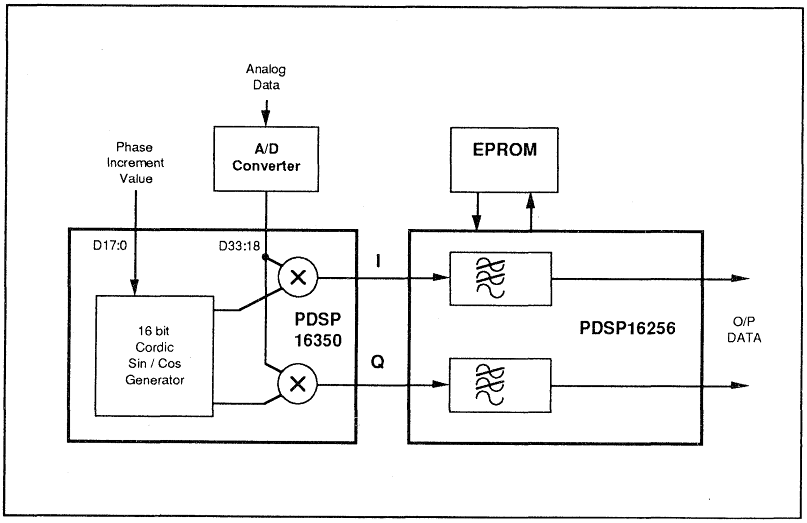

Fig. 5 IQ Split Function

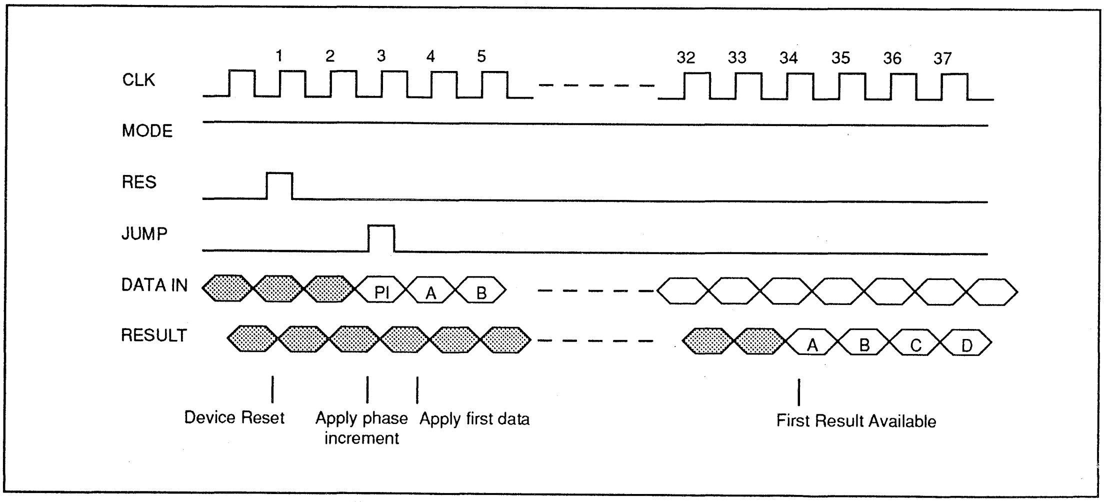

Fig. 6 Amplitude Modulation (34bit frequency accuracy)

Modulated Frequency

The output frequency can be modulated very simply, see Fig 8. Since the phase increment value can be loaded as a complete word every cycle, there is no need to provide internal double buffering to prevent spurious frequencies being generated during the load operation. Binary Frequency Shift Keyed (BFSK) modulation can easily be implemented by externally multiplexing between two phase increment values representing the two frequencies to be used. The value to be used can be instantaneously changed, thus maintaining phase coherence, whilst the bit to be transmitted changes from a mark to a space. Frequency hopping could also be simply effected by clocking a new random number into the DIN port once every thousand cycles, for instance. The output will reflect any change in the frequency after 31 system clock cycles.

If the phase increment value on the DIN port is changed on each clock cycle, then the output frequency will change without introducing any discontinuities. Thus, a linear frequency sweep can be achieved by incrementing the value on the DIN port by a fixed amount each cycle. Alternatively, a logarithmic sweep could be implemented by 'walking' a one across the DIN port. Shifting the input one place to the left every hundred cycles, for example, would double the frequency every time.

Chirp generation for FM - CW Radar systems is a typical example of the need for linear frequency sweeps. This application requires the generation of quadrature chirp waveforms and is illustrated in simplified form by Fig. 7. One waveform is needed for the transmitter, and the other for the receiver. The phase increment value is supplied by the counter block which simply increments at a rate determined by dividing down the time base clock. The synthesised frequency thus increases during the sweep period.

A number of the more significant phase increment bits are used to supply the addresses to a PROM. The output of this PROM is used to amplitude modulate the sine and cosine waveforms. In this manner it is possible to compensate, at the source, for any poor frequency versus gain characteristics of analog circuits further along in the system.

The digital outputs directly drive two D/A converters. Once in the analog world, it is necessary to remove the alias frequencies with low pass filters. The phase linearity and pass band ripple characteristics of these filters are very important, if the correct phase relationships are to be maintained between the two waveforms.

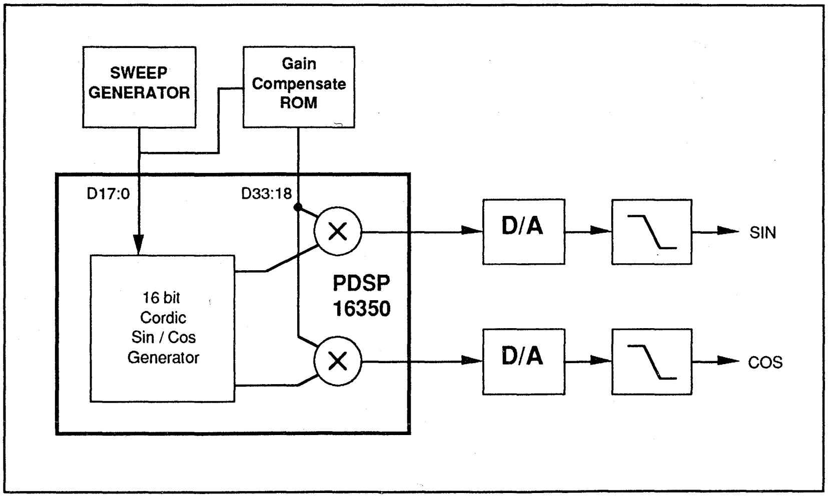

Fig. 7 Quadrature Chirp Generator

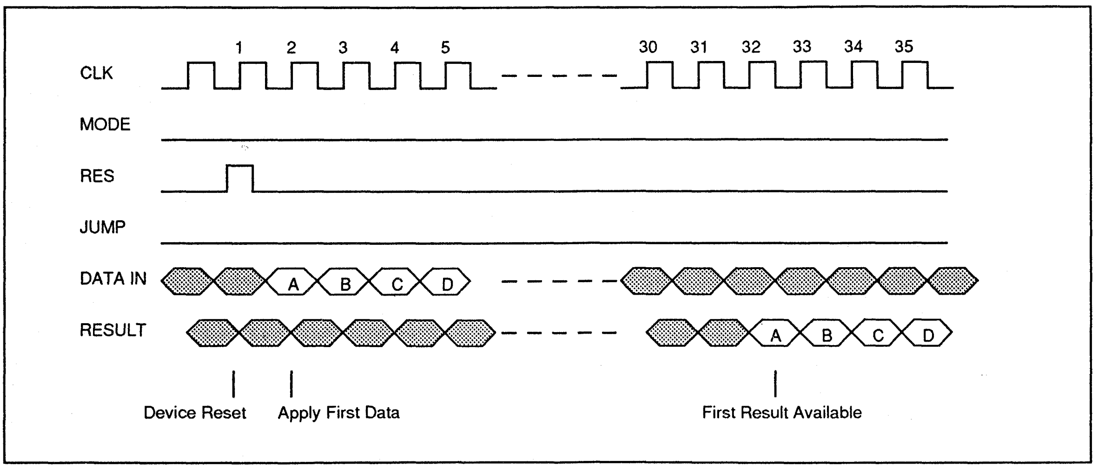

Fig. 8 Frequency Modulation Timing Diagram

Modulated Phase

Relative phase jumps may be made with or without amplitude modulation. For example, if a jump of 180 degrees is required, this can be done with a value of :

DIN33:0 = 10 0000 0000 0000 0000 0000 0000 0000 0000

This is loaded into the phase increment register for one cycle, then the normal increment value is re-loaded in the following cycle.

Alternatively, if no amplitude modulation is needed, an absolute jump to a phase value can be made, see Fig. 9. This can be done by activating the JUMP input during one cycle and also presenting the new phase value at the same time. For example, if a jump to 270 degrees is required :

DIN33:0 = 11 0000 0000 0000 0000 0000 0000 0000 0000

The RES (reset) input can alternatively be used if a jump to 0 degrees is needed. This avoids using the DIN inputs and can be used with or without amplitude modulation. The reset function is internally synchronised to the input clock.

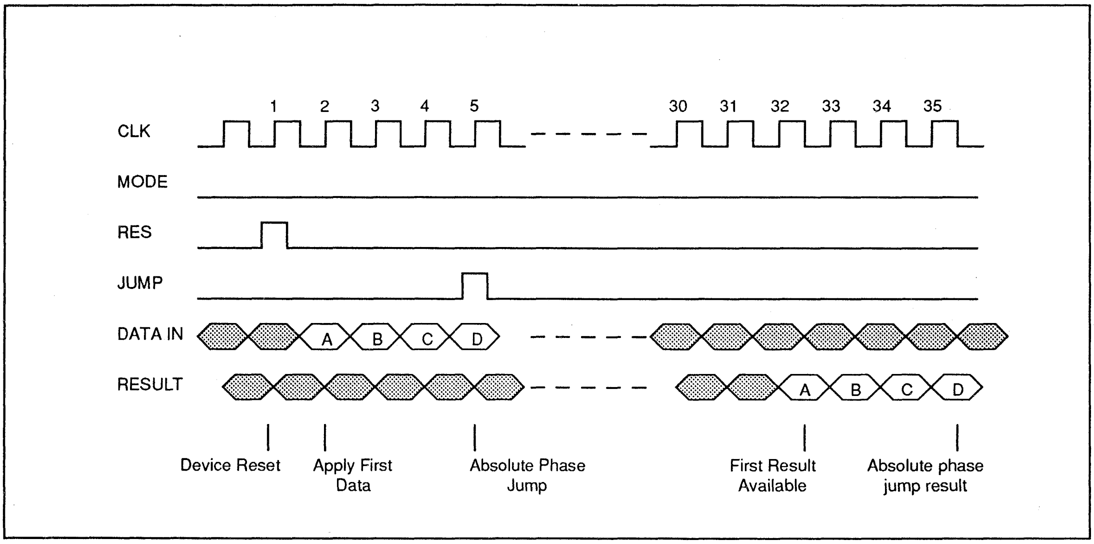

Fig. 9 Phase Modulation Timing Diagram

ABSOLUTE MAXIMUM RATINGS (Note 1)

Supply voltage Vcc	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to Vcc + 0.5V
Output voltage V_{OUT}	-0.5V to Vcc + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	
Military	-55°C to +125°C
Industrial	-40°C to 85°C
Junction temperature	150°C
Package power dissipation	3500mW
Thermal resistances	
Junction to Case $ø_{JC}$	5°C/W

NOTES
1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Vcc = Max, Outputs Unloaded, Clock Freq = Max.
5. CMOS levels are defined as
   VIH = VDD - 0.5v
   VIL = +0.5v
6. Current is defined as positive into the device.
7. The $ø_{JC}$ data assumes that heat is extracted from the top face of the package.

ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)

Commercial: T_{AMB} = 0°C to +70°C $T_{J(MAX)}$ = 95°C V_{CC} = 5.0V±5% Ground = 0V
Industrial: T_{AMB} = -40°C to +85°C $T_{J(MAX)}$ = 110°C V_{CC} = 5.0V±10% Ground = 0V
Military: T_{AMB} = -55°C to +125°C $T_{J(MAX)}$ = 150°C V_{CC} = 5.0V±10% Ground = 0V

Static Characteristics

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	3.0		-	V	
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{CC}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{CC}
Output S/C current	I_{SC}	40		250	mA	V_{CC} = Max

Switching Characteristics

Characteristic	Industrial			Military			Units	Conditions
	Min.	Typ.	Max.	Min.	Typ.	Max.		
D33:0 signal setup to clock rising edge	15		-	15		-	ns	
D33:0 signal hold after clock rising edge	4		-	4		-	ns	
CEN setup to clock rising edge	20		-	20		-	ns	
CEN hold after clock rising edge	0		-	0		-	ns	
JUMP, RES setup to clock rising edge	10		-	10		-	ns	
JUMP hold after clock rising edge	6		-	6		-	ns	
RES hold after clock rising edge	8		-	8		-	ns	
Clock rising edge to output valid	5		30	5		30	ns	30pF
Clock freq	DC		20	DC		20	MHz	
Clock High Time	15		-	15		-	ns	
Clock Low Time	20		-	20		-	ns	
$\overline{OES}$,$\overline{OEC}$ low to data valid	-		20	-		20	ns	30pF
$\overline{OES}$,$\overline{OEC}$ high to data high impedance	-		20	-		20	ns	30pF
Pipeline delay VIN to VOUT	31		31	31		31	CLKs	
Vcc Current (CMOS inputs)	-		430	-		450	mA	See Note 4
Vcc Current (TTL inputs)	-		460	-		500	mA	See Note 4

ORDERING INFORMATION

Commercial (0°C to +70°C)

PDSP16350 / C0 / AC	(20MHz - PGA)
PDSP16350 / C0/ GC	(20MHz - QFP)

Industrial (-40°C to +85°C)

PDSP16350 / B0 / AC	(20MHz - PGA)
PDSP16350 / B0/ GC	(20MHz - QFP)

Military (-55°C to +125°C)

PDSP16350 / A0 / AC	(20MHz - PGA)
PDSP16350 / A0/ GC	(20MHz - QFP)

DS3713 - 2.0

PDSP16488

SINGLE CHIP 2D CONVOLVER WITH INTEGRAL LINE DELAYS

The PDSP16488 is a fully integrated, application specific, image processing device. It performs a two dimensional convolution between the pixels within a video window and a set of stored coefficients. An internal multiplier accumulator array can be multi-cycled at double or quadruple the pixel clock rate. This then gives the window size options listed in Table 1.

An internal 32k bit RAM can be configured to provide either four or eight line delays. The length of each delay can be programmed to the users requirement, up to a maximum of 1024 pixels per line. The line delays are arranged in two groups,which may be internally connected in series or may be configured to accept separate pixel inputs. This allows interlaced video or frame to frame operations to be supported.

The 8 bit coefficients are also stored internally and can be downloaded from a host computer or from an EPROM. No additional logic is required to support the EPROM and a single device can support up to 16 convolvers.

The PDSP16488 contains an expansion adder and delay network which allows several devices to be cascaded. Convolvers with larger windows can then be fabricated as shown in Table 2.

Intermediate 32 bit precision is provided to avoid any danger of overflow, but the final result will not normally occupy all bits. The PDSP16488 thus provides a multiplier in the output path, which allows the user to align the result to the most significant end of the 32 bit word.

FEATURES

- 8 or 16 bit pixels with rates up to 40 MHz
- Window sizes up to 8 x 8 with a single device
- Eight internal line delays
- Supports interlace and frame to frame operations
- Coefficients supplied from an EPROM or remote host
- Expandable in both X and Y for larger windows
- Gain control and pixel output manipulation
- 84 pin PGA or 132 pin QFP

Data Size	Window Size Width X Depth		Max Pixel Rate	Line Delays
8	4	4	40MHz	4x1024
8	8	4	20MHz	4x1024
8	8	8	10MHz	8x512
16	4	4	20MHz	4x512
16	8	4	10MHz	4x512

Table 1 Single Device Configurations

Max Pixel Rate	Pixel Size	Window size						
		3x3	5x5	7x7	9x9	11x11	15x15	23x23
10MHz	8	1	1	1	4	4	4	9
10MHz	16	1	2	2	-	-	-	-
20MHz	8	1	2	2	6	6	8	-
20MHz	16	1	4	4	-	-	-	-
40MHz	8	1	4 *	4 *	-	-	-	-
40MHz	16	2	-	-	-	-	-	-

* Maximum rate is limited to 30 MHz by line store expansion delays

Table 2 Devices needed to implement typical window sizes

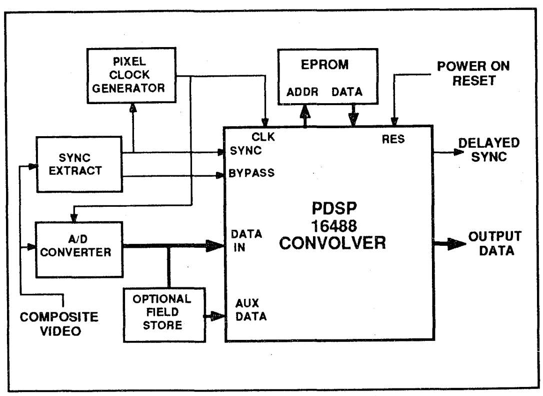

Fig. 1 Typical , Stand Alone, Real Time System

GC	SIG	GC	SIG	GC	SIG	GC	SIG
1	N/C	34	N/C	67	N/C	100	N/C
2	D0	35	X2	68	IP1	101	VDD
3	$\overline{\text{OEN}}$	36	X3	69	GND	102	F0
4	BIN	37	X4	70	IP2	103	D15
5	$\overline{\text{PC1}}$	38	N/C	71	N/C	104	N/C
6	VDD	39	X5	72	VDD	105	D14
7	GND	40	GND	73	UP3	106	D13
8	OVER	41	X6	74	VDD	107	GND
9	N/C	42	X7	75	IP4	108	D12
10	HRES	43	N/C	76	GND	109	GND
11	R/$\overline{\text{W}}$	44	X8	77	IP5	110	VDD
12	$\overline{\text{CE}}$	45	X9	78	GND	111	VDD
13	N/C	46	VDD	79	IP6	112	D11
14	N/C	47	VDD	80	VDD	113	D10
15	GND	48	VDD	81	IP7	114	D9
16	N/C	49	X10	82	VDD	115	GND
17	$\overline{\text{DS}}$	50	$\overline{\text{MASTER}}$	83	N/C	116	CLK
18	GND	51	N/C	84	L7	117	CLK
19	VDD	52	X11	85	GND	118	CLK
20	$\overline{\text{PROG}}$	53	X12	86	L6	119	GND
21	GND	54	$\overline{\text{SINGLE}}$	87	GND	120	GND
22	CS3	55	GND	88	L5	121	D8
23	CS2	56	GND	89	VDD	122	VDD
24	CS1	57	N/C	90	L4	123	D7
25	CS0	58	X13	91	VDD	124	D6
26	VDD	59	X14	92	L3	125	D5
27	$\overline{\text{RES}}$	60	N/C	93	VDD	126	D4
28	$\overline{\text{PC0}}$	61	X15	94	L2	127	GND
29	N/C	62	VDD	95	GND	128	D3
30	DELEOP	63	BYPASS	96	L1	129	N/C
31	X0	64	IP0	97	F1	130	D1
32	X1	65	VDD	98	L0	131	D2
33	N/C	66	N/C	99	N/C	132	N/C

Pin out Table (132 pin ceramic QFP - GC132

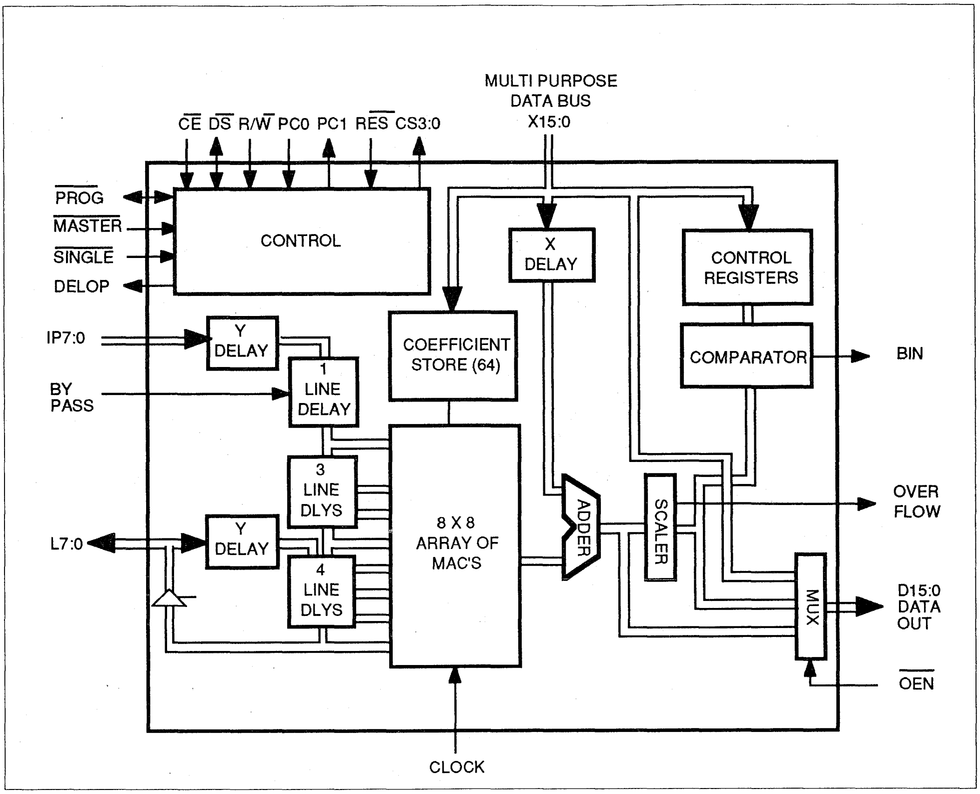

Fig. 2 Functional Block Diagram

PIN NO AC PACKAGE	FUNCTION	PIN NO AC PACKAGE	FUNCTION	PIN NO AC PACKAGE	FUNCTION	PIN NO AC PACKAGE	FUNCTION
A1	L0	M3	X15	K12	RES	B9	D7
B1	F1	N3	X14	K13	CS0	A9	D8
C2	L1	M4	X13	J12	CS1	B8	CLK
C1	L2	N4	SPARE	J13	CS2	B7	SPARE
D2	L3	M5	SINGLE	H12	CS3	A7	D9
D1	SPARE	N5	X12	G12	PROG	B6	D10
E2	L4	M6	X11	G13	DS	A5	D11
E1	L5	M7	MASTER	F12	CE	B5	SPARE
F2	L6	N7	X10	E13	R/W	A4	D12
G2	L7	M8	X9	E12	HRES	B4	D13
G1	IP7	N9	X8	D13	OV	A3	D14
H2	SPARE	M9	X7	D12	PC1	B3	D15
J1	IP6	N10	X6	C13	BIN	A2	F0
J2	IP5	M10	X5	C12	OEN	F1	VDD
K1	IP4	N11	X4	B13	D0	N6	VDD
K2	SPARE	M11	X3	A13	D1	F13	VDD
L1	IP3	N12	X2	A12	D2	A6	VDD
L2	IP2	N13	X1	B11	D3	H1	GND
M1	IP1	M13	X0	A11	D4	N8	GND
N1	IP0	L12	DELOP	B10	D5	H13	GND
N2	BYPASS	L13	PC0	A10	D6	A8	GND

Pin out Table (84 pin PGA - AC84)

NAME	TYPE	DESCRIPTION
IP7:0	INPUT	Pixel data input to the first line delay. [most significant byte in 16 bit mode]
L7:0	I/O	Pixel data input to the second group of line delays. [least significant byte in 16bit mode]. Alternatively an output from the last line delay when the appropriate mode bit is set.
BYPASS	INPUT	The first line delay in the first group is bypassed when this input is active. (High)
HRES	INPUT	Resets the line delay address pointers when high. Normally the composite sync signal in real time applications. In non real time systems it defines a frame store update period, when low.
X15:0	DUAL FUNCTION	Address/data connections from a MASTER or SINGLE device to the external coefficient source, with X15 defining EPROM or Host support. Otherwise they provide the expansion data input.
D15:0	OUTPUT	Signed 16 bit scaled data or multiplexed 32 bit intermediate data. During intermediate transfers the most significant half is valid when the clock is low, and the least significant half when clock is high.
PC1	OUTPUT	During programming a MASTER device outputs a timing strobe on this pin. This is passed down the chain in a multiple device system, using the $\overline{\text{PC0}}$ input on the next device.
PC0	INPUT	This pin is used in conjunction with $\overline{\text{PC1}}$ in multiple device systems. It terminates the write strobe from a MASTER device which is EPROM supported.
DELOP	OUTPUT	This output provides a version of the HRES input which has been delayed by an amount defined by the user.
$\overline{\text{DS}}$	I/O	The data strobe from a host computer. Active low. This pin will be an output from an EPROM supported MASTER device which provides strobes to the remaining devices.
$\overline{\text{CE}}$	INPUT	An active low enable which is internally gated with R/ $\overline{\text{W}}$ and $\overline{\text{DS}}$ to perform reads or writes to the internal registers. In a SINGLE or MASTER device, which is supported from an EPROM, the bottom 72 addresses are always used and CE is not needed. $\overline{\text{CE}}$ can then be used to initiate a new register load sequence after the power on load sequence.
R/ $\overline{\text{W}}$	INPUT	Read / not write line from the host CPU. When an EPROM is used this pin should be tied low.
$\overline{\text{PROG}}$	I/O	This pin is normally an input which signifies that registers are to be changed or examined. It is, however, an output from an EPROM supported SINGLE or MASTER device indicating to the rest of the system that registers are being updated.
CLK	INPUT	Clock. All events are triggered on the rising edge of the clock, except the latching of least significant expansion inputs . Internally the clock can be multiplied by two or four in order to increase the effective number of multipliers.
BIN	OUTPUT	This output indicates the result from the internal comparison. A high value indicates that the pixel was greater than the internal threshold. The output is only valid from the last device in a chain.
OV	OUTPUT	When high this output indicates that there has been a gain control overflow.
$\overline{\text{RES}}$	INPUT	Active low power on reset signal.
$\overline{\text{SINGLE}}$	INPUT	Tied to ground to indicate a SINGLE device system. Internal pull up resistor.
$\overline{\text{MASTER}}$	INPUT	Tied to ground to indicate the MASTER device in a multiple device system. Must be left open circuit in a SINGLE device system. Internal pull up.
$\overline{\text{OEN}}$	INPUT	Output enable signal. Active low.
CS3:0	OUTPUTS	Four address bits from a MASTER specifying one of sixteen devices in a multiple device system. Must be externally decoded to provide chip enables for the additional devices.
F1:0	OUTPUTS	These bits indicate the field selection given by the auto select logic. The same coding as that used for Control Register bits C5:4 is used.
VCC / GND	SUPPLY	Four Power and ground pairs. All must be connected.

BASIC OPERATION

The PDSP16488 convolver performs a weighted sum of all the pixels within an N x N two dimensional window. Each pixel value is multiplied by a signed coefficient, or weight, and the products are summed together. In practice positive weights would be used to produce averaging effects, with various distribution laws, and negative weights would be used for edge enhancement. The window is moved continuously over the video frame, and for real time operation a new result must be obtained for every pixel clock. In most applications odd sized windows will be used, resulting in a centre pixel whose value is modified by the surrounding pixels.

OUTPUT ACCURACY

With 8 bit pixels, and an 8 x 8 window, it is possible for the accumulated sum to grow to 22 bits within a single device. With 16 bit pixels, and an 8 x 4 window (the maximum possible), the sum can grow to 29 bits. The PDSP16488 actually allows for word growth up to 32 bits, and thus allows several devices to be cascaded without any danger of overflow. Since coefficients can be negative, the final result is a 32 bit signed two's complement number.

In a particular application the desired output will lie somewhere within these 32 bits, the actual position being dependent on the coefficient values used. This causes problems in physically choosing which output pins to connect to the rest of the system. To overcome this problem the PDSP16488 contains an output multiplier, or gain control, which allows the final result to be aligned to the most significant end of the 32 bit internal result.The provision of a multiplier, rather than a simple shifter, allows the gain to be defined more accurately.

The sixteen most significant bits of the adjusted result are available on output pins, and contain a sign bit.

OUTPUT SATURATION

If the output from the convolver is driving a display, negative pixels will give erroneous results. An option is thus provided which forces all negative results to zero, which are then interpreted as black by the display. At the same time positive results, which overflow the gain control, are forced to saturate at the most positive number ie peak white. In this mode the output sign bit is always zero,and should not be connected to an A/D converter.

A separate option forces both negative and positive overflows to saturate at their respective maximum values, but in scale negative results remain valid. A gain control overflow warning flag is also available, which can be used in a host CPU supported system to change the gain parameters if overflows are not acceptable.

BINARY OUTPUT

The PDSP16488 contains a 16 bit arithmetic comparator which allows the output from the gain control to be compared with a previously programmed value. An output flag allows the user to detemine if the result was above or below a value contained within an internal register.

MULTIPLIER ARRAY

The PDSP16488 contains sixteen 8x8 multipliers each producing a 16 bit result. Internally the pixel clock supplied by the user can be multiplied by two or four, which together with the proprietary architecture, allows each multiplier to be used several times within a pixel clock period. This increases the effective number of multipliers, which are available to the user, from 16 to 32 or 64 respectively. This architecture produces a very efficient utilization of chip area, and allows the line delays to be accommodated on the same device.

The sixteen multipliers are arranged in a 4 deep by 4 wide array, resulting in effective arrays of 4 by 8 or 8 by 8 with the multi-cycling options. The multiplier array can also be configured to handle 16 bit signed pixels; the effective number of available multipliers is then halved.

LINE DELAY OPERATION

Internal RAM is arranged in two separate groups, and can be configured to provide line delays to match the chosen size of the convolver. When a four deep arrangement is used, with 8 bit pixels, four line delays are available, and each can be programmed to contain up to 1024 pixels. In an eight deep array, or if16 bit pixels are needed, each line can contain up to 512 pixels. Figure 4 illustrates the options available.

The first line delay in one of the groups can optionally be switched in or out under the control of an input pin. It is used to delay the pixel input when data is obtained from another convolver in a multiple device system, or it is used to support interlaced video.

Signals L7:0 may be used as pixel inputs or outputs. They are configured as inputs at power-on to avoid possible bus conflicts, but by setting a mode control bit can become outputs. They can then be used to drive another device when multiple PDSP16488's are required.

INTERLACED VIDEO

When using real time interlaced video, a picture or frame is composed from two fields, with odd lines in one field and even lines in the other. An external field delay is thus required to gather information from adjacent lines, and the convolver needs two input busses. The bus providing the delayed pixels has an extra internal line delay. This is only used in the field containing the upper line in any pair of lines, and must be bypassed in the other field. It ensures that data from the previous field always corresponds to the line above the present active line, and avoids the need to change the position of the coefficients from one field to the next.

Figure 3 shows the translation from physical to internal line positions, for single device interlaced systems. Line N is the line presently being convolved, which is either one or two lines previous to the line presently being produced.

When windows requiring four or more lines are to be implemented, the first line delay, in the group supplied from the L7:0 pins, must always be by-passed. This by-pass option is controlled by Register B, bit 7 and is not effected by the BYPASS input pin.. The coefficients must be loaded into the locations shown, which match the translated line positions, with unused coefficients, shown shaded, loaded with zero's.

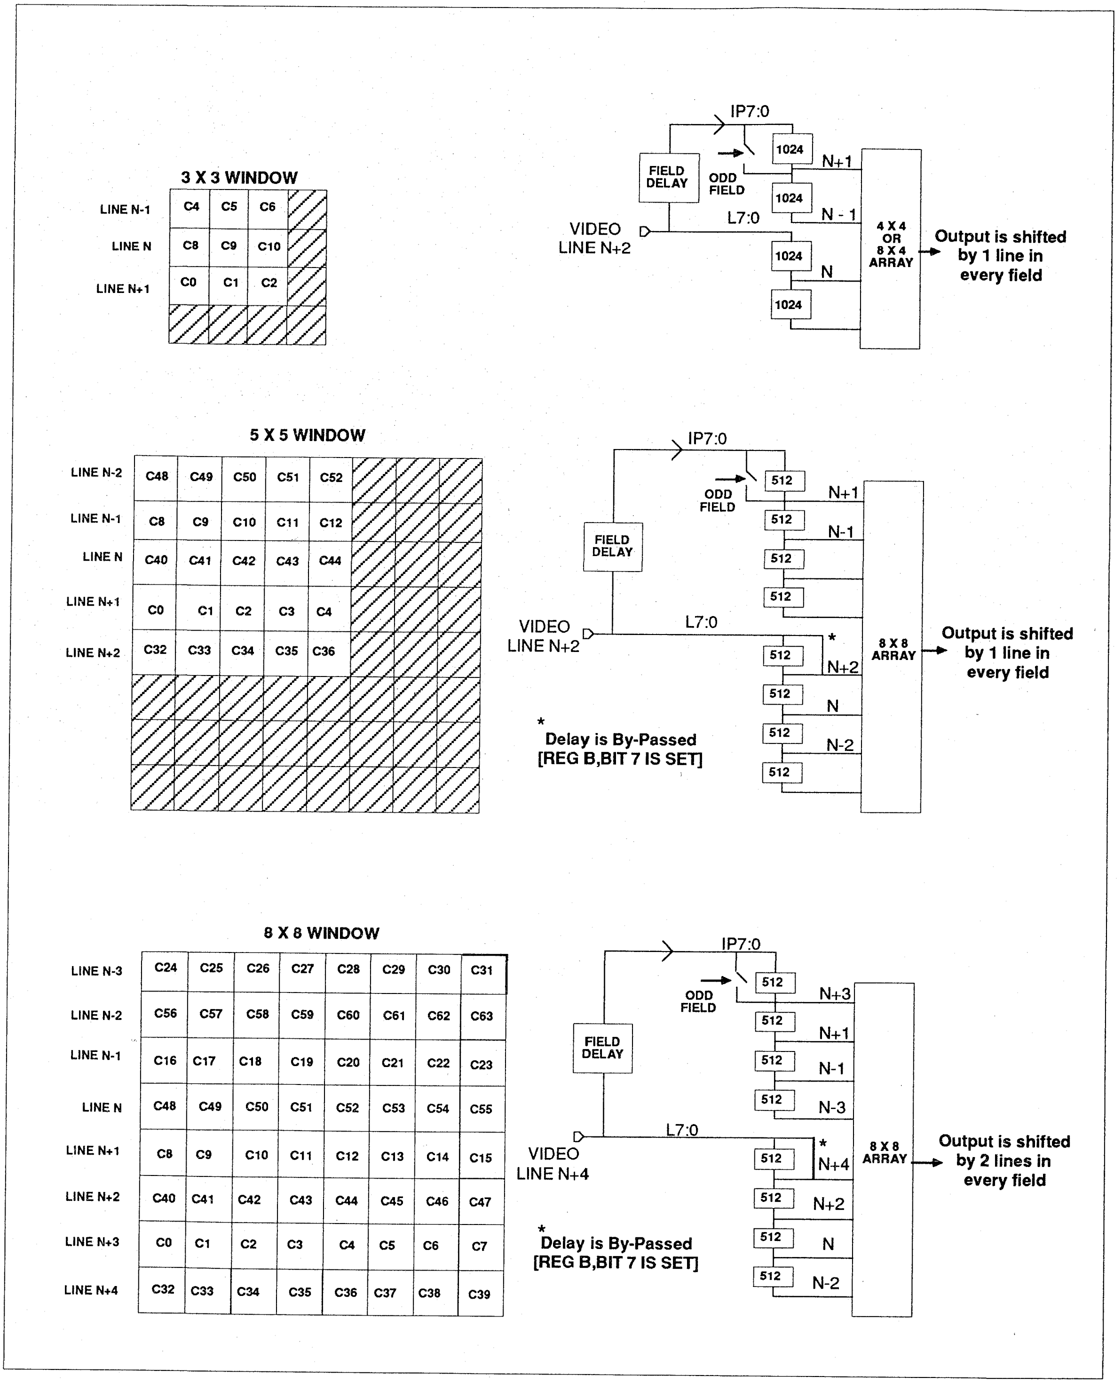

Figure 3. Line Delay Allocations in Single Device Interlaced Systems

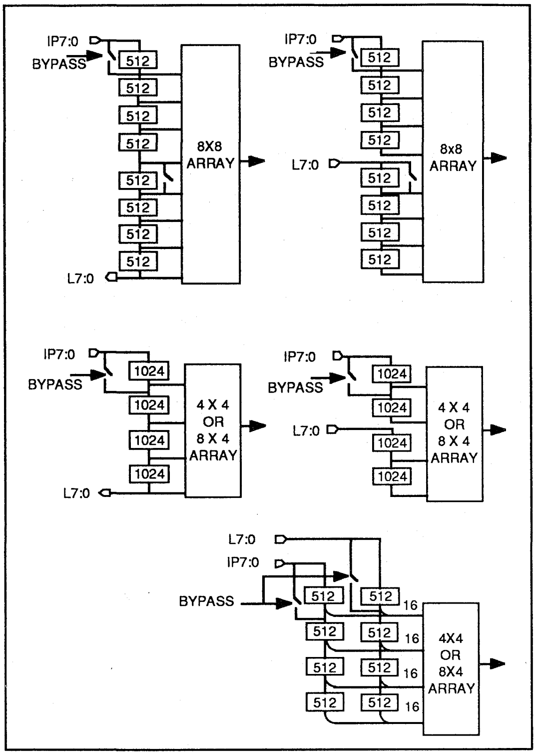

Fig. 4. Line Delay Configurations

DEFINING THE LENGTH OF THE LINE DELAY

Figure 4 defines the maximum line lengths available in each of the window size options. The actual line lengths can be defined in one of three ways, to support both real time applications, taking pixels directly from a camera, and also use in systems supported by a frame store. In the former case the line delays must be referenced to video synchronization pulses. In the latter case the line lengths are well defined, and the horizontal flyback 'dead times' will have been removed.

To support real time applications an option is provided in which the length of the line delay is defined by the number of clocks obtained whilst an input pin (HRES) is in-active. HRES would normally be composite sync when the convolver is directly attached to an NTSC or PAL video camera.

Conceptually, the line delay is achieved by reading the previous contents of a RAM based line store, and then writing new information to the same address. When HRES is active write operations are inhibited, and the address counter is reset. During an active line the counter is incremented by the pixel clock. If the maximum count is reached before the end of a line, then write operations are terminated and wrap-around effects avoided.

The active going edge of HRES, marking the end of a line, is normally asynchronous to the pixel clock, and it is possible for an additional pixel to be stored on some lines. This has no effect on the convolver operation, and will not cause a cumulative shift in the pixel position from line to line.

An alternative means of defining the line length is, however, provided when an exact number of pixels is needed. HRES going in-active then starts the delay operation for every line, but it ceases when the 10 bit value contained in two registers is reached. This method can avoid the need to store blank pixels at the end of a line before sync goes active. With this method the line must contain an even number of pixels, but the value loaded into the control registers defining the line length, must be one less than the even number needed.

In an image processing system, the pixel clock is often re-synchronized, or even inhibited, during blanking or sync. The next line is then started with a precise time interval from the end of sync to the first pixel clock edge. This avoids any visible pixel jitter at the beginning of the line, which would otherwise be present since pixel clock is asynchronous with respect to video sync pulses.

When using the PDSP16488 the pixel clock should not be inhibited, or re-synchronized, until the delayed version of the HRES input goes active. This is present on the DELOP output pin. This will ensure that no pixels on the right hand edge are lost due to the internal pipeline delay.

If the pixel clock is a continuous signal, the user must ensure that the HRES in-active transition meets the timing requirements defined in Figure 10. The active going edge at the end of a line need not be synchronized.

When pixels are read/written to a frame store, an alternative line delay configuration is needed. Within the frame store lines would be stored in contiguous locations, with no gaps caused by the flyback period between the lines. This method of use makes the HRES defined line delay operation difficult to use, and an alternative mode of operation is provided. The HRES input is then driven by a system provided signal, which defines a complete frame store update period. It is not a line defining signal. The high to low transition of this signal will initiate the line store update sequence and allow the internal address pointers to increment. These pointers will be synchronously reset at the end of a line, when they reach the pre-programmed value. They will then immediately start a new operation using address zero. The actual line delay must be pre-loaded into two control registers as described previously.

Write operations back to the frame store must allow for the total pipeline delay. This can be achieved by inhibiting write operations until the delayed version of HRES goes low at the DELOP output pin. Write operations then continue until it goes back high. The PDSP16488 assumes that data is valid when a clock signal is applied, and that it also meets the set up and hold requirements given in Figure 10. If data is not valid, due for example to a frame store DRAM refresh cycle, then the user must externally inhibit the clock. The clock supplied to the convolver will in this mode be a signal which defines a frame store cycle time.

The use of the convolver in a line scan system is similar to its use with a frame store. These systems have no flyback period, and the address counter must be synchronously reset at the end of the line and then allowed to continue.

GAIN CONTROL

The gain control is provided as an aid to locating the bits of interest in the 32 bit internal result. The magnitude of the largest convolved output will depend on the size of the window, and the coefficient values used. The function of the gain control is then to produce an output, which is accurate to 16 bits, and which is aligned to the most significant end of this 32 bit word. The sixteen most significant bits of the word are available on output pins, and the largest number need only have one sign bit if the gain control is correctly adjusted.

Fiigure 5 indicates the mechanism employed with the required function implemented in two steps. Two mode control bits allow one of four 20 bit fields to be selected from the final 32 bit value. These four fields are positioned with the first at the most significant end, and then at four bit displacements down to the least significant end.

By setting an enabling bit, the field selection can optionally be done automatically. This feature should only be used in the real time operating mode, when HRES defines video lines. Internal logic examines the most significant 13, 9, or 5 bits from the 32 bit result, and makes a field selection dependent on which group does not contain identical sign bits. If less than five sign bits are obtained, the logic will select the field containing the most significant 20 bits.

The automatic selection is particularly useful when a fixed scene is being processed. The selection is reset when any internal register is updated (ie PROG has been active) and is then held in-active for ten further occurances of the HRES input. This allows the internal multiplier/ accumulator array to be completely flushed before a field selection is made. As convolver outputs of greater magnitude are produced the field selection logic will respond by selecting a more significant field. The most significant field found necessary remains selected until $\overline{\text{PROG}}$ again goes active. Even if the automatic field selection is not enabled, two outputs, F1:0, will still indicate which field would have been selected. These are coded in the same way as Register C, bits 5:4.

Having chosen a field, either manually or automatically, it is then multiplied by a 4 bit unsigned integer. This is contained within a user programmed register, and the multiplication will produce a 24 bit result. The middle 16 bits of this result contain the required output bits. The gain control multiplier can overflow in to the unused most significant four bits if the parameters are chosen wrongly. This condition is indicated by an overflow flag .

By setting appropriate mode control bits, further manipulation of the gain control output is possible. One option allows all negative outputs to be forced to zero, and at the same time positive gain control overflows will saturate at the maximum positive number. A different option will saturate positive and negative overflows at their respective maximum values, but otherwise leaves them unchanged. Occasional overflows can be tolerated in some systems, and this option prevents any gross errors.

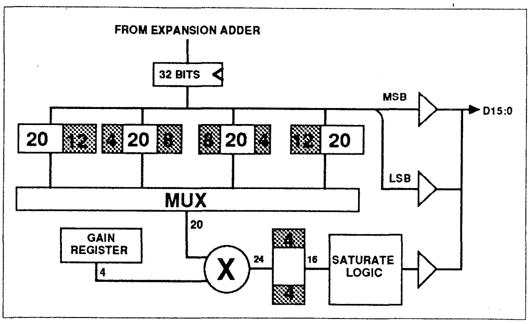

Fig. 5. Gain Control Operation

EXPANSION

Multiple devices can be connected in cascade in order to fabricate window sizes larger than those provided by a single device. This requires an additional adder in each device which is fed from expansion data inputs. This adder is not used by a single device or the first device in a cascaded system, and can be disabled by a mode control bit.

The first device in the cascaded system must be designated as a MASTER device by tying an input pin low. Its expansion input bus is then used as the source of data for the coefficient and control registers in all devices in the system.

In order to reduce the pin count required for 32 bit busses, both expansion in and data out are time multiplexed with the phases of the pixel clock. When the clock is high the least significant half will be valid, and when the clock is low the most significant half will be valid.

In practice this multiplexing is only possible with pixel clocks up to 20MHz. Above these frequencies the multiplexing must be inhibited by setting a Mode Control bit (Register A, Bit 7). The intermediate data accuracy will then be reduced, since only the lower 16 bits of the internal 32 bit intermediate sum are available on the output pins. In such systems the coefficients must be scaled down in order to keep the intermediate and final results down to 16 bits. The final device should not use the gain control, and instead should simply output the non-multiplexed 16 bit result. The overflow flag and pixel saturation options will not be available.

PIXEL INPUT AND OUTPUT DELAYS

In a real time system, when line delays are referenced to video sync pulses present on the HRES input, the first pixel from the last line delay does not appear on the L7:0 pins until the fifth active pixel clock edge after HRES has gone low. This is illustrated in Figure 7. In a vertically expanded system, this output provides the input to the first line delays in the vertically displaced devices. The internal logic is thus designed to always expect this five clock delay. Compensation must thus be applied to the devices which are directly connected to the video source, such that the first pixel is not valid until the fifth clock edge.

For this reason the PDSP16488 contains an optional four clock pipeline delay on each of the pixel data inputs. When the delay is used the first pixel in a video line must be available on the input pins after the first pixel clock edge. This would be so if the device were connected to an A/D converter, since that would introduce a one pixel pipeline delay. If the system introduces any further external pipeline delays, then the internal delay should be bypassed, and the user should ensure that the first pixel is valid after the fifth clock edge.

The use of this four clock delay is controlled by Bit 3, in Control Register B. This delay is in addition to the delays which are provided to support expansion in both the X and Y directions, and are controlled by Register D, Bits 3:2. Both delays are in fact simply added together in the device, but are provided for conceptually different reasons.

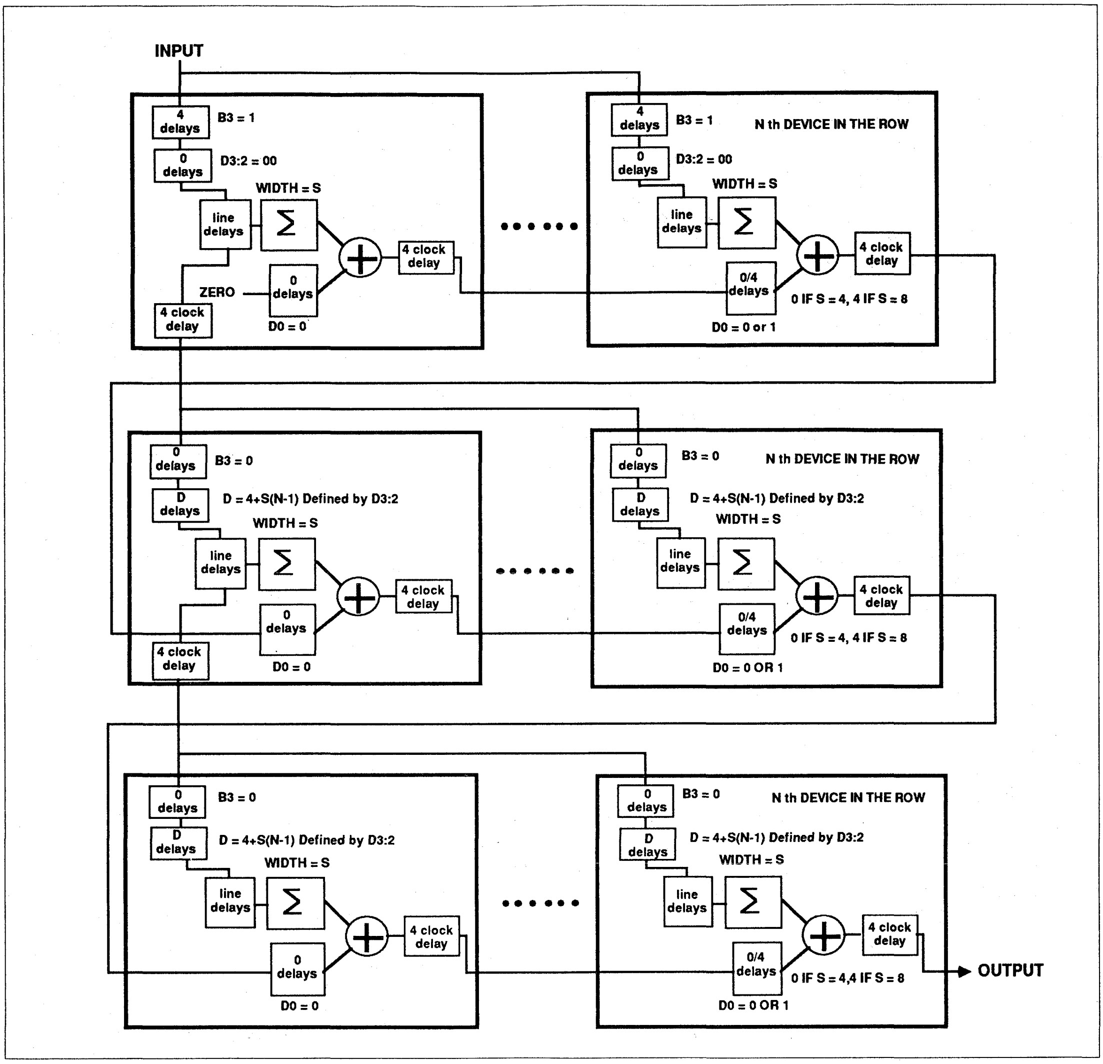

Fig. 6. Multi-Device Delay Paths

DELAY COMPENSATION FOR LARGE WINDOWS

A large window is composed of several partial windows each of which is implemented in an individual device. If necessary the partial window must be padded with zero coefficients to become one of the standard sizes. When constructing a large window it is necessary to delay the expansion data inputs in order to compensate for growth in the horizontal direction. Delays in the partial sums are also necessary to compensate for the total pipeline delay needed to produce the previous complete horizontal stripe.

Within each device in a horizontal stripe, apart from the first, the expansion input must be delayed by the width of the partial window, before it is added to the internal sum. Since partial windows can only be 4 or 8 pixels wide,a delay of 4 or 8 pixel clocks is needed. There is, however, an in-built delay of 4 pixels in the inter device connection, and the PDSP16488 thus only needs an option to delay the expansion input by an additional four pixels.

The data from the last device in a horizontal row of convolvers feeds the expansion input of the first device in the next row. This is shown in Figure 6. With this arrangement, the position of the partial window as illustrated, is the inverse of its vertical position on a normal TV screen. Thus the top, left hand, device corresponds to the bottom, left hand, portion of the complete window.

The output from the last device in the row is delayed with respect to the original data input by an amount given by the formula;

DELAY = 4 + [N-1].S where N is the number of devices in a row and S is the partial window width, ie 4 or 8.

The internal convolver sums, in each of the devices in the next row, must be delayed by this amount before they are added to results from the previous row. This is more conveniently achieved by delaying data going into the line stores. The required cumulative delay with respect to the first horizontal stripe is then automatically obtained when more than two rows of devices are needed.

Two bits in Control Register D are used to define one of four delay options. These delays have been selected to support systems needing from two to eight devices and are described in the applications section.

COEFFICIENTS

Sixty-four coefficients are stored internally and must be initially loaded from an external source. Table 3 gives the coefficient addresses within a device, with coefficent C0 specified by the least significant address and C63 by the most significant address. Table 5 shows the physical window position within the device which is allocated to each coefficient in the various modes of operation. Horizontally the coefficient positions correspond to the convolution process as if it were conceptually observed on a viewing screen, ie the left hand pixel is multiplied with C0. In the vertical direction the lines of coefficients are inverted with respect to a visual screen, ie the line starting with C0 is actually at the bottom of the visualized window.

The coefficients may be provided from a Host CPU using conventional addressing, a read/write line, data strobe, and a chip enable. Alternatively, in stand alone systems, an EPROM may be used. A single EPROM can support up to 16 devices with no additional hardware.

When windows are to be fabricated which are smaller than the maximum size that the device will provide in the required configuration, then the areas which are not to be used must contain zero coefficients. The pipeline delay will then be that of a completely filled window.

TOTAL PIPELINE DELAY

The total pipeline delay is dependent on the device configuration and the number of devices in the system. Table 4 gives the delays obtained with the various single device configurations when the gain control is used. These delays are the the internal processing delays and do not include the delays needed to move a given size window completely into a field of interest. When multiple devices are needed, additional delays are produced which must be calculated for the particular application. These delays are discussed in the applications section.

The PDSP16488 contains facilities for outputing a delayed version of HRES to match any processing delay. Control register bits allow this delay to be selected from any value between 29 and 92 pixel clocks.

Function	Hex. Addr
Mode Reg A	00
Mode Reg B	01
Mode Reg C	02
Mode Reg D	03
Comparator LSB	04
Comparator MSB	05
Scale Value	06
Pixels / Line LSB	07
Pixels / Line MSB	08
C0 - C15	40 - 4F
C16 - C31	50 - 5F
C32 - C47	60 - 6F
C48 - C63	70 - 7F
Unused	09 - 3F

Table 3 Internal Register Addressing

Data size	Window Size	Pipeline Delay
8	4x4	34
8	8x4	30
8	8x8	26
16	4x4	28
16	8x4	26

Table 4 Pipe line dalays

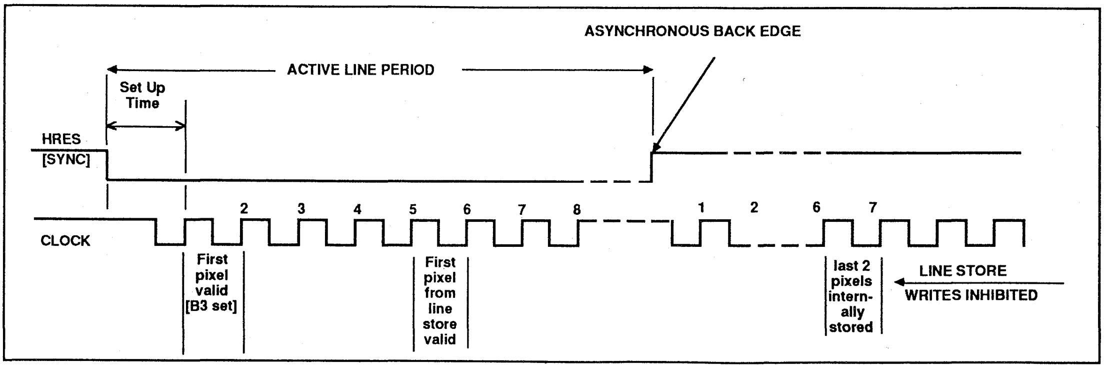

Fig. 7 Pixel Input Delays

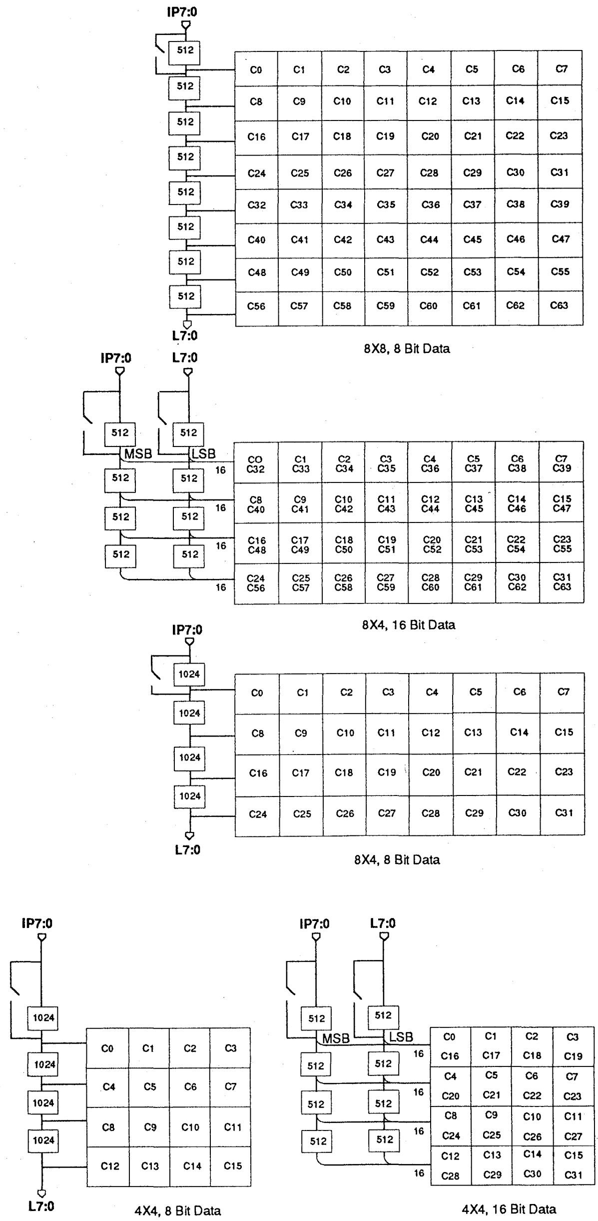

NOTE
Two coefficients occuring in the same box have identical values

Table 5 Physical Coefficient Position

LOADING REGISTERS FROM A HOST CPU

The expansion data inputs [X14:0] on a single or master device are connected to the host bus to provide address and data for the internal registers. In a multiple device system the remaining devices receive addresses and data which have been passed through the expansion connection between earlier devices in the cascade chain. Each device needs an individual chip enable plus a global data strobe, read/write line, and $\overline{\mathrm{PROG}}$ signal from the host.

Registers are individually addressed and can be loaded in any sequence once the global $\overline{\mathrm{PROG}}$ signal has been produced by the host. The latter would normally be produced from an address decode encompassing all the necessary device addresses.

If a self timed system is to be implemented, a timing strobe must be passed down the expansion chain through the $\overline{\mathrm{PC1}}/\overline{\mathrm{PC0}}$ connections. The $\overline{\mathrm{PC0}}$ output from the final device is used as a host REPLY signal, and indicates that the last device has received data after the propogation delay of previous devices. The timing strobe is produced in the MASTER device from the host data strobe, and will appear on the $\overline{\mathrm{PC0}}$ output. This feature allows the user to cascade any number of devices without knowing the propogation delay through each device. The timing information for this mode of operation is given in Figure 8.

The host can also read the data contained in the internal registers. The required device is selected using chip enable with the R/$\overline{\mathrm{W}}$ line indicating a read operation. Single device systems output the data read on X7:0, but in multiple device systems data is read from the D7:0 outputs on the final device in the chain. These must be connected back to the host data bus through three-state drivers. When earlier devices in the chain are addressed, the register contents are transferred through the expansion connections down to the final device. In the self timed configuration the data will be valid when the $\overline{\mathrm{REPLY}}$ goes active, as shown in Figure 8.

If the $\overline{\mathrm{REPLY}}$ signal is not to be used , the $\overline{\mathrm{PC0}}/\overline{\mathrm{PC1}}$ connections are not necessary, and the host data strobe for a write operation must be wide enough to allow for the worst case propogation delay through all the devices (TDEL). If the data or address from the host does not meet the set up time given in Fugure 8, the width of the data strobe can be simply extended to compensate for the additional delay. When reading data the access time required is: TACC + (N - 1).TDEL using the maximum times obtained from Figure 8.

HOST CONTROL LINES

X7:0 — 8 bit data bus. In a single device system this bus is bi-directional; in other configurations it is an input. Only a SINGLE or MASTER device is connected directly to the host. Other devices receive data from the output of the previous device in the chain.

X14:8 — 7 bit address bus which is used to identify one of the 73 internal registers. Connected in the same manner as X7:0.

X15 — X15 must be open circuit on the MASTER device

$\overline{\mathrm{PC0}}$ — An input from the previous PC1 output in a multiple device chain. Not needed on a SINGLE device or if the self timed feature is not used.

$\overline{\mathrm{PC1}}$ — Reply to the host from a SINGLE device or from the last device in a cascade chain. It indicates that the write strobe can be terminated. Connected to PC0 input of the next device at intermediate points in the chain if the self timed feature is used.

R/$\overline{\mathrm{W}}$ — Read/Not Write line from the host CPU which is connected to all devices in the system.

$\overline{\mathrm{CE}}$ — An active low enable which is normally produced from a global address decode for the particular device. This must encompass all internal register addresses.

$\overline{\mathrm{DS}}$ — An active low host data strobe which is connected to all devices. in the system.

$\overline{\mathrm{PROG}}$ — An active low global signal, produced by the host, which is connected to all devices in the system. Together with a unique chip enable for every device, it allows the internal registers to be updated or examined by the host. PROG and CE should be tied together in a single device system.

LOADING REGISTERS FROM AN EPROM

In the EPROM supported mode, one device has to assume the role of a host computer. If more than one device is present, this must be the first component in the chain, which must have its $\overline{\mathrm{MASTER}}$ pin tied low.

The MASTER device contains internal address counters which allow the registers in up to 16 cascaded devices to be specified. It also generates the $\overline{\mathrm{PROG}}$ signal and a data strobe on the pins which were previously inputs. These outputs must be connected to the other devices in the system, which still use them as inputs. The R/W input should be tied low on all devices.

The width of the data strobe is determined by the feedback connection from the PC1 output on the last device to the $\overline{\mathrm{PC0}}$ input on the MASTER. The $\overline{\mathrm{PC0}}/\overline{\mathrm{PC1}}$ connections must be made between devices in a multiple device system; in a single device system the connection is made internally.

The available EPROM access time is determined by an internal oscillator and does not require the pixel clock to be present during the programming sequence. Any pixel clock re-synchronization in a real time system will thus not effect the coefficient load operation. The relevent EPROM timing information is shown in figure 9.

The load procedure will commence after reset has gone from active to in-active, and will be indicated by the $\overline{\mathrm{PROG}}$ output going active. The data from 73 EPROM locations will be loaded into the internal registers using addresses corresponding to those in Table 3. Within a particular page of 128 EPROM locations, the first nine locations supply control register information, and the top 64 supply coefficients. The middle 55 locations are not used. If the window size is 8 x 4, the top 32 locations will also contain redundant data, and if the size is 4 x 4 the top 48 will be redundant.

In a multiple device system the load sequence will be repeated for every device, and four additional address bits will be generated on the CS3:0 pins. These address bits provide the EPROM with a page address, with one page allocated to each device in the system. Within each page only 73 locations provide data for a convolver, the remainder are redundant as in the single device system. The CS3:0 outputs must also be decoded in order to provide individual chip enables for each device. These can readily be derived by using an AS138 TTL decoder. Bits in an internal control register determine the number of times that the sequence is repeated.

If changes to the convolver operation are to be made after power-on, activating the $\overline{CE}$ input on the MASTER or SINGLE device will instigate the load procedure. Additional EPROM address bits supplied from the system will allow different filter coefficients to be used.

EPROM CONTROL LINES

X7:0 — 8 bit data from the EPROM to the MASTER or SINGLE device. Otherwise data is received from the previous device in the chain.

X14:8 — Lower 7 address bits to the EPROM from a MASTER or SINGLE device. Otherwise an input from the data outs of the previous device.

X15 — Tied to ground on a MASTER device to indicate the EPROM mode.

R/$\overline{W}$ — Tied low on all devices.

$\overline{DS}$ — An output from a MASTER or SINGLE device which provides a data strobe for the other devices.

CS3: 0 — Four additional address bits for the EPROM which are provided by the MASTER device. They allow 16 additional devices to be used and must be externally decoded to provide chip enables.

$\overline{PC0}$ — An input on the MASTER device which is driven from the $\overline{PC1}$ output of the last device in the chain. Used internally to terminate the write strobe. Connected to previous $\overline{PC1}$ outputs at intermediate points in the chain. Not needed for a SINGLE device.

$\overline{PC1}$ — An output connected to the $\overline{PC0}$ input of the next device in the chain. The last device feeds back to the MASTER. Not needed for a SINGLE device.

$\overline{CE}$ — An enable which is produced by decoding CS3:0 from the MASTER. It is not needed for a MASTER or SINGLE device which will always use the bottom block of addresses with internally generated write strobes. It can however be used on these devices to initiate a new load procedure after the initial power on sequence.

$\overline{PROG}$ — An active low going signal produced by an EPROM supported MASTER or SINGLE device. An input to all other devices. It indicates that a register load sequence is occuring, either after power on, or as the result of CE as explained above. It remains active until register 73 in the final device has been loaded. Four bits in a control register define the number of cascaded devices.

SYSTEM CONFIGURATION

The device is configured using a combination of the state of the $\overline{SINGLE}$ and $\overline{MASTER}$ pins, and the contents of the four Mode Control registers. In a MASTER or SINGLE device the state of the X15 pin is used to define whether the system is EPROM or host supported.

MODE CONTROL REGISTERS

REGISTER A Bit Allocation

BIT	CODE	FUNCTION
3:0	XXXX	Number of extra devices from1-15
6:4	000	8 bit, 8x8 window, 10MHz max, 8x512 line delays.
6:4	001	16 bit, 8x4 window, 10MHz max, 4x512 line delays.
6:4	010	16 bit, 4x4 window, 20MHz max, 4x512 line delays.
6:4	011	8 bit, 8x4 window, 20MHz max, 4x1024 line delays.
6:4	101	8 bit, 4x4 window, 40MHz max, 4x1024 line delays
7	0	Multiplexed exp. data
7	1	Non-mux. exp. data

BITS 3:0 — These bits are 'don't care' when using a host computer but to a MASTER device, in an EPROM supported system, they define the number of interconnected chips. The EPROM must contain contiguous 128 byte blocks for each of the devices in the system and a 4 bit counter in the MASTER device will sequence through up to 16 block reads. An internal comparator in the MASTER causes the loading of the internal registers to cease when the value in the counter equals that contained in these bits. The bits are redundant in a SINGLE device which only uses one 128 byte block.

BITS 6:4 — These bits define one of the five basic configurations. The line delays will automatically be configured to match the chosen window size and pixel accuracy. The maximum clock rate that is available to the user reflects the internal mutiplication factor.

BIT 7 This bit must be set if the pixel clock is greater than 20MHz. It disables the output time multiplexing, and instead outputs the least significant half of the 32 bit intermediate sum for the complete clock cycle. When the gain control is used, the output multiplexing will automatically be disabled.

REGISTER B Bit Allocation

BIT	CODE	FUNCTION
0	0	Second line delay group fed from the first group
0	1	Second line delay group fed from L7:0 which become inputs
2:1	00	Store pixels to end of line
2:1	01	Store pixels till count is reached
2:1	10	Frame store operation
2:1	11	Not Used
3	0	No delays on pixel inputs
3	1	4 delays on both pixel inputs
4	0	Use expansion adder
4	1	Expansion adder disabled
6:5		Not used
7	0	Use first delay in second group
7	1	Bypass first delay in second group

BIT 0 This bit defines the input for the second group of line delays. It must be set in the 16 bit pixel modes, and is set by power on reset.

BIT 2:1 These bits control the mode of operation of the line stores. In real time systems pixels can be stored either until HRES [SYNC] goes active , or until a pre-determined count is reached. In the frame store mode line store operations are continuous, with a pre-determined line length.

BIT 3 When this bit is set four pipeline delays are added to the pixel inputs to compensate for the internal/ external delays between line stores. The extra delay is only necessary when a device supplied with system video in which the first pixel in a line is valid in the period following the first active clock edge. See Fig 7. The delay is not necessary if the device is fed from the output of another convolver. When set this bit will add four additional delays to those defined by Register D, bits 4: 2.

BIT 4 When this bit is set the expansion adder will not be used. It is automatically set in a MASTER or SINGLE device.

BIT 7 This bit controls the bypass option on the first line delay on the L7:0 inputs. It is only effective when an 8 bit pixel mode is selected, which also needs more than four line delays. When L7:0 are used as outputs it should always be reset. In the 16 bit modes the bypass function is only controlled by the BYPASS pin, and the bit is redundant.

REGISTER C Bit Allocation

BIT	CODE	FUNCTION
0	0	Field selection defined by C5:4
0	1	Automatic field selection
3:1	000	DELOP = 29 + 0 clks
3:1	001	DELOP = 29 + 8 clks
3:1	010	DELOP = 29 + 16 clks
3:1	011	DELOP = 29 + 24 clks
3:1	100	DELOP = 29 + 32 clks
3:1	101	DELOP = 29 + 40 clks
3:1	110	DELOP = 29 + 48 clks
3:1	111	DELOP = 29 + 56 clks
5:4	00	Select upper 20 bits
5:4	01	Select next 20 bits
5:4	10	Select next 20 bits
5:4	11	Select bottom 20 bits
7:6	00	By-pass the gain control
7:6	01	Normal gain control O/P
7:6	10	Saturate at max + and -ve values.
7:6	11	Force -ve to zero.Sat.+ve values.

BIT 0 If this bit is set, the 20 bit field selected from the 32 bit result, is defined automatically by internal logic.

BITS 3:1 These bits are in conjunction with Register D, bits 7:5 to define the pixel delay from the HRES input to the DELOP pin. They are used to match the appropriate processing delay in a particular system. The minimum delay is 29 pixel clocks.

BITS 5:4 These bits define which of the four 20 bit fields out of the 32 bit final result is selected as the input to the gain control. They are redundant when the gain control is not used, or if Register C, bit0, is set.

BITS 7:6 These bits define the use of the gain control as given in the table. Intermediate devices in a multiple device system MUST by-pass the gain control, otherwise the additional pipeline delays will effect the result.

REGISTER D Bit Allocation

BIT	CODE	FUNCTION
0	0	X15:0 Not delayed
0	1	X15:0 Delayed
1	0	Internal sum not shifted
1	1	Internal sum multiplied by 256
3:2	00	I/P to line stores not delayed
3:2	01	I/P to line stores delayed by 4
3:2	10	I/P to line stores delayed by 8
3:2	11	I/P to line stores delayed by 12
4	0	Un-signed pixel data input
4	1	2's complement pixel data input
7:5	XXX	Add 0 to 7 clock delays to DELOP output.

BIT 0 If this bit is set the expansion data input is delayed by four pixel clocks before it is added to the present convolver output. It is used in multiple device systems when the partial window width is 8 pixels.

BIT 1 When this bit is set the internal sum is shifted to the left by 8 places before being added to the expansion input. It is used when two devices are used, each in an 8 bit pixel mode, to fabricate a 16 bit pixel mode.

BITS 3::2 These bits define the delays on both sets of pixel inputs before entering the line stores. The delays are always identical on both sets.

BIT 4 When this bit is set the convolver interprets 8 or 16 bit pixels as 2's complement signed numbers

BIT 7:5 These bits add 0 to 7 additional clock delays to those selected by Register C, bits 3:1.

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage Vcc	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to Vcc + 0.5V
Output voltage V_{OUT}	-0.5V to Vcc + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Max. junction temperature	
commercial	100°C
industrial	110°C
Package power dissipation	3000mW
Thermal resistances, junction to case $ø_{JC}$	5°C/W

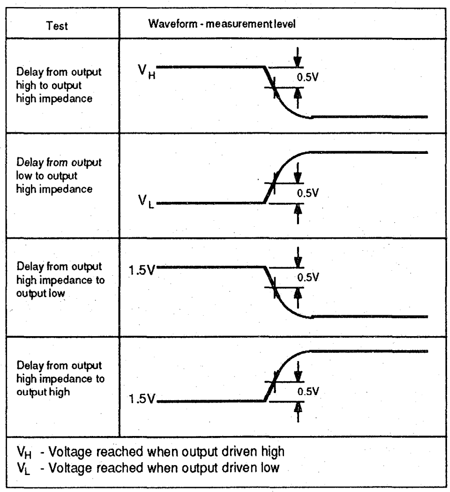

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as negative into the device.

NOTE: Signal pins $\overline{PROG}$, $\overline{PC0}$, X15, $\overline{MASTER}$, $\overline{SINGLE}$, $\overline{DS}$, BYPASS ans 0V have pull-up resistors in the range 15kΩ to 200kΩ

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.0		-	V	
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{CC}.no internal pull up
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{CC}.no internal pull up
Output S/C current	I_{SC}	10		300	mA	V_{CC} = Max

Characteristic	Symbol	Value Min.	Value Max.	Units	Notes
DS Hold Time after REPLY active	T_{DSH}	20		ns	Only applicable for read ops & if $\overline{REPLY}$ is used.
Host Address/data Set Up Time	T_{HSU}	0		ns	Only applicable if $\overline{REPLY}$ is used. Otherwise time is referenced to risng edge of strobe when set up must be N xTDEL, for N devices
Read Set UpTime to prevent Write	T_{RA}	5		ns	
Host Signal Hold Time	T_{HH}	5		ns	Must always be guaranteed.
Expansion in to Data Out in $\overline{PROG}$ mode	T_{DEL}		30	ns	No clocks are needed in $\overline{PROG}$ mode
Delay from strobe to $\overline{PC1}$ [Equivalent to $\overline{PC0}$ to $\overline{PC1}$ delay]	T_{EXP}		50	ns	Greater than TDEL under all conditions
Chip Enable Set Up Time	T_{CSU}	0		ns	
$\overline{PROG}$ Set Up Time	T_{PSU}	0		ns	
$\overline{PROG}$ Hold Time	T_{PH}	0			
Chip Enable Hold Time	T_{CH}	0			
$\overline{PC1}$ In-active Delay after $\overline{DS}$ in-active	T_{PCH}		50	ns	Defines Data Strobe in-active time
Coefficient Read Time	T_{ACC}		50	ns	From MASTER or SINGLE device
Coefficients valid Time before $\overline{REPLY}$	T_{RSU}	5		ns	

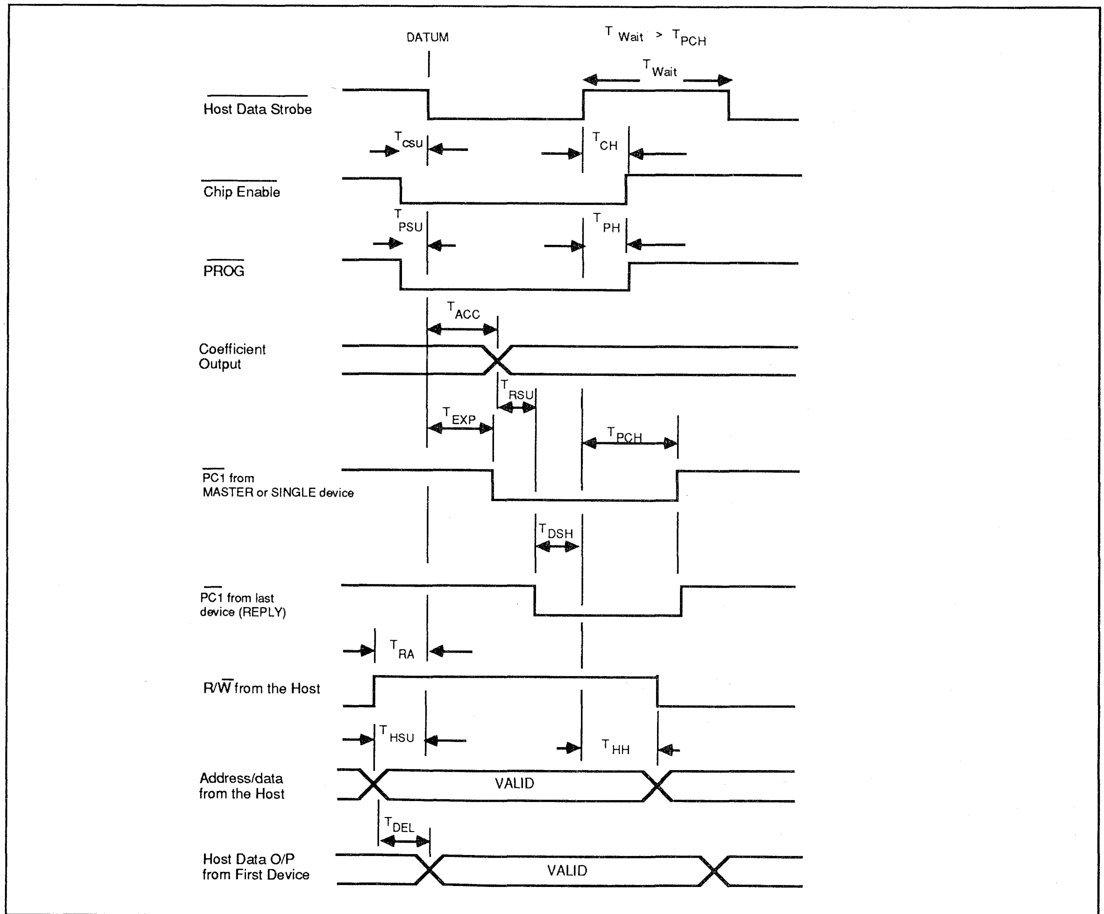

FIG. 8. Host Timing

Characteristic	Symbol	Value		Units	Notes
		Min.	Max.		
Delay from Data Strobe to MASTER $\overline{PC1}$	T_{PCD}		50	ns	
Delay from $\overline{PC0}$ Input to Write in-active	T_{WH}	5		ns	
$\overline{PC1}$ In-Active Delay	T_{PCH}		50	ns	
Write from $\overline{MASTER}$ In-Active	T_{WW}	250		ns	
Write In-Active to new Address	T_{AD}		30	ns	
EPROM Data Set Up Time	T_{DS}	20		ns	
Data Strobe from MASTER	T_{RW}	10		ns	Single device
Chip Enable Set Up Time	T_{CSU}	0		ns	
Chip Enable Hold Time	T_{CH}	0		ns	
Availible EPROM Access Time	T_{DA}	200		ns	
Expansion In to Data Out	T_{DEL}		30	ns	
$\overline{PC0}$ to $\overline{PC1}$ Delay	T_{EXP}		50	ns	Greater than T_{DEL} at all temps

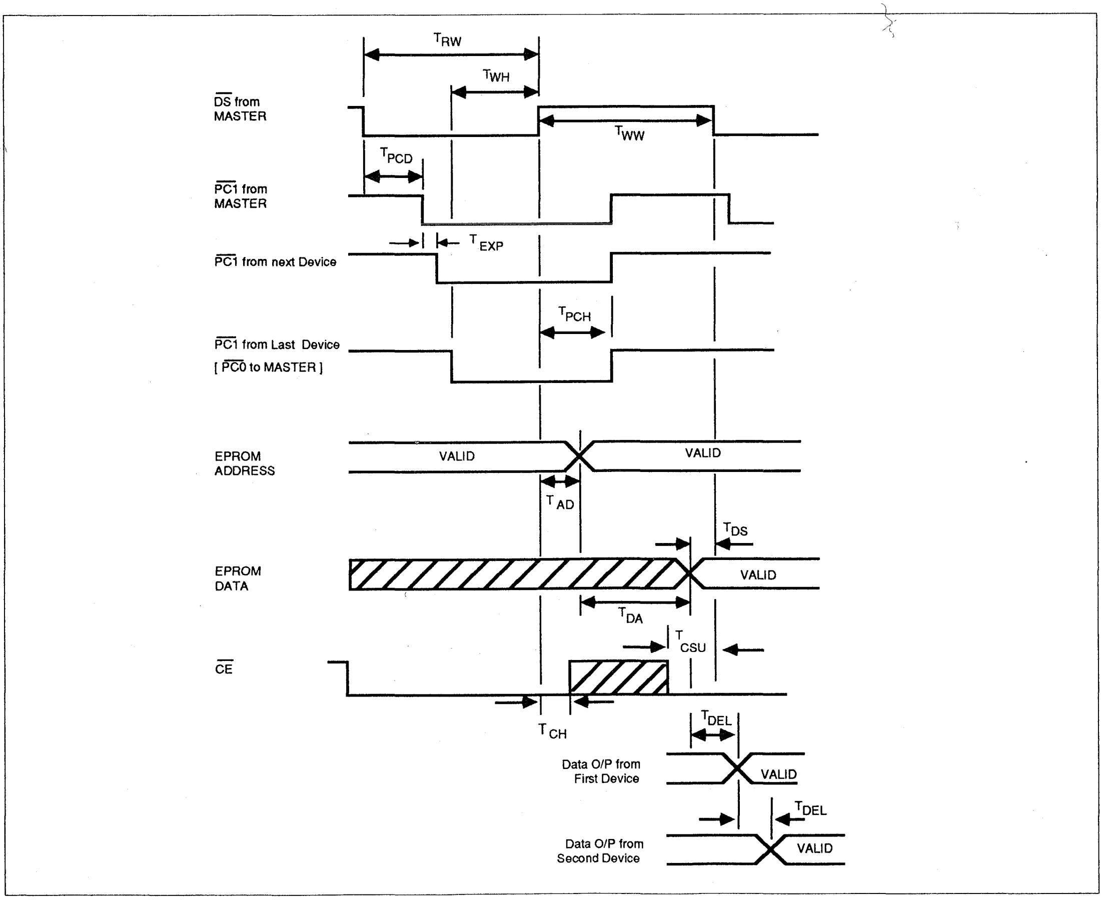

Fig. 9. EPROM Timing

Characteristic	Symbol	Value		Units	Notes
		Min.	Max.		
Pixel Clock Low Time	T_{CL}	25 (a) 10 (b)		ns	(a) 32 Bit Muxed Output (b) 16 Bit Output
Pixel Clock High Time	T_{CH}	25 (a) 10 (b)		ns	(a) 32 Bit Muxed Output (b) 16 Bit Output
Data in Set Up Time	T_{DSU}	10		ns	
Data in Hold Time	T_{DH}	0		ns	
CLK rising to Output delay	T_{RD}		21	ns	Increase to 40ns for DELOP output
Line Store Output Delay	T_{LD}		20	ns	
HRES In-active Set Up Time	T_{RSU}	10		ns	
$\overline{\text{Output Enable}}$ Time	T_{DLZ}		15	ns	Measured with a 15kΩ series resistor and 30pF load capacitance
Output Disable Time	T_{DHZ}		15	ns	

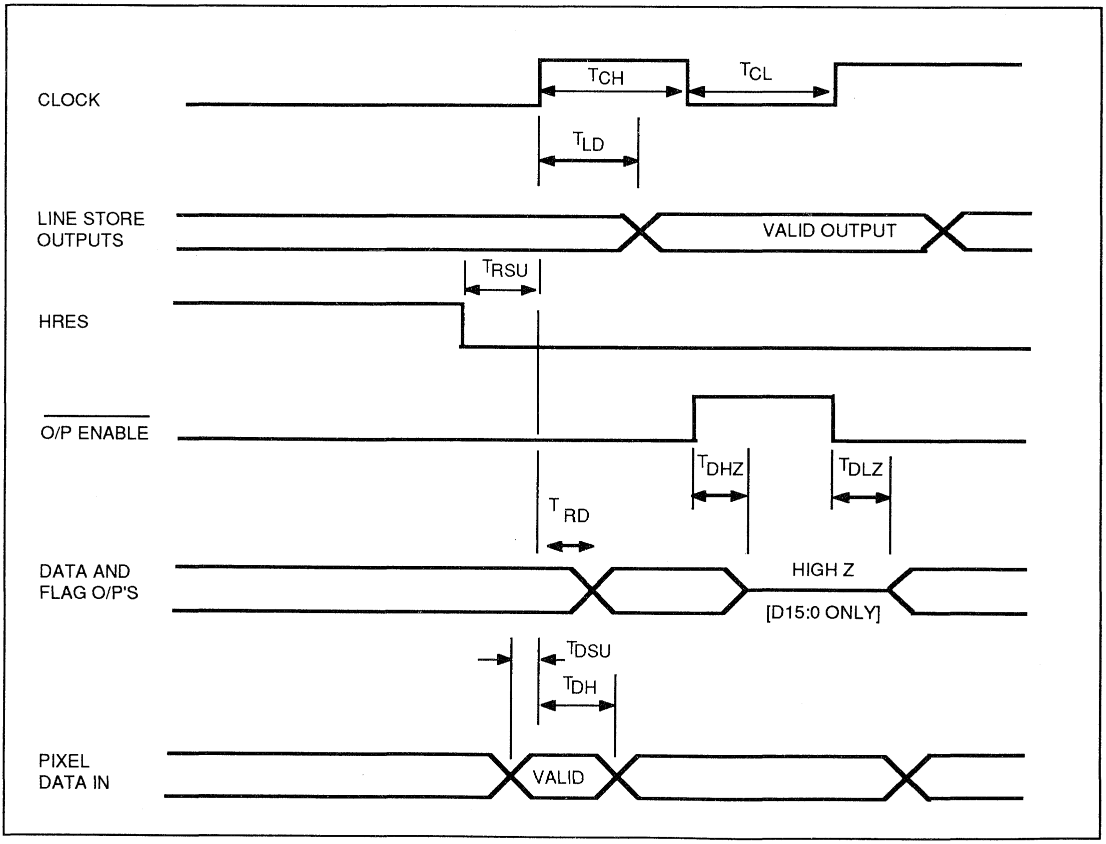

Fig. 10. I/O Timing

APPLICATIONS INFORMATION

DEVICE REQUIREMENTS

The number of devices required to implement a given convolver window depends on the size of the window, the required pixel rate, and whether the pixel accuracy is to be 8 or 16 bits. In practice the PDSP16488 supports windows requiring one, two, four, six, or eight devices without additional logic. Table 2 gives typical window sizes which may be obtained with the above number of devices.

Figures 11 through 18 show system interconnections for these arrangements. Other configurations are possible but may need the support of additional pixel/line delays and/or expansion adders. Although not necessarily shown, all configurations can be supported by either an EPROM or a Host Computer . Interlaced or non-interlaced video may also be used, unless explicitly stated otherwise in the text.

Expansion with 8 bit pixels is a straightforward process and the number of devices needed is easily deduced from the window sizes available in a single device. At pixel rates above 20MHz it may not be practical to use more than four devices. since the full 32 bit intermediate precision is not available. The lack of expansion multiplexing reduces the intermediate precision to 16 bits. The partial sum outputs must thus not overflow these 16 bits; this will require the coefficients to be scaled down appropriately with a resulting loss in accuracy.

Expansion with 16 bit pixels can be achieved in several ways. The simplest way is to use two devices, each working with 8 bit pixels. One device handles the least significant part of the data, and its output feeds the expansion input of a second device. This performs the most significant half of the calculation. The least significant half is then added to the most significant sum, after the latter has been multiplied by 256 ie shifted by eight places. This shift is done internally and controlled by Register D, bit 1. The internal 32 bit accuracy prevents any loss in precision due the shift and add operation.

The window size with this arrangement is restricted to that available in a single device, at the required pixel rate but with 8 bit pixels. Thus two devices can be used , for example, to provide an 8 x 8 window with 16 bit pixels and 10 MHz rates.

If a larger extended precision window is needed, it is possible to use four devices. Each device is then programmed to be in a 16 bit data mode, but should be restricted to rates below 20 MHz, if the 32 bit intermediate precision is to be maintained. In the 16 bit modes, however, the output from the last line delay is not available due to pin limitations. This is not a problem in a four device interlaced system, since half of the devices will be fed from an external field delay. In non interlaced systems additional external line delays would be needed. An alternative approach would be to configure all the devices in the appropriate 8 bit mode, do separate least significant and most significant calculations, and then combine the results in an external adder after a wired in shift.

SINGLE DEVICE SYSTEMS

Figures 11 illustrates both EPROM and Host supported single device systems, with or without interlaced video. In both cases the $\overline{\text{SINGLE}}$ and X15 pins must be tied tied low, and the $\overline{\text{PC0}}$, $\overline{\text{PC1}}$, and $\overline{\text{DS}}$ pins are redundant. The $\overline{\text{PROG}}$ pin becomes an output and indicates that a register load sequence is occuring. The first line delay must always be bypassed in a non interlaced system, however, since an internal pull up is provided, the BYPASS pin can be left open circuit for the correct operation. With interlaced video the BYPASS input is used to distinguish between the odd and even fields.

The CE input may be left open circuit if coefficients are to be simply loaded after a power on reset signal; the latter being applied to the $\overline{\text{RES}}$ input. Alternatively the $\overline{\text{CE}}$ input may be used to change the coefficients at any time after power on reset; the EPROM would then need additional address bits for the extra sets of coefficients that are to be stored.

In an interlaced system the pixels from the previous field must use the IP7:0 inputs, and the live pixels must use the L7:0 inputs. Interlaced sysytems requiring extended precision pixels are non supported with a single device, since the L7:0 inputs are then use for the least significant 8 bits, and the IP7:0 inputs for any more significant bits.

If the X15 pin is left open circuit, an internal pull up will configure the device in the host supported mode. The host must then supply a data strobe and a R/$\overline{\text{W}}$ control line. The X7:0 pins must be connected to the host data bus, and are used to both load and read back register values. The $\overline{\text{PROG}}$ and CE pins may be connected together, and then driven by a host address decode. The output on $\overline{\text{PC1}}$, which provides a $\overline{\text{REPLY}}$ to the host, need not be used if the width of the data strobe is greater than the maximum TEXP value given in Figure 7.

The configuration bits 6:4 in REGISTER A define the window size, maximum pixel rate, and pixel resolution. Window sizes smaller than the maximum in any configuration are implemented by filling in the window with `zero' coefficients. Bits 3:0 are irrelevent in the SINGLE mode, as is bit 7 if the gain contol is used.

The result would be expected to lie in either the bottom 20 bits of the 32 bit result , or possibly in the next 20 bit field displaced by four bits. Register C, bits 5:4, must thus select one of these fields for subsequent use by the gain control. The gain is then adjusted such that the 16 outputs available on pins are in fact the 16 most significant bits of the result. The gain needed is application specific, but if too much gain is used the OV pin will indicate an overflow.

Register B, bits 2:1, must be set to select the required method of defining the length of the line delays, and the use of bit 3 is dependent on any external pixel delays before the convolver input. No additional delays are needed on the pixel inputs in a single device system, and REGISTER D, bits 4:2, should be reset. The pipeline delay in the DELOP output path should match one of those in Table 4, and is window size dependent.

DUAL DEVICE CONFIGURATIONS

Two devices, each configured with 8 bit pixels and 8W x 4D windows, can be used to provide an 8 x 8 window at up to 20 MHz pixel rates. Figure 12 shows both the non interlaced and interlaced arrangements.

Video lines containing up to 1024 pixels are possible in both configurations, since each device only needs four line delays. One device is configured as the MASTER by grounding the $\overline{\text{MASTER}}$ pin; the other then receives control signals in the normal way and has its $\overline{MASTER}$ and $\overline{SINGLE}$ pins left open circuit.

The internal convolver sum, in the device producing the final result, must be delayed by 4 pixels to match the inherent delay in the expansion output from the other device. This is actually achieved by delaying the pixel inputs to the line stores [Register D bits 3:2 = 01]. No additional delay in the expansion input is needed, but the pipeline delay used to produce DELOP must be four clocks greater than that given in Table 4 for a single device. The DELOP output is redundant in one of the two devices.

Two devices can also be used to support systems requiring 16 bit pixels. With this approach the 16 x 8 multiplication is mechanized as two 8 x 8 operations, with the results added together after the most significant half has been shifted by 8 places to the most significant end. This shift operation is controlled by Register D, Bit 1. Both convolvers are programmed to contain the same coefficients. The convolved output can theoretically grow to 30 bits, and the appropriate field must be selected before using the gain control.

Examples of this operating mode are shown in Figure 13. Each device must be configured in the same 8 bit pixel operating mode, but the device producing the final result must use the 8 place shift option on its internal sum.

The least significant 8 bits of the pixel are connected to the MASTER device and the most significant 8 bits are connected to the device producing the final result.. The internal sum in this device must be delayed by four pixels to match the delay in the expansion output from the first device. This is actually achieved by delaying the pixel inputs to the line stores(Register D, bits 4:2, = 001]. The expansion input needs no additional delay [Register D bits 1:0 = 10].

The actual pixel precision can be any number of pixels between 8 and 16, and may be a signed or unsigned number. Any unused, more significant bits, must respectively be either sign extended or be tied low.

DELOP must have four additional pipeline delays in order to match the total processing delay. This output can be obtained from either device.

FOUR DEVICE SYSTEMS

Four devices, each in the 8x8 mode, can be used to provide a 16 x 16 window, with 8 bit pixel resolution and 10 MHz clock rates. The partial sum from the first device in each row must be delayed by eight pixel clocks before it is added to the result from the next device. This provides the eight pixel displacement to match the width of the window. The delay is actually provided by four additional delays in the expansion input to the next device, plus the inherent four clock delays in outputing results from the first device. Register D, Bit 0 controls the additional delay.

The internal convolver sums, in the two devices in the second row, must be delayed by 12 clocks before they are added to the result from the first row. This twelve clock delay is necessary because of the combination of the eight pixel horizontal displacement delay , and the four clock delay in outputing the result from the last device in the top row. It is actually achieved by delaying the pixel inputs to the line stores. (Register D, bits 3:2 = 11].

The DELOP output must have 20 delays additional to those in a single device. This compensates for the twelve delays added to the convolver sums in the second row, plus an additional eight delays to compensate for the partial width of the first device in the secind row.

Four devices can also be used to give an 8x8 window, but with a 30 MHz pixel clock. Each device is configured to provide a 4x4 partial window, but the maximum pixel rate is reduced from 40 to 30 MHz because of the response of the line delay expansion circuitry. Intermediate precision is restricted to 16 bits, since time multiplexed data outputs cannot be used above 20 MHz.

This configuration requires no additional delay in the expansion inputs, and the inputs to the line stores in both devices in the second row must be delayed by 8 clock cycles [Register D bits 3:2 = 10]. The DELOP output needs twelve additional clock delays to match the processing delay.

Figures 14 and 15 show non-interlaced and interlaced versions of the above 8 x 8 and 4 x 4 arrangements

Figure 16 shows how four devices can also be used to provide an 8x8 window, with 16 bit pixels and 20MHz clock rates. The expansion data from a previous device needs no additional delay since the partial window size in each device is only 4x4. The internal convolver sums from each device in the second row must be delayed by 8 Clks and the DELOP output must have 12 additional delays. If this arrangement is to be used in a non-interlaced application, the field store must be replaced by four line delays.

SIX DEVICE SYSTEMS

As shown in figure 17, six devices, each in an 8Wx4D mode using 8 bit pixels, can provide a 16W x 12D window at 20MHz clock rates. Expansion inputs from previous devices in a row [but not the first device in each row] need an extra 4 Clks of delay since the partial window is eight pixels wide. Internal convolver sums need a differential delay of 12 Clk cycles from row to row [Register D bits 3:2 = 11].

The DELOP output must have 32 additional delays to match the total processing delay.

EIGHT DEVICE SYSTEMS

Two additional chips will extend the above six device configuration to a 16 x 16 window. Internal convolver sums must have differential delays of 12 clock cycles between rows, as in the six device system. The DELOP output needs 44 additional clock delays.

NINE DEVICE SYSTEMS

Nine devices each in the 8 x 8 mode will provide a 24 x 24 window with 8 bit data and 10 MHz pixel clocks. This is shown in Figure 18. Expansion data inputs from previous devices in a row [but not the first device in each row] need an extra 4 Clks of delay. The internal convolver sums need differential delays of 20 Clk cycles between rows. Sixteen of the latter delays can be provided internally by setting Register B, bit3, and also Register D, bits 3:2. The four extra delays must be provided externally.

The DELOP output needs 56 clock delays in addition to the 29 required for the 8 x 8 single device configuration.

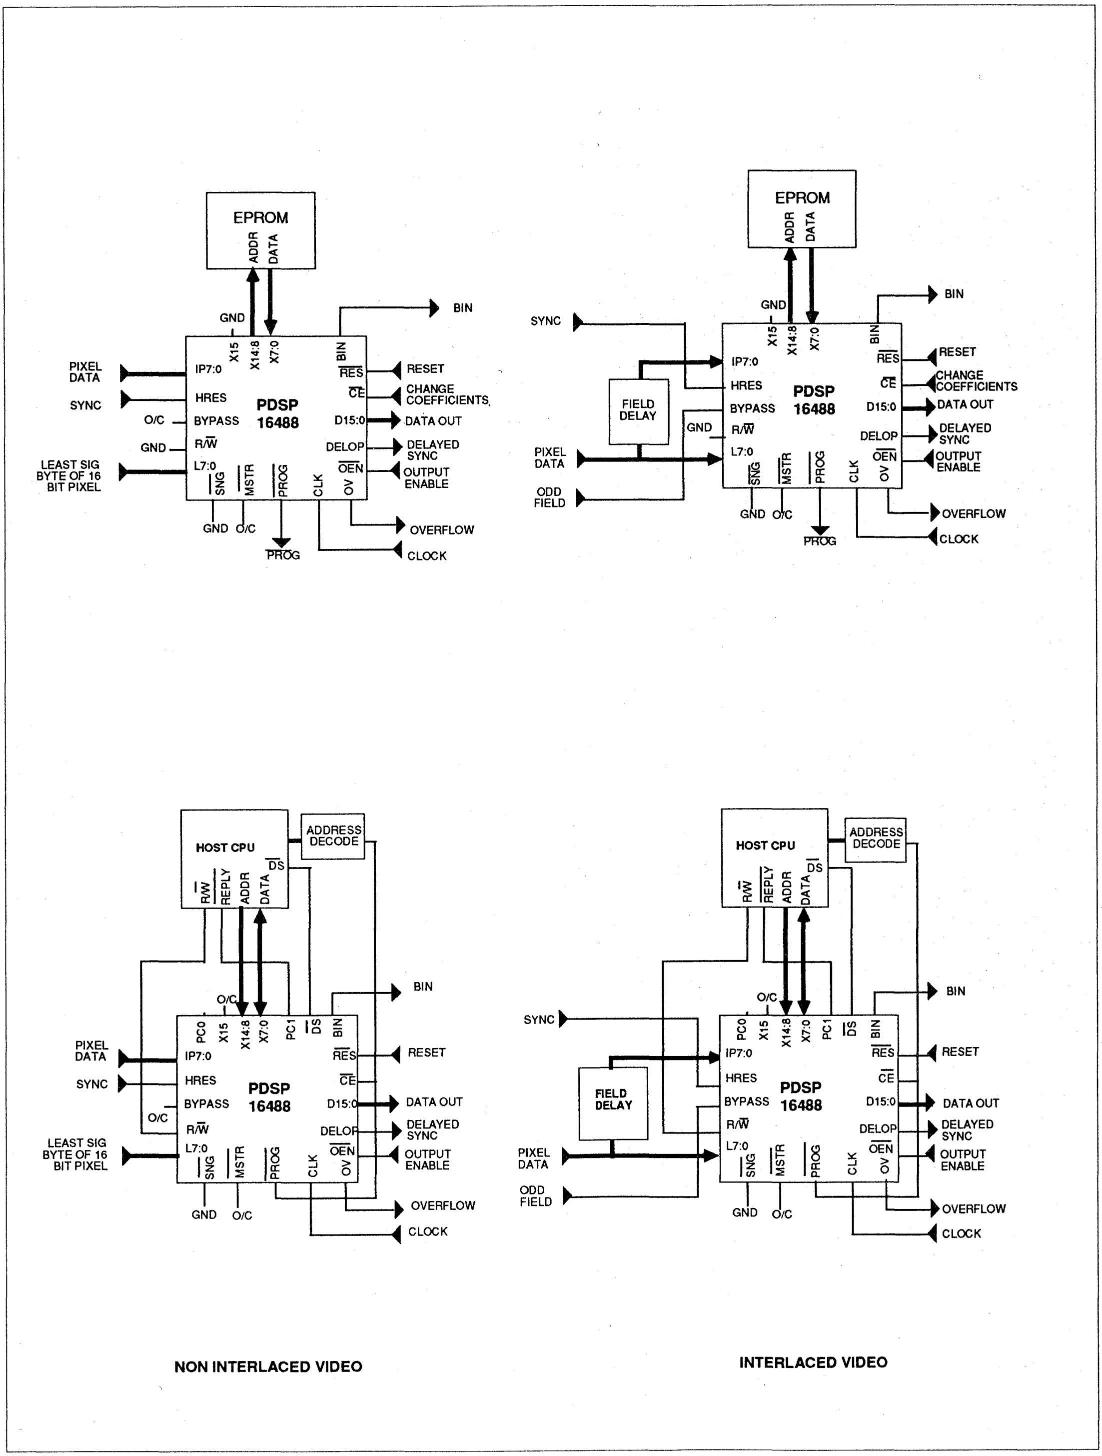

Figure 11 Single Device Systems

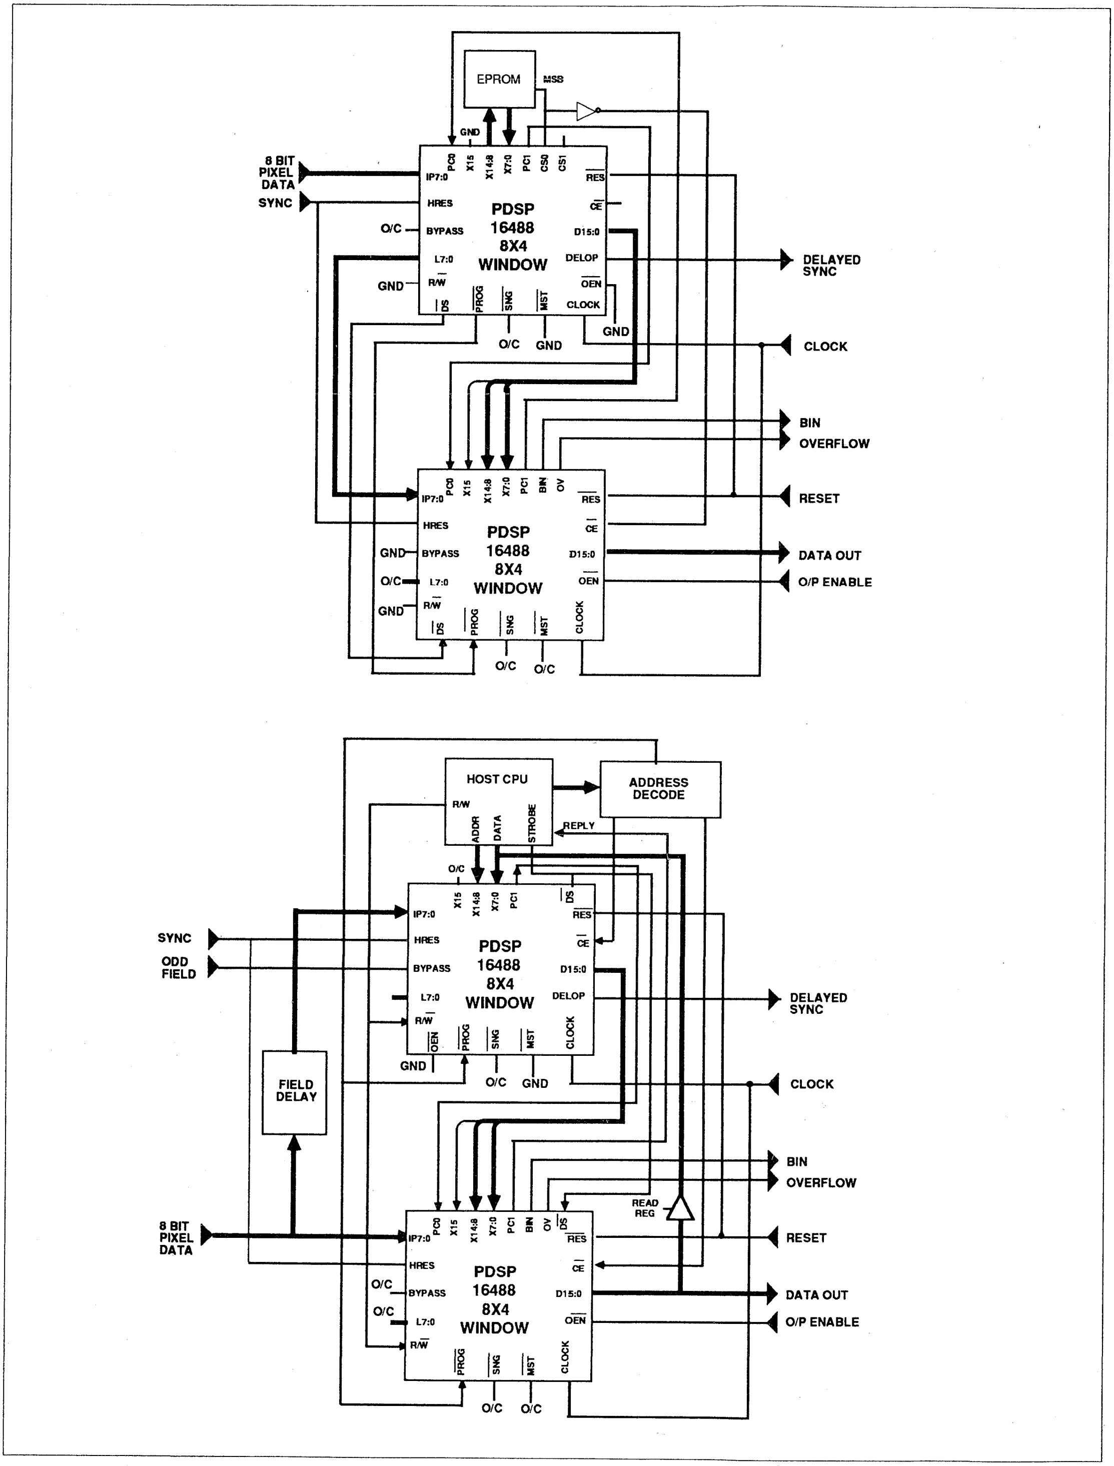

Figure 12. 8 Bit Dual Device Systems

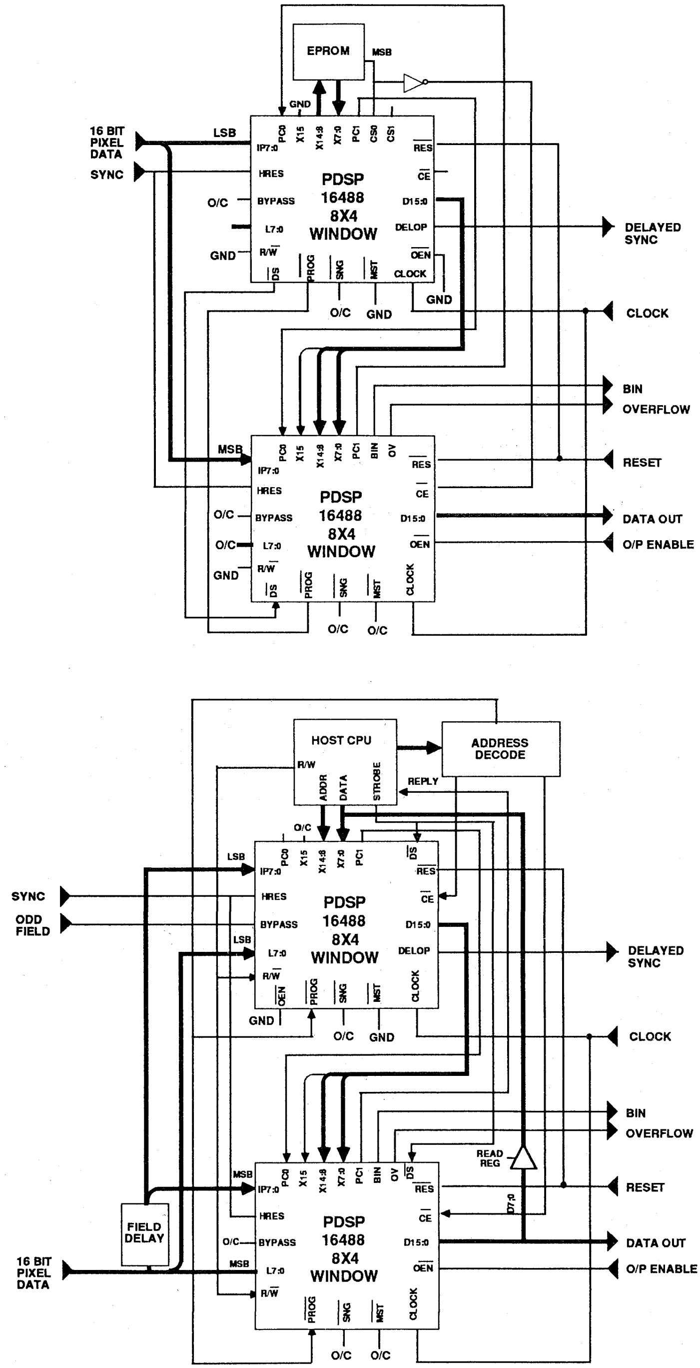

Figure 13. Dual Device 16 Bit Systems.

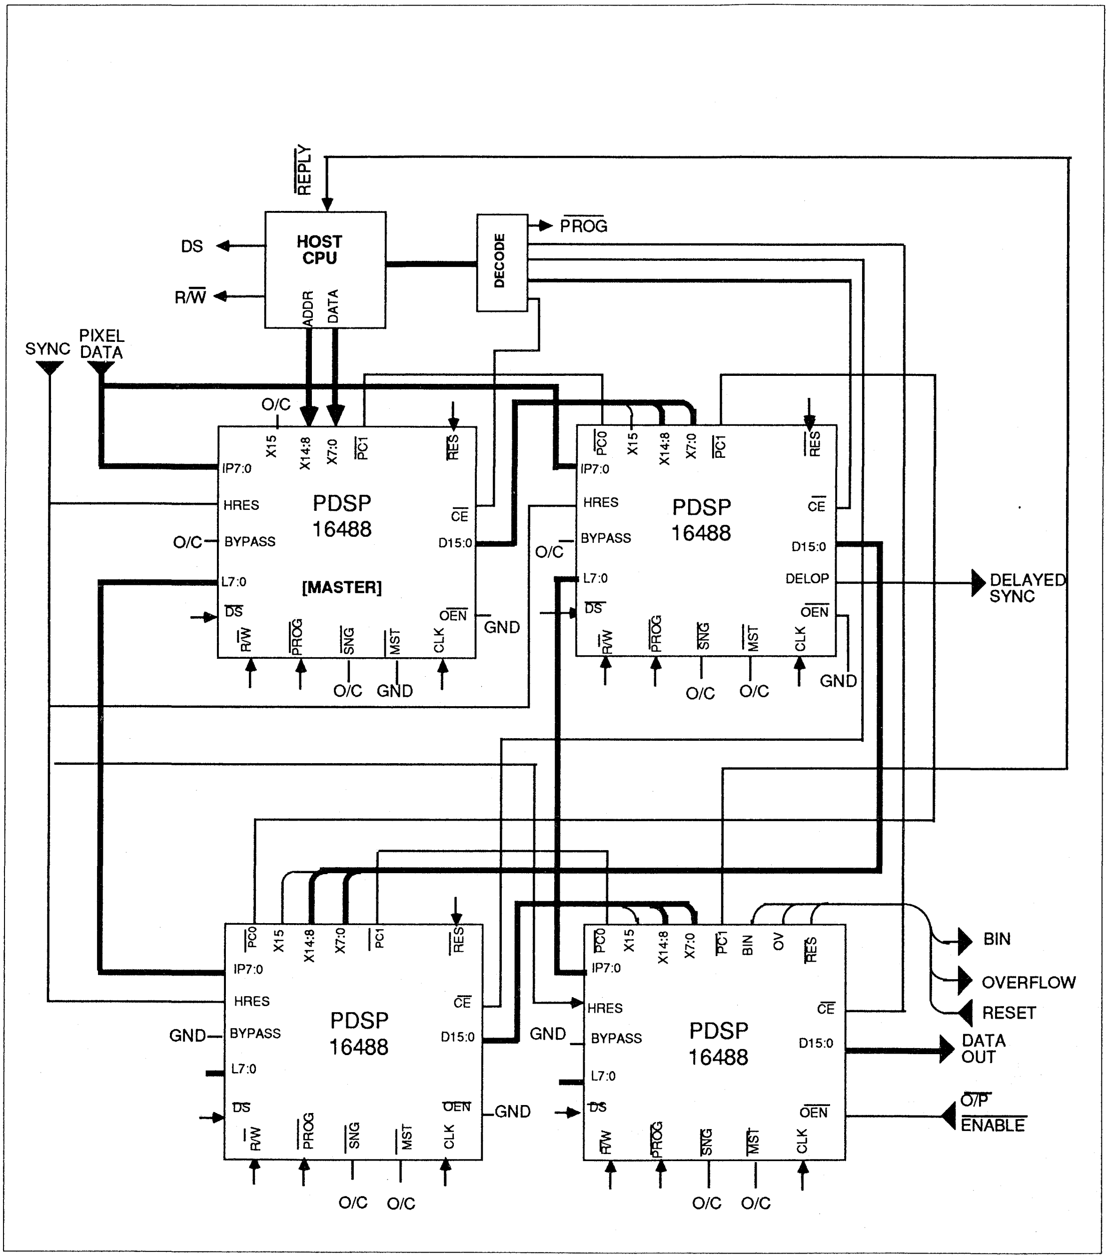

Figure 14. Four Devive Non Interlaced System.

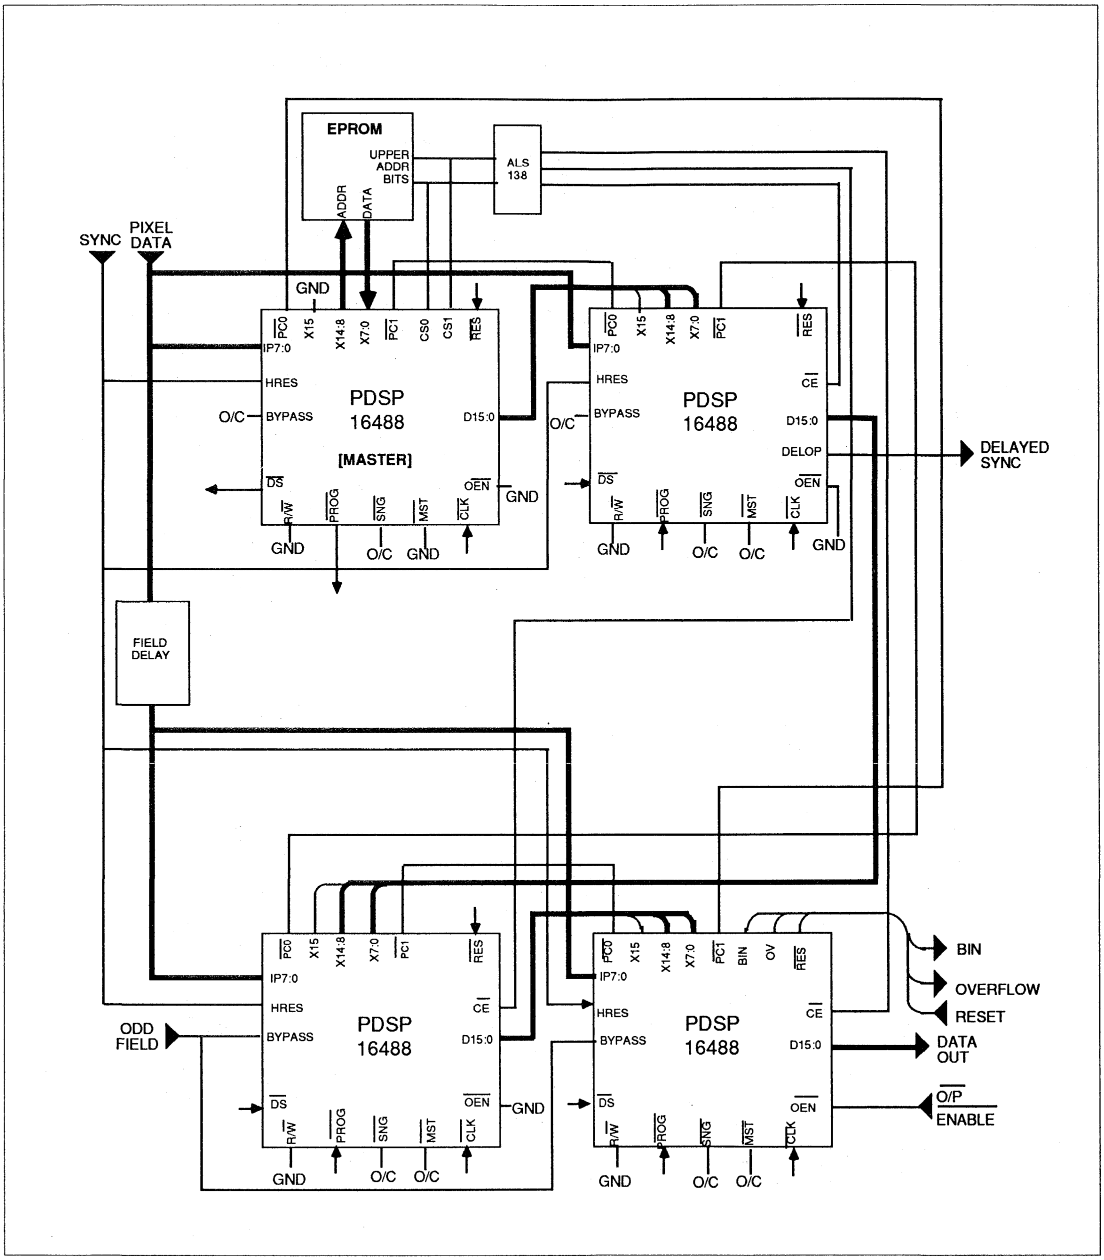

Figure 15. Four Deviice Interlaced System.

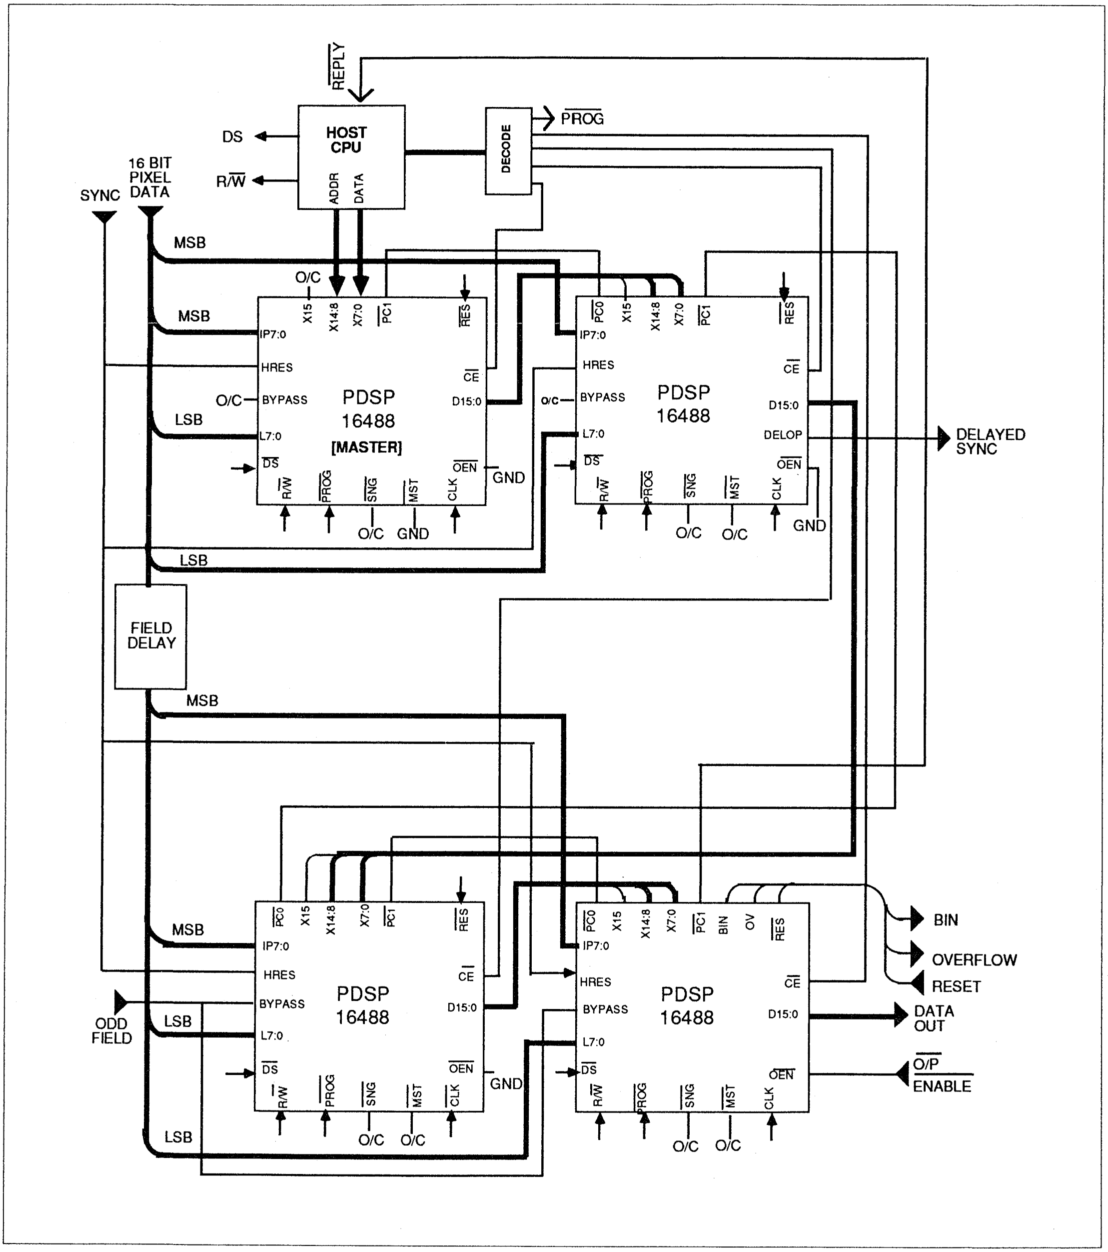

Figure 16. Four Device System with 16 Bit Pixels

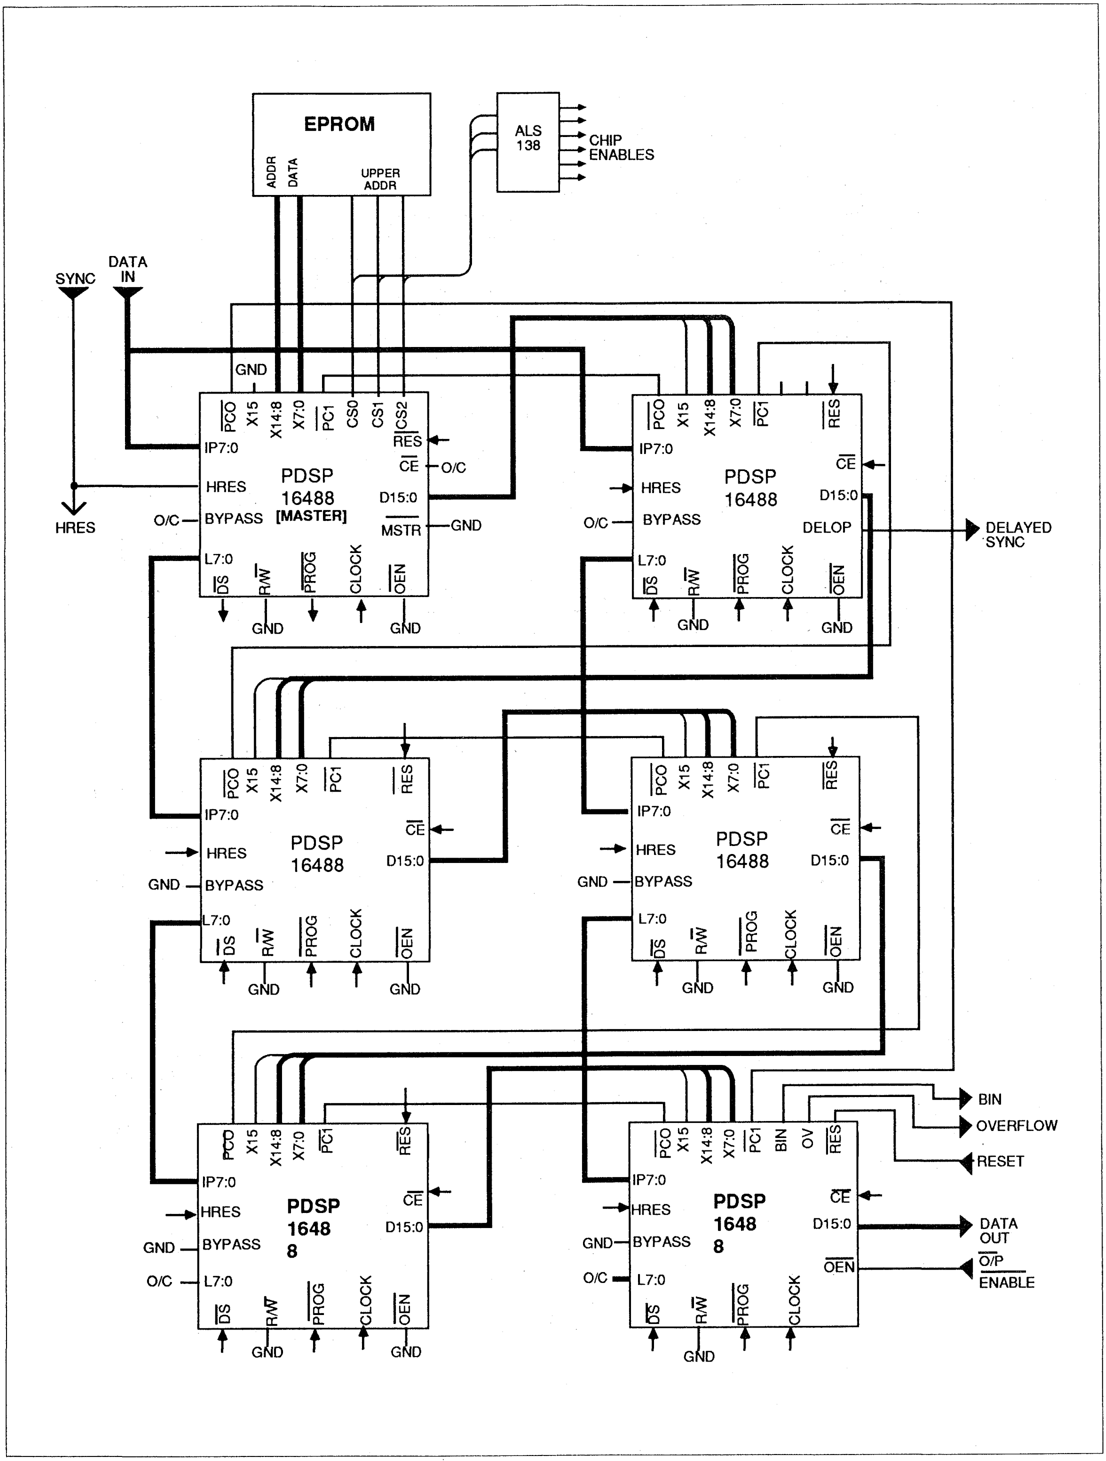

Figure 17. Six Device Non Interlaced System.

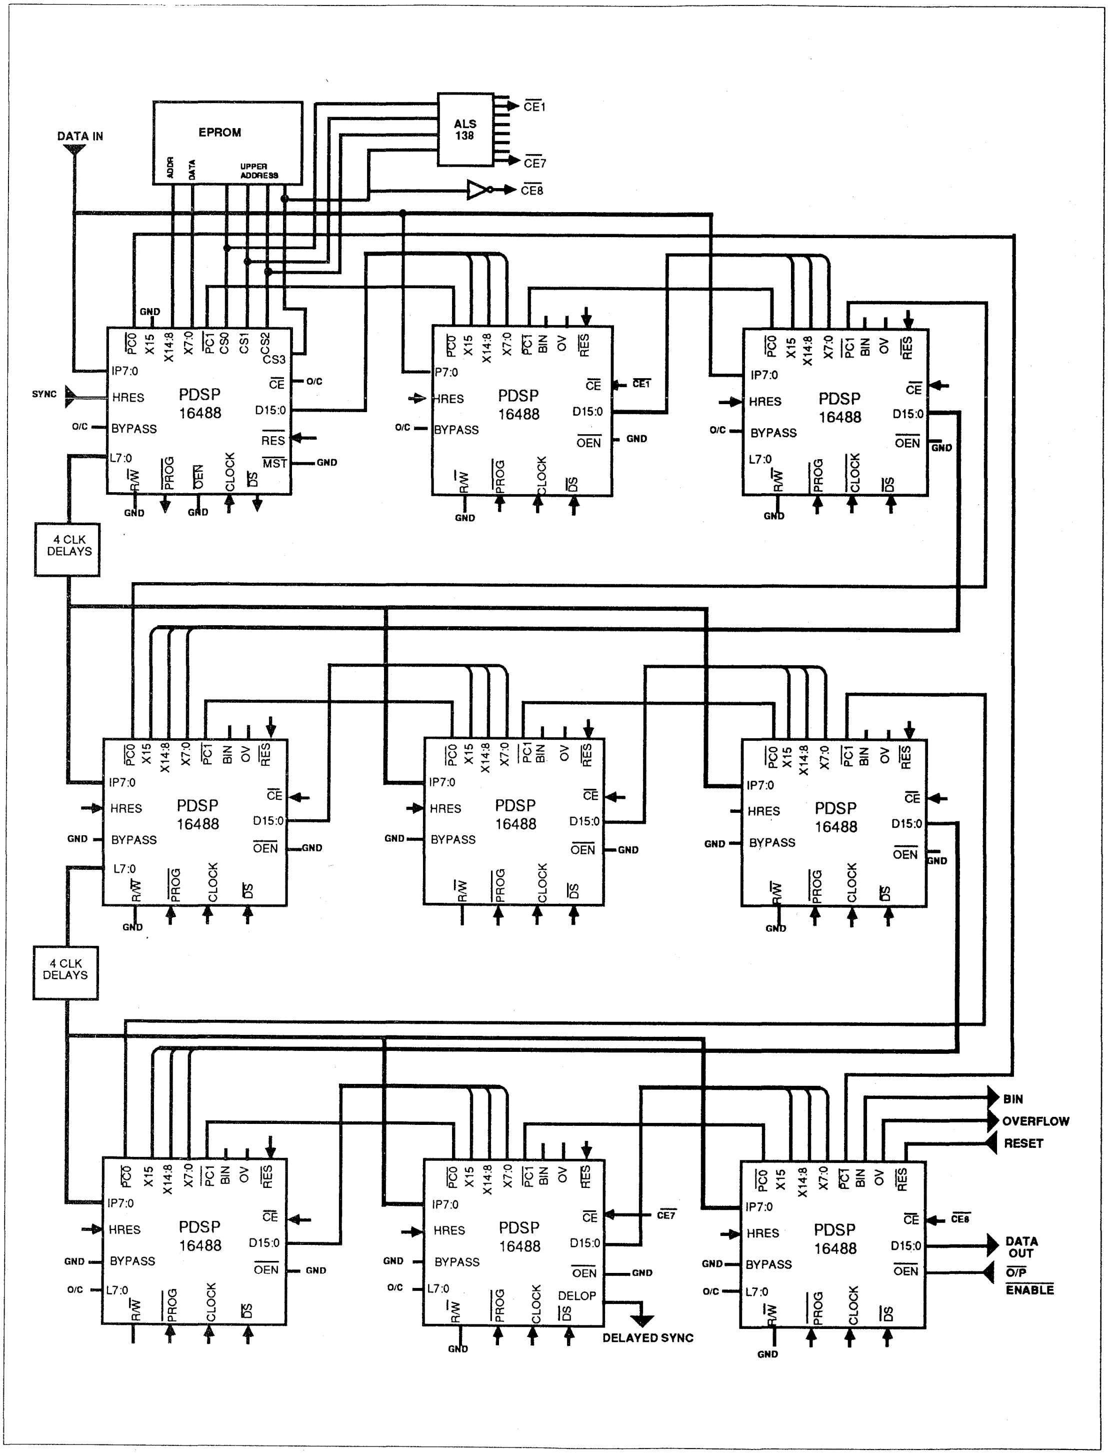

Figure 18. Nine Device Non Interlaced System.

ELECTRICAL CHARACTERISTICS

Operating Conditions

Commercial T_{amb} = 0°C to +70°C, T_{JMAX} = 95°C, V_{CC} = 5.0V ± 5%, fround = 0V

Industrial: T_{amb} = -40°C to +85C, T_{JMAX} = 110°C, V_{CC} = 5.0V ± 10, fround = 0V
Military : T_{amb} = -55°C to +125°C, T_{JMAX} = 150°C, V_{CC} = 5.0V ± 10%, fround = 0V

ORDERING INFORMATION

Commercial (0°C to +70°C)
PDSP1488 / C0 / AC (PGA
PDSP1488 / C0 / GC (QFP)

Industrial (-40°C to +85°C)
PDSP1488 / B0 / AC (PGA
PDSP1488 / B0 / GC (QFP)

Military (-55°C to +125°C)
PDSP1488 / A0 / AC (PGA
PDSP1488 / A0 / GC (QFP)

Call for availability of high rel and MIL883 screened parts

GEC PLESSEY
SEMICONDUCTORS

PRELIMINARY INFORMATION

DS3475 - 4.1

PDSP16510A

STAND ALONE FFT PROCESSOR

(Supersedes DS3475 - 3.1 March 1992)

The PDSP16510 performs Forward or Inverse Fast Fourier Transforms on complex or real data sets containing up to 1024 points. Data and coefficients are each represented by 16 bits, with block floating point arithmetic for increased dynamic range.

An internal RAM is provided which can hold up to 1024 complex data points. This removes the memory transfer bottleneck, inherent in building block solutions. Its organisation allows the PDSP16510 to simultaneously input new data, transform data stored in the RAM, and to output previous results. No external buffering is needed for transforms containing up to 256 points, and the PDSP16510 can be directly connected to an A/D converter to perform continuous transforms. The user can choose to overlap data blocks by either 0%, 50%, or 75%. Inputs and outputs are synchronous to the 40MHz system clock used for internal operations.

A 1024 point complex transform can be completed in some 98µs, which is equivalent to throughput rates of 450 million operations per second. Multiple devices can be connected in parallel in order to increase the sampling rate up to the 40MHz system clock. Six devices are needed to give the maximum performance with 1024 point transforms.

Either a Hamming or a Blackman-Harris window operator can be internally applied to the incoming real or complex data. The latter gives 67dB side lobe attenuation. The operator values are calculated internally and do not require an external ROM nor do they incur any time penalty.

The device outputs the real and imaginary components of the frequency bins. These can be directly connected to the PDSP16330 in order to produce magnitude and phase values from the complex data.

ASSOCIATED PRODUCTS

PDSP16540 Bucket Buffer

PDSP16330 Pythagoras Processor.

PDSP16256 Programmable FIR Filter.

PDSP16350 I/Q Splitter / NCO

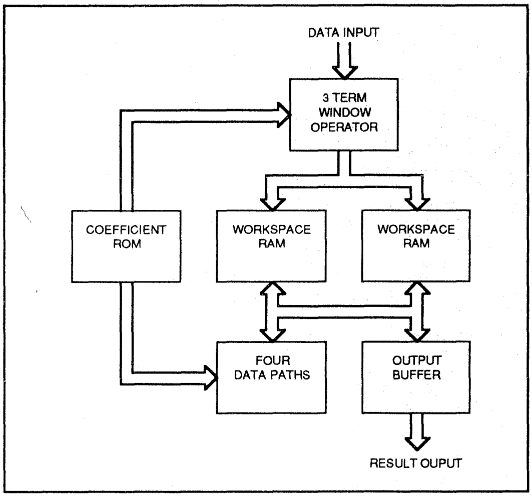

Fig. 1. Block Diagram

FEATURES

- Completely self contained FFT Processor
- Internal RAM supports up to1024 complex points
- 16 bit data and coefficients plus block floating point for increased dynamic range
- 450 MIP operation gives 98 microsecond transformation times for 1024 points
- Up to 40MHz sampling rates with multiple devices.
- Internal window operator gives 67dB side lobe attenuation and needs no external ROM.
- 84 pin PGA or 132 surface mount package

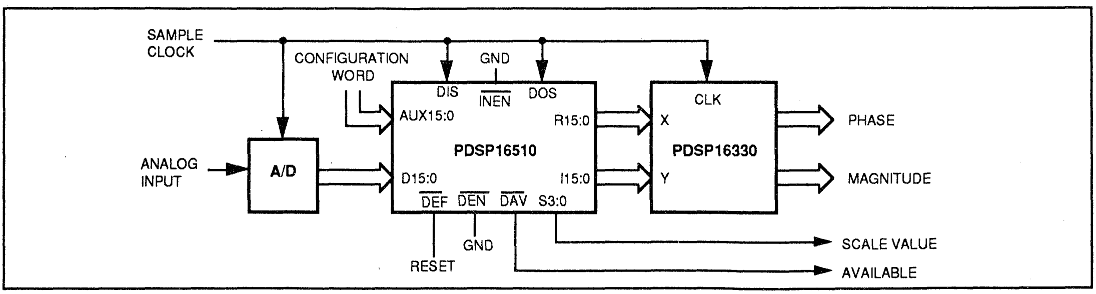

Fig. 2. Typical 256 Point Real Only System Performing Continuous Transforms

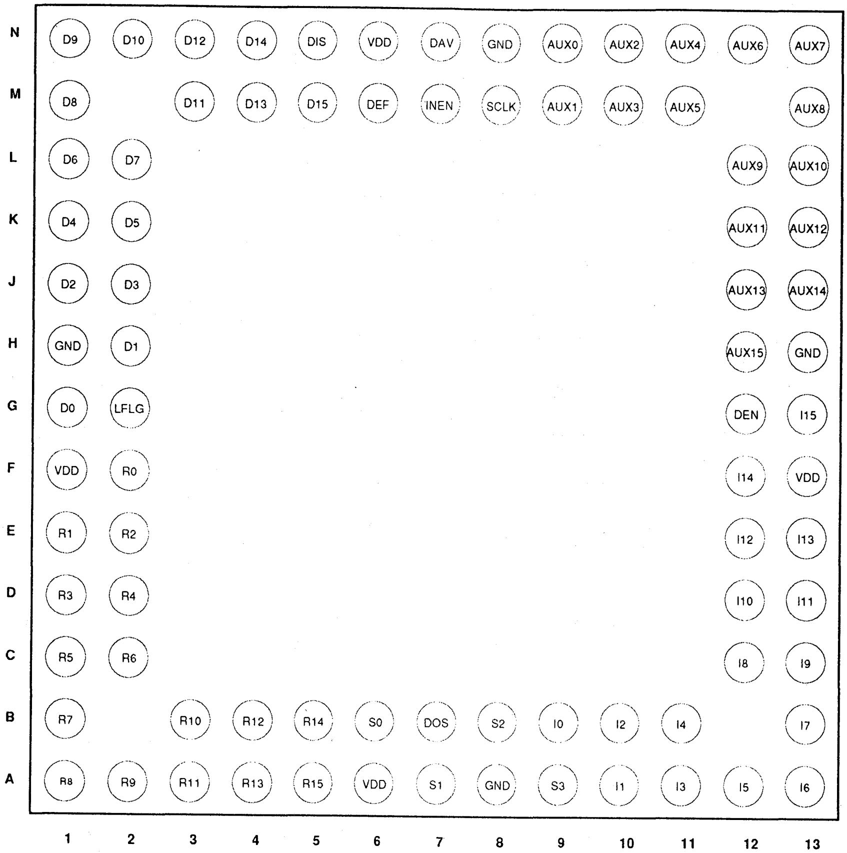

Pin Out for 84 PGA Package (AC84) - bottom view

PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC
1	VDD	23	AUX13	45	GND	67	D8	89	GND	111	GND
2	GND	24	VDD	46	VDD	68	D7	90	R3	112	S1
3	I7	25	AUX12	47	SCLK	69	D6	91	VDD	113	GND
4	I8	26	GND	48	GND	70	D5	92	R4	114	DOS
5	I9	27	AUX11	49	GND	71	GND	93	GND	115	DOS
6	I10	28	VDD	50	DAV	72	VDD	94	R5	116	VDD
7	VDD	29	GND	51	GND	73	D4	95	R6	117	S2
8	I11	30	AUX10	52	INEN	74	GND	96	R7	118	GND
9	GND	31	AUX9	53	VDD	75	D3	97	R8	119	S3
10	I12	32	AUX8	54	DEF	76	VDD	98	GND	120	GND
11	VDD	33	AUX7	55	GND	77	D2	99	VDD	121	VDD
12	I13	34	VDD	56	DIS	78	GND	100	R9	122	I0
13	GND	35	AUX6	57	VDD	79	D1	101	VDD	123	I1
14	I14	36	VDD	58	D15	80	VDD	102	R10	124	GND
15	VDD	37	AUX5	59	D14	81	D0	103	R11	125	I2
16	I15	38	GND	60	GND	82	LFLG	104	R12	126	I3
17	GND	39	AUX4	61	D13	83	GND	105	R13	127	I4
18	DEN	40	AUX3	62	D12	84	R0	106	GND	128	GND
19	AUX15	41	AUX2	63	D11	85	GND	107	R14	129	VDD
20	GND	42	VDD	64	D10	86	R1	108	R15	130	I5
21	AUX14	43	AUX1	65	VDD	87	VDD	109	DISAB	131	I6
22	GND	44	AUX0	66	D9	88	R2	110	S0	132	VDD

Pin Out for 132 Leaded Chip Carrier (GC132)

SIGNAL	TYPE	DESCRIPTION
D15:0	I	Data input during real only mode. The real component in complex data mode.
AUX15:0	I	When DEF is active AUX15:0 are used to define the operating mode as defined in Table 3. When DEF is in-active AUX15:0 either provide the 16 bit imaginary component of complex input data, or a second set of real only inputs.
R15:0	O	These pins output the real component of the transformed data when DAV and DEN are active. Otherwise they are high impedance.
I15:0	O	These pins output the imaginary component of the transformed data when DAV and DEN are active. Otherwise they are high impedance.
$\overline{\text{DEF}}$	I	The high going edge of DEF is used to internally latch the contents of AUX15:0, which then define the operating mode. In the simplest system DEF is a power on reset. When DEF is low the internal control logic is reset.
SCLK	I	System clock used for internal computations.
S3:0	O	These pins indicate the number of shifts towards the binary point which have occurred as the result of the conditional scaling logic. When the data path right shift is restricted to 2 places per pass, state 15 is used to indicate an overflow and only a total of 14 shifts is possible.
LFLG	O	This flag indicates that data is being loaded into the device. It goes active in response to an INEN input, and may be programmed to go in-active after the complete, one quarter, or one half a data block has been loaded.
$\overline{\text{INEN}}$	I	The use of this input is mode dependent. It is either used as an active low, load enabling, signal for the DIS strobe, or it is used to initiate a new block load operation.
DIS	I	The rising edge of this input is used to load data into the device.
DOS	I	The rising edge of this input is used to dump data from the device. In most applications it may be tied to the DIS input, even if the output rate must be higher than the input rate because of overlapped data blocks. The DIS input is then internally divided down.
$\overline{\text{DAV}}$	O	An active low signal that indicates that a transform is complete. Transformed data will then be output in normal sequential order using DOS. It may be optionally programmed to be delayed by 24 DOS strobes to match the delay through a PDSP16330.
$\overline{\text{DEN}}$	I	This input is used to enable the data dump operation when DAV has gone active. If it is tied low the device will automatically dump data when DAV goes active. Otherwise the device will wait for the enabling signal to go low before the dump operation commences.
DISAB	I	Only available in the 132 pin GC package. When high the block floating logic is disabled.
VDD	P	+5V pins
GND	P	Ground pins

NOTE. All references to DEF, INEN, DAV, and DEN within the text do not contain the bar designator, signifying an active low signal. This is considered to be implied by the signal name and is not meant to imply a change in the signal function.

FUNCTIONAL OPERATION

The PDSP16510 performs decimation in time, radix 4, forward or inverse Fast Fourier Transforms. Data is loaded into an internal workspace RAM in normal sequential order, processed, and then dumped in the correct order. With real only input data the processing time can approximately be halved for a given transform size. Two real inputs then replace a single complex input, and are processed in parallel.

Either a Blackman-Harris or a Hamming window can be generated internally, and applied to the incoming real or complex data with no time penalty. No external ROM is needed to support these windows. The Blackman-Harris window gives improved dynamic range over the Hamming window when two closely spaced frequencies are to be detected, and one is of smaller magnitude than the other. It does, however, reduce the actual frequency resolution, and the Hamming window may then be preferable.

Data in and out of the device is represented by 16 bit real and imaginary components, with 16 bit sine and cosine values contained in an internal ROM. Conditional scaling, coupled with word growth through the butterfly data path, gives increased dynamic range. Transforms can be computed with sample sizes of either 256 or 1024 data points. The 256 point option can alternatively be used to simultaneously execute either four 64 point transforms, or sixteen 16 point transforms. The 16 point mode can only be used with a rectangular window, and no overlapping of data blocks is possible.

The device can be configured, either, to perform continuous transforms in a real time application, or as slave processor to a more general purpose signal processing system. In the continuous mode, with transform sizes of 256 points or less, it contains three internal control units which simultaneously allow new data to be loaded, present data to be transformed, and previous results to be dumped. Additional, external, input/output buffering is not needed. The internal input buffer also allows data blocks to be overlapped by either 50% or 75%, apart from the mode with no overlaps.

When 1024 point transforms are to be calculated, without loss of incoming data during the transform time, it is necessary to use an input buffer. This requirement is satisfied by a single PDSP16540 support device.

In any of the real or complex modes it is possible to obtain higher performance by connecting devices in parallel. It is then possible to increase the sampling rate to that of the system clock used for internal operations.

The mode of operation of the device is controlled by 16 bits in a control register. These are loaded through the AUX15:0 port when a control signal DEF is active low. This port is also used to provide the imaginary component of complex input data, and, if complex transforms are to be performed, an external tristate buffer will be needed to isolate the control information. This should only be enabled when DEF is active. DEF is also used to initialise the internal circuitry, and can be a simple power on reset if control parameters need not be subsequently changed.

DATA PRECISION

During each pass of a radix-4 fast Fourier transform it is possible for either component of a particular result to grow by a factor of up to four in the first pass, and 5.242 in subsequent passes. This is between two and three bits in each pass and the data path must allow for this word growth to avoid any possibility of overflow. At the end of the data path the word is again reduced to 16 bits by discarding least significant bits. Any un-necessary word growth to prevent overflow thus results in loss of arithmetic precision, and has a detrimental effect on the dynamic range achievable.

In practice these large word growths only occur when bipolar complex square waves are transformed, and even then will not occur on every pass. The PDSP16510 compromises by allowing a 2 bit word growth during the butterfly calculation in the first pass. This is equivalent to ignoring the most significant bit of the 19 bit final result, which is assumed to be an extra sign bit, and then selecting the next 16 bits for

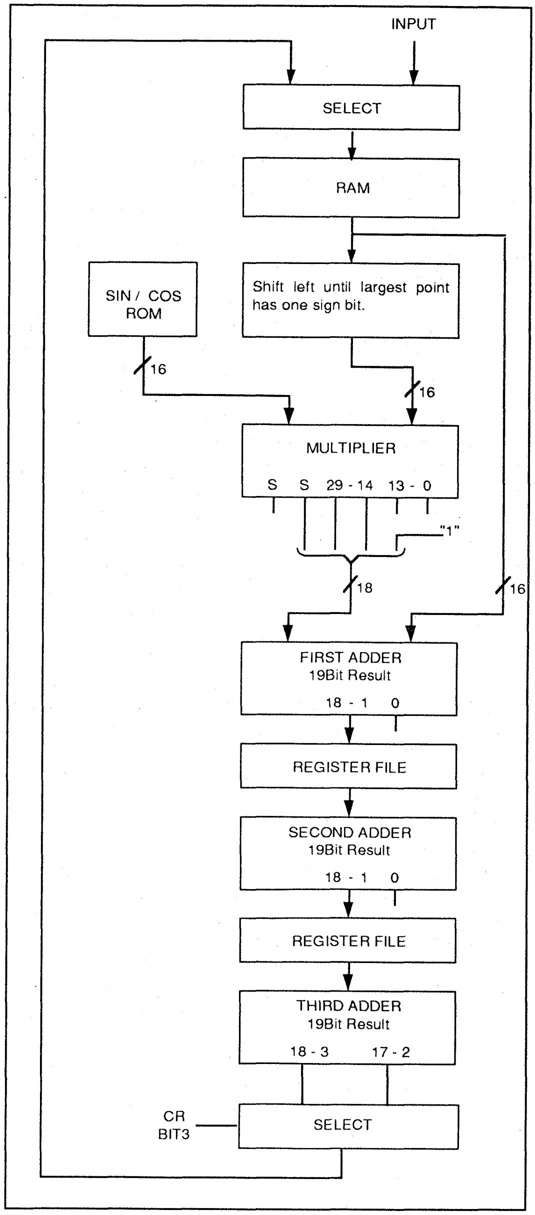

Fig. 3 One of Four Data Paths

storage. In subsequent passes a Control Register Bit allows the user to continue to select these 16 bits, or instead to use the 16 most significant bits. The latter option is equivalent to a 3 bit word growth. The 2 or 3 bit word growth option applies to ALL subsequent passes and is not a per pass option.

If the 2 bit option is selected there is a possibility of overflow occurring in one of the passes. The prediction of overflow is mathematically difficult, and only occurs with specific complex square waves. Scaling down the inputs cannot be guaranteed to prevent overflow because of the block floating point shifting scheme, which is discussed later. Overflow can NEVER occur if the 3 bit option is chosen, but at the expense of worse dynamic range.

When overflow does occur a flag is raised which can be read by the user (see later discussion on scale tag bits), and the results ignored. In addition all frequency bins are forced to zero to prevent any erroneous system response.

Even with only 2 bit word growth poor dynamic range will be obtained if the data is simply reduced to 16 bits, and becomes worse when the incoming data does not fully occupy all the bits in the word. These problems are overcome in the PDSP16510, however, by a block floating point scheme which compensates for any unnecessary word growth.

During each pass the number of sign bits in the largest result is recorded. Before the next pass, data is shifted left [multiplied by 2], once for every extra sign bit in this recorded sample. At least one component in the block then fully occupies the 16 bit word, and maximum data accuracy is preserved

Up to four shifts are possible before every pass after the first, with a total of fifteen for the complete transform. At the end of the transform the number of left shifts that have occurred is indicated on S3:0. Lack of pins prevents a separate output being available to indicate that overflow has occurred in the 2 bit word growth option. For this reason the maximum number of compensating left shifts in this mode is restricted to 14. State 15 is then used to indicate that overflow has occurred.

The first step in the butterfly calculation multiplies 16 bit data values with 16 bit sine/cosine values, to give 18 bit results. This increased word length preserves accuracy through the following adder network, and has been shown through simulations to be an optimum size for transform sizes up to 1024 points. This is particularly true when the input data is restricted to below 16 bits, as is necessary with practical A/D converters with very high sampling rates. The bottom bit of this 18 bit word is forced to logical one and as such is a compromise between truncation and true rounding. It gives a lower noise floor in the outputs compared to simple truncation.

To prevent any possibility of overflow during the butterfly calculation the word length is allowed to grow by one bit through each of the three adders. The least significant bit is always discarded in the first two adders . Sixteen bits are then chosen from the final adder in the manner discussed earlier, and the number of sign bits in the largest result is recorded for use in the following pass.

Fig. 3 shows one of the four internal data paths which can compute a radix-4 butterfly in twelve system clock cycles. This equates to completing the butterfly in 3 cycles for the complete device.

DATA TRANSFERS

The data transfer mechanism to and from the internal

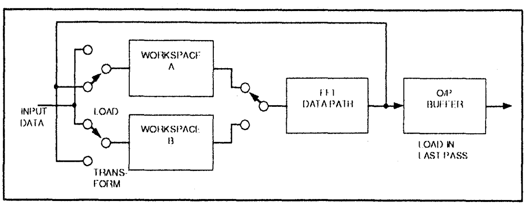

Fig. 4. RAM Organization with 256 Data Points

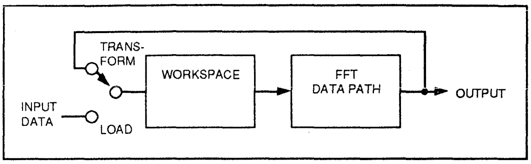

Fig. 5. RAM Organization with 1024 Point Transforms

RAM has been designed for use in a wide variety of applications. The provision of an independent input strobe (DIS), allows data to be loaded without the need for additional external buffering. An independent output strobe (DOS) is also provided. DIS and DOS can thus be tied together, this being particularly useful when the device is performing the inverse transform back to the time domain. Transfer of data occurs internally from DIS to SCLK, so although thay can be of different frequencies, they must be synchronous to each other. In the same way transfer of data also occurs from SCLK to DOS, so while DOS can also be independent of SCLK it must also be synchronous to it. Inputs and outputs are both supported by flag and enabling signals which allow transfers to be properly co-ordinated with the internal transform operation.

In many applications the DIS and DOS inputs can be tied together and fed by the sampling clock. If the output rate must be higher than the input rate, as with multiple devices supporting overlapped data samples, both strobes can still be connected together. The clock supplied should then be twice or four times the sampling clock, and an internal divider can be used to provide the correctly reduced input rate. The provision of a separate DOS pin does, however, allow the output rate to be different to the input rate, and therefore faster than strictly needed. Further output processing at higher rates is then possible if this is advantageous to system requirements.

The internal workspace is double buffered when 256 point transforms are to be performed. A separate output buffer is also provided. These resources, together with separate input and output buses, allow new data to be loaded and old results to be dumped, whilst the present transform is being computed. Additional, external, input buffering is not needed to prevent loss of incoming data whilst a transform is being performed.

When block overlapping is required, internally stored data will be re-used, and a proportionally smaller number of new samples need be loaded. Note that the internal window operator still functions correctly since it is actually applied during the first pass, and not whilst data is being loaded. The internal RAM organisation is shown in Fig. 4. It should be

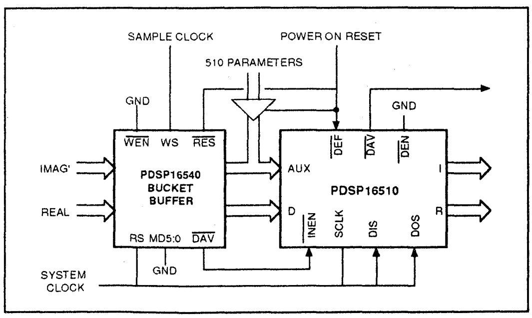

Fig. 6. 1024 Point Transforms with I/P Buffer

noted that the amount of overlap between I/O transfers and transforms is completely under the control of the system, since an input enable signal (INEN) and an output enable (DEN) can be used to initiate transfers.

In the 1024 point mode there is insufficient workspace for input and output buffering in addition to working memory. The device is then configured in a mode with separate load, transform and dump operations. The internal arrangement is shown in Fig. 5. The support of an external input buffer is needed if incoming samples are not to be lost whilst a transform is in progress. This is loaded at the sample clock rate and transferred to the FFT processor as quickly as possible. In this mode the PDSP16510 always expects to receive 1024 words, regardless of the amount of block overlapping. Data stored internally cannot be re-used when block overlapping is required, and data from the external buffer must be re-read as necessary.

Fig. 6 illustrates a typical 1024 point system with an input buffer which supports complex input data. The input buffer can be provided by a PDSP16540 Bucket Buffer without the need for any external control logic. It supplies RAM for 1024 x 32 complex words, and allows transfers to the FFT Processor at the full system clock rate. The PDSP16540 also supports the standard 50% and 75% data block overlapping, but in addition allows the user to define the amount of overlap to within 32 words.

If no incoming data is to remain un-processed, the user must ensure that the time taken to acquire sufficient data to instigate a new transform is greater than or equal to the transformation time itself. The latter can be calculated from Table 4, once the system clock rate has been defined. When 1024 point transforms are performed, both the time to read data from the input buffer, and also the time to dump data, must be included in the calculation to determine the minimum time in which data can be loaded into the external buffer.

The peak transfer rate is limited by the characteristics of the I/O circuits, but can be greater than the sampling rate which is determined by the transform time. When load and dump operations are not concurrent with transform operations (as in the 1024 point modes), then the maximum I/O rate is equal to the system clock rate, Ø. When other transform sizes are specified, the sampling rate, S, is reduced by a factor F. This is defined below where Ø is in MHz and L is the system clock low time in nanoseconds :

$$S = F\varnothing, \quad \text{where } F = 4 / (6+0.001\varnothing L)$$

F is typically 0.66 and applies to all transforms except for those of 1024 points, even if INEN is driven such that concurrent operations do not actually occur (Note also that S must be

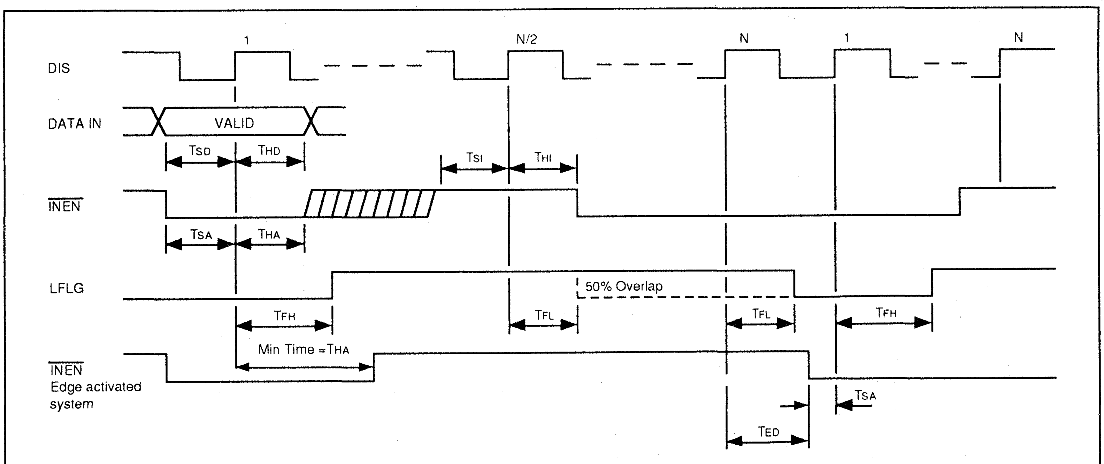

Characteristic	Symbol	16510A,A0,B0,C0 Min	Max		Units
Data In set up Time	T_{SD}	10			ns
Data In Hold Time	T_{HD}	0			ns
$\overline{INEN}$ active going set up	T_{SA}	8			ns
$\overline{INEN}$ active Hold Time	T_{HA}	0			ns
$\overline{INEN}$ in-active Hold Time to ensure no load	T_{HI}	2			ns
$\overline{INEN}$ in-active going set up for no load operation	T_{SI}	8			ns
Delay to LFLG going active (30 pf load)	T_{FH}		10		ns
Delay to LFLG going in-active (30 pf load)	T_{FL}		10		ns
Min time to $\overline{INEN}$ low in edge mode	T_{ED}	15			ns

Table 1. Advanced Timing Information with Continuous Inputs.

synchronous to SCLK). If this causes a system limitation in a single device application, then the device can be configured for pseudo, Mode 2, multiple device operation. Separate load, transform, and then dump operations will then always occur, but DEN must be low when a transform is complete or DAV will never go active. See the section on multiple device operation.

LOADING DATA

Data loading is controlled by three signals; DIS an input strobe, INEN a load enable, and LFLG an output flag. Detailed timing information is given in Table 1. Once sufficient data has been acquired, a transform will automatically commence. This is normally after a complete block has been loaded, except when a single device is performing overlapped transforms of 256 points or less. With 75% overlapping, transforms will commence after 25% of a new block has been loaded, and with 50% overlapping transforms commence after 50% of the data has been loaded. The remainder of the block is provided by data already stored in the internal RAM.

The data strobe is used to load data into the internal workspace RAM, and data must meet the specified set up and hold times with respect to its rising edge. DIS can be a continuous input since the device only loads data when an input enabling signal is active.

An internal synchronisation interval is necessary between the last sample being loaded with the DIS strobe and transforms being started with the system clock. This can be up to twelve system clock periods when data transfers and transforms are overlapped. The transform times given later in Table 4 are maximum values, and include these twelve periods.

The way in which the INEN signal controls data loading is dependent on whether a single or multiple device is to be implemented, and the status of Control Register Bit 12.

When Bit12 is set in a SINGLE device system the INEN signal is simply used as an enable for the DIS strobes. When INEN is low, and provided the relevant set up and hold times have been satisfied, data will be loaded with the rising edge of the DIS strobe. If no gaps occur within the incoming data, INEN can be tied permanently low, provided that the sampling rate has been chosen such that transforms are completed before a new block of data is loaded. For transforms of less than 1024 points, data will then be continually processed without any loss of information. In the 1024 point modes the device will cease loading data when 1024 samples have been loaded, and even if INEN remains low no more data will be accepted until the previous results have been dumped.

In a multiple device system an edge is ALWAYS needed to commence a load operation, and Bit 12 has a different purpose. The edge is provided by INEN going low. Loading will cease when a complete block (or group of blocks with multiple concurrent transforms) of data has been loaded, even if INEN remains low. INEN must go high at some point after the minimum hold time has been satisfied, and then return low AFTER ALL DATA HAS BEEN LOADED, before a new load operation can commence. Low going edges which occur before all data has been loaded will be ignored.

The INEN edge mode is actually provided for the correct operation of multiple device systems, but if Bit 12 in the Control Register is reset in the SINGLE device mode, the edge activated operation will still be possible. With all but 256 point complex transforms, the single device edge mode of operation is identical to that of a multiple device system. With 256 point transforms, and their concurrent derivatives, the location of the low going edge in the data stream is dependent on the amount of block overlapping. The low going edge transition must be provided after 64 samples have been loaded with 75% overlapping, and after 128 samples have been loaded with 50% overlapping. With no overlapping the edge must be provided after 256 samples have been loaded.

In a single device system with Bit 12 set, INEN can be taken high to inhibit the load operation when gaps occur in the data stream. In the INEN edge activated mode gaps in the data stream can only be accommodated if the DIS clock is externally inhibited. Taking INEN high will not inhibit the loading of data in this mode.

With gaps in the data stream the peak sampling rates can be higher than continuous sampling rates. When data loading is not coincident with transform operations the peak rate can equal that of the system clock, otherwise it is reduced by the factor, F, given on the opposite page.

When Control Register Bit 12 is set in any multiple device mode, the DEF high going edge will also initiate a load operation after it has been internally synchronised to the rising DIS edge. If the first device in a multiple device system is programmed in this manner, the transform sequence will automatically start when DEF goes in-active. The other devices need the INEN edge as usual, and must have Bit 12 reset. A fuller explanation of the use of Bit 12 in a multiple device mode is given in the section on I/O In Multiple Device Systems. Note that the use of Bit 12 in a single device system (Control Register Bits 10:9 = 00) is completely different to its use in a multiple device mode.

The LFLG output goes active in response to the DIS rising edge used to load the first data sample, and indicates that a load operation is occurring. In an edge activated system the LFLG output will go high as the result of the first high going DIS edge after INEN has gone low. In the simple INEN enabling mode, internal logic counts the number of valid inputs and detects when the programmed block length has been reached. LFLG then goes low and will go high again in response to the next valid DIS strobe. LFLG will go low when DEF is active and will go high in response to the first INEN enabled DIS edge after DEF has gone in- active.

The active going LFLG edge does not normally have any system significance, but in the block overlapping modes the in-active going edge will occur when 50% or 75% of the data has been loaded. By driving the INEN input on one device with the LFLG output from a previous device, this edge can be used to partition data between several devices in a multiple device system. It can also be used to provide an address marker for a user defined input buffer, when executing 1024 point transforms with a single device. It is not needed, however, when the input buffer is provided by the PDSP16540.

DUMPING DATA

Data output is controlled by an output strobe [DOS], a dump enable signal [DEN], and a Data Available signal [DAV]. The DAV signal is used to indicate that the internal output buffer contains transformed data, and the DEN input is used to control the outputting of that data. The output buffer within the device is clocked by the DOS input, and must be primed with a number of DOS strobes (see "user notes - stopping DOS") once a transform is complete in order to transfer data to the output pins. DAV will not go active until this priming has occurred.

The state of the DEN input at the end of a transform is used to control the transition of the active going edge of the DAV output with respect to the DOS strobes. The latter are then used to transfer data from the device to the next system component. If the DEN input is tied low in a single device system, the active going DAV transition will be internally synchronised to the rising edge of a DOS clock. If DEN is not tied low it must be guaranteed to be low at the end of the internal transform operation for this synchronization to occur. Since there is no external indication of this event, the user must take care to only allow DEN to go high whilst DAV is active, if this DAV synchronous mode is needed.

SYNCHRONIZED DAV OPERATION

In the DAV synchronised mode the first rising edge of the DOS clock, after DAV has gone active, must be used to transfer the first transformed sample from the output pins to the next system component. It should be noted that the output buffer will have been primed before the active DAV transition, since DOS must be a continuous clock, and there is then no delay before the first output becomes valid. The DAV output can be used as a clock enable for this next device, and transfers will continue in normal sequential order until the required data has been dumped. DAV will then go inactive in response to the last DOS edge which was used to transfer data to the next device.

This mode of automatically dumping data when it is ready finds applications in real time data flow systems, and detailed timing is given in Table 2. It should be noted that the DOS input MUST be continually present before DAV goes active. If this is not the case the DAV output will not go active at the correct time, and the internal output circuitry will not be primed. Once DAV is active, however, it is possible for DOS to be irregular, and DEN can be used to inhibit the action of the output strobe as discussed previously. For the correct operation of the device the user must ensure that DOS becomes continuous and DEN remains low once DAV goes in-active.

When continuously transforming data such that new outputs are internally available before the previous block has been completely dumped, then DAV would normally stay active and give no indication that one block dump had been finished and another block started. Additional internal circuitry is, however, provided to ensure that DAV goes inactive for one DOS high time, thus supplying an inter block marker.

ASYNCHRONOUS DAV MODE

If DEN is not active in a single device when the transform is complete, then the device will wait for DEN to go active before any data is dumped. This mode is suitable for applications in which output processing is under the control of a remote host, such as a general purpose digital signal processor. The DAV output will then go active as soon as the output buffer is full, and will not be synchronised to the DOS edge. In such systems the DOS strobe may not necessarily be present at this time. Table 3 gives the relevant timing information.

In this host controlled dump mode the PDSP16510 waits for the host to activate the DEN input after DAV has gone active. DEN then functions as an enable for the host produced data strobes on the DOS pin. DEN may either stay active for the complete transfer, or may be used to enable each DOS

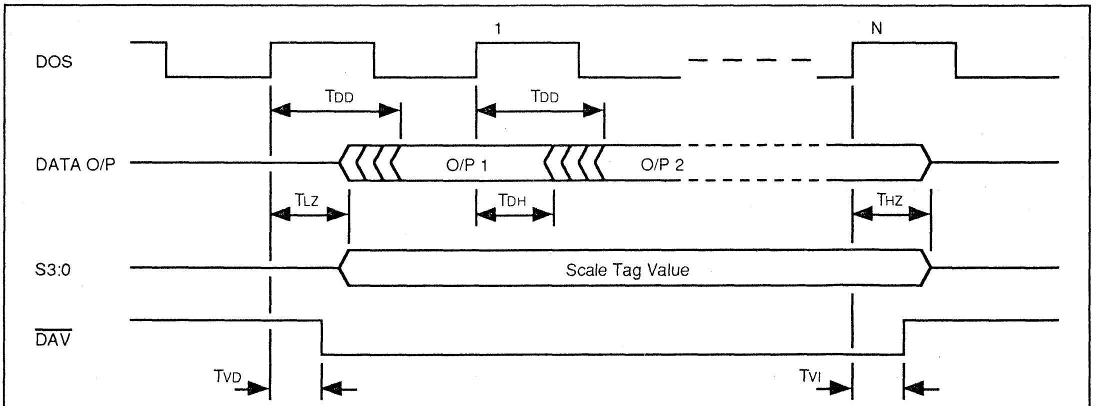

Characteristic	Symbol	16510A,A0,B0,C0 Min	Max		Units
Output Enable Time	T_{LZ}		15		ns
Output Disable Time	T_{HZ}		15		ns
Data Delay Time (30 pf load)	T_{DD}		15		ns
Data Hold Time	T_{DH}	2			ns
$\overline{\text{DAV}}$ active Delay Time (30 pf load)	T_{VD}	1	10		ns
$\overline{\text{DAV}}$ in active Delay Time (30 pf load)	T_{VI}	1	10		ns

Table 2. Output Timing with DEN tied low. (Advanced Data)

input. When DEN and DOS are both active an internal read operation occurs, and an address generator is incremented. DAV goes in-active in response to the DOS edge needed to read the last output, unless Bit 15 in the Control Register is set. In this case DAV goes in-active when the next INEN edge is received for reasons given later.

In host controlled systems the time to dump data could be longer than the transform time. The dump time in such a system will dictate the maximum sampling rate that can be used without the loss of incoming data. In the 1024 point mode, when the loss of data is not important, the PDSP16510 is designed to not accept new data until the previous results have been dumped. Such a system needs no input buffer, and INEN can be permanently tied low if the edge activated mode is not in use. If the loss of data is to be avoided an input buffer is needed and the host must have received all the results before a new block of data has been loaded into the buffer.

For 256 point transforms, with host controlled dumping, it is still possible to overlap load and dump operations. The maximum dump times, however, must be less than the load times to avoid data corruption. Previously converted outputs will be actually corrupted, rather than inputs simply not being used.

If the loss of incoming data is not important, the device can be forced to do separate load, transform, and then dump operations. The corruption of results will then never occur, no matter what dump time is taken. This can be achieved by ensuring that INEN is not active between loading a block of data and completing the dump of the results from that data. The same ends can be achieved if the INEN edge activated mode (Bit 12 reset) is used, and the inverted DAV edge is used to drive the INEN input. This then initializes a new load operation only when the previous dump has been completed.

Results are transferred from the device with the rising edge of the DOS strobe when DEN is active. This is consistent with using the device in a data flow architecture, as is commonly employed in data processing systems. In a typical microprocessor based system, however, data is normally expected to become valid before the end of the data strobe produced by the processor. It is thus necessary for the user to provide a 'dummy' data strobe in order to transfer data to the outputs which can then be read by the host during the next data strobe. In addition further ' dummy ' strobes are needed each time DAV goes active in order to prime the output circuitry. The actual output sequence is given in Table 3 for a single device systemand is described more fully in "user notes - stopping DOS".

GENERAL DUMP CONSIDERATIONS

The tri-state drivers on the output buses are only enabled when both DAV and DEN are active. When DEN is tied permanently low the output bus will start to become valid from the DOS edge which also generates the DAV output. The next DOS edge can then be used to transfer the first output to the next device. When DEN is driven low in response to the DAV output, the outputs start to become valid when DEN goes low. The Scale Tag outputs become valid at the same time as data, and when enabled will continue to indicate the correct value until all frequency bins have been dumped. If at any time during the dump operation DEN goes in- active, then both the

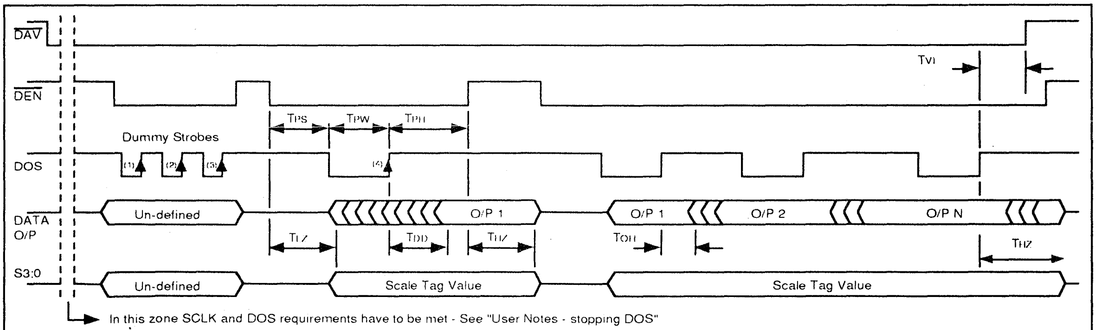

Characteristic	Symbol	16510A,A0,B0,C0 Min	Max		Units
DEN Set Up Time	T_{PS}	10			ns
Host Strobe Width	T_{PW}	10			ns
DEN Hold Time	T_{PH}	5			ns
DAV in-active going Delay (30 pf load)	T_{VI}		10		ns
Output Enable Time (see Fig 13)	T_{LZ}		10		ns
Output Data Delay Time (30 pf load)	T_{DD}		15		ns
Output Disable Time (see Fig 13)	T_{HZ}		10		ns
Read Cycle Time	T_{RC}	25			ns
Old Data Hold Time	T_{OH}	2			ns

Table 3. Host Controlled Output Timing. (Advanced Data)

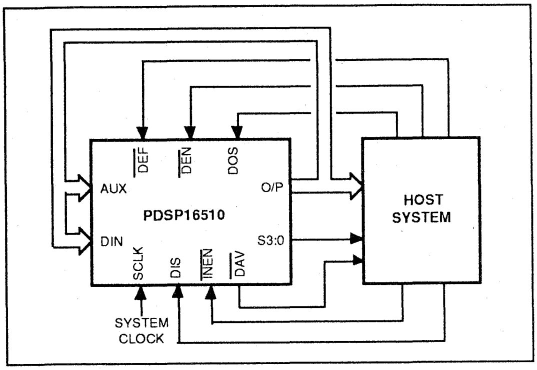

Fig. 7. Host Controlled System

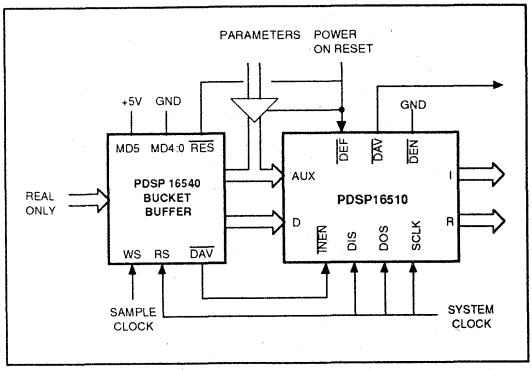

Figure 8. 1024 Point Real Transforms

data and scale tag outputs will go high impedance after the delay shown in Table 3.

Valid transformed data is actually available within the device from DAV going active until INEN again goes active, and a new set of data is loaded. The output tristate drivers, however, normally go high impedance when DAV goes inactive once a dump operation has been completed. In order to support systems in which it may be necessary to read the transformed data more than once, a Control Register Bit is provided which keeps the DAV output active until a further INEN edge is received. The user must then keep track of how many outputs have been dumped before INEN is generated to start a new load operation.

The DAV output can be delayed by an amount equivalent to the pipeline delay through the PDSP16330. This option is invoked by setting a control bit, and allows DAV to indicate that polar data is available at the output of the PDSP16330. When the option is used the tri-state outputs will be enabled when data is actually available and DEN is active, and not when DAV eventually goes active.

Two Control Register Bits allow a range of dump size options to be supported. In some applications the results of interest may only lie in the lower 25 or 50% of the frequency bins, the sampling rate having been chosen to prevent aliasing, and the transform size having been selected to give the required frequency resolution. In other systems it is only necessary to output the second half of a given sized transform. This is useful when filtering is to be performed in the frequency domain using Overlap /Discard Fast Convolutions. With this method FIR filters with N taps can be implemented in the frequency domain using 50% overlapped transforms on 2N samples. After multiplication in the frequency domain with the required frequency response, the inverse transform is performed and the first half of each output is discarded. Since only half the results are dumped, the dump clock need not be twice the rate of the clock used to load data.

FULL CO - PROCESSOR OPERATION

A single device can be configured as a co-processor to a host system in which both the loading and dumping of data is under the control of the host. Such a system is shown in Figure 7, in which DEN is a host provided enable for host read operations, and INEN is an enable for host write operations. DIS and DOS are host data strobes.

The host loads a block of data into the PDSP16510, using DIS enabled by INEN, which is then automatically transformed. The DAV output provides a flag indicating that the transform is complete, and results are then read by the host using DOS enabled by DEN. A new set of inputs is not normally loaded until the previous results are complete. If, however, 1024 point transforms are not to be performed, loading new data could coincide with dumping previous results. This, however, would require a host system with separate input and output buses, and which also allowed coincident transfers. As discussed previously, transferring results must take no longer than loading new data to prevent corruption of the outputs.

In the system illustrated by Figure 7, the host also controls the mode of operation of the FFT processor. The DEF signal is produced from an address decode, and the control parameters are loaded from the host bus by connecting the AUX inputs to the data outputs.

REAL ONLY TRANSFORMS WITH A SINGLE DEVICE

In the simplest case real transforms can, of course, be computed by forcing zero levels on the imaginary input pins. The device can, however, be configured to internally perform two simultaneous real transforms instead of a single complex transform. The block floating point logic will then use data from both blocks when it determines the number of shifts to be applied. This dual transform technique is used to increase the maximum permissible sampling rates, but since an additional data pass is required in order to un-scramble the transformed data, the actual performance is not quite double that possible with a complex transform of the same size. The 4 x 64 point complex mode becomes an 8 x 64 real mode, but the change from 16 x 16 complex transforms to 32 x 16 real transforms is not supported.

When a real transform is performed the algorithm produces complex results for each of the incoming data blocks, but each result only represents the first half of the frequency domain data. This does not cause any loss of information since the two halves are mirror images of each other. As with complex transforms, it is necessary for a different system configuration to be used when 1024 point transforms are required. These are considered later, and the following only applies to 256 or 64 point transforms.

given previously.

The time taken to dump the transformed data must be no more than the load time, if continuous inputs are to be supported and I/O operations are concurrent with transforms. With block overlapping the dump time must be reduced to the time taken to load the partial block. This dump time must include four extra DOS strobes needed to prime the output circuitry when a transform is complete. These, in effect, can be added to the transform time such that with concurrent I/O and 0%, 50%, or 75% overlapping;

nS or (nS)/2 or (nS)/4 must be gtr than or equal to PK + 4W

where n is the transform size, S is the input DIS period, P is the number of clock periods given in Table 4, K is the system clock period, and W is the DOS period which can be less than S if necessary. Note also that S must be synchronous to SCLK, and if an asynchronous ratio is required then a pdsp16540 input buffer should be used.

When DIS and DOS are produced from a common source the minimum allowable sampling period must be increased to allow for the extra dumping time. Thus when DIS and DOS have equal periods and, for example, there is no overlapping;

(n - 4)S must be greater than or equal to PK

The maximum sampling rates given in Table 5 allow for the extra dumping time.

The load and dump operations are not concurrent with transforms in the 1024 point modes, and an external input buffer will be needed if loss of incoming data is to be avoided.

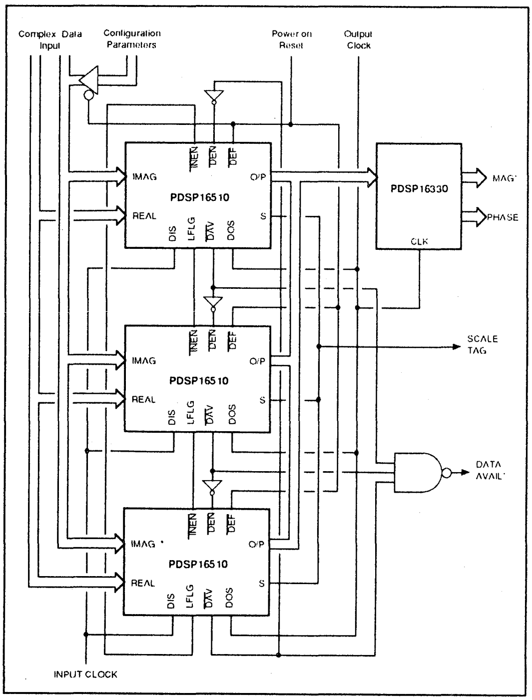

Figure 9. Multiple Device Configuration

This is loaded at the sampling rate and then data is transferred to the PDSP16510 at a user defined rate. The time taken to load this external buffer must be at least equal to the sum of the time to transfer data in and out of the FFT processor and the transform time itself. When data blocks are overlapped by 50% or 75%, no more than one half or one quarter of the block, respectively, must have been loaded in the same time. In the 1024 point modes the dump time can be any user defined value, and need not be increased to allow for block overlapping. The dump time, however , does directly effect the maximum sampling rates which can be accommodated without loss of incoming data.

The maximum sampling rates for 1024 point transforms at any load and dump rate can be calculated from the following relationship:

1024S or 512S or 256S > 1024B + PK + D

for 0%, 50%, or 75% overlapping respectively. S, P, and K were defined opposite. B is the clock period in which data is read from the input buffer and loaded into the device, D is the total dump time allowing for the four extra DOS periods. The periods of the load and dump clocks cannot be less than the system clock period. The maximum sampling rates given in Table 5 assume that a 40 MHz I/O rate is used, and that all results are dumped.

MULTIPLE DEVICE SYSTEMS

In real time applications several devices may be used in parallel in order to increase the sampling rate, but not to increase the transform size. When all outputs are commoned together, and feed a single output processor, then the data dump time must always be less than or equal to the time taken to load the data block (or 50% or 25% of the time with block overlapping). In most configurations with block overlapping the dump rate requirements will limit the maximum input rate, if only one output processor is provided. This can be avoided if the system provides separate output processors for every device. The system clock used for internal calculations then ultimately imposes a limit on the maximum sampling rate possible.

A multiple device system performing complex transforms with a single output processor is shown in Figure 9. The INEN/LFLG signals are used to co-ordinate the segmentation of data between devices. The in-active going edge of LFLG instigates the load procedure in the next device, and, since this edge can be programmed to occur either 25%, 50%, or 100% through the load operation, it can cause the next device to commence loading before the previous one has finished. In this manner data block overlapping is achieved. When multiple concurrent transforms are performed (for example 4 x 64 or 8 x 64) two LFLG transitions are sometimes needed to support block overlapping. This is fully explained in the section on Mode 1 sampling rates.

In any of the multiple device modes an INEN edge transition is needed to start a new load procedure when the previous one has finished. When the LFLG output from the last device is fed back to the INEN input of the first device, continuous transforms will be executed. This continuous sequence can be started by the rising edge of DEF if Control Register Bit 12 is set in the first device (see section on Loading

Data). This bit must not be set in the other devices. Since all devices are supplied from a common input bus and have a common source of control parameters, this Bit 12 inversion is best mechanized with an Exclusive OR gate in the AUX12 input line of the first device. The input can then be inverted when DEF is active but otherwise not be effected. Once the first device has been started with the DEF edge, the sequence will continue automatically using the LFLG /INEN connection between devices.

In many applications data is transformed continuously after power on, and the concept of a first data sample does not exist. If, however , the opposite is true, the first data sample must be present on the input pins such that it can be loaded with the second rising DIS edge after DEF has gone in-active. The data must meet the set up and hold times given in Table 1, and DEF itself must meet the parameters normally met by the INEN rising edge. The latter requirement is necessary to avoid a possible one DIS cycle variance, due the internal DEF synchronization logic. If the position of the first data sample is not important, it is not necessary for DEF to have any set up specification.

Without the feedback from the last device, the first device would wait for another externally supplied initialising pulse. In such a system with N devices in parallel, then N continuous transforms must be executed before the first device can wait for a new INEN input.

When only one output processor is provided the data outputs from all devices are connected together, and internal logic will enable the tri-state outputs when a device is ready to output data i.e. DAV goes active. When data blocks are overlapped it is possible that the output rate requirements will limit the input sampling rate (see section on Multiple Device Sampling Rates). Additional output processors will remove this restriction, and the correct choice of multiple device operating mode will optimise the sampling rates that can be achieved with a given number of devices.

The synchronisation intervals, necessary to co-ordinate input and output operations with the transform operation, lead, in effect, to some uncertainty in the time needed to complete a transform. Thus a particular device in a multiple device system can effectively complete a transform in less system clock periods than another device in the same system. To prevent one device turning on its output bus before the previous one has finished, it is either necessary to use a faster output rate than would otherwise be required, or to use the inverted DAV output from one device to drive the DEN input of the next. The latter option allows DIS and DOS to be connected together, and ensures that the second device will not output data until the first device has finished.

This method of driving the DEN input from the inverted DAV output from a previous device requires a change to the single device DAV and DEN operation. If DEN is active at the end of a transform in a multiple device system, the DAV output will go active when the output circuit has been primed by the DOS strobes. This operation is identical to that provided for a single device system, and is transparent to the user as long as DEN and DOS are active . If DEN is not active, however, the DAV output will not asynchronously go active as happens in a single device system. Instead DAV will only go active when DEN eventually goes active. Since DEN is the inverted DAV output from a previous device, it is thus never possible for two devices to be actively outputtng data. The DAV active going edge remains synchronised to the DOS strobe since the DEN input will only go active when a previous DAV goes in-active. A further change to the output circuitry ensures that the output buffer is primed even though DEN is not active. The first word, however, only progresses as far as the final output latch. The output bus is not enabled, and address increments do not

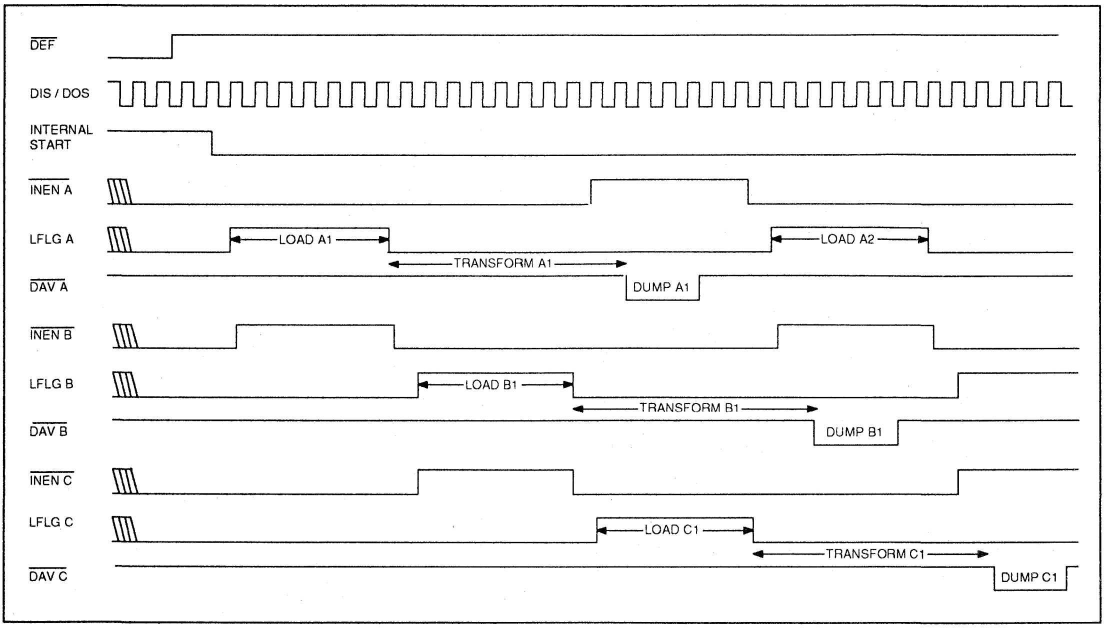

Figure 10. Three Device System with Separate Load, Transform, and Dump Operations

occur, until DEN is finally received. This modification to the internal control logic ensures that the output buffer does not impose unnecessary gaps between consecutive transforms. These gaps would, in turn, force the required DOS frequency to be greater than the DIS frequency (or greater than twice or four times the frequency with 50% and 75% overlaps).

The system illustrated by Figure 9 produces a common DAV output by OR'ing together all the individual, active low, DAV outputs. This is not guaranteed to give an indication when one transform has finished, and the next one has started, since it may simply glitch as one DAV goes in-active and the next one goes active after some delay. This glitch will not cause system problems since it occurs at a point clear of the high going edge of the DOS strobe. To provide a marker for the end of a transform each in-active going DAV edge should set its own latch, which is then reset by a subsequent DOS edge. The output of the latches can then be OR'd together if necessary.

Three multiple device operating modes are actually provided, and are selected with Control Register Bits 10:9. The choice of a particular mode is application dependent, and will effect the maximum sampling rate achievable with a given number of devices.

MULTIPLE DEVICE SAMPLING RATES

MODE 1. (BITS 10:9 = 01)

In this mode transfers in and out of the device are concurrent with transform operations. This mode must not be used for 1024 point transforms due to internal memory size restrictions. When real transforms are performed in this mode, only the real data input is used, regardless of the amount of block overlapping.

The increase in performance is directly related to the number of devices provided, but the input and output rates are limited to FØ where F and Ø are as defined previously. Within this restriction the theoretical performance is given by;

$$NnS > PK+4W, \quad \text{or } 0.5NnS > PK+4W, \quad \text{or } 0.25NnS > PK+4W$$

for 0%, 50%, or 75% overlapping. N is the number of devices, n is the transform size, S is the DIS strobe period, P is the number of system clock periods given in Table 4, K is the system clock period, and W is the DOS strobe period. Note that DIS should be synchronous to SCLK, and also that DOS should be synchronous to SCLK.

If an output processor is provided for every device, two devices with 50% block overlapping or four devices with 75% block overlapping will give the same sampling rates as a single device with no overlapping. If only one output processor is provided, the two or four times increase needed in the output rate over the input rate, usually imposes a limit on the input rate, since the output rate is limited to a factor, F, of SCLK.

In this operating mode the DIS and DOS strobes can often be tied together, since a faster DOS strobe gives no improvement in the sampling rates possible. This remains true even when the output rate must be twice or four times the input rate due to block overlapping. Options can then be used which internally divide the DIS strobe by two or four, and thus allow the input to be driven by the faster DOS strobe.

In this mode the LFLG goes in-active after 25%, 50%, or 100% of the block has been loaded. When multiple transforms are performed concurrently (for example 4 x 64) a LFLG transition occurs at the relevant point whilst the first block in the group is being loaded. LFLG then goes high again and returns low at the overlap point in the last block. This double LFLG transition allows two devices to support 50% block overlapping, since the first transition from the first device can be used to initiate the load procedure in the second device. The second transition from the second device then initiates a new load procedure in the first device. The additional edges from each device have no effect since they occur when the device they are driving is already doing a load operation.

In such a two device system supporting 50% overlaps the inverted DAV from the first device must drive the DEN input of the second device. The data dumping time is then shared equally between both devices. The second device only outputs data when the first has finished, but both dumps must be finished in the time taken to load the group of blocks if only one output processor is provided. Without the DAV/DEN connection one device would only have had the time needed to load half of one sub block in which to dump its data.

In a similar manner four devices will handle 75% overlaps when concurrent multiple transforms are to be computed. The second, third, and fourth devices make use of the first transition, and ignore the second. The first device uses the second transition from the last device, and ignores the first. With the DAV/DEN connection each device will have one quarter of the load time to dump its data when a single output processor is provided .

More than two devices will provide increased performance for multiple transforms with 50% overlapping, and more than four devices will increase the performance with 75% overlapping. External logic is then needed to ensure that each device only uses the correct LFLG transition. Any device should only use the negative LFLG transition from a previous device if its own LFLG is low, and the LFLG output from the previous device plus one is low.

MODE 2 (BITS 10:9 = 10)

This mode is suitable for all transform sizes, since separate load, transform, and then dump operations occur. More devices than required by Mode 1 are necessary to achieve a given sampling rate, but the input and output rates can be any value up to the full system clock rate with the A grade part. As with Mode 1, additional output processors are needed to avoid the sampling rate restriction imposed by block overlapping.

The number of devices, N, needed to achieve a given sample rate can be derived from the following formula:

$$NnS > nS + PK + D \quad \text{for no overlapping}$$
$$NnS > 2 \times [nS + PK + D] \quad \text{for 50\% overlapping}$$
$$NnS > 4 \times [nS + PK + D] \quad \text{for 75\% overlapping}$$

N is the number of devices, n is the transform size, S is the DIS strobe period, P is the number of system clock periods given in Table 4, K is the system clock period, and D is the total dump time including 4 extra DOS periods as discussed previously. The DIS and DOS periods are any value defined by the user, down to the system clock period with the A grade part. Note that DIS should be synchronous to SCLK, and also DOS should be synchronous to SCLK.

In this mode increasing the output clock frequency will allow a greater continuous input rate. The provision of separate DIS and DOS pins allows this to be mechanized, and the DOS frequency can be increased to that of the system clock used internally. When the sum of the dump time (including four extra DOS periods for output priming) plus 12 system clock periods (the transform time variation caused by input synchronization) is less than the load time, one device will be guaranteed to have finished dumping before the next one starts. The inverted DAV to DEN connection between devices is then not needed, and all DEN inputs can be grounded.

The LFLG transitions occur at the same times as Mode 1, except that the double transition does not occur with multiple concurrent transforms. Fig. 10 illustrates a timing sequence with three devices. Real transforms still only use the real inputs regardless of the amount of block overlapping.

MODE 3 (BITS 10:9 = 11)

Multiple device Mode 3 is provided in order to improve the performance when block overlapping is needed, and separate output processors are provided. In this mode transfers in and out of the device are never concurrent with transform operations. The device will actually load extra data such that the required data to perform two overlapped transforms is stored internally. The amount of internal RAM prohibits the use of this mode when performing overlapped 1024 point transforms.
LFLG will go in-active after a normal data block have been loaded, regardless of the overlap selected. The device, however, continues to load more data. Thus, for example, in the 4 x 64 mode, five 64 point blocks will be loaded. This technique allows each device in the system to complete two or four overlapped transforms (depending on the amount of overlap) before any new data is needed. When doing a straightforward 256 point transform the device will load 256 + 128 data points.

The full benefits are only obtained if more than one output processor is provided, but an extra processor is not always necessary for every device. Sampling rates up to the system clock rate are possible. The equations defining the sampling rates become:

$$(N - 1)L > 2PK + 2D \text{ for 50\% overlaps}$$
$$(N - 1)L > 4PK + 4D \text{ for 75\% overlaps}$$

where L is the time needed to load a normal block of data but not including the extra data, P is the number of system clock periods given in Table 4, K is the system clock period, and D is the total dump time including 4 extra DOS periods. As before, both DIS and DOS must be synchronous to SCLK.

When real transforms are to be performed on single sourced data, an external FIFO is needed to provide pairs of data blocks. These are loaded simultaneously into the real and imaginary inputs. See the section on real transforms.

OPERATING MODES

The operating mode of the PDSP16510 is determined by the condition of 16 bits in an internal Control Register. The status of these bits is defined by the inputs present on the AUX15:0 pins when the DEF input is active. The DEF input can be a simple power on reset if the operating mode is fixed once power is supplied. The AUX pins are also used to provide the imaginary component of the complex input data. Thus, if complex inputs are needed, the mode definition must be implemented through a tri-state buffer which is only enabled when DEF is active. The imaginary input data must be disabled during this time.

Table 6 lists the functionality of each of the bits in the mode control register, and further explanations are as follows:-

BITS 2:0

These bits define one of 7 options for the sample size and type of data. In the 1024 point options the device will assume the non concurrent operating mode, regardless of whether a single or multiple device system is specified. The internal control logic will then ensure that data is loaded, transformed, and dumped in sequential operations.

For other data set sizes, loading, transforming, and dumping, can all occur simultaneously with a single device; the actual overlap will be dependent on the relative occurrences of the INEN input. Only in Mode 1 can concurrent operations be done with multiple devices.

BIT 3

This bit determines the number of right shifts built into the data path. In either condition only two right shifts occur during the first pass. If the bit is reset, three shifts occur in subsequent passes and the block floating point scheme allows up to fifteen compensating left shifts. If it is set, two shifts occur in every pass and overflow is possible. This is indicated by reducing the number of compensating left shifts to fourteen, and using scale tag value fifteen to indicate that overflow has occurred.

BITS 5:4

These bits define the choice of window operator. If other windows are needed they must be applied externally. The fourth option is used to specify the inverse transform, which does not require the use of a window operator. When 16 x 16 complex transforms are specified by Bits 2:0, only the rectangular window can be used. The use of any of the other options will cause the device to enter an internal test mode.

BITS 8:6

These bits define 0%, 50%, or 75% data block overlapping, and the division factor on the DIS input. Overlapping must not be specified with 16 x 16 complex transforms.
Two decodes allow the DIS input to be divided by two or four, when 50% and 75% overlapping is respectively needed. These options allow the DOS and DIS input pins to be still supplied from a common source, even though the output rate must be faster than the input rate. The frequency of this source would be dictated by the output rate requirement, with the input rate internally reduced by the correct amount.

Special decodes are provided to support real only transforms from dual sources, using both the real and auxiliary inputs. When data is from a single source, and no overlaps are needed, only the real input should be used. If 50% or 75% overlaps are needed from a single source of real data, the device always expects blocks to be simultaneously loaded. An external FIFO is then needed to supply data to the real inputs after a delay of one block. Each block is thus loaded twice, firstly through the Auxiliary inputs and then through the Real inputs.

BIT 10:9

These bits define a single device system, or one of three multiple device possibilities. The choice between the first and second multiple device mode is dependent on the transform size and the sampling rate needed. The third mode should only be used when overlapped multiple transforms with less than 1024 points are to be performed simultaneously. It changes the LFLG logic and allows sampling rates up to the system clock rate to be achieved with multiple output processors.

BIT 11

BITS	Dec'	OPTION
2:0	000	16 x 16 COMPLEX
	001	4 x 64 COMPLEX
	010	256 COMPLEX
	011	1024 COMPLEX
	100	8 X 64 REAL
	101	2 X 256 REAL
	110	2 X 1024 REAL
	111	NOT USED
3	0	SHIFT 3 PLACES AFTER PASS1
	1	ALWAYS SHIFT 2 PLACES
5:4	00	RECTANGULAR
	01	HAMMING WINDOW
	10	BLACKMAN-HARRIS
	11	INVERSE TRANSFORM
8:6	000	NO OVERLAP
	001	50% OVERLAP
	010	50% OVERLAP AND DIS ÷ 2
	011	75% OVERLAP
	100	75% OVERLAP AND DIS ÷ 4
	101	DUAL SOURCE, NO OVERLAP
	110	DUAL SOURCE, 50% OVERLAP
	111	DUAL SOURCE, 75% OVERLAP
10:9	00	SINGLE DEVICE
	01	N DEVICES, CONCURRENT I/O
	10	N DEVICES, LOAD-TRANS-DUMP
	11	SPECIAL MULTIPLE TRANSFORM
11	00	DAV NOT DELAYED
	01	24 CLK DAV DELAY
12	0	INEN EDGE ACTIVATED
	1	INEN IS SIMPLE ENABLE
14:13	00	O/P FIRST QUARTER
	01	O/P FIRST HALF
	10	O/P LAST HALF
	11	O/P ALL RESULTS
15	0	NORMAL DAV
	1	KEEP DAV ACTIVE TILL INEN

Table 6. Mode Control Bit Allocations

When this bit is set the PDSP16510 will not generate DAV until 24 DOS clocks after data was actually valid. In this case the output tri-state drivers will be enabled at the correct time, even though the DAV signal was not externally valid. Host controlled dumping should not be used.

BIT 12

When this bit is set in the single device mode, the INEN input is a simple load enable signal. When it is reset an INEN edge is needed at the end of a load sequence before a new one can commence.

When it is reset in a multiple device mode it has no action, but when it is set it will cause the DEF high going edge to also initiate a load operation.

BIT 14:13

These bits allow four dump size options to be provided. Individual frequency bins are not accessible.

BIT 15

Under normal circumstances DAV would be expected to go invalid when a transform has been dumped. In some applications, however, it may be necessary to read the outputs more than once. When this bit is set, DAV will remain valid until the next INEN input, and will indicate that the transformed data still remains in the internal buffer. As soon as the next INEN is received the transformed data will be overwritten. Whilst DAV remains active the output tri-states will be enabled.

WINDOW OPERATORS

Since only a finite segment of a signal can be observed and processed at any one time, it is impossible to obtain pure spectral lines. Discontinuities are introduced at the boundaries of the observation interval which lead to spectral leakage. Windows are weighting functions applied to the data in order to reduce these discontinuities at the boundaries.

In the time domain the signal has to be observed through a finite window as a matter of accord. This is in fact equivalent to multiplying the signal with a set of uniform weights i.e. a rectangular window operator. In the frequency domain the spectrum of the data will be the spectrum of this weighting function shifted to the sinusoidal frequencies of the components in the data.

The rectangular window has a Fourier Transform which is a SINC(X) function. This has sidelobes which are only 13dB down from the main lobe. This severely limits the dynamic range of the system since a second sinusoid in close proximity would have its main lobe swamped by this side lobe. This would occur if its amplitude was a mere 13dB down from the first sinusoid.

Window operators are thus mathematically constructed to cancel these sidelobes as far as possible. Unfortunately this is normally done at the expense of making the main lobe spread over more frequency bins. This reduces the ability of the system to resolve two frequencies, and can only be overcome by using more data samples. This may not always be possible because of other system constraints.

A common rule of thumb defines the resolution of an FFT system as half the full width of the mainlobe. The width of the mainlobe for a rectangular window is two frequency bins; for the Hamming window it is four bins; for the Blackman-Harris

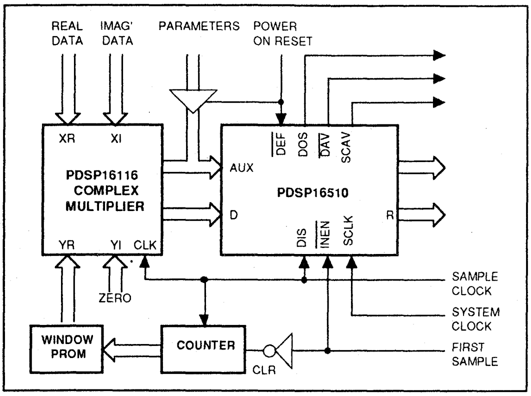

Fig. 11. External Window Generator

window it is six bins.

The latter two windows are actually supported by the PDSP16510. These are constructed on the fly as needed, and take the general form:

$$A - B\cos x + C\cos 2x \quad \text{where } x = (2\pi n)/N,\ n = 0 \text{ to } N-1$$

For Hamming, A = 0.54, B = 0.46, C = 0

For Blackman-Harris, A = 0.42323, B = 0.49755,C=0.07922

These windows can be applied to any of the transform size options, except the 16 x 16 complex variant. When the latter is specified the rectangular window option MUST be selected, or the device will be configured in an internal test mode.

If other operators are required these must be applied externally. This can be conveniently achieved with either a PDSP16112 or a PDSP16116, both of which are complex multipliers but with different accuracies. Fig. 11 shows how either one can be configured to perform two separate multiplications with one input common to both. This arrangement is necessary to perform the window function on complex inputs.

Important features of the windows generated by PDSP16510, and other commonly used windows, are illustrated in Table 7. The results are obtained from the reference quoted, which should be consulted for a full mathematical treatment. The significance of each parameter is outlined below :

Highest Side Lobe Level

The inherent rectangular window has sidelobes which are only 13dB down from the mainlobe. These severely limit the dynamic range. The object of the window is to improve this situation with better side load attenuation.

Mid-Point Loss

In line with the filter concept it is possible to conceive of an additional processing loss for a tone of frequency mid-way between two bins. This is defined as the ratio of the coherent gains of two tones, one at the mid-point and one at the sample point. It is expressed in dB in Table 8.

Overall loss

An overall figure for the reduction in signal to noise ratio can be obtained by adding the mid-point loss to the reciprocal of the equivalent noise power bandwidth in dB. It is a measure of the ability of the window to detect single tones in broadband noise. The variance between windows is less than 1dB.

6.0dB Bandwidth

This figure, expressed in bin widths, represents the ability of the window to resolve two tones and should be as close to unity as possible. As the highest sidelobe level is reduced, this parameter tends to get worse, and a compromise must be used when choosing a window.

Overlap Correlation

In many practical systems the squared magnitudes of successive transforms are averaged to reduce the variance of the measurements. If, however, a windowed FFT is applied to non overlapping partitions of the sequence, data near the boundaries will be ignored since the window exhibits small values at those points. To avoid this loss partitions are usually overlapped by 50% or 75%, which might, at first sight, remove the need to average successive transforms. If non-windowed

Window Operator	Highest Side Lobe	Mid-Point Loss dB	Overall Loss dB	6dB Bandwidth	Overlap Correlation	
					75%	50%
Rectangular	-13	3.92	3.92	1.21	75	50
Hamming	-43	1.78	3.1	1.81	70.7	23.5
Dolph-Chebyshev [C = 3.5]	-70	1.25	3.35	2.17	60.2	11.9
Kaiser-Bessel [C = 3]	-69	1.02	3.55	2.39	53.9	7.4
Blackman	-58	1.1	3.47	2.35	56.7	9
Blackman-Harris [3 term]	-67	1.13	3.45	1.81	57.2	9.6

Table 7. Window Performance (from The use of Windows for Harmonic Analysis. F J Harris. Proc IEEE Vol 66. Jan 1978)

Arithmetic Accuracy	Max Tone WRT Noise	Slot Noise Test	2 Tones with Freq Spread
16 bit,unconditional scaling	60	44	45
24 bit arithmetic with unconditional scaling, 16 bit inputs	88	67	65
16 bit inputs with PDSP16510 block FP	74	61	63
Full 32 bit Floating point with 16 bit inputs	93	82	67

Table 8. Comparative Dynamic Range Measurements

transforms are overlapped by 75% or 50%, then 75% or 50% of the data will be correlated. When windows are applied, however, the data common to both transforms will be operated upon by different portions of the window waveform. The difference in these portions will dictate the amount of correlation between overlapped data. At 50% overlap Table 7 shows that with all windows the data is virtually independent, and successive averaging would still be needed. At 75% overlap figures are obtained which are closer to the 75% correlation obtained with no window.

Examination of Table 7 shows that the Blackman-Harris window gives performance very similar to that of the Kaiser-Bessel and Dolph-Chebyshev windows. The latter two windows can not be computed as they are needed since they are mathematically too complicated. The values are normally pre-computed and stored in a ROM; this would need to contain 1M bits to match the accuracy of the rest of the system.

Use of the Hamming window gives worse dynamic range than the more complex windows, but it has less effect on the overlap correlation and it has a smaller main lobe width.

SPECTRAL PERFORMANCE

There are two important parameters in the measurement of spectral response: resolution and dynamic range. Resolution defines how closely two sinusoids can be spaced in frequency and still be identified; dynamic range defines how great the difference in the amplitudes of the sinusoids may be and yet the smaller one still identified. Resolution is determined by the observation time [i.e. the width of the frequency bin] and the window operator that is used. Dynamic range is also determined by the window operator, but in a hardware implementation it is also influenced by the number of bits used to represent the data throughout the calculation.

The hardware effects include the accuracy of the A/D converter, the number of bits representing the window operator and the twiddle factors, and the way the growth in word length is handled as the FFT calculation proceeds. The obvious way to overcome these limitations is to use floating point arithmetic; but in real life the accuracy of the A/D converter is fixed and the sample size is limited. Floating point arithmetic is thus an overkill solution for the majority of applications. This is especially true for transform sizes up to 1024 points, which is the intended application area.

Figures given for the dynamic range of a system must be carefully interpreted, since there is no exact definition of the measurement. Three different ways of measuring dynamic range have been investigated using 1024 point transforms.

The 'best' dynamic range figures will be obtained with single tone measurements, and these results are often quoted to indicate the need for greater bit accuracies. The measure is the ratio of a full scale sinusoid to the average noise level and the results will be essentially independent of the window operator. The results given by the PDSP16510 are compared to various other configurations in the first column of Table 8. With this method the dynamic range is bound to improve as more bits are used to represent the data. Theoretically 6 dB of dynamic range will be obtained for every bit representing the input data, if the internal arithmetic accuracy gives no degradation in performance. In practice this improvement has no significance since the incoming waveforms will be much more complex than a single sinusoid.

An alternative method of determining dynamic range is with a slot noise test. White noise is passed through a narrow-band notch filter, several frequency bins wide, and the FFT computed. There is no noise in the filtered slot at the input to the FFT, but there is noise in the frequency bins corresponding to the width of the notch. Dynamic range is measured as the difference in dB of the average signal power and the average noise power and can be considered to give more useful results. Comparative results from various configurations are also given in the second column of Table 8. The performance with 24 bit data is seen to be little better than that obtained with the PDSP16510. This can be attributed to the scaling scheme, word growth, and rounding method used within the device.

When two nearby tones are to be capable of detection, the window operator will dictate the performance of the system. The final column in Table 8 illustrates the results obtained using two sinusoids of different amplitudes, with the larger one residing mid-way between two frequency bins, and the smaller 5.5 bins away. The two frequencies are five bins apart to avoid the effects of the mainlobe widths. The dB figures given are the difference in amplitude between the two signals when the smaller one is still just detectable as a separate peak from the larger one.

This technique illustrates the performance of the window, since the amount by which sidelobe structure of the larger signal swamps the mainlobe of the smaller signal will determine if the smaller is detected. The theoretical attenuation of the highest sidelobe levels, with respect to the mainlobe, for the window options provided by the PDSP16510 have been given in Table 7, and represent the dynamic range that can be obtained if arithmetic effects are ignored. The results in the final column in Table 8 are the practical results given by the device, and as with the slot noise test indicate that the arithmetic scheme used by the PDSP16510 is equivalent to using 24 bit data. The Blackman Harris window was used in all cases.

USER NOTES - STOPPING DOS

(1) GENERAL DESCRIPTION

The transform is calculated internally fully synchronous to SCLK. However, as all outputs are referenced to DOS, a transfer has to be made between the two clocks. In addition, some dummy DOS strobes are needed to operate the internal control logic, and to advance data from the internal RAMs to the output pins.

The most simple configuration for the device is to have DOS running continuously and for DEN to be permanently active. When this happens the user will just be aware of data appearing on the output pins on the same DOS cycle when DAV goes active. However, there are many situations where either DOS is not continuously running, or DEN is not permanently active. To help explain how to operate the device in these situations, the internal operation of the output circuits must be described. For those who are not going to be interrupting DOS, the remainder of this section can be ignored.

(2) INTERNAL RAM - GENERAL DESCRIPTION.

For single device operation of transforms less than 1024 points, the internal RAM is shared between three separate operations which enable the device to output old transformed results, calculate the current transform, and input new data ready for the next transform. All these operations, along with the internal control logic, are controlled by a 12-cycle state machine. The RAM operations are :

(a) 2 cycles in every 12 are dedicated to reading new information in the input buffer and writing it to the RAM

(b) 2 cycles in every 12 are dedicated to reading the contents of the RAM and advancing that data to the output buffer.

(c) 8 cycles in every 12 are dedicated to the read and write operations of the transform currently being calculated.

(3) SEQUENCE OF EVENTS

The sequence of events relating to the output control and data flow is as follows :

(3.1) An SCLK rising edge :

(a) An internal flag is raised to indicate that the transform has finished and data is available to be dunped. Data will be present in the internal RAM, and the output address generator will be at the cortrect address. Access to the RAM at this moment, however, has not been made.

(b) If at this moment the device is programmed to be a single device, and DEN is inactive, then DAV will be made active - ie without the presence of DOS. If DEN is active at this point, or the device is programmed in any multiple device mode, then DAV will remain inactive.

(3.2) Accessing the RAM at this point

At this moment, when DAV has been made active before data appears on the output pins, data is not yet in the output buffer. Internally the precise SCLK cycle at which the RAMs are read and written to the output buffers now has to be waited for. This cycle, as described above ocurrs 2 in every 12 SCLK cycles, so at worst case 6 SCLK cycles have to elapse until data is guaranteed to be in the output buffer.

If the DOS rate is similar to the SCLK rate, and the user has been immediately applying DOS pulses (on seeing DAV go active) hoping to get data off the chip, then this will not actually happen.

The next internal flag raised is the one which indicates that the output data has been successfully read from the RAMs and is now in the output buffer.

(3.3) The next DOS rising edge (regardless of DEN status) :

The flag indicating that the RAMs have been read is transferred to circuitry operating on DOS. The output enable signal, DEN, does not have to be present at this point.

(3.4) The next DEN-Enabled DOS rising edge (ie the 1st one of this sequence)

The output state machine receives it's first edge.

(3.5) The next DEN-Enabled DSO rising edge (ie the 2nd)

Internal output address generators start to count (ready for fetching the next set of output data).

(3.6) The next DEN-Enabled DOS rising edge (ie the 3rd)

An enable signal is raised for the final data latch in the output buffer.

(3.7) The next DEN-Enabled DOS rising edge (ie the 4th)

(a) The final data in the output buffer latch clocks-through new data and presents it to the output pads.

(b) The output pads come oput of high impedance.

(c) If DAV was previously inactive, it is now made active.

(4) OUTPUT SCENARIOS

Considering the above sequence, therefore, some single device situations can now be explained :

(4.1) DOS is continuously present, but DEN is inactive (Transform size less than 1024)

In this case, when the transform is complete, as the device is programmed as a single device and DEN is inactive, DAV will be made active. Even though DOS is running, the status of DAV at this point does not rely on it.

The user can now monitor the status of DAV, and after at least 6 SCLK cycles can initiate some further action, eg by external control force DEN active at some later time when the rest of the system is ready to accept the transformed data. Independently of this external control, the next DOS pulse will start to operate the sequence of events as described above (ie point No. 3.3). When DEN is eventually made active, the remainder of the above sequence (points Nos 3.4 to 3.7) is executed, with 4 DEN-Enabled DOS pulses needed before data is observed on the output pins.

If however the user immediately forces DEN active upon monitoring DAV go active and waiting for the required 6 SCLK cycles, then 5 DOS pulses would have to be issued. The first of these 5 would start the sequence of events as described above (3.3), and the fact that it is enabled by DEN would be irrelevant. The required DEN enabled pulses in this situation would be the 2nd, 3rd, 4th and 5th pulses supplied.

(4.2) DOS is not running, and DEN is inactive. (Transform sizes less than 1024)

In this situation, again as the device is programmed to be a single device and DEN is inactive at the point where the transform is complete, DAV will be made active regardless of the state of DOS. The user can now monitor this event on DAV and after waiting a further 6 SCLK cycles, use it to switch on DOS and to make DEN active.

DOS can now be switched on for at least one pulse (but may be more), and the sequence of events as described earlier (from point No 3.3) will start. DEN can then be made active, whereby a further 4 DEN-Enabled DOS pulses will be required before data is seen on the output pins. This is the situation shown in table 3.

Alternatively, DEN and DOS could be made to operate on the same cycle. In this case data will appear on the output pins on the 5th DOS pulse (the first would not actually require the presence of DEN, but the 2nd, 3rd, 4th and 5th would)

(4.3) 1024 point transforms, single device mode.

In the case of 1024 point transforms, the internal RAM is no longer operated in the manner described in section 2. The RAM is instead totally dedicated to one operation at a time. Thus data for a transform will be loaded, and all 12 out of 12 SCLK cycles will be available for the transfer of input data to the RAMs. During the transfrom no transfers from the input to the RAM or from the RAM to the output are possible. This is why DIS and DOS can be equal to SCLK for 1024 point transforms.

If 1024 point transforms are being performed and the device is programmed as a single device, then "asynchronous" operation of DAV is possible as described earlier for transform sizes less than 1024 points. If DEN is inactive at the time the transform has finished calculating, then DAV will be made to go active regardless of the state of DOS. Although 6 SCLK cycles do not have to be waited for as in section 3.2, a transition has to be made from the transform controlling the internal RAM to the output circuits cntrolling it. This operation plus the time taken to advance data from the RAMs to the output buffer takes exactly 4 SCLK cycles.

Hence the sequence of events is exactly as described in section 3, except that section 3.3 should read 4 SCLK cycles rather than 6. The analysis of sections 4.1 nad 4.2 are also true if the 6 SCLK cycle time is substituted with 4 SCLK cycles.

(5) DUMMY DOS STROBES AFTER DEF

In addition to the dummy DOS strobes needed prior to dumping data, it is necessary to provide at least 4 DOS strobes after DEF has gone inactive, but before DAV goes active. These initialise the internal address counters and do not rely on DEN also being active. They are needed every time DEF has been used to change the operating mode.

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage Vcc	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to Vcc + 0.5V
Output voltage V_{OUT}	-0.5V to Vcc + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Junction Temperature, Commercial	100°C
Junction temperature, Industrial	115°C
Junction Temperature, Military	155°C
Package power dissipation	5000mW

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as positive into the device.

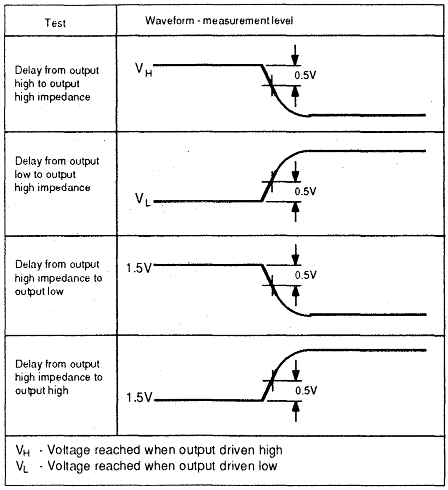

ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise state)

PDSP16510A C0 Tamb = 0 C to + 70°C. Vcc = 5.0v ± 5%
PDSP16510A B0 Tamb = -40 C to + 85°C. Vcc = 5.0v ± 10%
PDSP16510A A0 Tamb = -55 C to +125°C. Vcc = 5.0v ± 10%

Characteristic	Symbol	Value			Units	Notes
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.0		-	V	SCLK, DIS, DOS, DEN need 3V
Input low voltage	V_{IL}	-		0.8	V	DEN needs 0.7V max
Input leakage current	I_{IN}	-10		+10	µA	GND < V_{IN} < V_{CC}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	µA	GND < V_{OUT} < V_{CC}
Output S/C current	I_{SC}	10		300	mA	V_{CC} = Max

SWITCHING CHARACTERISTICS

Characteristic	Symbol	Min	Max	Conditions
Clock Frequency (MHz)	Ø	DC	40	Max Ø high time is 1msec
Clock High Period (ns)	T_{CH}	13		
Clock Low Period (ns)	T_{CL}	10		
Max DOS, DIS Frequency	$Ø_D$		FØ	Less than 1024 points or Mult Dev Mode 1. Note $F = \frac{4}{6 + 0.001ØT_{CL}}$
Max DIS Frequency	$Ø_D$		Ø	1024 points or Mult Dev Modes 2 and 3
Max DOS Frequency	$Ø_D$		Ø	

SCLK to DIS/DOS RELATIONSHIP
Both DIS and DOS must be synchronous to SCLK. Ideally they should both be produced from SCLK, in which case the SCLK rising edge would either be first or coincident with the DIS and DOS rising edges. In any event, the rising edge of SCLK must not fall between 2ns and 10ns after the rising edge of either DIS or DOS

ORDERING INFORMATION

PDSP16510A C0 AC	(Commercial -	PGA Package)
PDSP16510A C0 GC	(Commercial -	Leaded Chip Carrier)
PDSP16510A B0 AC	(Industrial -	PGA Package)
PDSP16510A B0 GC	(Industrial -	Leaded Chip Carrier)
PDSP16510A A0 AC	(Military -	PGA Package)
PDSP16510A A0 GC	(Military -	Leaded Chip Carrier)
PDSP16510A/MA/GCPR	(Military -	Screened Leaded Chip Carrier. See separate datasheet for details)

PDSP16540

32K BUCKET BUFFER

The PDSP16540 Bucket Buffer is for use in systems which require a reservoir in which a block of data is accumulated, whilst previous data is being transferred to other system elements and then processed. It thus prevents the loss of incoming data whilst the previous block is being processed. Like a FIFO all address are generated internally.

It differs from a normal FIFO, however, by allowing the user to define both the length of the data block and also the amount of the old data to be re-read before the new data is added. The latter feature supports the block overlapping requirements of Digital Signal Processing Systems performing Fast Fourier Transforms. It also provides wide, 32 bit, input and output buses, unlike normal byte wide FIFO's. This wide configuration supports the16 bit real and imaginary components of the complex data found in many DSP systems.

In particular, the device can be directly connected to the PDSP16510 FFT Processor without any external logic. The FFT Processor requires the support of an input buffer when 1024 point transforms are to be continuously performed and no incoming data is to remain un-processed.

The number of words, which are read as a complete block, can be programmed in multiples of 32 up to a maximum of 1024. The amount of new data in this block can separately be programmed in multiples of 32 words. In this manner the percentage of new data in a complete block is under the control of the user, and the device is not restricted to only supporting the requirements of the PDSP16510.

A Read Me Flag is raised at a user defined point during the loading of new data. This allows the next system component to prepare itself to accept data. Data is not actually transferred, however, until all the user defined amount of new data has been loaded, and a Data Available Flag goes active. The gap between the two flags can be programmed to provide sufficient time to prepare the device which is to accept data from the buffer. This provide a much more flexible solution than the simple Full Flag offered by a standard FIFO.

ASSOCIATED PRODUCTS

PDSP16510 FFT Processor

PDSP16520 Quad Port Synchronous RAM

PDSP16116 Complex Multiplier

PDSP16318 Complex Accumulator

PDSP16330 Cartesian to Polar Converter

PDSP16340 Polar to Cartesian Converter

FEATURES

- 1K x 32 bit dual port RAM for use as a reservoir in data flow systems
- Up to 40 MHz read rates and 16 MHz write rates
- Buffer size user programmable up to 1k words
- A user programmble amount of old data can be re-read before new data is added
- Provides the input buffer requirements for the PDSP16510 FFT Processor when 1024 point continuous transforms are performed
- User programmable get ready to Read Me Flag
- Data Available Flag indicates the required amount of new data has been acquired
- 84 Pin PGA or 132 Pin QFP

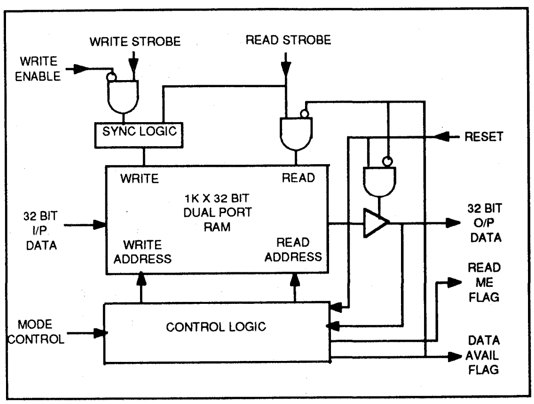

Figure 1. Simplified Block Diagram

NAME	TYPE	SIGNAL DESCRIPTION
IP31:0	I/P	32 bit input bus. If MD5 is high, pins IP16:31 are redundant
D31:0	O/P	32 bit output bus. This bus will be high impedance until the Data Available Flag is active. It then remains low impedance until the required amount of data has been read. D15:0 become inputs during reset, and may be used to define the operating conditions.
RS	I/P	The read strobe must be continuous, and the rising edge transfers data to the output pins.
WS	I/P	Write strobe used to load data into the internal RAM. This strobe may be asynchronous to the read strobe, and may be continuous or intermittent.
$\overline{\text{WEN}}$	I/P	Write enable which when low allows the write strobe to load data.
$\overline{\text{DAV}}$	O/P	Data Available Flag. This signal goes active low when the required amount of new data has been written to the RAM. The complete block of data will then be read from the RAM in sequence using the read strobe. The next system component must be ready to accept the information, which will consist of both new and old data, in amounts defined by MD2:1. The flag will go in-active for one read strobe period every time new data is written to the RAM, and stays in-active when the complete block has been transferred.
RMF	O/P	Read Me Flag. This signal goes active high when a user defined amount of new data has been written to the RAM. It can go active before DAV goes active, and thus allows the system to prepare itself for data when it becomes available. It stays active until the complete block has been read.
MD0	I/P	When MD0 is low the block length is 1024 words. When it is high the block length is defined in groups of 32 words by the data on D4:0 during reset.
MD2:1		MD2:1 define the amount of new data within the block length as defined above. The options are 1024 (00), 512 (01), 256 (10), or the number defined in groups of 32 words by D9:5 during reset (11). When the number of new words is less than the block length defined by MD0, the first words read from the RAM will be data previously stored.
MD4:3	I/P	MD4:3 define the number of new words which are written before the Read Me Flag goes active. The options are 1024, 512, 256 or the number defined in groups of 16 words by D15:10 during reset.
MD5	I/P	When this pin is high the device will support the real transform mode of the PDSP16510. Only IP15:0 input pins are then used and 2 blocks are acquired before the flags go active. Both blocks are then read in parallel using the 32 output pins.
RES	I/P	When this pin is low outputs D15:0 become inputs, which are used to define the operating mode if the internal options have not been selected. The input can be power on reset.
GND	I/P	Four ground pins. All must be connected
VCC	I/P	Four +5 volt pins.All must be connected

	1	2	3	4	5	6	7	8	9	10	11	12	13
N	D7	D8	D10	D12	D14	VDD	D17	GND	D19	D21	D23	D25	D26
M	D6		D9	D11	D13	D15	D16	D18	D20	D22	D24		D27
L	D4	D5										D28	D29
K	D2	D3										D30	D31
J	D0	D1										MD0	MD1
H	GND	RMF										MD2	GND
G	RS	WEN										MD3	MD4
F	VDD	DAV										MD5	VDD
E	WS	IP 0										IP 31	RES
D	IP 1	IP 2										IP 29	IP 30
C	IP 3	IP 4										IP 27	IP 28
B	IP 5		IP 8	IP 10	IP 12	IP 14	IP 16	IP 17	IP 19	IP 21	IP 23		IP 26
A	IP 6	IP 7	IP 9	IP 11	IP 13	VDD	IP 15	GND	IP 18	IP 20	IP 22	IP 24	IP 25

Pin Out Diagram - Bottom View (84pin PGA - AC84)

FUNCTIONAL DESCRIPTION

The PDSP16540 is designed for use in synchronous data flow systems in which the transfer between system elements is contolled by a continuously available system clock. This system clock is usually at the maximum rate that the system elements will allow, since it is governing the rate at which processing can be performed on the acquired data. The rate at which external data is actually inputed to the system (the sampling rate in DSP terminology) is usually much slower than the internal system, or computational, rate. The PDSP16540 then provides a reservoir for data which is acquired at the sampling rate and then processed with the higher speed system clock rate.

Data is written to the RAM using an asynchronous write strobe when a write enable input is active. The enbling signal must meet the set up and hold times given in Table 1. Data is read from the RAM using a read strobe which is expected to be continuously availble and not to just go active when read operations are actually needed. It is normally the high speed system clock discussed earlier. All RAM addresses are generated internally since the device is partitioning consecutive data inputs into pre-defined blocks, which are then transferred to the rest of the system at the system clock rate.

All internal read and write operations are actually performed by the continuous read strobe. When a write strobe is received, internal syncronization occurs and the write operation is actually done with the read strobe. If data is being read from the RAM when a write operation is requested, the read sequence will be interupted for one read strobe period. The flag indicating that data is available goes in-active for this strobe period and the next system element should not accept data during this perfiod.

The correct operation of the write synchronization circuit requires that write operations occur at a slower rate than that of the read strobe. In fact the write strobe period must be at least twice the read strobe period plus some internal delays. Table 1 gives the actual maximum writing rates, and shows that the rate must be reduced when the block of data which is read from the RAM is not completely composed of new data. The maximum writing rate is limited by the need to have read a complete block before the requested amount of new data has been loaded.

A Data Available Flag is provided which goes active when the pre-defined number of words have been written to the RAM. The data read sequence then automatically starts and the flag will go in-actice when the pre-programmed amount of data has been read. An additional get ready to Read Me Flag is provided which can separately be programmed to occur at any point during the block write operation. This flag has no internal action but can be used to warn the next system element that data is to be expected.

DEFINING THE LENGTH OF THE BLOCKS

The amount of new data written to the RAM before the Data Availble Flag is raised, and the amount of data which is then read from the RAM are separately definable. In this way the user can define the amount of old data which is re-read before the new data will be accessed. These overlapping data blocks are required in systems performing frequency domain transforms, when a window operator is applied to prevent frequency discontinuities between the blocks. The resulting loss of information, caused by de-emphasizing the data near the edges, is recovered by overlapping the blocks.

The mode control input MD0 is used to define the block length during the read operation. When MD0 is tied low the read block length will be 1024 words. When MD0 is tied high the block length is defined by the state of pins D4:0, which become inputs whilst the RESET input is active. A tri-state buffer is needed on the outputs which is only enabled during RESET, and whose inputs define the block length. These five inputs allow the block length to be defined in multiples of 32 words, from a minimum of 32 up to the maximum of 1024. The decode of the five bits (0 - 31) should be considered as defining additional blocks of 32 words above the 32 word minimum.

The mode control inputs MD2:1 are used to define the number of new words in the total block defined as above. Decodes 0 through 2 define 1024, 512, and 256 new words respectively. Decode 3 is used when a finer definition is needed, and makes use of the states of pins D9:5 during reset. The decodes of the five bits (0 - 31) then define additional groups of 32 words above a 32 word minimum.

USING THE FLAGS

The data available flag (DAV) always goes active when the required number of new words have been written to the buffer, and the first word to be read is available at the output pins. The rising edges of the read strobes must then be used by the system to transfer the complete block of data to the next system component. The minimum write periods given in Table 1 ensure that the first word will have been read before it is replaced with new data.

Internal logic will increment the read address counter and DAV will go in-active when the complete block has been read. The DAV output will also go in-active for one read strobe period every time a new word is written to the buffer. Write operations to the next system component should be inhibited for that cycle, and the DAV ouput must be used as write enable for the next device. All DAV transitions are produced by the rising edge of the read strobe.

An additional flag is provided which can be used to warn the next system component that data is to be expected. This get ready to read me flag (RMF) can be programmed to occur at any point (within 16 words) during the write operation. Decodes 0 through 2, from mode control inputs MD4:3, will cause the flag to go active after 1024, 512, or 256 words respectively have been loaded. Decode 3 allows the state of pins D15:10 during RESET to be used to define the transition point. Decodes 0 through 63 define form 0 to 63 additional groups of 16 words after the minimum 16 words have been loaded. The RMF flag goes in-active at the same time DAV goes in-active.

The gap between the RMF and DAV outputs should be sufficient to ensure that the next system component can immediately accept data once DAV goes active. The RMF flag has no internal action within the PDSP16540.

SUPPORTING THE PDSP16510

The PDSP16510 FFT Processor does not contain sufficient RAM to allow it to perform continuous 1024 point transforms without ignoring some of the incoming data. When the PDSP16540 is used as an input buffer, continuous transforms can be executed without any loss of information.

When block overlapping is not needed, or if the amount is restricted to either 50% or 75%, the mode control inputs can be directly used to define the operation of the PDSP16540. The D15:0 pins need not be used to define the block lengths.It should be noted, however, that the reset input is still needed to initialise the device, even though the state of the D15:0 pins is irelevent at that time. Figure 1 shows such a system.

Tying MD0 low defines the block length to be 1024 words, and tying MD2:1 appropriately high or low will produce the required decodes to provide 0%, 50%, or 75% overlaps. With 50% overlapping 512 new words are loaded, and with 75% overlapping 256 new words are needed. MD5 should be tied low unless real only transformsd are to be done (See the next section).

The DAV output is used to drive the INEN input on the PDSP16510 and the RMF flag is not used. The PDSDP16510 must be used in the mode in which INEN is an enabling signal, rather than its edge activated mode (Control Register Bit 12 must be set). The LFLG transition produced by the PDSP16510 is not used by the PDSP16540, since internal logic computes the starting address for the read operation.

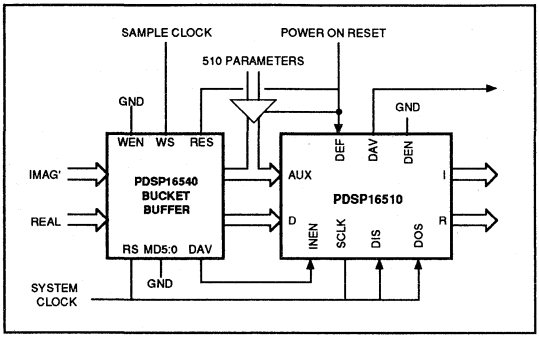

Figure 1. Typical 1024 Point FFT System

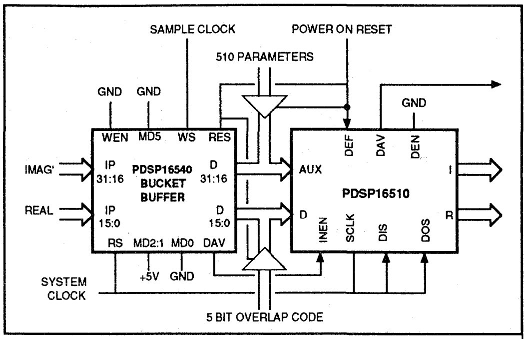

Figure 2. System with Non Standard Overlaps

Figure 2 shows a 1024 point system which allows the amount of overlap to be any value within 32 words. The 5 bit overlap code defines groups of 32 new words which are written to the buffer, in addition to the minimum number of 32 words. The smaller the number of new words written, the greater is the overlap with the previous block.

During reset the D31:0 outputs from the PDSP16540 will be high impedance and the 5 bit code is inputed on D9:5. This high impedance state also allows the PDSP16510 control parameters to be inputed on its AUX 15:0 bus without any conflicts.

The rate at which data is written to the PDSP16540 must be such that 1024 words can be transferred between the devices, transformed , and then moved to the output circuit for analysis before the DAV flag goes active again. Since the read operation is interrupted for one cycle every time a write operation occurs, the equation controlling the minimum writing period is given by;

$$NS > 1024B + \frac{1024BS}{S-B} + T + D$$

where N is the amount of new data written to the buffer, S is the period of the write strobe, B is the read strobe period, T is the transform time as given in the data sheet for the PDSP16510, and D is the time to transfer data from the PDSP16510 to the next system device.

It must be noted that the above minimum write period only applies if continuous inputs are to transformed without the loss of any incoming information. Peak writing rates can be much higher if gaps occur within the incoming data stream. The minimum periods given in Table 1 then limit the writing rate.

When the PDSP16510 uses a 40 MHz clock, dumps its transformed data with a 40MHz strobe, and the PDSP16540 uses a 40 MHz read strobe, then the minimum S period is149 ns. This equates to a 6.7 MHz writing rate when blocks are not overlapped, 3.35 MHz with 50% overlaps (512 new words), or 1.675MHz with 75% overlaps (256 new words).

Characteristic	Min	Max	Notes
RS Period,Tp	25ns		
RS Low Time	8ns		
RS High Time	8ns		
WS Period	2Tp+10ns		Both conditions must be satisfied
WS Period	$\frac{Tp \times L}{N}$		L = Block length, N = amount of new data written
WS Low Time	10ns		
WS High Time	10ns		
$\overline{\text{WEN}}$ set up wrt WS going high	2 ns		$\overline{\text{WEN}}$ going active or in-active
$\overline{\text{WEN}}$ Hold wrt WS going high	8 ns		
Data In Set Up wrt RS going high	8 ns		
Data In Hold Time wrt RS going high	0 ns		
Delay from RS going high to O/P Data		19ns	All output delays are with 30pf loads
$\overline{\text{DAV}}$,RMF transition wrt to RS going high	10ns	18ns	Going active or in-active
Time to go Low Z wrt to RS going high		19ns	Occurs when DAV also goes active
Time to go High Z wrt to RS going high		12ns	Occurs when DAV also goes in-active

Table 1. Timing Information

The amount of overlapping is dependent on the needs of a particular application, and is usually subject to some compromise. If the above maximum writing rates are marginally not adequate, the amount of overlap can possibly be reduced to achieve the required performance. Mode control inputs MD2:1 should then all be tied high, and outputs D9:5 used as inputs during reset to define the number of new words to be written.

SUPPORTING REAL ONLY TRANSFORMS

If MD5 is tied high the PDSP16540 will support the PDSP16510 when two concurrent 1024 point real transaforms are to be performed. It does not suppoprt block overlapping in this mode.

ABSOLUTE MAXIMUM RATINGS [See Notes]

Supply voltage Vcc	-0.5V to 7.0V
Input voltage V_{IN}	-0.5V to Vcc + 0.5V
Output voltage V_{OUT}	-0.5V to Vcc + 0.5V
Clamp diode current per pin I_K (see note 2)	18mA
Static discharge voltage (HMB)	500V
Storage temperature T_S	-65°C to 150°C
Ambient temperature with power applied T_{AMB}	0°C to 70°C
Junction temperature	150°C
Package power dissipation	3000mW
Thermal resistances	
Junction to case $ø_{JC}$	5°C/W

NOTES ON MAXIMUM RATINGS

1. Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
2. Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
3. Exposure to absolute maximum ratings for extended periods may affect device reliablity.
4. Current is defined as positive into the device.

STATIC ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise state)

Tamb = 0 C to +70°C.

Vcc = 5.0v ± 10%

Characteristic	Symbol	Value Min.	Typ.	Max.	Units	Conditions
Output high voltage	V_{OH}	2.4		-	V	I_{OH} = 4mA
Output low voltage	V_{OL}	-		0.4	V	I_{OL} = -4mA
Input high voltage	V_{IH}	2.8			V	
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	μA	GND < V_{IN} < V_{CC}
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	μA	GND < V_{OUT} < V_{CC}
Output S/C current	I_{SC}	10		300	mA	V_{CC} = Max

Real only data is written to the buffer using the IP15:0 inputs, and the IP31:16 inputs are redundant. Two blocks of data are acquired before DAV goes active, and both blocks are then read in parallel using all thirty two outputs.

MD0,1, and 2 must be tied low in order to define blocks of 1024 words which totally consist of new data. The RMF flag is not needed by the PDSP16510, but will actually go active after the defined number of words in the second block have been loaded. Control Register Bits 8:6 in the PDSP16510 must be set to 101 in order to expect data on both its real and imaginary inputs.

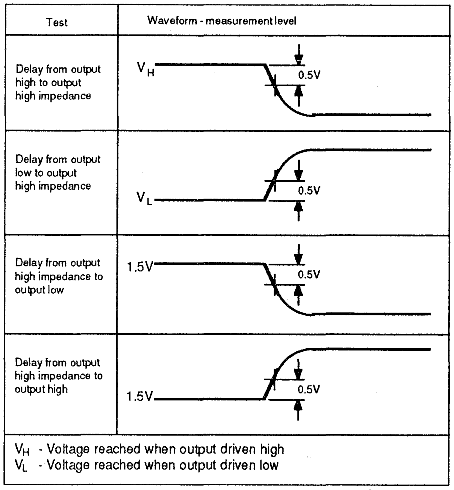

ORDERING INFORMATION

PDSP16540 C0 AG	Commercial	- PGA package
PDSP16540 C0 GC	Commercial	- Ceramic QFP
PDSP16540 B0 AG	Industrial	- PGA package
PDSP16540 B0 GC	Industrial	- Ceramic QFP
PDSP16540 A0 GC	Military	- Ceramic QFP

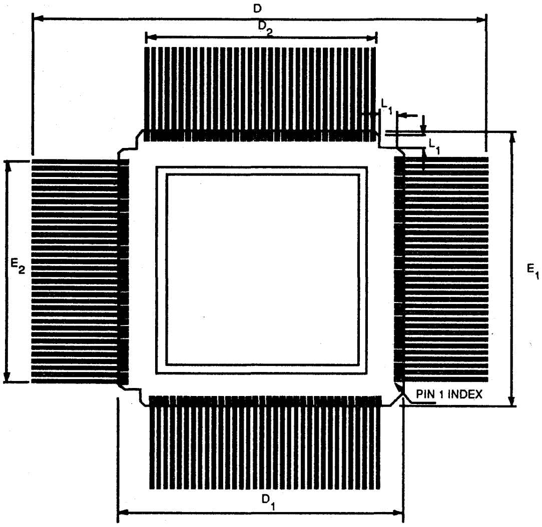

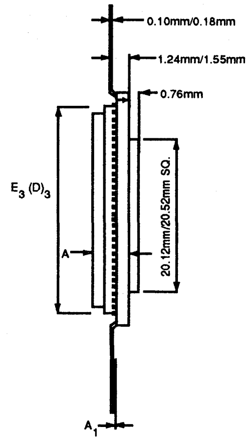

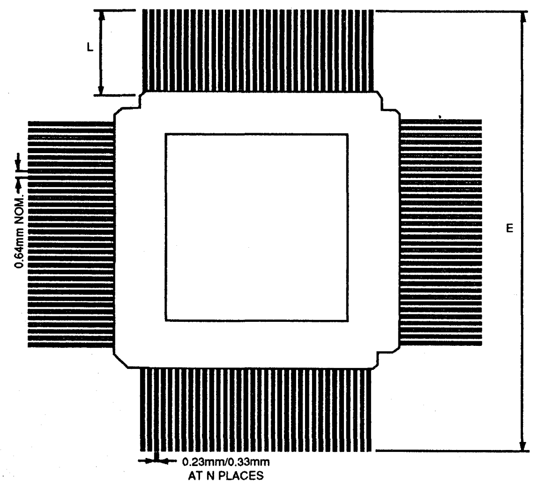

Sym -bol	MILLIMETRES			Note
	MIN	Nominal	MAX	
A	2.08		2.59	
A_1		0.20		
D	36.96		38.35	
D_1	24.00		24.38	
D_2		20.32		
D_3	20.62		21.30	
E	36.96		38.35	
E_1	24.00		24.38	
E_2		20.32		
E_3	20.62		21.03	
L	6.48		6.99	
L_1		1.27		
N		132		
NOTE				

Sym -bol	INCHES			Note
	MIN	Nominal	MAX	
A	0.082		0.102	
A_1		0.008		
D	1.455		1.510	
D_1	0.945		0.960	
D_2		0.800		
D_3	0.812		0.828	
E	1.455		1.510	
E_1	0.945		0.960	
E_2		0.800		
E_3	0.812		0.828	
L	0.255		0.275	
L_1		0.050		
N		132		
NOTE				

132 - Pin Power Ceramic QFP - GC132
(used for PDSP16340, PDSP16350, PDSP16488,
PDSP16510, PDSP16540)

ORDERING INFORMATION

PDSP16540 C0 AG	Commercial	-	PGA package
PDSP16540 C0 GC	Commercial	-	Ceramic QFP
PDSP16540 B0 AG	Industrial	-	PGA package
PDSP16540 B0 GC	Industrial	-	Ceramic QFP
PDSP16540 A0 GC	Military	-	Ceramic QFP

Application Notes

GEC PLESSEY
SEMICONDUCTORS

A FAST FFT PROCESSOR USING THE PDSP16000 FAMILY

THE PROBLEM

The current industry standard benchmark for the execution of a 1024 point complex FFT is of the order of 2ms. This article describes how an FFT processor almost an order of magnitude faster may be built using GEC Plessey Semiconductors' PDSP16000 family of highly integrated CMOS DSP building blocks.

Fig.1 shows diagrammatically the fundamental operation of the FFT algorithm, the Butterfly.

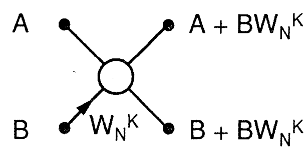

Fig.1 The Butterfly

This operation requires a complex mutiply, an addition and a subtraction to complete. If we set a target of 256µs for a 1024 point FFT (roughly ten times faster than can be achieved with an FFT processor designed around a 100ns MAC), then the time allowed to calculate the Butterfly is:

$256/(N/_2 \log_2 N)$ µs (where N=1024)

=50ns.

The complex multiplication operation itself requires four real multiplications, an addition and a subtraction, so that the complex Butterfly requires four multiplication, three additions and three subtractions to be executed in 50ns.

The 28-bit results from these operations are rounded to 16 bits before being passed to the adder and subtractor. The subtractor calculates:

$$(X_R.Y_R - X_I.Y_I) = P_R$$

to form a 17-bit real result P_R.

The adder calculates:

$$(X_R.Y_I + X_I.Y_R) = P_I$$

to give a 17-bit imaginary result P_I

The add and subtract operations may, depending on the data, cause the results to grow by one bit (hence the 17-bit wide outputs). The PDSP16112 operates using 2's complement arithmetic, hence if fractional 2's complement is used, the outputs will lie in the range:

$$-2 \leq P < 2$$

for inputs in the range:

$$-1 \leq X \text{ or } Y < 1.$$

For outputs in the range:

$$-1 \leq P < 1$$

the 17th bit (MSB) will duplicate the 16th bit (the sign bit).

Both inputs and outputs are registered. On the rising edge of the clock, data is clocked into the input registers. At the same time a new result is clocked into the output registers. The maximum clock frequency is 20MHz, giving a full complex multiply in 50ns. The final operations required in calculating the Butterfly are the addition and subtraction. The PDSP16318 Complex Accumulator (Fig. 3) can be configured for these operations.

Using ECL logic and the fastest available ECL array multiplier (BIT's ECL MAC will operate at 100MHz) it is possible to construct a 50ns Butterfly Processor. The problems with this, though, are quite severe. The ECL Butterfly Processor consumes a great deal of power, requires its I/O bus to operate at 100MHz and occupies a large amount of board space. Such a processor is far from simple to design and places horrendous access time requirements on external memory.

THE SOLUTION

GEC Plessey Semiconductors' PDSP16000 family solution takes a very different approach to that above; at the heart of the Butterfly Processor is the PDSP16112 Complex Multiplier (Fig. 2)

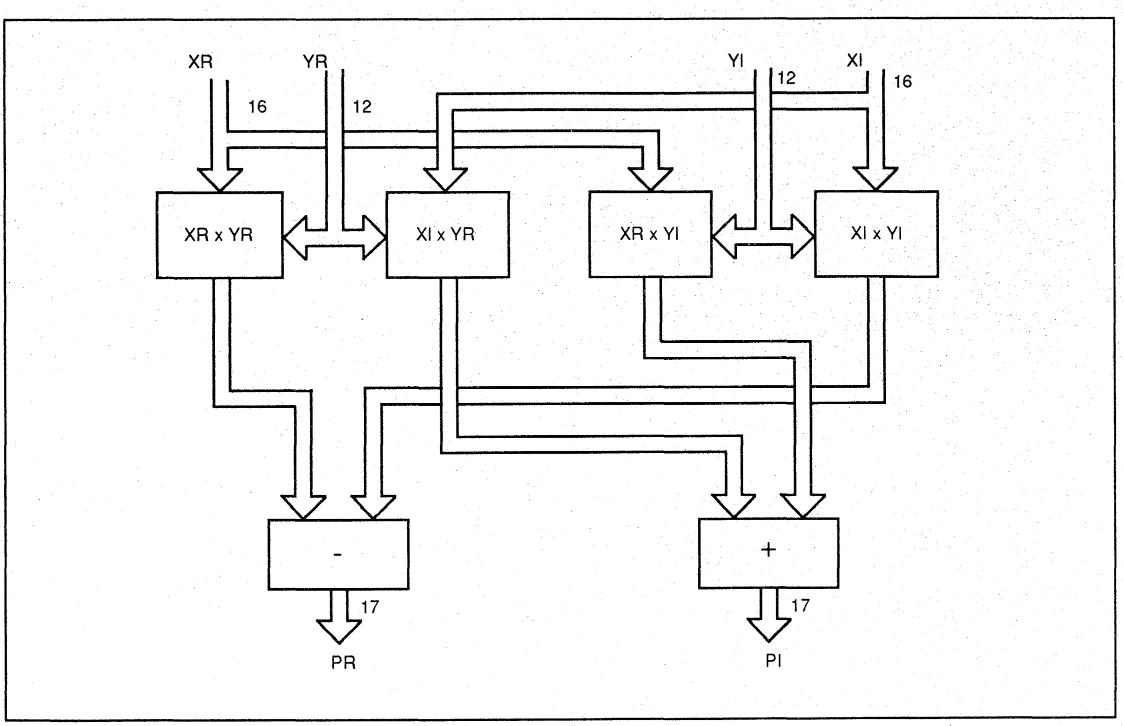

Fig.2 PDSP16112A 20MHz Complex Multiplier

This device contains four pipelined 16x12 Array multipliers, a 17-bit Adder and a 17-bit subtractor. The multipliers accept data from the X_R, X_I, Y_R, Y_I inputs and perform the four multiplies necessary to implement a complex multiplication:

$$X_R.Y_R;\ X_R.Y_I;\ X_I.Y_R;\ X_I.Y_I$$

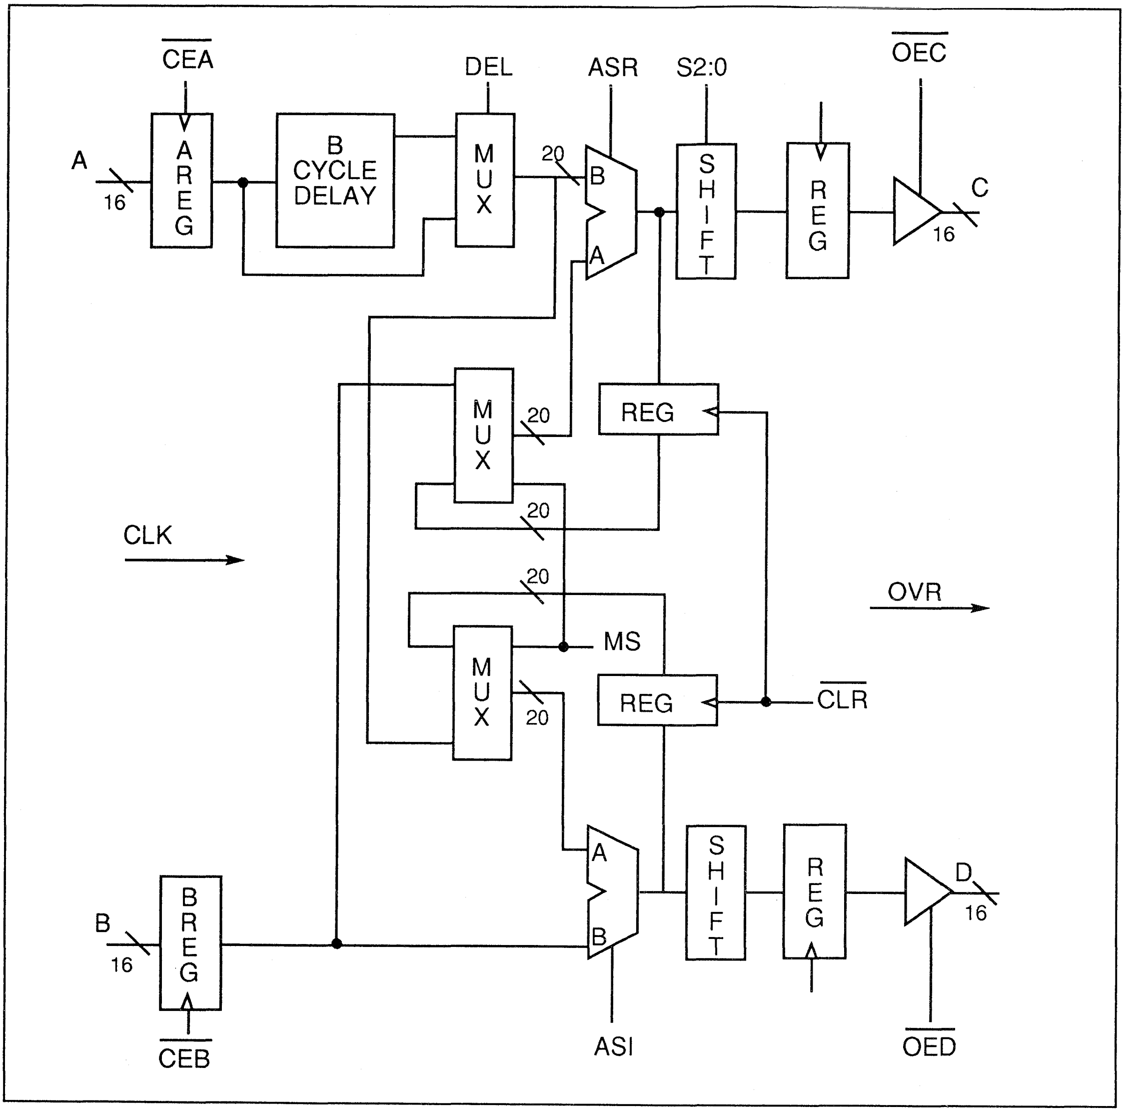

Fig.3 PDSP16318A 20MHz Complex Accumulator

This device has a variety of applications in filtering, correlation and FFT. In filtering and correlation applications, a single PDSP16318A is used in conjunction with a PDSP16112A to form a complex MAC. When used in FFT applications, a pair of PDSP16318As are used with a PDSP16112A to form a Butterfly Processor capable of executing a Radix 2 DFT Butterfly every 50ns, using 16 bit data and 12 bit twiddle factors.

Fig 4 illustrates the connections between the devices.

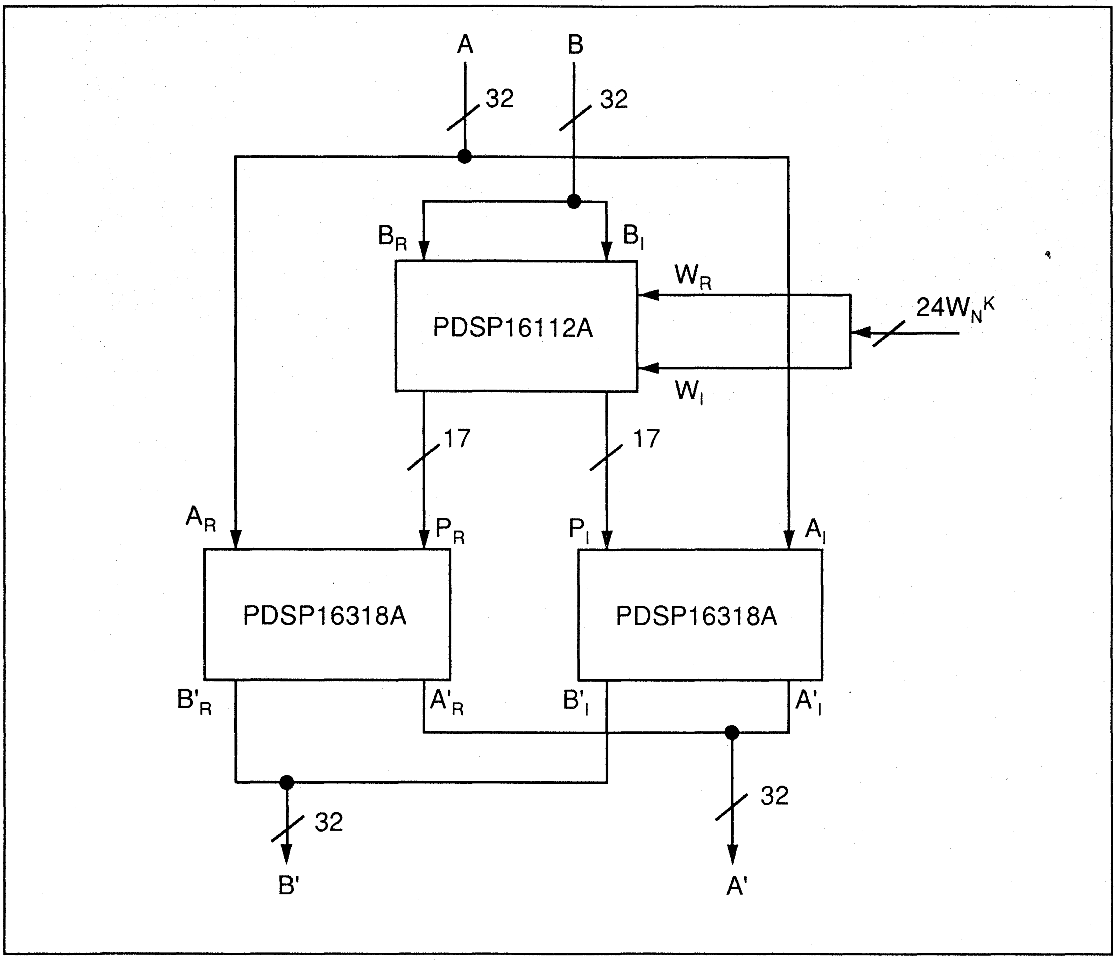

Fig.4 Radix 2 Butterfly Processor using PDSP16112A & PDSP16318A

The PDSP16112A provides the real and imaginary parts of BW_N^K to the two PDSP16318As. One of the PDSP16318As calculates the real parts of $A+BW_N^K$ and $A-BW_N^K$, the other, the imaginary parts of $A + BW_N^K$ and $A-BW_N^K$.

For even greater throughput, one chip-set (16112+2 x 16318) may be allocated to each column of the FFT. As an example, 10 chip-sets will allow the execution of a 1024 point complex FFT in a mere 26μs!

Application Note AN47 'A RADIX 2 BUTTERFLY PROCESSOR' describes the Butterfly hardware in greater detail.

MEMORY REQUIREMENTS

In the 1960s, when the FFT algorithms were first being developed, memory was an expensive commodity. This led to the invention of 'In-Place' FFT algorithms in which the results A', B' of a Butterfly are put back into the locations from where the inputs A, B are read. With a Butterfly Processor as fast as the one described above, In-Place Algorithms pose a nasty problem.

Examination of Fig. 4 shows that in every 50ns cycle, two reads from and two writes into memory have to be accomplished. The obvious implication is that the RAM has to have an access time less than 12.5ns. Such RAM is expensive as these speeds are right at the limits of that achievable for CMOS RAM - clearly an alternative arrangement must be found.

THE CONSTANT GEOMETRY ALGORITHM

The Constant Geometry Algorithm is illustrated in Fig. 5. On each pass of the FFT, the read/write address sequence is the same, but the addresses written to after each Butterfly are different to those from where the input data is read. This requires twice as much RAM as for In-Place algorithms.

The key to the use of the Constant Geometry Algorithm is the recognition of the order in which data points are addressed. Fig. 5 illustrates the Butterfly structure of the Constant Geometry Algorithm. For an N point transform, the read addressing sequence is:

A		**B**
0	and	N/2 + 0
1	and	N/2 + 1
2	and	N/2 + 2
3	and	N/2 + 3

or in general for n = 0 (N/2-1)

the addresses are n and N/2 + n.

For the same N point transform the write address sequence is:

A'		**B'**
0	and	1
2	and	3
4	and	5
6	and	7

or in general for n = 0 (N/2-1),

the write addresses are 2n and 2n + 1.

MEMORY CONFIGURATION

The implication so far is that four memory accesses are required every cycle, which is only 50ns long hence needing 12.5ns memory cycles! The reality is that four separate blocks of RAM may be configured as one in such a way that any given device is only accessed twice for each four access Butterfly cycle.

This reduces the required access time to only 25ns which is feasible with current RAM devices. The required RAM bandwidth may be reduced by a further factor of two to a more comfortable 50ns by 'double buffering', ie using two banks of storage, one for reading and one for writing. After each pass these storage banks swap roles, data being passed back and forward between them via the Butterfly Processor. This step doubles the amount of RAM required, but reduces each RAM device's required I/O bandwidth by a factor of two. to one cycle every 50ns.

Half of the required memory configuration is as shown in Fig. 6. This structure is duplicated, each half alternating between sourcing and receiving data to and from the Butterfly Processor.

Each memory is divided up into four quadrants each quadrant being a separate 32-bit block of RAM with separate input and output ports. The two left hand quadrants are configured to accommodate data points with even valued addresses, the right quadrants accommodating data points with odd valued addresses. The upper two quadrants accommodate data points with address values greater than (N/2-1), where N is the transform size, the lower two quadrants accommodate data points with addresses with values less than or equal to (N/2-1).

The left hand quadrants have their inputs commoned to become the "A" input bus, the right hand quadrants have their inputs commoned to become the "B" input bus. The upper quadrants have their outputs commoned to become the "B" output bus, and the lower quadrants have their outputs commoned to become the "A" output bus. These buses are connected to the Butterfly Processor "A', B', B, A" output and input buses respectively.

The mode of addressing the composite RAM is surprisingly simple as each quadrant is supplied with exactly the same address.

The example of Fig. 7 shows the storage of data points for a 16 point transform according to the Even-Odd (N/2-1) rules. These 16 data points are addressed two at a time by the 8 term address sequence 0,0,1,1,2,2,3,3 or the address sequence 0,1,2,3,0,1,2,3 depending upon whether reading or writing is required. Unlike all other FFT algorithms, this address sequence remains unchanged throughout each pass of the transform.

Reading Mode

When reading data from the RAM, quadrants enabled alternate between the left and the right pairs of quadrants and the address sequence used is 0,0,1,1,2,2,3,3. As can be seen from the example in Fig. 7 this simple count sequence on the address port will result in the data points read onto the A and B buses being:

A		B
0	and	8 on the 1st cycle with the left quadrants enabled
1	and	9 on the 2nd cycle with the right quadrants enabled
2	and	10 on the 3rd cycle with the left quadrants enabled
3	and	11 on the 4th cycle with the right quadrants enabled
4	and	12 on the 5th cycle with the left quadrants enabled
5	and	13 on the 6th cycle with the right quadrants enabled
6	and	14 on the 7th cycle with the left quadrants enabled
7	and	15 on the 8th cycle with the right quadrants enabled

Write Mode

When writing data to the RAM the quadrants enabled are the lower pair for the first half of the operation and the upper pair for the second half of the operation. The address sequence used is 0,1,2,3,0,1,2,3. As can be seen from the example in Fig. 7, this simple count sequence on the address port will result in the data points written into the RAM from the A' and B' buses in the following manner:

A'		B'
0	and	1 on the 1st cycle with the lower blocks enabled
2	and	3 on the 2nd cycle with the lower blocks enabled
4	and	5 on the 3rd cycle with the lower blocks enabled
6	and	7 on the 4th cycle with the lower blocks enabled
8	and	9 on the 5th cycle with the upper blocks enabled
10	and	11 on the 6th cycle with the upper blocks enabled
12	and	13 on the 7th cycle with the upper blocks enabled
14	and	15 on the 8th cycle with the upper blocks enabled

Reference to the Constant Geometry diagram of Fig.5 will show that this address sequence is as required by the algorithm. Each separate RAM package is only accessed once each cycle resulting in 50ns access time RAM being sufficiently fast for this speed application.

COEFFICIENT ADDRESSING

The coefficient addressing sequence required by the Constant Geometry Algorithm is as simple as the Data Addressing sequence. The correct sequence for a normally ordered input, bit reversed output Forward Transform is as illustrated in Fig. 5.

To generate the sequence of values of K for the coefficients W_N^K in Fig. 5 use the following routine:

$$K = \text{Bit Reversed } [0,\ldots.(2^{(m-1)}-1)]$$

this sequence is repeated $^1/_2\, N.2^{m-1}$ times where m is the column number of the FFT and N is the number of points.

Thus for 16 points, for example, the sequence on the 4th pass is given by the count sequence of 0 to (2 to the power 3) - 1 = 7 is repeated (16/2) / (2 to the power 3) = once.

This count from 0 to 7 is then bit reversed to give 0, 4, 2, 6, 1, 5, 3, 7 as shown in Fig.5.

The coefficients need to be accessed on the same 50ns cycle as data, requiring the use of RAM as the storage medium.

FFT PROCESSOR CONFIGURATION

The architecture of the complete FFT processor is as shown in Fig 9. A configuration using 50ns access 256 word RAM devices will in conjunction with PDSP16112A and PDSP16318A Arithmetic processors execute a 1024 point Complex FFT in just 256µs, a solution that is realised entirely with CMOS Logic, fits on a single board yet delivers a benchmark eight times faster than the Industry Standard.

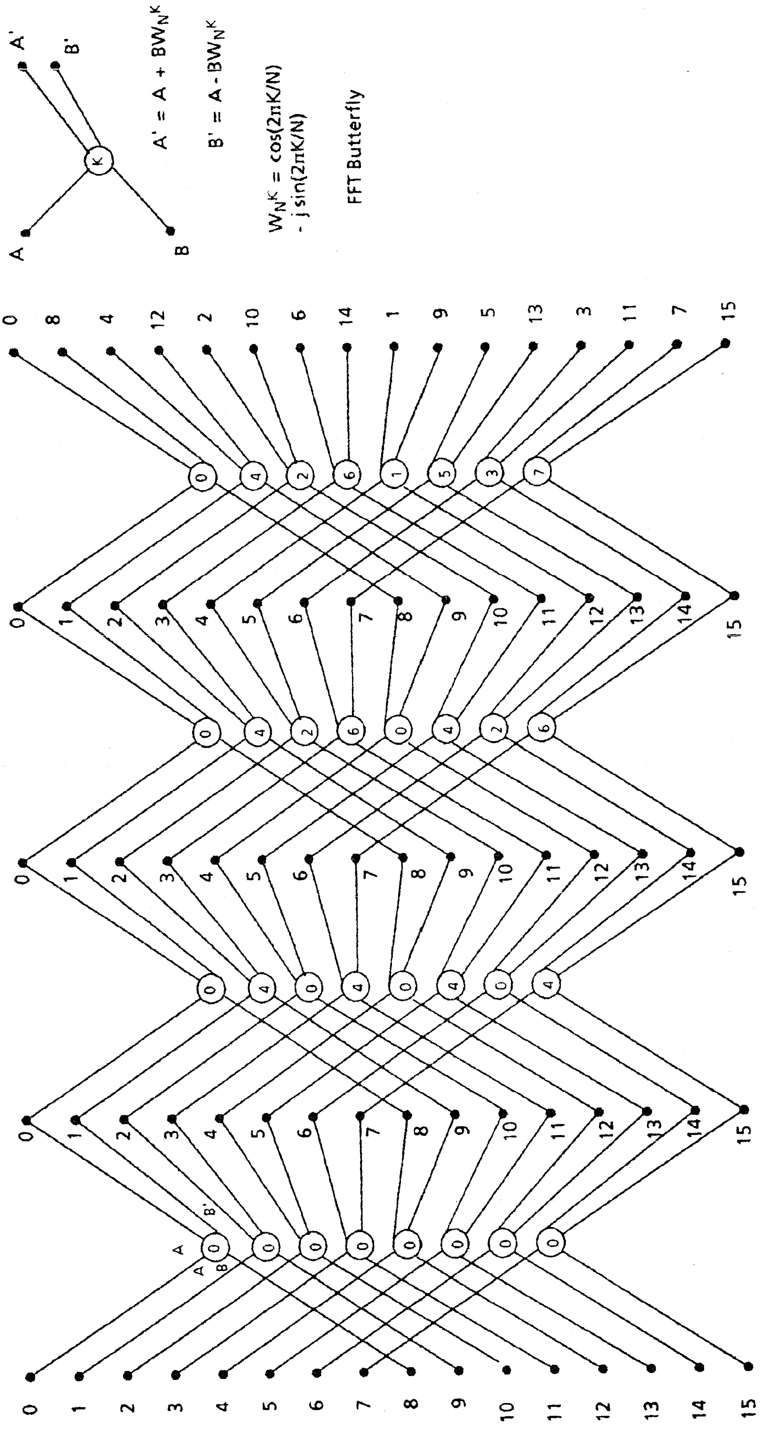

Fig.5 16 point constant geometry DIT Radix 2 algorithm

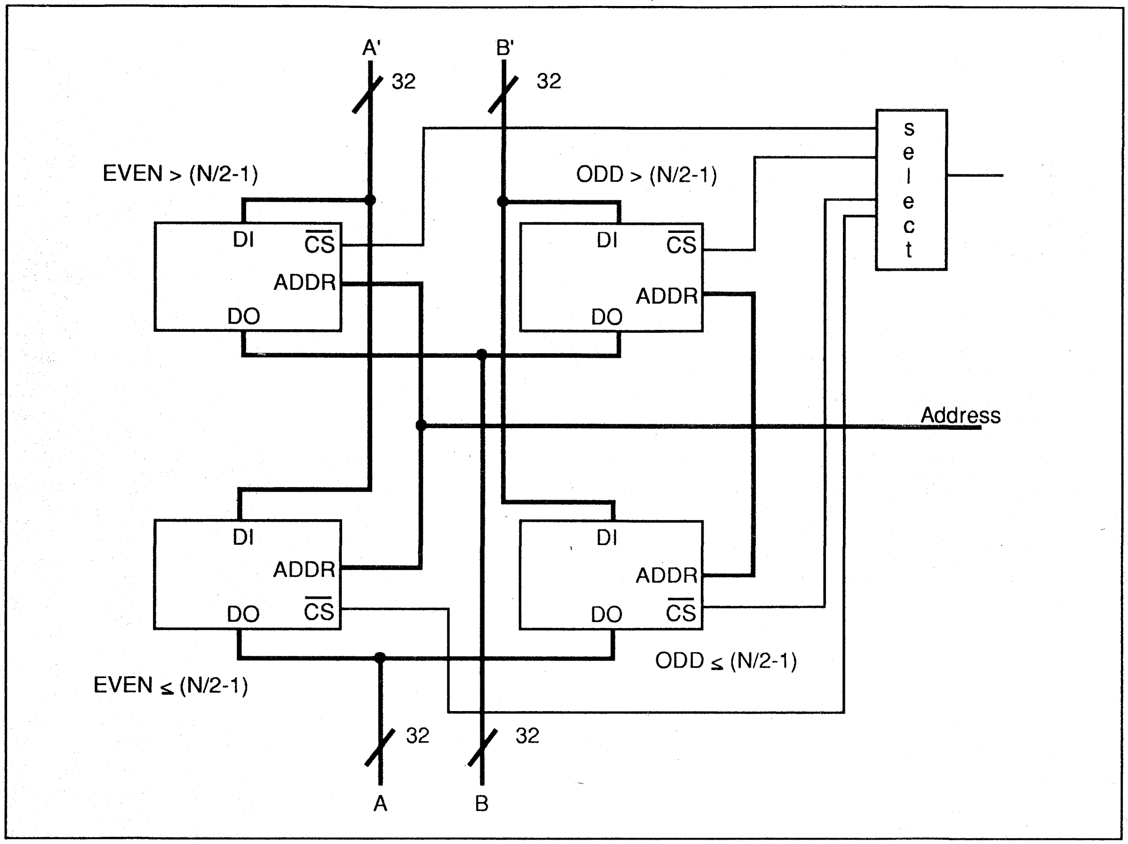

Fig.6 Memory Configuration For Constant Geometry Algorithm

The memory is configured in four quadrants each comprising 32 bit wide blocks of RAM. All four quadrants share the same Address Bus and Read/Write select lines. Each quadrant is independently enabled via its chip select control.

The blocks of RAM have separate Read and Write Data ports which are connected to the A, B, A', B' Butterfly Processor ports as indicated.

EVEN > {(N/2) - 1}

DATA POINT	8	10	12	14
RAM ADDRESS	0	1	2	3

ODD > {(N/2) - 1}

DATA POINT	9	11	13	15
RAM ADDRESS	0	1	2	3

EVEN ≤ {(N/2) - 1}

DATA POINT	0	2	4	6
RAM ADDRESS	0	1	2	3

ODD ≤ {(N/2) - 1}

DATA POINT	1	3	5	7
RAM ADDRESS	0	1	2	3

Fig.7 16 point Transform Example

The data points of a 16 point transform are stored in the four RAM quadrants at the addresses and in the quadrants indicated.

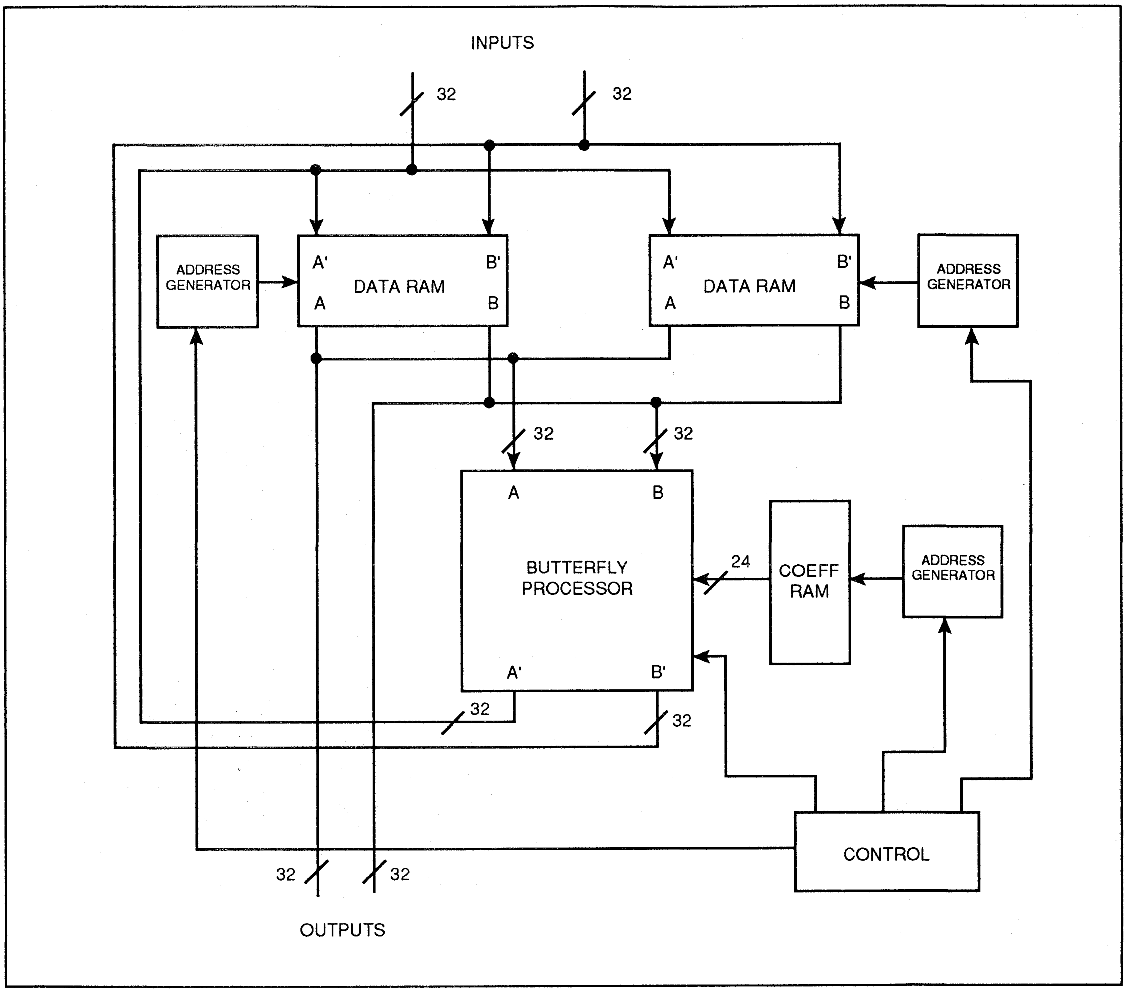

Fig.8 FFT Processor architecture

INTERFACING THE PDSP FAMILY

INTRODUCTION

GEC Plessey Semiconductors' PDSP family of DSP functional blocks are fabricated on a high speed CMOS process, and incorporate several design features to ease interfacing and board layout. However there are a few precautions which should be taken which will ensure trouble-free board design and operation.

All parts in the PDSP family are designed with the generic structure of Fig 1

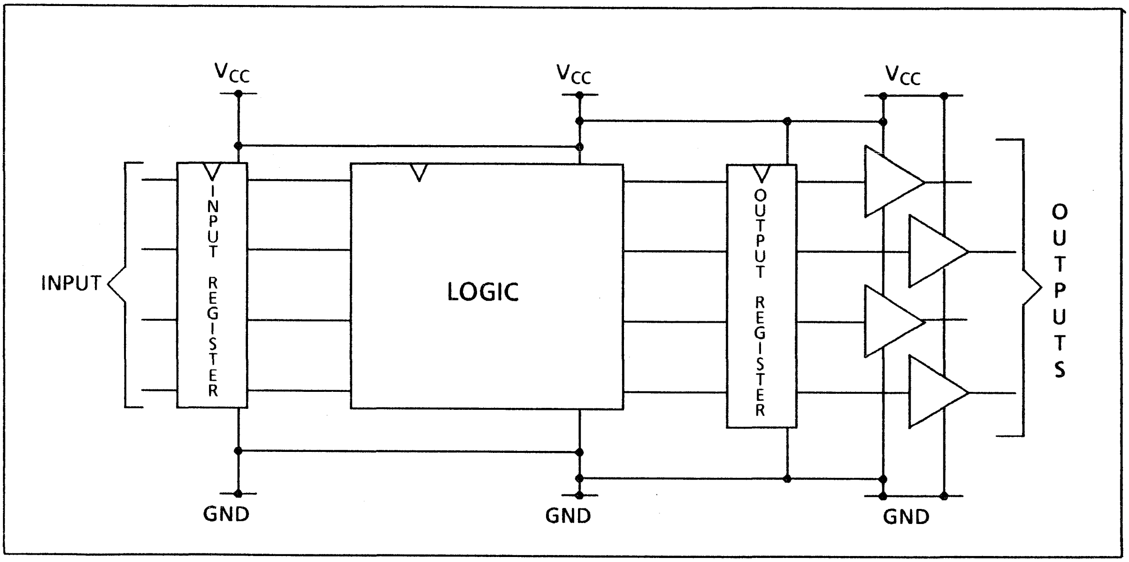

Fig. 1 PDSP structure

The registered input is designed with a positive set-up time (ie data must be presented before the rising edge of the clock) and zero hold time (ie the data is allowed to change anytime after the rising edge of the clock). The input levels are designed for compatability with LSTTL outputs (V_{IH} = 2.2V, V_{IL} = 0.8 V), and the output, although conventional CMOS stages, are specified into a load of 2 standard LSTTL inputs + 20pF for track loading .

All PDSP devices (with the exception of the PDSP16112/A) have tri-state output buffers preceded by an output register, this ensures that the output data is valid for a whole clock cycle. To simplify timing requirements further, the clock to output valid delay is generally less than half a cycle at the maximum specified clock rate.

NOISE

The operating margins of all devices on a board of high-speed logic can best be maintained by providing a quiet environment free of noise spikes, undershoot, and ringing. The key elements in creating such an environment are good supply decoupling and termination of interconnections.

POWER DISTRIBUTION

To maintain wide operating margins across all devices on a board, the supply impedance at each device must be kept to a minimum. The internal design of PDSP devices is such that the input registers, main logic, and output buffers have separate supply pins. This arrangement is designed to ensure that current spikes generated in the output drivers do not modulate the supply to the input gates, hence altering the thresholds. Although these multiple supply pins are internally connected, the internal paths are not particularly low impedance, and therefore each individual V_{CC} pin should be separately decoupled.

The total supply impedance at a device is a function of the supply line impedance and the decoupling capacitors. In practice, the effect of local decoupling does not extend very far, because of the very fast edges of the current spikes generated by CMOS output stages and the inductive nature of the PCB tracks. In order to minimise the effects of these transients, the decoupling capacitors should be high quality, low inductance parts mounted as close as possible to the device pins, with as short a track length as is practical. Capacitor values should be in the 0.1 to 0.47µF region, too small and there will be insufficient decoupling, too large and the equivalent inductance will reduce decoupling efficiency. The quality of the ground connection is also important, this should be either a solid plane or a grid to minimise inductance and prevent loss of noise margin due to differential ground noise between devices.

Low frequency current transients can best be handled by tantalum capacitors mounted close to the edge connector where the panel tracks meet the backplane power distribution system. Such large capacitors provide bulk energy storage which prevents voltage drops due to the long inductive path between the logic board and the system power supply.

TRACK TERMINATION

On a large board PCB tracks look like shorted transmission lines to the signals they are carrying . This causes reflection of the signal resulting in undershoot, overshoot or ringing. Particular cases which can cause difficulty are large RAM arrays being addressed by the PDSP1601- the long track lengths and heavy capacitative loading can store and reflect amounts of energy leading to severe ringing - and LSTTL to CMOS interface via long tracks which can suffer severe undershoot. In both cases track termination is best effected by a series resistor at the driving end (typically 10 or 18 ohms). Parallel termination is not recommended since it reduces the voltage swing at the input (making the noise margin even worse), consumes DC power (hardly desirable in a CMOS system) and doesn't work very well in any case.

VERIFICATION

When a board design is complete and the prototype built, it is good practice to check the power supplies to each device and the signals on the buses with a wideband 'scope to ensure that excessive noise, ringing or undershoot is not present. A board which works on the bench but which is marginal because of noise problems will almost certainly exhibit gremlins in the field.

OPTIMISING THE ACCURACY OF AN FFT SYSTEM

The two major design specifications of an FFT system are speed and accuracy. Matching the processing speeds of each of the different sections of the FFT system is a straightforward design task and optimises the design by ensuring that each piece of hardware operates at its maximum processing rate. The equivalent design task that ensures that each section of the FFT processor has the optimum dynamic range is far more complex, but may be simply determined for five example FFT processor systems described below.

LIMITS TO SYSTEM ACCURACY

Arithmetic accuracy relates directly to the achievable dynamic range of the FFT processor or its ability to discriminate low amplitude signals in the presence of large signals. There are several limiting factors to the overall system accuracy from the A to D converters through arithmetic bit widths to the FFT algorithm itself. An optimum design maintains a constant dynamic range throughout all sections of the FFT. A surprising result is that the A to D converter and the Arithmetic processor have a far greater influence over the achievable dynamic range than the stored constants used for windowing and as 'Twiddle Factors' in the FFT calculation itself.

The following optimisation criteria formed the basis for the development of GEC Plessey Semiconductors' PDSP family of Complex Digital Signal Processing Building Blocks, and the bit widths selected for the data and coefficent paths. GEC Plessey Semiconductors offers optimum FFT implementations for five A to D converter sizes, simplifying the algorithm required to allow high speed for 8 and 10 bit systems, optimal bit widths for 12 and 14 bit systems and offering a new highly integrated solution for 16 bit systems.

FFT OPTIMISATION CRITERIA

Brigham and Cecchini (Ref.1) addressed this optimisation problem by developing a nomogram for determining the maximum bit widths required for each stage of a 1k FFT given an input A to D converter bit accuracy. The following tables are derived from this nomogram and relate to the specifications achievable with the GEC Plessey Semiconductors' PDSP family of complex signal processing building blocks.

The relevant sections of an FFT system that determine the dynamic rage are:

1. The input A to D Converter

The error factors contributed by the input A to D converter that limit dynamic range are quantisation, saturation and aperture jitter.

Quantisation is simply the number of bits in the A to D including sign bits; Saturation simply refers to errors caused when the input signal becomes larger than the maximum output value of the A to D converter. Aperture jitter refers to the difference between the point in time that the sample was meant to be taken, and when it actually was taken.

Assuming saturation can be avoided, and that quantisation is the selection criterion, then to maintain the dynamic range offered by the selected bit width, the maximum tolerable aperture jitter is given in Table 1.

These numbers assume a 2:1 overlap between consecutive transforms, and 50ns butterflies as offered by the PDSP 3-Chip Butterfly solution.

Overlap refers to the number of new samples written into the FFT between each FFT calculation. For an N point FFT if N new points are written into the processor between each calculation, then there is no overlap; if N/2 new points are written then there is a 2:1 overlap; N/4 then 4:1 etc.

If the overlap is increased to 4:1, then the tolerable jitter time doubles. If there is no overlap, then the tolerable aperture jitter times halves. Similarly if the butterfly time is doubled to 100ns, the tolerable jitter time doubles.

2. Number of Bits in Weighting Function Lookup Table

The weighting function applies a window to the input data to de-emphasise points at each end of the sample sequence. This operation minimises the distortions that result from the FFT's assumption that the input samples are part of a periodic signal. There are many different window functions used for different applications, all simply multiply each sample by a number whose value is related to the position of the sample within the input sequence.

The optimum number of bits in the weighting function lookup table, including the sign bit and assuming rounding, can be determined from Table 2. The limitation to dynamic range arising from this operation is caused by quantisation errors in the weighting function values which will be comparable with the A to D quantisation errors for comparable weighting function bit widths.

A to D bit width	Jitter as % of sample period	FFT Size: 1024	512	256	128	64
		Max. tolerable aperture jitter in ns				
8	0.5%	2.51	2.26	2.01	1.75	1.50
10	0.15%	0.79	0.71	0.63	0.56	0.48
12	0.03%	0.16	0.14	0.13	0.11	0.09
14	0.005%	0.03	0.02	0.02	0.02	0.02
16	0.0025%	0.01	0.01	0.01	0.01	0.01

Table 1

A to D bit width	Minimum weighting function bit width	Acceptable bit widths marginal addtional error
8	7	(6)
10	9	(8)
12	11	(10)
14	13	(12)
16	15	(14)

Table 2

3. Number of Bits in Sin-Cos Lookup Table

The Sin-Cos lookup table is the source of the "twiddle factors' used within the FFT calculations. Despite the fact that these values are used many times during the FFT, the required accuracy is no more stringent than that of the A to D converter even for large transforms (these numbers are derived from a nomogram for 1024 point transforms).

The optimum number of bits in the Sin-Cos lookup table, including the sign bit and assuming rounding, can be determined from Table 3.

A to D bit width	Minimum Sin-Cos lookup bit width
8	7
10	9
12	11
14	13
16	15

Table 3

4. Number of Bits in Arithmetic Section

The arithmetic section has the most profound effect upon the overall FFT accuracy, with bit width and chosen algorithm contributing to the total errors introduced and resulting limitation to dynamic range. The bit width contribution is easy to understand; more bits mean smaller errors. The algorithm differences depend upon the scheme used to allow for the inevitable word growth experienced within an FFT processor. The two schemes considered are termed Unconditional Scaling, and Conditional Scaling.

For unconditional scaling, the worst case word growth is assumed to occur during every pass of the FFT and fixed shifts are introduced to eliminate the possibility of overflow. This algorithm is easy to implement and allows very fast processors to be constructed, though as worst case word growth does not in practice always occur, the accuracy and hence dynamic range of this algorithm is reduced.

For conditional scaling, the results of each pass are examined and shifts are onyl used if word growth has actually occurred. This data-dependent algorithm is slower and requires more hardware to execute but offers greater dynamic range as illustrated in Table 4. The arithmetic bit widths shown include sign bits and maintain the dynamic range offered by the input A to D word width.

DYNAMIC RANGE ACHIEVED

Brigham and Cecchini used a rigorous measure of dynamic range in developing their nomogram. This measure gives the absolute worst case estimate of dynamic range, but is not directly related to the dynamic range observed from real signals in the systems in practice. Equally the simple test of resolving a small signal in the presence of a large signal does not directly relate to the worst case dynamic range experience in practice. The true result lies somewhere between the two. Table 5 collects the specifications set out in the previous tables and adds the worst case dynamic range estimates as calculated by Brigham and Cecchini.

A to D bit width	Minimum arithmetic section bit width unconditional scaling	Minimum arithmetic section bit width conditional scaling
8	14	9
10	16	11
12	18	13
14	21	16
16	23	18

Table 4

A to D bit width	Weighting bit width	Sin-Cos bit width	Unconditional arithmetic bit width	Conditional arithmetic bit width	Dynamic range
8	7	7	14	9	38dB
10	9	9	16	11	50dB
12	11	11	18	13	62dB
14	13	13	21	16	74dB
16	15	15	23	18	86dB

Table 5

Table 5 shows clearly that for a given A to D ocnverter bit width, the accuracy required for the 'twiddle factors' within the FFT is always less than the accuracy required for the arithmetic section even when conditional shift algorithms are employed.

THE PDSP FAMILY AND OPTIMUM FFTs

The PDSP Family of Complex DSP Building Blocks implements a Radix 2 DIT Butterfly using just three CMOS devices with an execution time of 50ns per butterfly. This Butterfly processor is optimised for FFTs with a 16 bit data path, and 12 bit coefficients. The optimal arithmetic format is supported by the PDSP16112A Complex Multiplier and two PDSP16318A Complex Accumulators. The Butterfly processor may be configured within Conditional or Unconditional Shift architectures, with all the shifting logic required for either algorithm contained within the PDSP16318 Complex Accumulator.

Optimum FFT configurations using this processor are:

1. 8 or 10 bit A to D based systems using Unconditional Shifting.
2. 12 or 14 bit A to D based systems using Conditional Shifting

Systems that wish to make use of the dynamic range offered by 16 bit A to D converters require more sophisticated shifting algorithms such as Block Floating Point. This algorithm is automatically supported by another PDSP product the PDSP16116. The PDSP16116 Complex Multiplier, together with two PDSP16318 Complex Accumulators and two PDSP1601 ALUs form a five chip 100ns FFT Butterly that supports block floating point arithmetic automatically allowing FFT systems dynamic ranges to exceed 74dB.

EXAMPLE OF PDSP FFT DYNAMIC RANGE

The following plots (Figs. 1-7) demonstrate the actual dynamic range achieved using a GEC Plessey Semiconductors PDSP16112 and PDSP16318 Butterfly processor on a 64 point complex transform using Unconditional Scaling.

The test signal was a combination of a full scale complex sinusoid (samples such that it accumulated within a single frequency bin, see Fig.1) combined with another complex sinusoid 60dB down and sampled such that it was spread across several frequency bins (see Fig.2). The effect of using complex sinusoids is to eliminate the negative frequency components of the input signals, and hence remove the image signal from the FFT result. This test waveform is a more severe test than a simple sinusoid as both real and imaginary components of the FFT inputs are used, giving greater opportunity for word growth and error accumulation. The smaller signal is deliberately chosen to have a non-integer number of cycles within the window so that the energy will be spread to adjacent bins. Eliminating the windowing function from the calculation ensures a worst case result as the effect of this sampling is not reduced.

The plot in Fig.3 illustrates the composite input to the FFT showing both real and imaginary components of the input. This input waveform is (of course) made up from four sinusoids, but the low amplitude signals are not perceptible visually. The plot in Fig.4 illustrates the 'perfect' transform output calculated with 32 bit floating point arithmetic, and the plot in Fig.5 shows the actual transform output obtained from the PDSP FFT system. Brigham and Cecchini predict a worst case dynamic range of 50dB based on the limitation of the 16 bit data path using unconditional scaling.

The actual dynamic range observed between the two sinusoid components is 62dB with approximately 64dB between the large and the noise floor. This value demonstrates the difference between the two methods of estimating dynamic range. The plot in Fig.6 shows a close up view of the difference between the small signal and the noise floor with the large signal component removed. With windowing the small signal would be further separated from the noise floor as its energy would be more concentrated into one frequency bin. The final plot in Fig.7 illustrates the result achieved by a system using 16 bit arithmetic for both data and Sin-Cos values and using the same unconditional scaling. This plot is superimposed upon the result from the PDSP system that uses 12 bit coefficients and unconditional scaling. This composite plot shows clearly that no significant improvement in dynamic range is gained despite the more accurate Sin-Cos values used.

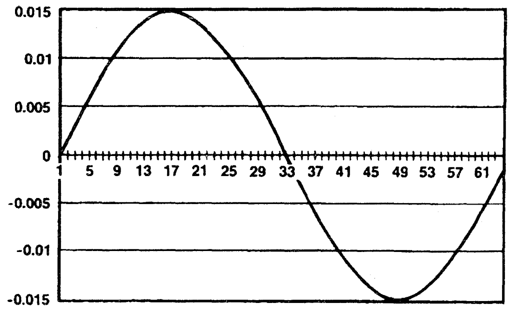

Fig.1 Imaginary portion of large input signal

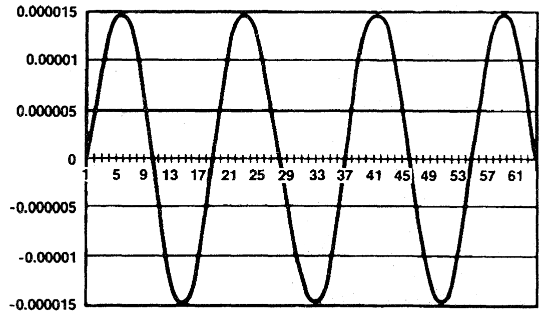

Fig.2 Imaginary portion of small input signal

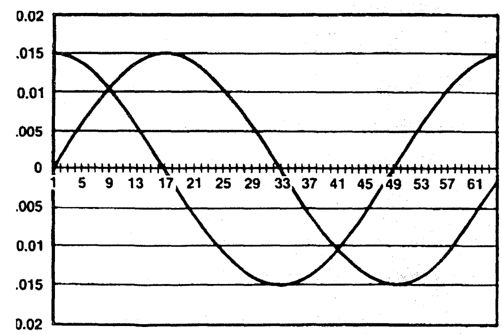

Fig.3 Complex composite input signal

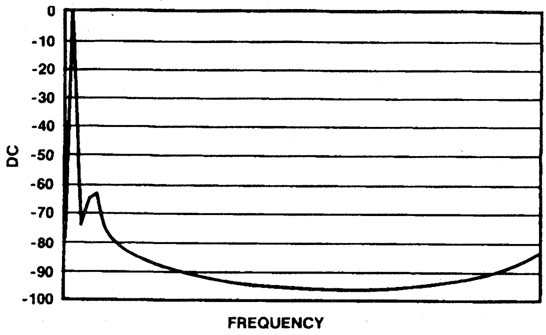

Fig.4 Perfect 64 point FFT output magnitude

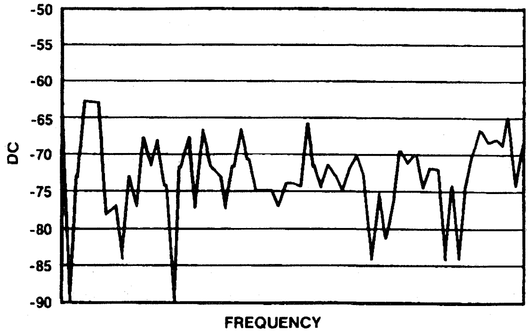

Fig.6 PDSP 64 point FFT detail of noise floor

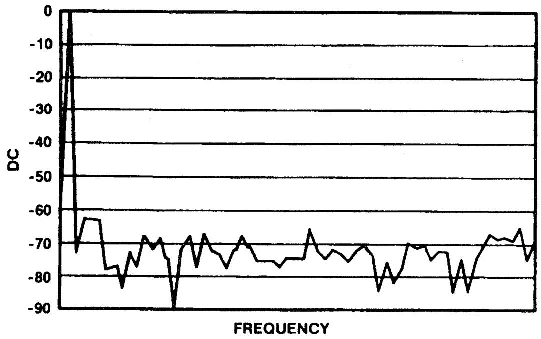

Fig.5 PDSP 64 point FFT output magnitude

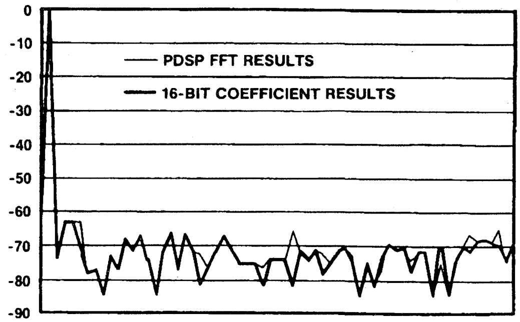

Fig.7 PDSP 64 point FFT output magnitude vs FFT using 16 bit coefficients

REFERENCES

1. A Nomogram for determining FFT system Dynamic Range - E.O. Brigham and L.R. Cecchini E-Systems, Inc. Melphar Division 7700 Arlington Blvd, Falls Church, VA 22046, U.S.A.

A 50ns BUTTERFLY PROCESSOR

GEC Plessey Semiconductors' PDSP16112A Complex Multiplier and PDSP16318A Complex Accumulator have been designed to allow the calculation of Radix 2 Decimation in Time Butterfly operations at very high speeds. One PDSP16112A in conjunction with two PDSP16318As is capable of generating a new result every 50ns whilst dissipating less than 1.5W, giving a 1024 point complex Fast Fourier Transform in just 256µs - almost an order of magnitude faster than the current norm.

Fig 1 shows the Butterfly operation diagrammatically, each Butterfly operation requires one Complex Multiplication, one Complex addition and one Complex subtraction.

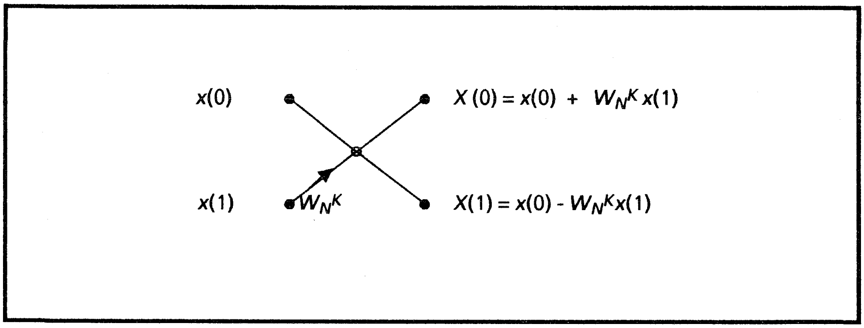

Fig 1 The Butterfly

Fig. 2 shows how the devices are connected to form a hardware Butterfly Processor. The PDSP16112A Complex Multiplier calculates $x(1)W_N^K$, the two PDSP16318As calculate respectively the real and imaginary parts of $x(0) + x(1)W_N^K$ and $x(0) - x(1)W_N^K$. The data format employed is fractional 2's complement, the data entering the PDSP16318s has had the binary point shifted right one place since the number range of the 17 bit output, P, from the PDSP16112 is $-2 \leq P < 2$. Three shift control lines allow control of the overall scaling factor, output overflow is indicated by the OVR flag which is active if the MSB goes above bit 15 of the output.

The PDSP16318s contain delay registers which compensate for the pipeline delay through the PDSP16112 Complex Multiplier in order to simplify addressing. The X(0) and X(1) outputs occur 9 cycles after the corresponding x(0) and x(1) inputs.

Application Note AN47 describes the Butterfly Processor in greater detail, while Application Note AN50 describes a complete FFT Processor built around the 50ns Butterfly Processor.

For those applications where very high speed is required, multiple Butterfly Processors may be used. Ten processors in a pipeline array can execute a 1024 point Complex FFT every 26µs, one Processor handling each column of the transform.

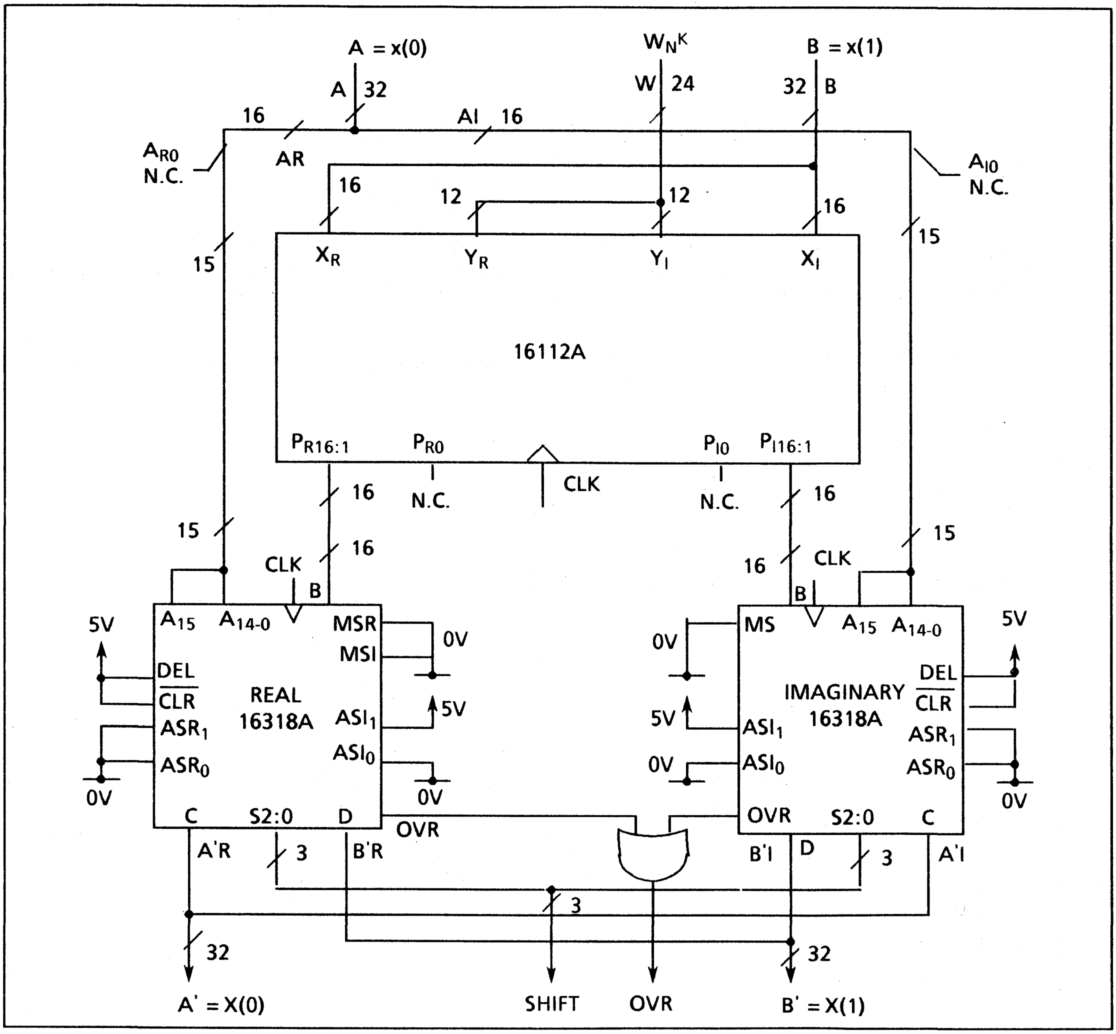

A = x(0)
A
32
W_N^K
W
24
B = x(1)
32
B
16
AR
AI
16
A_R0
N.C.
A_I0
N.C.
16
12
12
16
15
15
X_R
Y_R
Y_I
X_I
16112A
P_R16:1
P_R0
P_I0
P_I16:1
N.C.
CLK
N.C.
16
16
15
16
16
15
CLK
CLK
B
B
5V
A_15
A_14-0
MSR
MSI
0V
0V
MS
A_15
A_14-0
5V
DEL
CLR
REAL
16318A
ASI_1
5V
5V
ASI_1
IMAGINARY
16318A
DEL
CLR
ASR_1
ASR_0
ASI_0
0V
ASI_0
ASR_1
ASR_0
0V
0V
0V
C
S2:0
D
OVR
OVR
S2:0
C
A'R
3
B'R
B'I
D
3
A'I
3
32
32
A' = X(0)
SHIFT
OVR
B' = X(1)

Fig 2

A 50 ns COMPLEX MULTIPLIER/ACCUMULATOR

The applications of complex multiplier/accumulators include digital demodulation, image rejection mixers, adaptive equalization, digital filtering, discrete Fourier Transforms and convolution. The PDSP16112A and PDSP16318A together form a complex multiplier/accumulator capable of generating a new result every 50ns in a system operating on a 20 MHz clock.

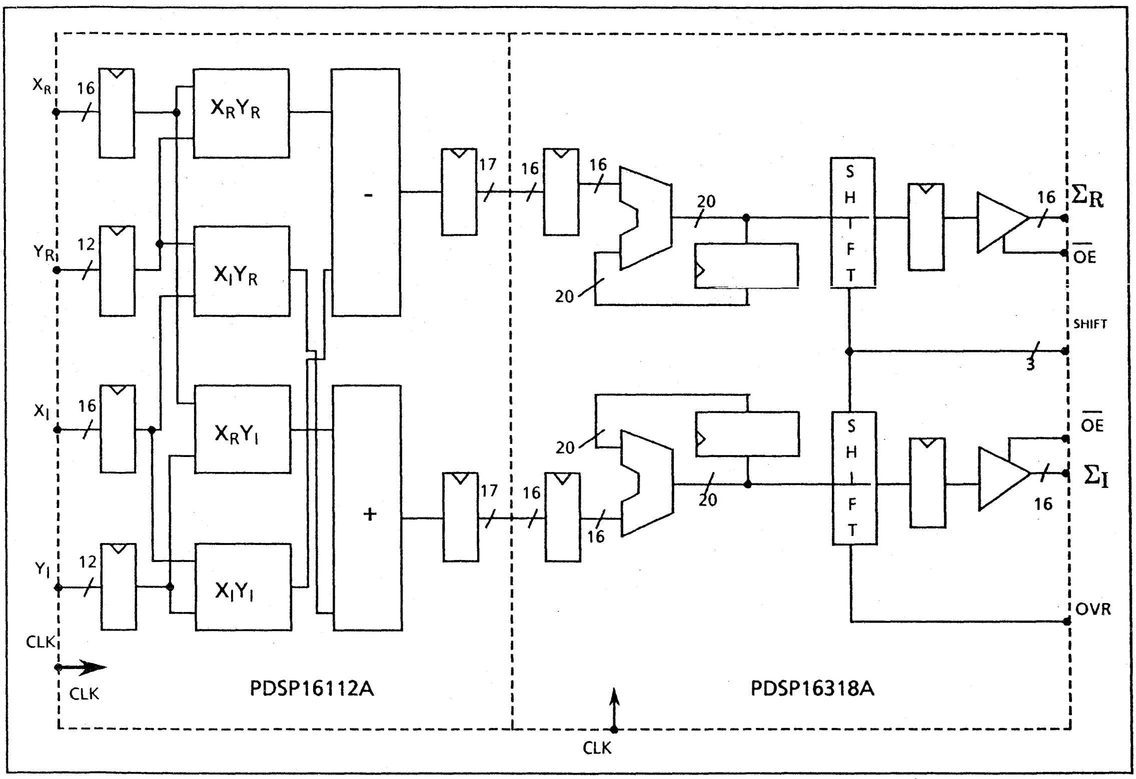

Fig.1

Fig 1 shows the block diagram of the CMAC. The PDSP16112A contains four 16x12 array multipliers, an adder, and a subtractor. The PDSP16318A provides two independent 20-bit-wide add-latch loops for accumulation, followed by output scaling shifters to generate a 16-bit output. If the MSB of the accumulator output goes outside the selected 16-bit output field the overflow flag (OVR) becomes active.

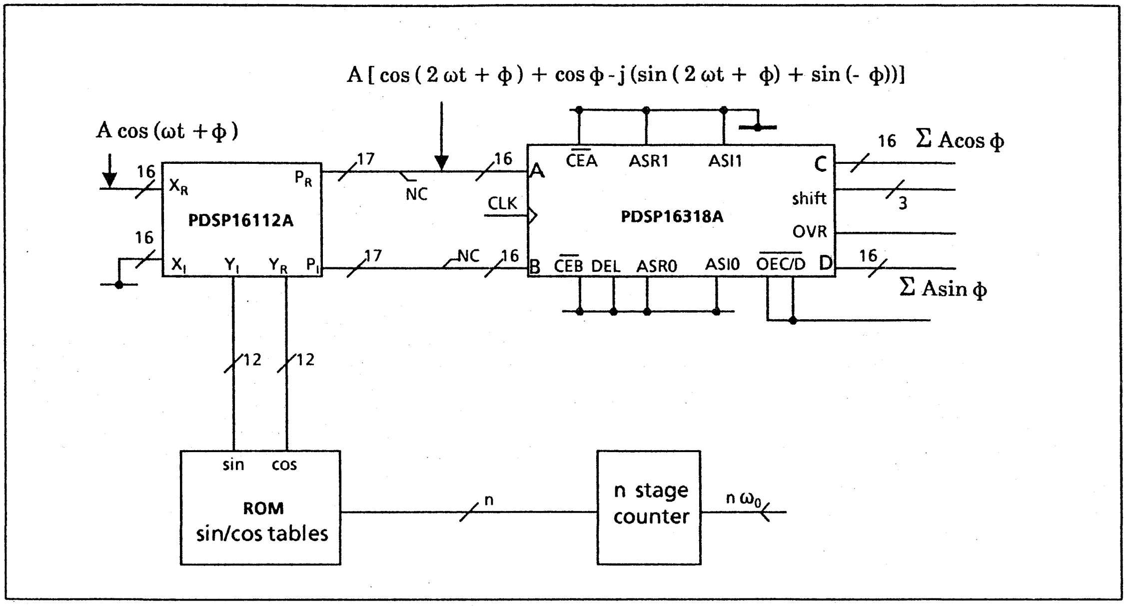

Fig.2 Digital demodulator

Fig 2 illustrates an application of the complex MAC as a digital demodulator. The input signal A cos ($\omega t + \phi$) may be phase and/or amplitude modulated with ϕ or A having a finite number of values which are held constant over the symbol period. The input is multiplied by $2e^{-j\omega t}$ (= cos ωt - j sin ωt) to give an output which, when integrated over the signal period in the accumulator, gives an output Σ A cos ϕ + j Σ A sin ϕ from which the original modulation is easily extracted (the PDSP16330 Pythagoras Processor is an obvious choice for this purpose).

For further details on Complex Signal Processing with these devices see Application Note AN49 'Complex Signal Processing with the PDSP16000 Family', Applications Brief AB04 - 'The Pythagoras Processor', and Application Brief AB10 'FIR Filtering with the PDSP16112 and PDSP16318'.

FIR FILTERING WITH THE PDSP16112 AND PDSP16318

The GEC Plessey Semiconductors' PDSP16112 Complex Multiplier and PDSP16318 Complex Accumulator are designed to perform very fast calculations on complex digital data for many signal processing applications and as such are ideally suited to the area of digital filtering. Digital filters fall into two groups, those with infinite impulse response (IIR) and those with finite impulse response (FIR). The main difference between these two types is that the output from a FIR filter may be calculated from only current and previous inputs, whereas the output from an IIR filter depends on previous output states as well. Although IIR filters may be designed to be more efficient than a FIR for a given order of filter, consideration must always be given to the stability of any design. FIR filters, on the other hand, are inherently stable, are generally easier to design and implement in hardware and have the additional advantage that they may be designed such that they are free of phase distortion (i.e. constant group delay).

The output, y_n, of a FIR may be calculated as the convolution of the input samples with the filter impulse response and can be represented by a difference equation such as:

$$y_n = b_0x_n + b_1x_{n-1} + \ldots\ldots + b_{N-1}x_{n-N+1}$$

or more generally:

$$y(n) = \sum_{k=0}^{N-1} h(k).x(n-k)$$

where coefficients bk represent the N samples of the impulse response, h(k), of the desired filter.

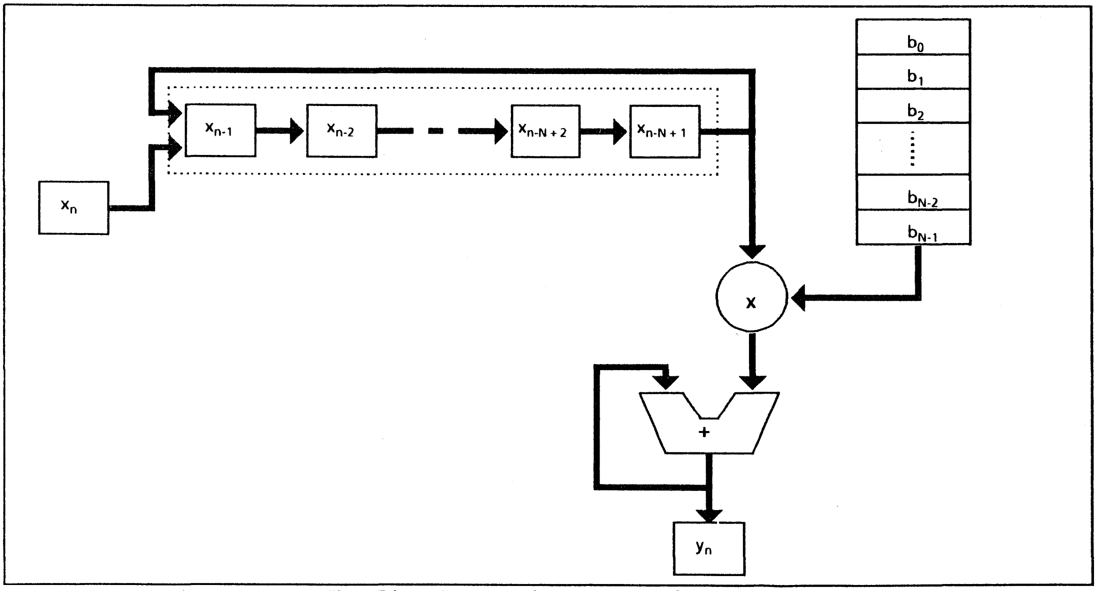

Fig 1 Direct Form Implementation of FIR Filter

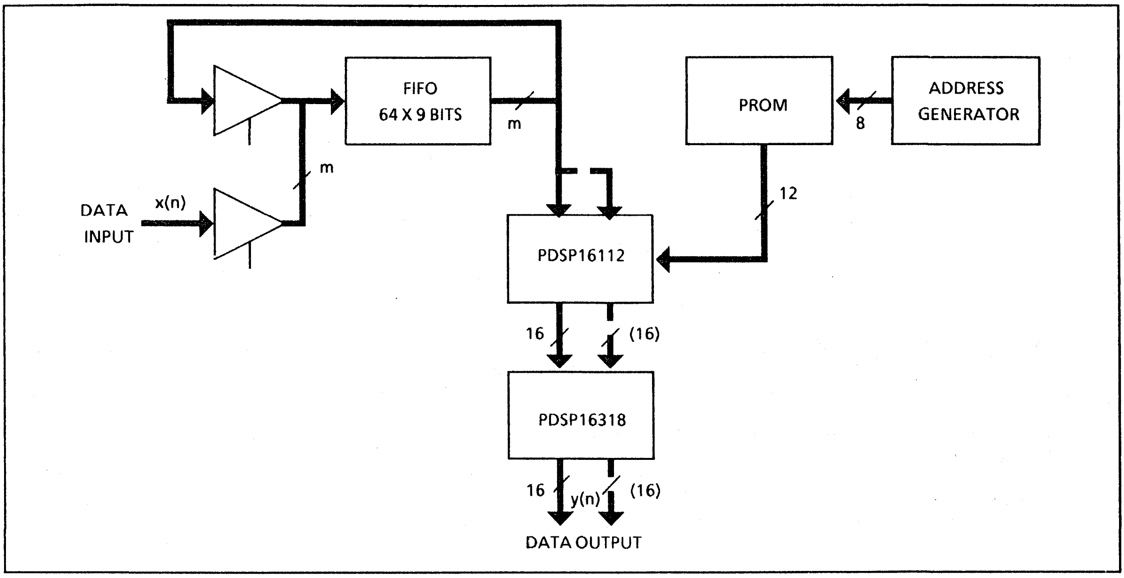

Fig 2 FIR Filter Block Diagram using PDSP16112/16318

A FIR filter may be implemented in a number of ways, but the simplest, with regard to hardware, is termed the direct form, as shown in Fig 1. This consists of just one multiplier and one accumulator plus memory to hold the input data and filter coefficients. The data memory is in the form of a shift register or FIFO, whereas the coefficients may reside in PROM. At each new sample, the data is rotated through the shift register, with the newest sample replacing the oldest within the data store. As each sample is rotated around the registers, it is multiplied by its relevant filter coefficient and the total sum accumulated to calculate the new output value.

The PDSP16112 and PDSP16318 allow filtering of either real or complex data at very high speed, being capable of accumulating a new multiplier result every 50ns. This, for example, enables a very simple 128 tap FIR filter to be implemented with an input signal bandwidth of 78kHz. A block diagram of the circuit for a FIR filter using the PDSP16112 and PDSP16318 is shown in Fig 2. This circuit also shows the use of a FIFO, as the data store, which may be cascaded to produce the depth and width of registers required, and an Address Generator to provide the correct address sequence for the coefficient PROM.

A RADIX 2 BUTTERFLY PROCESSOR

1 INTRODUCTION

The Fast Fourier Transform is a set of algorithms providing short - cuts for the computation of a Discrete Fourier Transform (DFT). FFT techniques can result in calculation times that are shorter than direct DFTs by a factor of 100 or more.

The commonest algorithm used for FFT is the Radix 2 Decimation in Time algorithm. This Application Note illustrates the use of the GEC Plessey Semiconductors PDSP16112 and PDSP16318 in the evaluation of this algorithm.

1.1 THE DFT ALGORITHM

The DFT of a limited sequence of values $\{x(n)\}$, $0 \le n \le (N - 1)$ is defined as:-

$$X(K) = \sum_{n=0}^{(N-1)} x(n) e^{-j(2\pi/N)nK} \quad , K = \{0,1,2,3.......(N-1)\} \qquad 1$$

that is, for N samples of data in the time domain $\{x(n)\}$ we can calculate a sequence of N values representing the signal in the frequency domain $\{X(K)\}$.

The difficulty with this direct evaluation is that $(N-1)^2$ multiplications and $N^2 - N$ additions must be performed. Clearly, large values of N require huge amounts of computation - a 1024 point DFT requires 2,094,081 arithmetic operations on the data.

1.2 THE FFT ALGORITHM

Equation 1 can be re-written as

$$X(K) = \sum_{n=0}^{(N-1)} x(n)\, W_N^{nK}, \quad where W_N = e^{-j2\pi/N} \qquad 2$$

If we split the sequence $\{x(n)\}$ into its even and odd numbered points then :

$$X(K) = \underbrace{\sum_{n=0}^{(N-1)} x(n) W_N^{nK}}_{\text{even only}} + \underbrace{\sum_{n=0}^{(N-1)} x(n)\, W_N^{nK}}_{\text{odd only}} \qquad 3$$

or

$$X(K) = \sum_{n=0}^{(N/2-1)} x(2n)\, W_N^{2nK} + \sum_{n=0}^{(N/2-1)} x(2n+1)\, W_N^{(2n+1)K} \qquad 4$$

Now,

$$W_N^2 = [e^{-j(2\pi/N)}]^2 = e^{-j2\pi/(N/2)} = W_{N/2} \quad 5$$

and $x_{even}(n) = x(2n)$, $x_{odd}(n) = x(2n+1)$

then

$$X(K) = \sum_{n=0}^{(N/2-1)} x_{even}(n) W_{N/2}^{nK} + W_N^K \sum_{n=0}^{(N/2-1)} x_{odd}(n) W_{N/2}^{nK}$$

the original transform has now divided into two separate smaller transforms combined in the following way:

$$X(K) = Xe(K) + W_N^K Xo(K) \quad 6$$

where Xe(K) and Xo(K) are the N/2 point DFTs of $x_{even}(n)$ and $x_{odd}(n)$ respectively.

Each of the sub transforms of equation 6 can be split into two, each of these shorter DFTs can then be divided in turn, and so on . If the number of points $N = 2^r$ where r is an integer, then this decimation can be continued until only 2 point DFTs remain.

The 2 point DFT X(K) , K = 0,1 can be evaluated as

$$X(0) = x(0) + x(1)$$
$$X(1) = x(0) - x(1) \quad 7$$

Note that there are no multiplications involved in a 2 point DFT as the values of $W_{N/2}{}^K$ for K=0,1are ±1. Non trivial multiplications by $w_N K$ are necessary in combining together the sub-DFTs, see Equation 6. These multipliers are often referred to as ' twiddle factors'.

Figure 1 shows the decomposition of an N point DFT into two N/2 point DFTs and twiddles.

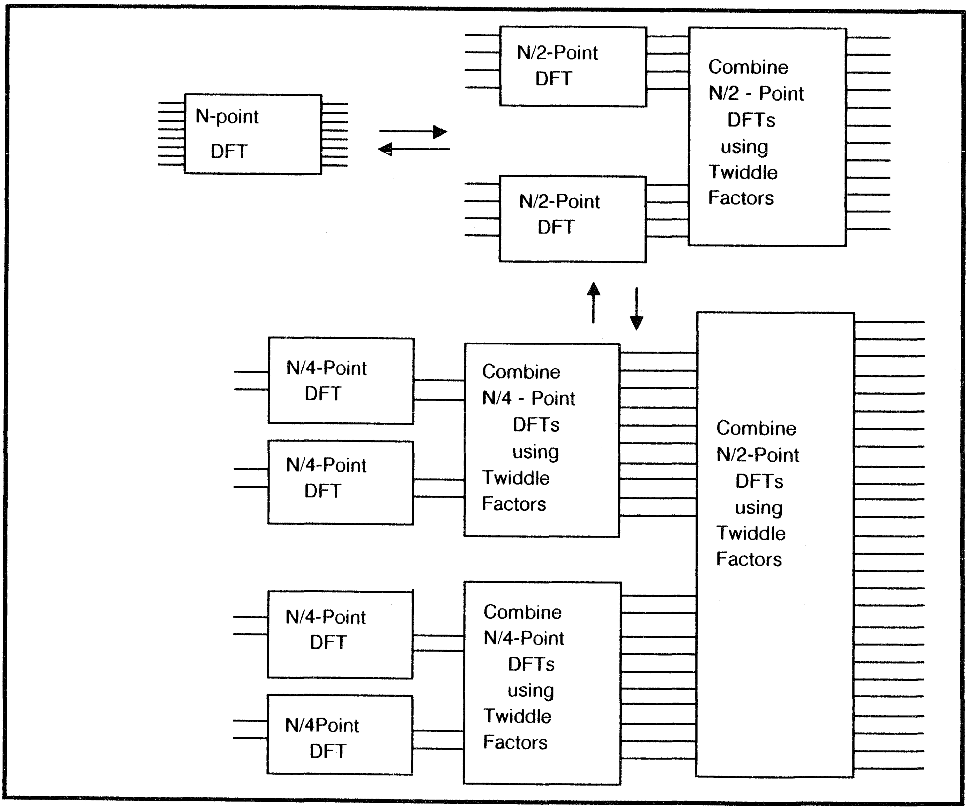

Fig 1 Typical decomposition for radix 2 FFTs

As an example, Figure 2 illustrates the splitting of an 8-point DFT into 2-point DFTs and twiddle factors.Fig 3 shows the arithmetic operations of the combined twiddle and 2-point DFT - the Butterfly

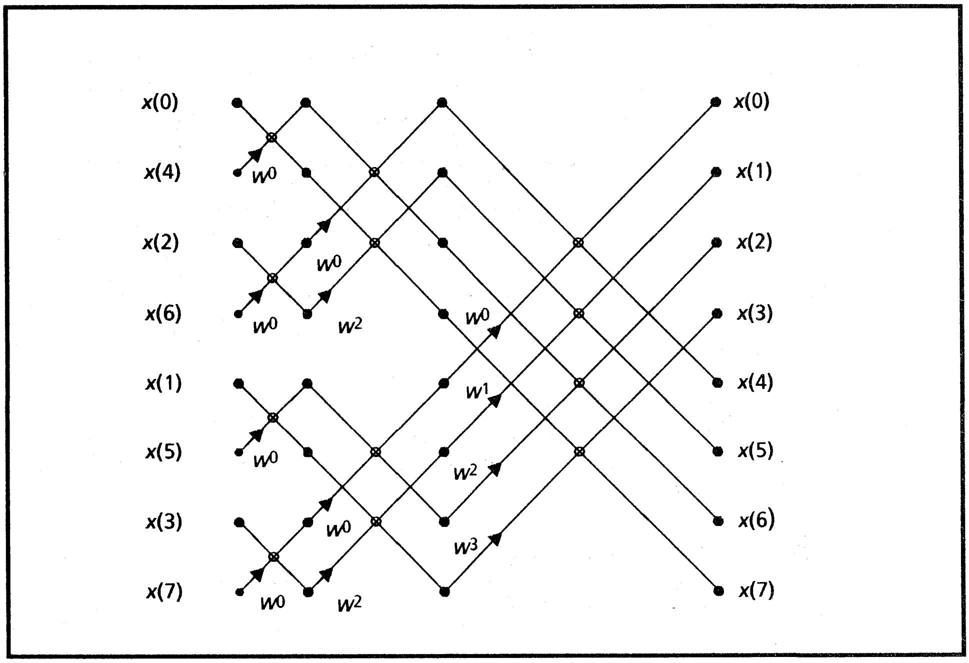

Fig 2 Eight -point FFT obtained by successive splitting into twos.

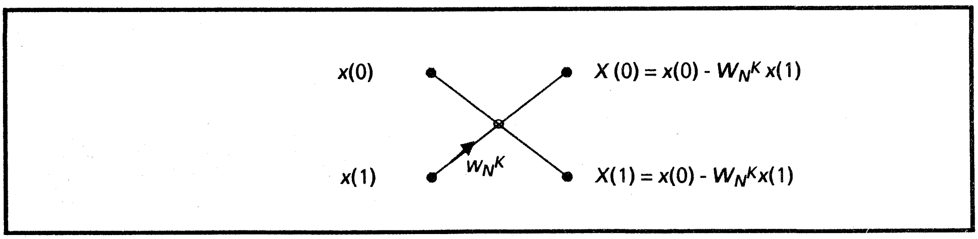

Fig 3 The Butterfly

It can be seen that as a result of this successive splitting the number of Butterfly operations is $^N/_2\log_2N$, each Butterfly requiring only one multiplication($W_N^K x$ can be calculated and stored for use twice). Contrast this with the $(N-1)^2$ multiplications required by the direct DFT:

N	$(N-1)^2$	$^N/_2\log_2N$ (FFT)
16	225	32
128	16129	448
256	65025	1024
1024	1046529	5120

Table 1 Number of multiplications required

1.3 REALISATION

The PDSP16112 and PDSP16318 have been designed to allow the calculation of Radix 2 DIT Butterflies at very high speed. A PDSP16112A in conjunction with a pair of PDSP16318s is capable of calculating a new Butterfly every 50ns.

2. ARCHITECTURE AND ARITHMETIC

Figure 4 shows the basic hardware architecture of the Butterfly processor. The PDSP16112 complex multiplier calculates BW_N^K, the two PDSP16318s calculate respectively the real and imaginary parts of $A+BW_N^K$ and $A-BW_N^K$. The 12 bit input ports on the complex multiplier are used for the twiddle factors, the 16 bit ports are used for data.

2.1 ARITHMETIC CONVENTIONS

The PDSP16112 and PDSP16318 operate on 2's complement fractional data. The form of an n-bit 2's complement fractional number is:

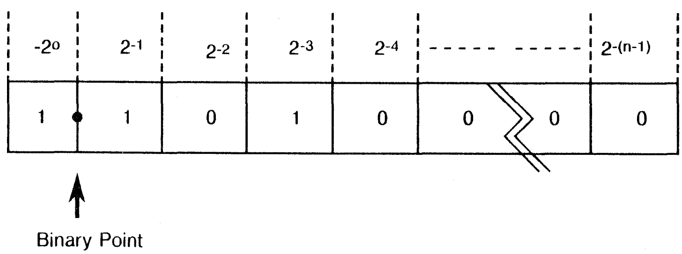

hence 1.1010 is - 0.3750 decimal.

The number range for 2's complement fractional numbers is:

$$-1 \leq n < 1.$$

In the PDSP16112 the 28-bit multiplier results are rounded to 16 bits before entering the adders (see Fig 5). The adder result is a seventeen bit number with two places ahead of the binary point, hence the output P has the range:-

$$-2 \leq P < 2$$

Fig 4 shows how the position of the binary point moves as data proceeds through the processor . The final position results from an unconditional shift at the output of the PDSP16318s, though more complex scaling strategies may be required to optimise dynamic range (see Section 5).

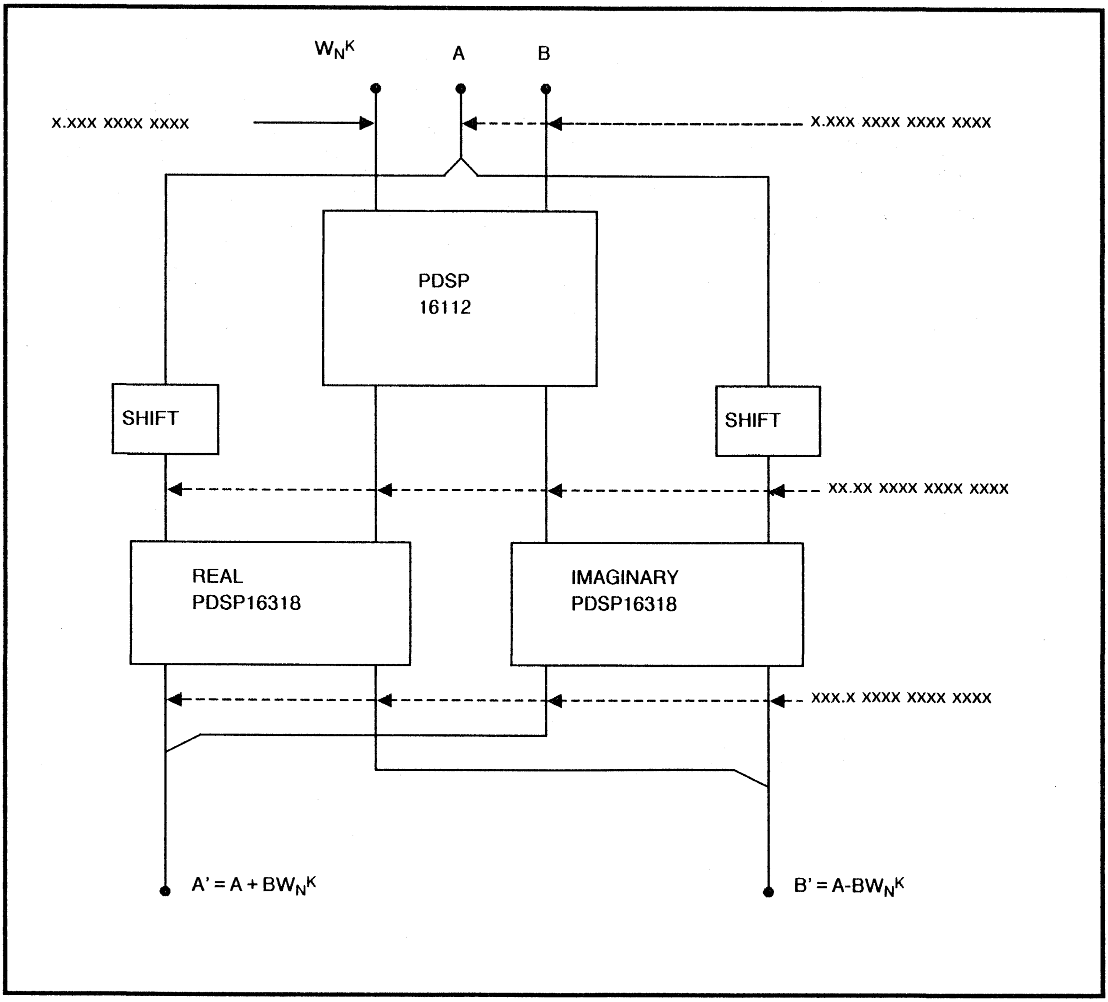

Fig. 4 Basic arrangement of butterfly processor

3. HARDWARE

Figures 5 and 6 show the block diagrams of the PDSP16112 Complex Multiplier and PDSP16318 Complex Accumulator. As can be seen in Fig. 5, there are a total of seven register delays in the data path through the complex multiplier. This pipeline delay would normally cause difficulties with addressing since the A and BW_N^K being presented to the 16318s would be eight cycles apart. This difficulty is avoided by the optional eight cycle delay on the 'A' port of the PDSP16318 which ensures that A and BW_N^K are presented to the adders together.

The structure of the PDSP16318 has been arranged such that a single PDSP16318 can be used with a PDSP16112 to form a complex MAC for filtering or correlation applications, or as in this case, a pair of PDSP16318s are used to handle real and imaginary data, with the internal adders performing complementary operations of A + B and A - B.

4. CIRCUIT DETAILS

Fig. 4 shows the detailed circuit of the Butterfly Processor. The magnitude range at the output of the Complex Multiplier is ± 2 represented as a 17-bit word. The top 16 bits of this word are routed to the B inputs of the Real and Imaginary 16318s, the LSBs being left unconnected. A corresponding shift must be applied to the A inputs to the 16318s in order that data words of the same weighting are presented to the adders. This is achieved by routing the most significant 15 bits of the real and imaginary components of the A' data into the least significant 15 bits of the A inputs on the Complex Accumulators. Input A_{15} must be connected to input A_{14} to provide sign extension for the shift.

Table 2 shows the functions of the various control lines on the 16318: DEL is active so that the 8-register delay is present in the 'A' data path. ASR 1:0 are set so that the 'A' adder gives A + B at its C' output and ASI are set to give A-B at the 'D' output. MS is set low to disable the accumulator feedback paths.

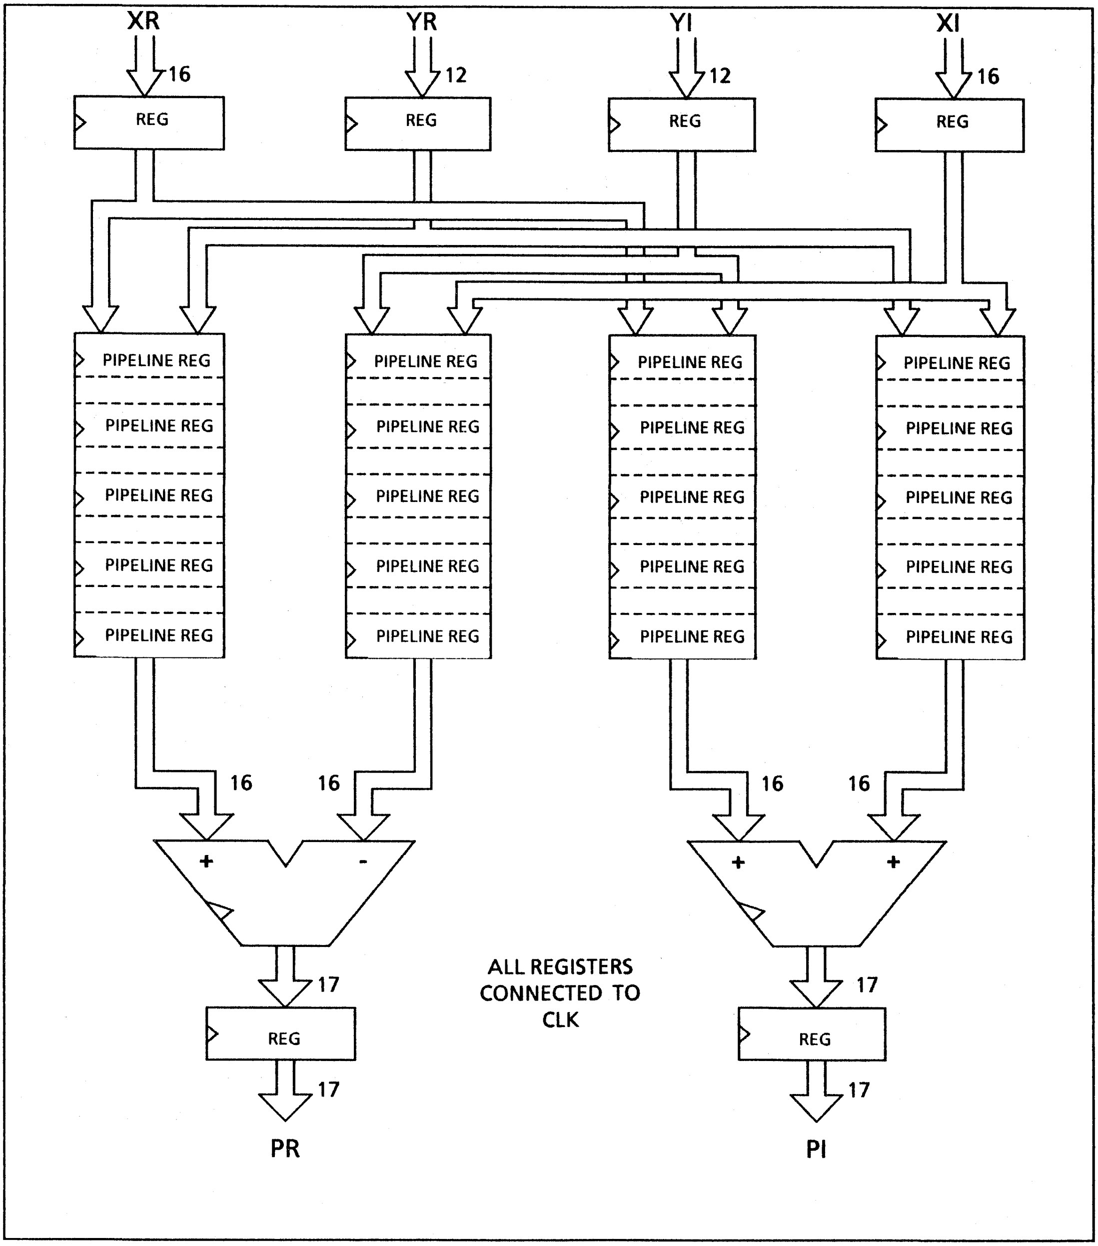

Fig.5 Pipelined multiplier structure of complex multiplier

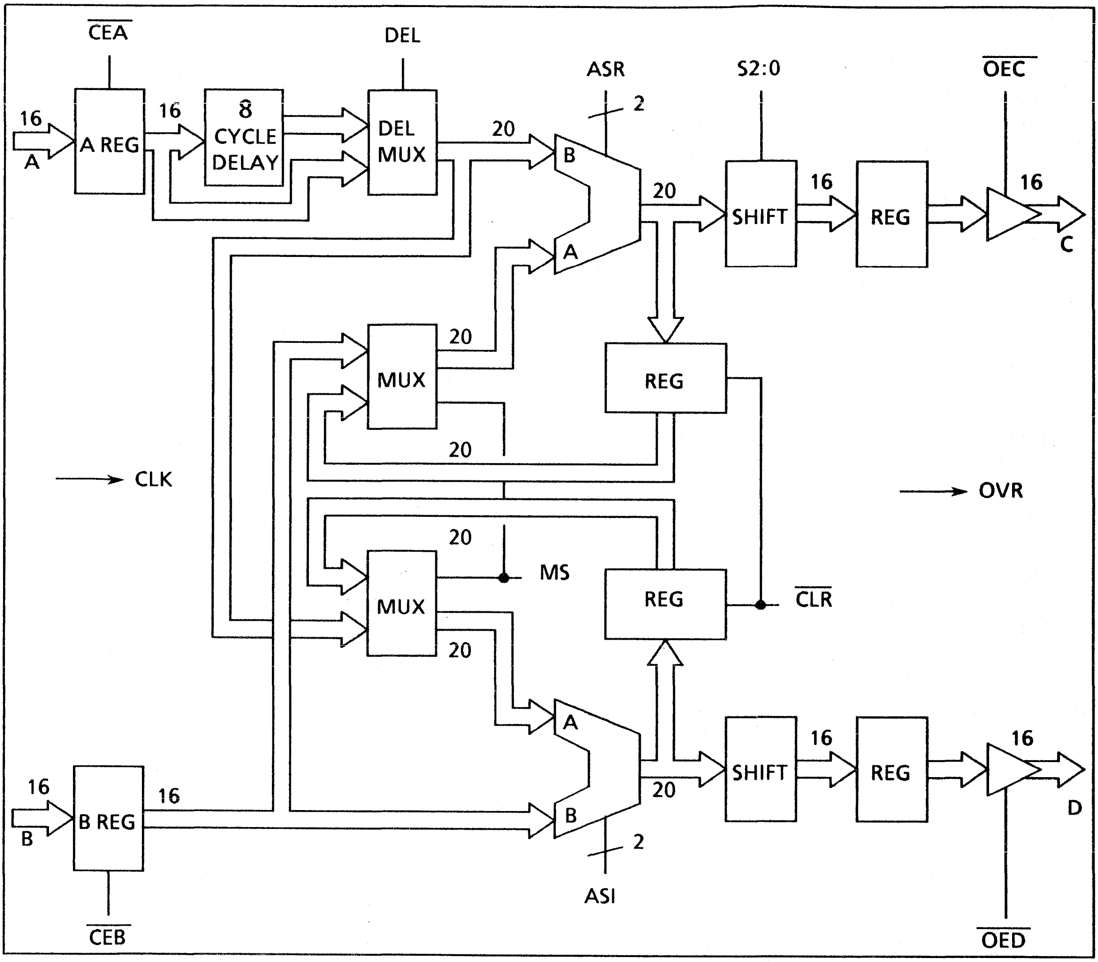

Fig.6 Block Diagram of Complex Accumulator

ASR or ASI ASX1	ASXO	ALU Function
0	0	A + B
0	1	A
1	0	A-B
1	1	B-A

DEL	Delay Mux Control
0	A port input
1	Delayed A port input

MS	Real and Imag' Mux Control
0	B port input / Del mux output
1	C accumulator / D accumulator

Table 2

5. WORD GROWTH

Word growth in the accumulators can be accommodated by the use of the Output Shifters. Since overflows can occur independently in either the real or imaginary 16318s, the two OVR overflow flags are ORed to generate a composite overflow warning. Table 3 shows the operation of the output shifter, note than an overflow will be flagged if a shift is selected which results in the MSB of the data being outside the output range. The Shift Control lines of the Real and Imaginary Complex Accumulators must be connected in parallel, otherwise the real and imaginary data components will have different weightings, causing invalid results on subsequent operations.

The simplest scaling scenario is to select the least significant seventeen bits of the adder result (shift code 011 in Table 3). This has the effect of adding another fixed shift of one place, and will prevent any possibility of an overflow in the adder output. This unconditional shifting will produce acceptable results most of the time, though in some situations where there is significant wideband noise the signal data may be scaled down by an excessive amount. Each pass through the Butterfly Processor will cause data to be scaled down by 2 bits.

Another method of scaling is to apply no shift at all at the output of the accumulator. This option is vulnerable to an overflow occuring within the accumulator. If overflow occurs an incorrect result is output from the Complex Accumulator, however the overflow flag of the Complex Accumuator flags the invalid data which may then be corrected by external circuitry, or discarded. Further overflow risk may be minimised by scaling down the input data before passing it to the FFT processor. This will globally reduce the data magnitude and hence reduce the probability of an overflow occurring.

A third solution is a compromise between the first two. A large FFT involves several passes through the data, for example a 1K Complex FFT requires ten passes. Situations that require scaling after every pass are rare, as are situations that require no scaling at all. The compromise solution is to select different shifts of the output data from the Complex Accumulator on alternate cycles, so that on the first, third, fifth etc passes data is not shifted on the way out of the Complex Accumulator, and on the second, fourth, sixth etc data is shifted down by one place. Experimentation with real data as opposed to test signals will reveal the optimum solution for each application, the overflow flags from the Complex Accumulators warning when overflows have occurred.

There are even scenarios when the scaling introduced by the Complex Multiplier is too much and output data from the Complex Accumulators needs to be scaled up. In these situations upward scaling of the Complex Accumulator outputs can be selected, indeed an adaptive scaling system could be constructed whereby the largest output from each FFT is monitored and the scaling is adjusted up or down accordingly.

S2:0			Adder Result																			
S2	S1	S0	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0				
0	0	1		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0			
0	1	0			15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
0	1	1				15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
1	0	0					15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	1						15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1	1	0							15	14	13	12	11	10	9	8	7	6	5	4	3	2
1	1	1								15	14	13	12	11	10	9	8	7	6	5	4	3

Table 3

NOTE: This table shows the portion of the adder result passed to the D15:0 and C15:0 outputs. Where fewer than 16 adder bits are selected, the ouput data is padded with zeros

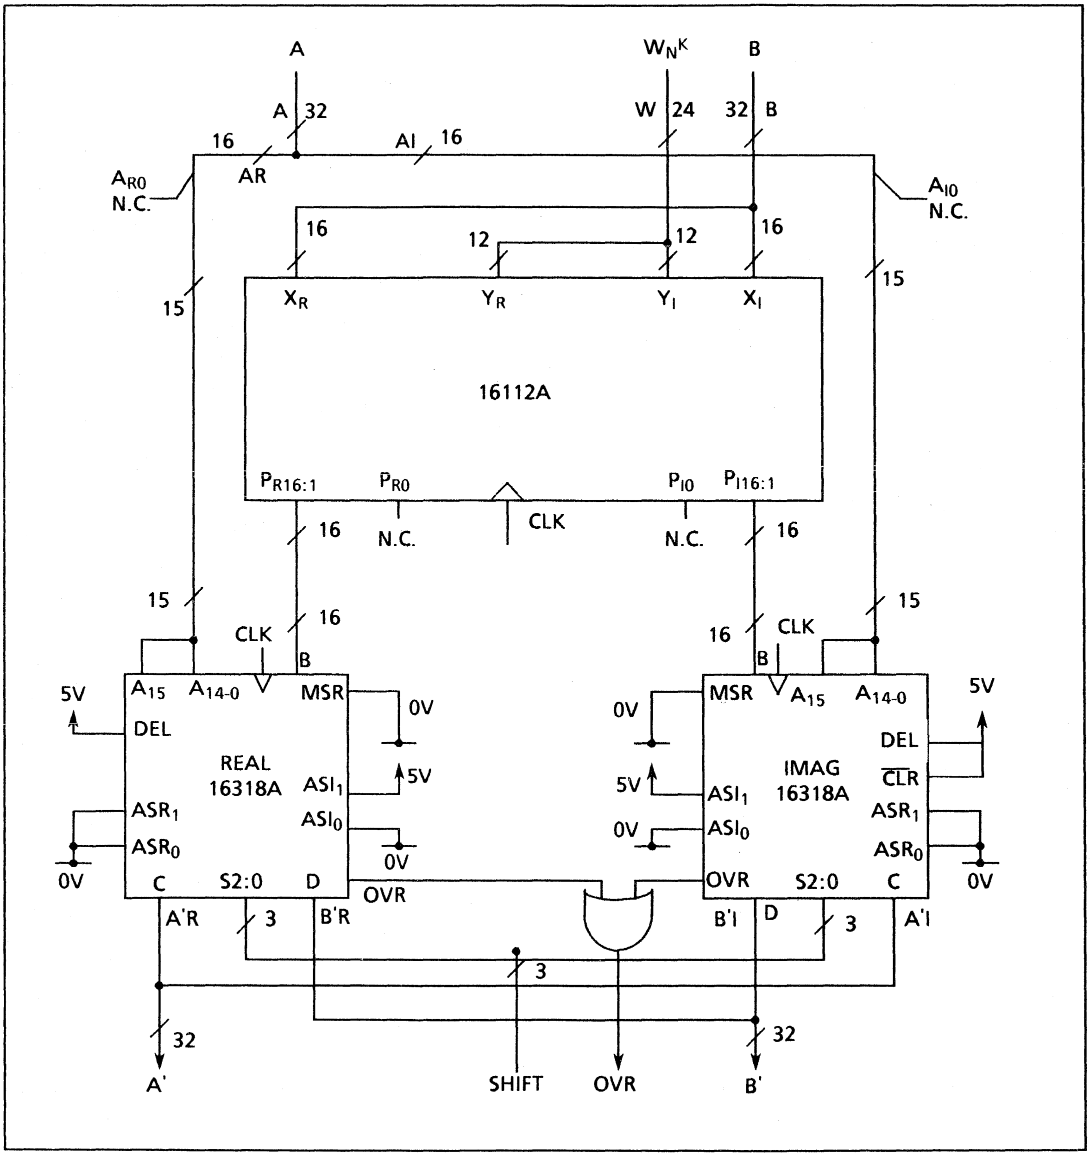
A
W_N^K
B
A 32
W 24
32 B
16
AR
AI 16
A_{R0} N.C.
A_{I0} N.C.
16
12
12
16
15
15
X_R
Y_R
Y_I
X_I
16112A
$P_{R16:1}$
P_{R0}
P_{I0}
$P_{I16:1}$
16
N.C.
CLK
N.C.
16
15
16
CLK
16
CLK
15
B
B
5V
A_{15}
$A_{14\text{-}0}$
MSR
0V
0V
MSR
A_{15}
$A_{14\text{-}0}$
5V
DEL
REAL 16318A
ASI_1
5V
5V
ASI_1
IMAG 16318A
DEL
$\overline{CLR}$
ASR_1
ASI_0
0V
ASI_0
ASR_1
ASR_0
0V
ASR_0
0V
C
S2:0
D
OVR
OVR
S2:0
C
0V
A'R
3
B'R
B'I
D
3
A'I
3
32
32
A'
SHIFT
OVR
B'

Fig.7

6. PROCESSOR TIMING

The circuit of Fig. 7 will operate with clock frequencies up to 20MHz if /A version parts used (10MHz for normal grade devices). The total latency is 9 cycles from inputs to outputs. The I/O timing of the processor is given in Table 4 (/A version devices, normal figures in brackets)

V_{cc} = 5v ± 10%, Tamb -40°C to 85°C

PARAMETER	VALUE MIN	MAX	UNITS	CONDITIONS
Input to CLK set up time	20(30)		ns	
CLK to Input hold time	5(8)		ns	
CLK to Output delay time		25(40)	ns	2 x LSTTL + 20pF
CLK MARK/SPACE ratio	40	60	%	

Table 4

A HIGH RESOLUTION FFT PROCESSOR USING THE PDSP16116/A

The PDSP16116A has been designed with an integral Block Floating Point system which can be used, in conjunction with other GEC Plessey Semiconductors' PDSP parts, to process FFTs with a combination of speed and accuracy previously unobtainable. All the functionality of this BFP system is contained within the PDSP parts, which are designed to interface easily to achieve a powerful FFT solution.

A butterfly processor based on the 20MHz PDSP16116A will allow the following FFT benchmarks:

1024 point complex radix-2 transform in 259us
512 point complex radix-2 transform in 118us
256 point complex radix-2 transform in 53us

This compares favourably with the current industry standard benchmark of around 2ms for a 1024 point complex FFT, but if speed is all important for a particular application, then the GEC Plessey Semiconductors PDSP16112/A 16x12 Complex Multiplier can double the PDSP16116/A performance with up to 70dB of dynamic range.

The FFT Algorithm

The Fast Fourier Transform is essentially a computationally efficient algorithm for extracting spectral information from signal waveforms, which may be in real time or recorded form (i.e. a transformation from the time domain to the frequency domain). It is often used to dramatic effect in a growing range of applications including radar and sonar processing, speech recognition and image processing. It is no less accurate than the related Discrete Fourier Transform (DFT), but it enjoys a vastly improved performance due to the 'divide and conquer' approach of its algorithm.

There are several variations of the FFT algorithm, each with their own merits. For high throughput, hardware implemented solutions, a variant of the Radix-2 Decimation-in-Time algorithm is most suitable. The 'Constant Geometry' algorithm (Fig.1) is easier to implement whereas the 'In-Place' algorithm (Fig.2) halves the amount of memory required.

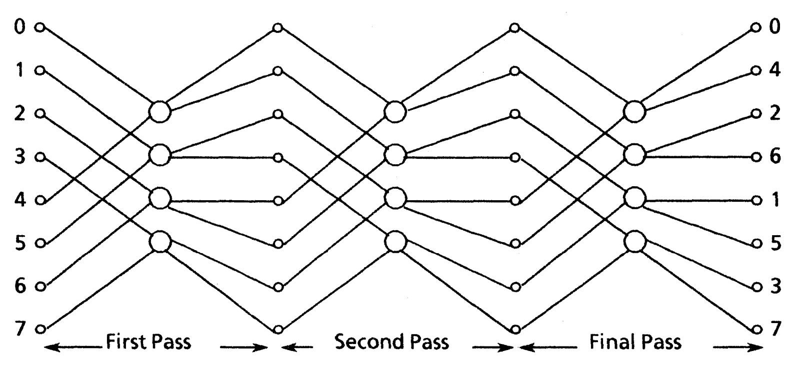

Fig 1 8 point constant geometry DIT radix 2 algorithm with normally ordered inputs and bit-reversed outputs

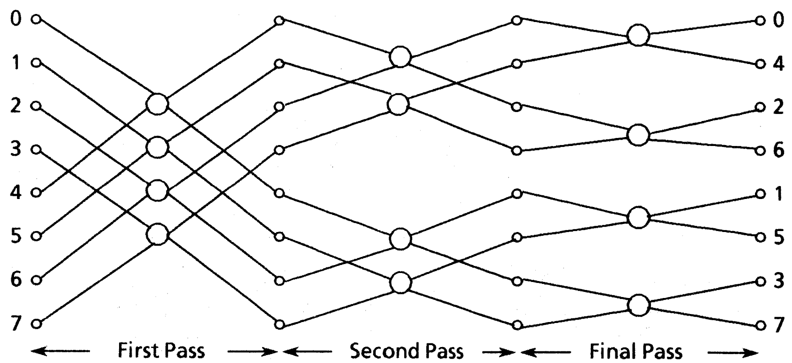

Fig 2 8 point in- place DIT radix 2 algorithm with normally ordered inputs and bit-reversed outputs

Both these variations are split vertically into a number of 'passes' (log_2N passes for an N-point transform), each pass consisting of N/2 'butterfly' operations:

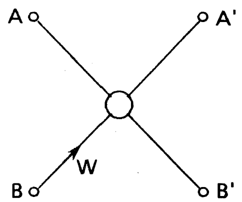

$$A' = A + B.W$$
$$B' = A - B.W$$

W is the complex coefficient and A and B are, for the first pass, the sampled data and then, in the second and subsequent passes, the values of A' and B' from the previous pass. The results of the FFT are the values of A' and B' from the butterflies of the final pass. In order to be compatible with previous FFT results, all points must be normalised to a universal format. These final complex number values (cartesian co-ordinates) may then be converted into magnitude and phase components (polar co-ordinates).

Defeating the Wordgrowth Problem

One of the most difficult problems to overcome when implementing an FFT algorithm in fixed point arithmetic is that of wordgrowth. The power of the PDSP16116s BFP system lies in its flexible and effective response to this problem. Before looking into the operation of this BFP system, the wordgrowth problem and some of the other solutions available are explained.

FFTs are implemented by means of successive multiplications and additions. Each time data is processed by an ALU (i.e. twice in each butterfly) there is the possibility of wordgrowth occuring: i.e. when two 16 bit words are added, they may produce a sum of 17 bits. The safe way to deal with this is to always pick the 16MSBs of the result. However, this will cause sign extension, i.e.repetition of the sign bit in the MSBs of the data. These two cases are illustrated in the examples below.

Wordgrowth occurs:	Wordgrowth does not occur:
0101110101110100	0101010111011010
+ 0110100110101001	+ 0010100010101110
01100011100011101	0011111101000100
	↑ sign extension

Sign extension can cause severe problems when the next multiplication occurs, as it is likely to lead to a product with a further extended sign bit. For example:

```
                           1 bit sign extension
                  0011010001010110
                * 0010001101000101
----------------------------------
00000111001101011101110100101110
    4 bit sign extension
```

After a few passes of the FFT, there is a danger that the data could become all sign bits and no information - not much use to anyone. The common alternative to this approach is to pick the 16 LSBs from the ALUs and hope that no wordgrowth occurs, as this will then lead to overflow. If overflow is flagged during the course of an FFT, then the calculation must be aborted. The input data is then scaled down and the calculation repeated. The hit and miss nature of this approach can be avoided by automatically scaling down the inputs and accepting the resulting penalty in accuracy. A 'conditional shift' system offers some degree of flexibility. Here, the 16 LSBs are selected from the second ALU in the butterfly hardware if no overflow occurs in any butterfly during that pass.

The PDSP16116 offers a superior solution to the problem by employing an intelligent control system which can monitor data magnitudes during the course of the FFT and adjust them as necessary so as to keep extended sign bits to a minimum, whilst eliminating the possibility of overflow. In fact, this system can not only deal with wordgrowth problems as they occur, but can also adjust underscaled input data in anticipation of these problems to ensure that a valid result is obtained at the end of the calculation.

A comparison of the data formats provided by each of the methods detailed above will clarify their differences. Given input data of the format:

X.XXX... (note the position of the binary point)

The UNCONDITIONAL SHIFT implementation will output all data at the end of a pass in the format:

XXX.X...

regardless of whether the data has increased in magnitude or not.

The CONDITIONAL SHIFT implementation will either output ALL data in the format :

XX.XX...

if the maximum wordgrowth was one bit in any butterfly; or, if two bits of wordgrowth occured in any butterfly, then ALL data will be output in the format :

XXX.X...

The BFP implementation can output EACH butterfly result in ANY of the following formats, according to the data magnitude:

If data is underscaled	.XXXX...
If no worgrowth occurs	X.XXX...
If wordgrowth occurs once	XX.XX...
If wordgrowth occurs twice	XXX.X...

The adaptability of the BFP system is clearly illustrated and it is this adaptability which allows the BFP system to defeat the wordgrowth problem.

How the BFP System Operates

A block floating point system is essentially an ordinary integer arithmetic system with some additional logic, the object of which is to lend the system some of the enormous dynamic range afforded by a true floating point system without suffering the corresponding loss in performance.

The initial data used by the FFT should all have the same binary weighting, i.e. the binary point should occupy the same position in every data word. This is normal in integer arithmetic. However, during the course of the FFT, a variety of weightings are used in the data words to increase the dynamic range available. This situation is similar to that within a true floating point system, though the range of numbers representable is more limited.

In the BFP system used in the PDSP16116, there are, within any one pass of the FFT, four possible positions of the binary point within the integer words. To record the position of its binary point, each word has a 2-bit word tag associated with it. By way of example, in a particular pass we may have the following four positions of binary point available, each denoted by a certain value of word tag:

XX.XXXXXXXXXXXXXX	word tag = 00
XXX.XXXXXXXXXXXXX	word tag = 01
XXXX.XXXXXXXXXXXX	word tag = 10
XXXXX.XXXXXXXXXXX	word tag = 11

At the end of each constituent pass of the FFT, the positions of the binary point supported may change to reflect the trend of data increases or decreases in magnitude. Hence, in the pass following that of the above example, the four positions of binary point supported may change to:

XXXX.XXXXXXXXXXXX	word tag = 00
XXXXX.XXXXXXXXXXX	word tag = 01
XXXXXX.XXXXXXXXXX	word tag = 10
XXXXXXX.XXXXXXXXX	word tag = 11

This variation in the range of binary points supported from pass to pass (i.e. the movement of the binary point relative to its position in the original data) is recorded in the Global Weighting Register (GWR). At the end of the final pass, the distance that the binary point has moved since the start of the FFT can be obtained by modifying the GWR according to the value of WTOUT of a particular word, as shown below:

WTOUT1:0	ADJUSTMENT TO GWR
00	SUBTRACT 1
01	NO ADJUSTMENT
10	ADD 1
11	ADD 2

For example, if the original data format was:

X.XXXXXXXXXXXXXXX

then, if the GWR = 01001 and the WTOUT = 10 for a particular word, the binary point has moved 10 places to the right of its original position and will be situated as shown below:

XXXXXXXXXXX.XXXXX

Using the GWR with Large FFTs

The Global Weighting Register represents the movement of the binary point in two's complement notation in a 5-bit field. An examination of FFT theory and the operation of the BFP system shows that, for an N-point transform, GWR will not exceed $(2+\log_2 N)$. This means that GWR can handle transforms as large as 8K by representing the movement of the binary point as a two's complement number. However, GWR can be used for much larger transforms by noting that GWR will never drop below -8, since with this degree of left shift, the rounding noise is amplified to fill the whole 16-bit data word. This fact allows GWR to be extended and represented as a six bit value simply by ANDing the two most significant bits to produce a new sign bit (Fig 3). This 6-bit field allows GWR to handle up to a 2097K transform.

Value of GWR	Decimal Equiv.	Meaning
00000	0	Binary point has not moved
00001	+1	Binary point has moved 1 place to the right
00010	+2	2
00011	+3	3
00100	+4	4
00101	+5	5
00110	+6	6
00111	+7	7
01000	+8	8
01001	+9	9
01010	+10	10
01011	+11	11
01100	+12	12
01101	+13	13
01110	+14	14
01111	+15	15
10000 *	+16	16
10001 *	+17	17
10010 *	+18	18
10011 *	+19	19
10100 *	+20	20
10101 *	+21	21
10110 *	+22	22
10111 *	+23	23
11000	-8	Binary point has moved 8 places to the left
11001	-7	7
11010	-6	6
11011	-5	5
11100	-4	4
11101	-3	3
11110	-2	2
11111	-1	1

* not in two's complement format

Table 1 GWR values and meanings

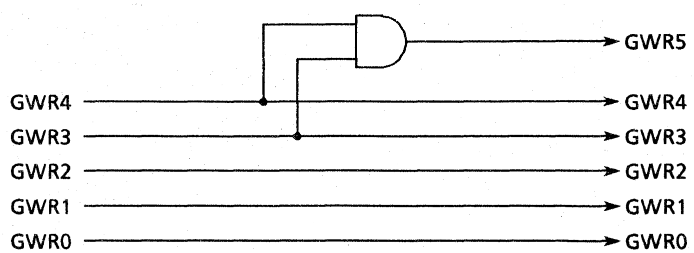

Fig 3 Extending GWR to 6 bits

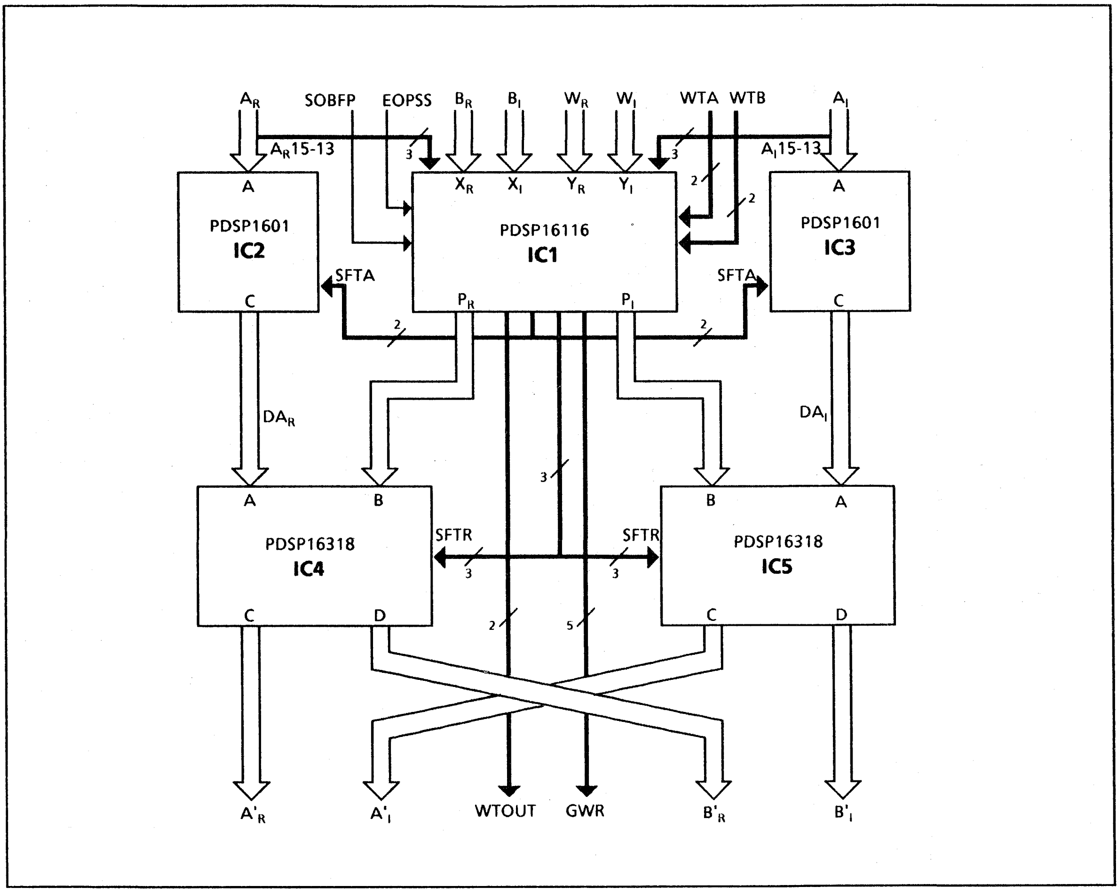

Fig. 4 Block Floating Point FFT butterfly

Construction of an FFT Butterfly Processor

As described earlier, the calculations A' = A + BW and B' = A-BW, forming a 'butterfly operation' must be carried out repeatedly in the course of an FFT. Fig.4 shows how a butterfly processor may be constructed using a single PDSP16116 in combination with two GEC Plessey Semiconductors PDSP16318s and two GEC Plessey Semiconductors PDSP1601s. The PDSP1601s are used to match the pipeline delay and shifting operations of the PDSP16116 to the datapath of the A word. The PDSP16318s are used to perform the complex addition and subtraction of the butterfly operation. Fig.5 details the underlying architecture of the processor.

A detailed list of the various connections required to combine these five chips into a butterfly processor appears in the Appendix. I/O connections are not specified as there are a number of I/O options that allow the butterfly processor to be interfaced with the rest of an FFT system.

A point to note is the hard-wired 1-bit right shift in the A-word data paths between the PDSP1601 outputs and the PDSP16318 inputs. This is to keep the A-word data format the same as the PDSP16116 output data format so that the two words may be added within the PDSP16318. The PDSP1601 applies a shift of 0 to 3 places to the right whereas data is output from the PDSP16116 with the binary point shifted from 1 to 4 places to the right. Hence an extra right shift of one place needs to be inserted in the PDSP1601 data path to keep the data formats compatible at the inputs to the PDSP16318 (data words must have their binary points in the same places before being added).

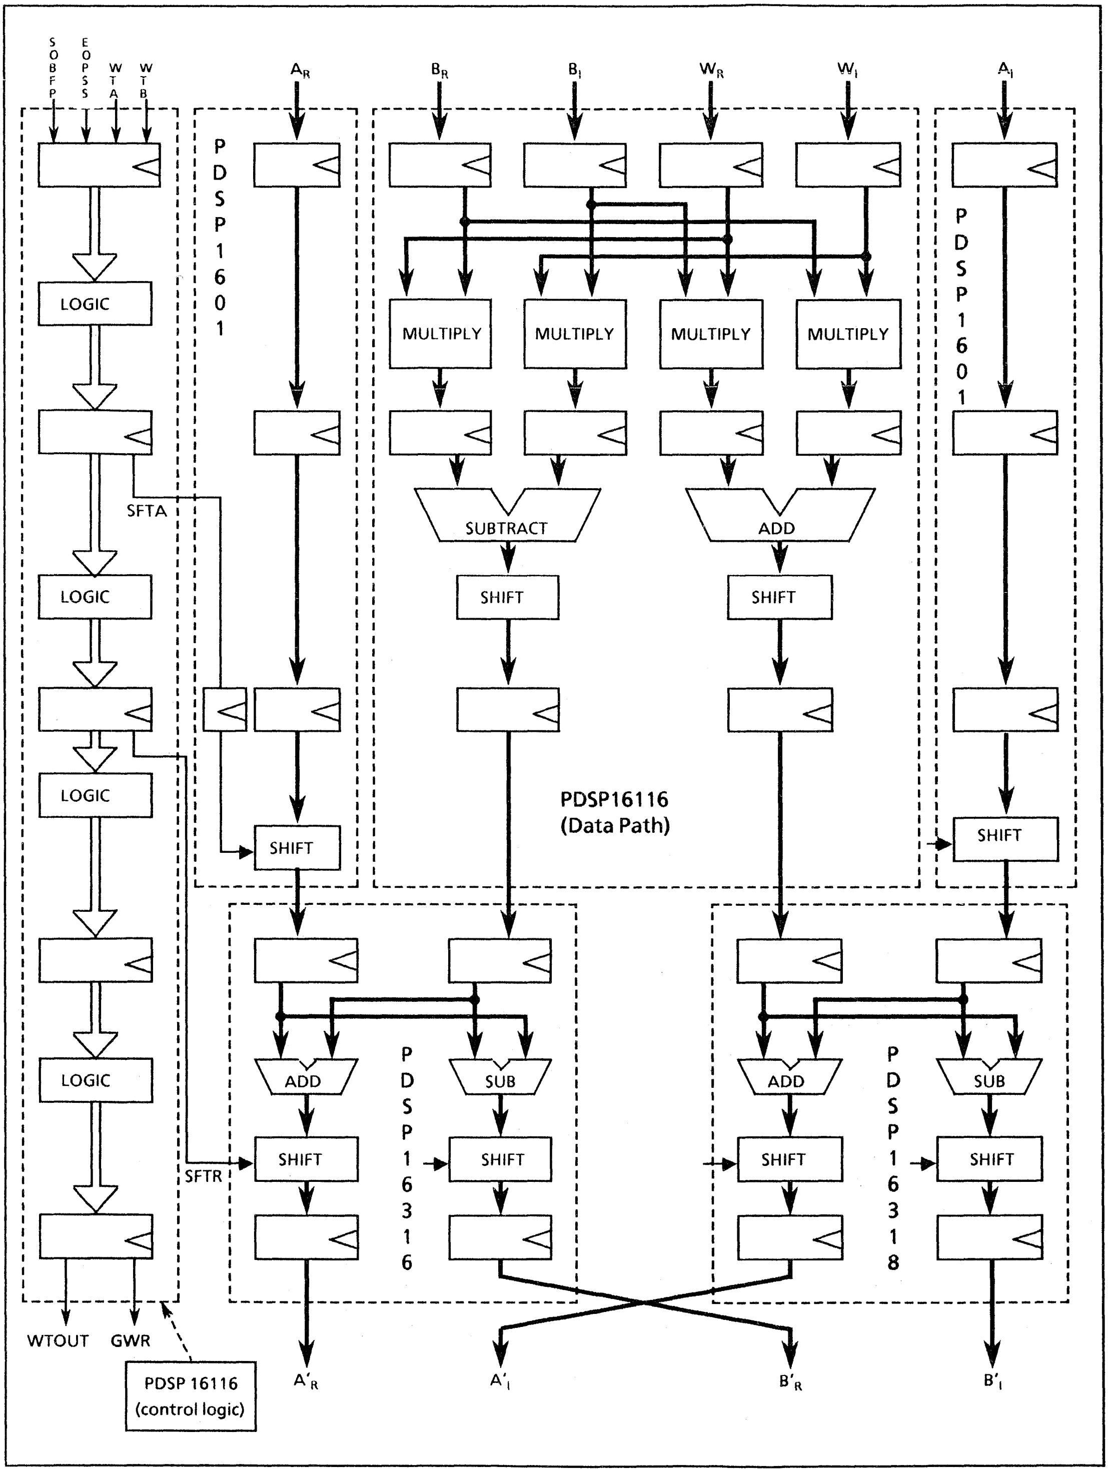

Fig 5 BFP Butterfly Detail

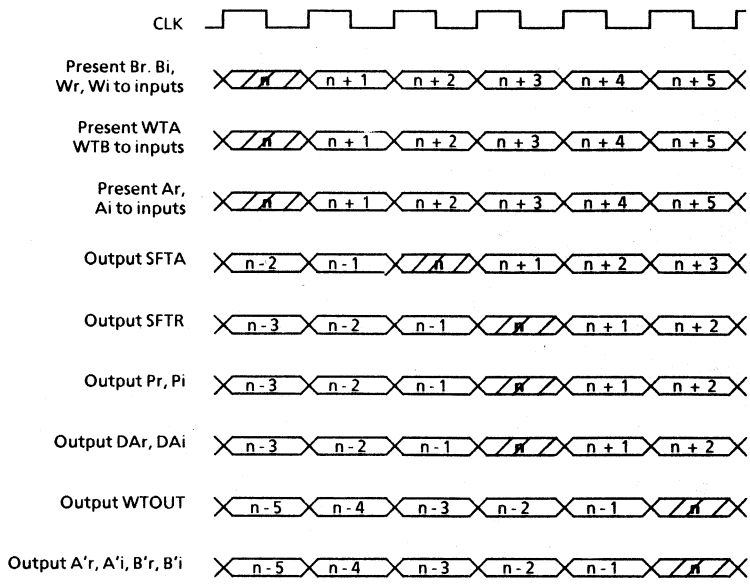

Fig 6 Data and Control Timing in the Butterfly

The Butterfly Operation

A new butterfly operation is commenced each cycle, requiring a new set of data for A, B, W, WTA and WTB. Five cycles later, the corresponding results A' and B' are produced along with their associated WTOUT. In between, the signals SFTA and SFTR are produced and acted upon by the shifters in the PDSP1601 and PDSP16318. The timing of the data and control signals is shown in Fig. 6.

The results (A' and B') of each butterfly calculation in a pass must be stored away to be used later as the input data (A and B) in the next pass. In every pass, each result must be stored together with its associated word tag, WTOUT. Although WTOUT is common to both A' and B', it must be stored separately with each word as the words are used on different cycles during the next pass. At the inputs, the word tag associated with the A word is known as WTA and the word tag associated with the B word is known as WTB. Hence the WTOUTs from one pass will become the WTAs and WTBs for the following pass. It should be noted that the first pass is unique in that word tags need not be input into the butterfly as all data must initially have the same weighting. Therefore, during the first pass alone, the inputs WTA and WTB are ignored.

Control of the FFT

To enable the block floating point hardware to keep track of the data, the following signals are provided:

SOBFP - start of the FFT
EOPSS - end of current pass

These inform the PDSP16116 when an FFT is starting and when each pass is complete. Fig. 7 shows the timing of these signals and an explanation of their use follows.

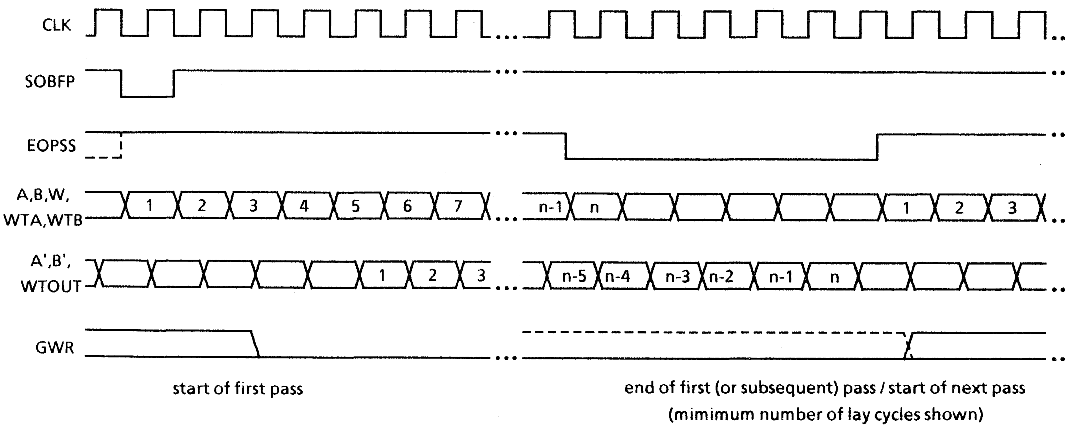

Fig 7 Use of BFP Control Signals

To commence the FFT, the signal EOPSS should be set high (where it will remain for the duration of the pass). SOBFP should be pulled low during the initial cycle, when the first data words A and B are presented to the inputs of the butterfly processor. The following cycle, SOBFP must be pulled high where it should remain for the duration of the FFT. New data is presented to the processor each successive cycle until the end of the first pass of the FFT. On the last cycle of the pass, the signal EOPSS should be pulled low and remain low for a minimum of five cycles, the time required to clear the pipeline of the butterfly processor so that all the results from one pass are obtained before commencing the following pass (should a longer pause be required between passes - to arrange the data for the next pass, for example - then EOPSS may be kept low for as long as necessary, the next pass cannot commence until it is brought high again). On the initial cycle of each new pass, the signal EOPSS should be pulled high and it should remain high until the final cycle of that pass, when it is pulled low again.

Building an FFT System

The Butterfly Processor is only one element of a complete FFT system. Also required are fast A/D converters at the front end of the system; a complex heterodyne filter to zoom-in on the frequencies of interest; fast memory and addressing circuits to store the data; additional fast memory and addressing circuits for the coefficients; an output normalisation circuit to make all data consistent; a Pythagoras Processor to extract magnitude and phase information from the results; finally, a D/A converter to allow the magnitude and phase information to be displayed on a video screen or oscilloscope. Fig. 8 shows how these blocks are connected. GEC Plessey Semiconductors make a range of high performance DSP devices which solve the more difficult problems outlined above. The complex heterodyne filter may be constructed from a combination of either a PDSP16116 or PDSP16112 complex multiplier and either a PDSP16318 complex accumulator or two PDSP1601 augmented arithmetic logic units. Output normalisation is a simple matter with the PDSP1601's adaptable barrel shifter and the PDSP16330 Pythagoras processor to convert Cartesian to polar coordinates.

Memory Requirements

Memory requirements differ according to whether the 'In-Place' or 'Constant Geometry' algorithms are used. In either case, two reads from memory (A and B) and two writes to memory (A' and B') have to be made each 100ns cycle.

For the In-Place algorithm, the results (A', B') of a butterfly are written to the same locations from which the inputs (A & B) were read. Hence, the memory must have an access time of 25ns to cope with the two reads and two writes.

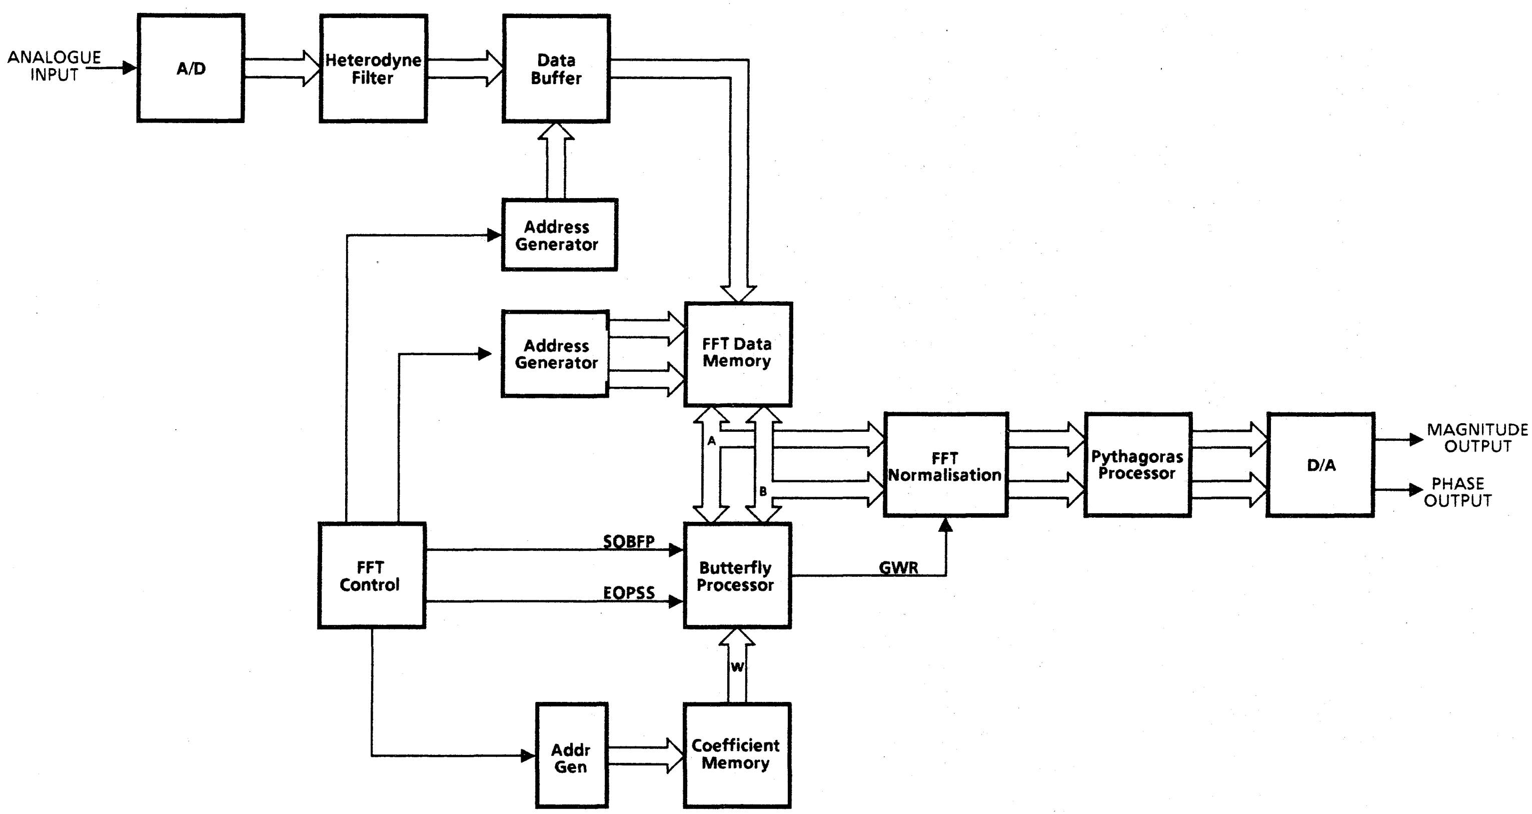

Fig 8 Typical FFT System

The Constant Geometry algorithm requires a memory access time of only 50ns, but the memory size must be double that of the In-Place algorithm. This is because the addresses written to after each butterfly are different from those from which the input data was read. This is possible due to the order in which data points are addressed.

The memory must be 32 bits wide to accommodate the real and imaginary parts of each word. Also, the 2 bit word tag must be stored with each word. This could be achieved by widening the memory to 34 bits or, alternatively, it could be stored in the LSB of the real and imaginary parts of the word, keeping the memory width at 32 bits. This would not affect the accuracy of the FFT, as the LSB is a rounded value in any case. There would be no problem in the initial pass when no word tags have been written to the memory as the PDSP16116 ignores the word tag inputs during the initial pass.

FFT Output Normalisation

In order to preserve the dynamic range of the data during the FFT calculation, the PDSP16116 employs a range of different weightings, however, at the end of the FFT, the data must be re-formatted to a pre-determined common weighting. This can be done by comparing the exponent of a given data word with the required universal exponent and then shifting the data word by the difference. The PDSP1601 ALU, with its multifunction 16-bit barrel shifter, is ideally suited to this task.

What value should the universal exponent take? Theoretically, the largest possible data result from an FFT is 1.27N times the largest input data, where N is the size of the FFT. This means that the binary point can move a maximum of $(1 + \log_2 N)$ places to the right. Hence, if the universal exponent is chosen to be $(1 + \log_2 N)$, this should give a sufficient range to represent all data points faithfully. In practice, the FFT output data may never approach the theoretical maximum, therefore it may be worthwhile trying various universal exponents and choosing the one best suited to the particular application.

Data is output from the butterfly processor with a two part exponent: the 5-bit GWR applicable to all data words from a given FFT and a 2-bit WTOUT associated with each individual data word. To find the complete exponent for a given word, the GWR for that FFT must be modified by the WTOUT value, the result being the number of places that the binary point has been shifted to the right during the course of the FFT. This value must be subtracted from the universal exponent, the difference being the shift required for that data word, which is input to the SV port of the PDSP1601.

As FFT data consists of real and imaginary parts, either two PDSP1601s must be used or a single PDSP1601 handling real and imaginary data on alternate cycles, the same shift being applied to both parts. An example of an output normalisation circuit is shown in Fig 9. Only 4-bit arithmetic is used in calculating the shift which means that very small (negative) values of GWR must be trapped and a forced 16-bit right shift applied. (N.B. It is easier to simply add the word tag value to the GWR to determine the shift rather than modifying it exactly. To compensate for this, the universal exponent should be increased by one)

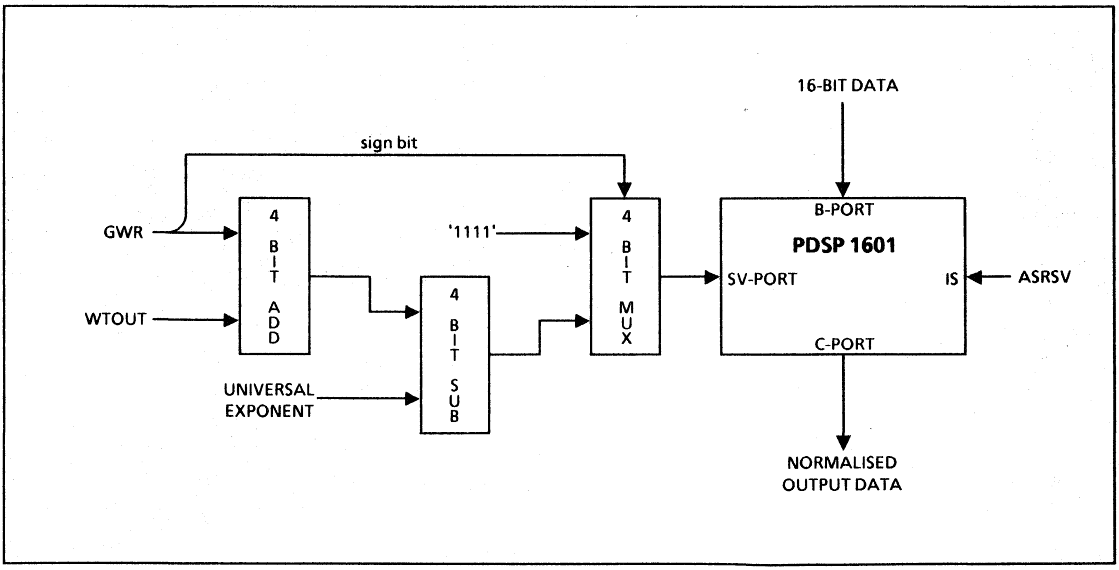

Fig 9 Output Normalisation Circuitry

Appendix A - Block Floating Point FFT Butterfly Net List

The following net lists give all the connections required for implementing the Block Floating Point FFT butterfly shown in Fig 4 :

IC 1 : PDSP16116 Complex Multiplier

Pin No.	Pin desc.	Net name	Connections
D3	PI14	PI14	IC5-C11
C2	PI15	PI15	IC5-D10
B1	WTOUT1	WTOUT1	external o/p
D2	WTOUT0	WTOUT0	external o/p
E3	SFTR0	SFTR0	IC4-L7 ; IC5-L7
C1	SFTR1	SFTR1	IC4-J7 ; IC5-J7
E2	SFTR2	SFTR2	IC4-J6 ; IC5-J6
D1	OEI		tie low
F3	CONX		tie low
F2	CONY		tie low
E1	ROUND		tie high
G2	AI13	AI13	external i/p ; IC3-H1
G3	AI14	AI14	external i/p ; IC3-F1
F1	AI15	AI15	external i/p ; IC3-G2
G1	AR13	AR13	external i/p ; IC2-H1
H2	AR14	AR14	external i/p ; IC2-F1
H1	AR15	AR15	external i/p ; IC2-G2
H3	YI15	WI15	external i/p
J3	YI14	WI14	external i/p
J1	YI13	WI13	external i/p
K1	YI12	WI12	external i/p
J2	YI11	WI11	external i/p
K2	YI10	WI10	external i/p
K3	YI9	WI9	external i/p
L1	YI8	WI8	external i/p
L2	YI7	WI7	external i/p
M1	YI6	WI6	external i/p
N1	YI5	WI5	external i/p
M2	YI4	WI4	external i/p
L3	YI3	WI3	external i/p
N2	YI2	WI2	external i/p
P1	YI1	WI1	external i/p
M3	YI0	WI0	external i/p
N3	XI0	BI0	external i/p
P2	GND	GND	0V supply rail
R1	VDD	VDD	+ 5V supply rail
N4	XI1	BI1	external i/p
P3	XI2	BI2	external i/p
R2	XI3	BI3	external i/p
P4	XI4	BI4	external i/p
N5	XI5	BI5	external i/p
R3	XI6	BI6	external i/p
P5	XI7	BI7	external i/p
R4	XI8	BI8	external i/p
N6	XI9	BI9	external i/p
P6	XI10	BI10	external i/p
R5	XI11	BI11	external i/p
P7	XI12	BI12	external i/p
N7	XI13	BI13	external i/p
R6	XI14	BI14	external i/p
R7	XI15	BI15	external i/p

IC 1 : PDSP16116 Complex Multiplier (continued)

Pin No.	Pin desc.	Net name	Connections
P8	CEY		tie low
R8	CEX		tie low
N8	XR15	BR15	external i/p
N9	XR14	BR14	external i/p
R9	XR13	BR13	external i/p
R10	XR12	BR12	external i/p
P9	XR11	BR11	external i/p
P10	XR10	BR10	external i/p
N10	XR9	BR9	external i/p
R11	XR8	BR8	external i/p
P11	XR7	BR7	external i/p
R12	XR6	BR6	external i/p
R13	XR5	BR5	external i/p
P12	XR4	BR4	external i/p
N11	XR3	BR3	external i/p
P13	XR2	BR2	external i/p
R14	XR1	BR1	external i/p
N12	XR0	BR0	external i/p
N13	YR15	WR15	external i/p
P14	YR14	WR14	external i/p
R15	YR13	WR13	external i/p
M13	GND	GND	0V supply rail
N14	VDD	VDD	+ 5V supply rail
P15	YR12	WR12	external i/p
M14	YR11	WR11	external i/p
L13	YR10	WR10	external i/p
N15	YR9	WR9	external i/p
L14	YR8	WR8	external i/p
M15	YR7	WR7	external i/p
K13	YR6	WR6	external i/p
K14	YR5	WR5	external i/p
L15	YR4	WR4	external i/p
J14	YR3	WR3	external i/p
J13	YR2	WR2	external i/p
K15	YR1	WR1	external i/p
J15	YR0	WR0	external i/p
H14	EOPSS	EOPSS	external i/p
H15	VDD	VDD	+ 5V supply rail
H13	SOBFP	SOBFP	external i/p
G13	WTB1	WTB1	external i/p
G15	WTB0	WTB0	external i/p
F15	WTA1	WTA1	external i/p
G14	WTA0	WTA0	external i/p
F14	MBFP		tie high
F13	CLK	CLK	external i/p - common to all ICs
E15	OSEL1		tie low
E14	OSEL0		tie low
D15	OER		tie low
C15	SFTA0	SFTA0	IC2-L6 ; IC3-L6
D14	SFTA1	SFTA1	IC2-L8 ; IC3-L8
E13	GWR0	GWR0	external o/p
C14	GWR1	GWR1	external o/p
B15	GWR2	GWR2	external o/p
D13	GWR3	GWR3	external o/p
C13	GWR4	GWR4	external o/p
B14	PR15	PR15	IC4-D10
A15	PR14	PR14	IC4-C11

IC 1 : PDSP16116 Complex Multiplier (continued)

Pin No.	Pin desc.	Net name	Connections
C12	VDD	VDD	+ 5V supply rail
B13	GND	GND	0V supply rail
A14	PR13	PR13	IC4-B11
B12	PR12	PR12	IC4-C10
C11	PR11	PR11	IC4-A11
A13	PR10	PR10	IC4-B10
B11	PR9	PR9	IC4-B9
A12	PR8	PR8	IC4-A10
C10	PR7	PR7	IC4-A9
B10	PR6	PR6	IC4-B8
A11	PR5	PR5	IC4-A8
B9	GND	GND	0V supply rail
C9	VDD	VDD	+ 5V supply rail
A10	PR4	PR4	IC4-B6
A9	PR3	PR3	IC4-B7
B8	PR2	PR2	IC4-A7
A8	PR1	PR1	IC4-C7
C8	PR0	PR0	IC4-C6
C7	PI0	PI0	IC5-C6
A7	PI1	PI1	IC5-C7
A6	PI2	PI2	IC5-A7
B7	PI3	PI3	IC5-B7
B6	PI4	PI4	IC5-B6
C6	VDD	VDD	+ 5V supply rail
A5	PI5	PI5	IC5-A8
B5	GND	GND	0V supply rail
A4	PI6	PI6	IC5-B8
A3	PI7	PI7	IC5-A9
B4	PI8	PI8	IC5-A10
C5	PI9	PI9	IC5-B9
B3	PI10	PI10	IC5-B10
A2	PI11	PI11	IC5-A11
C4	PI12	PI12	IC5-C10
C3	PI13	PI13	IC5-B11
B2	GND	GND	0V supply rail
A1	VDD	VDD	+ 5V supply rail

IC 2 : PDSP1601 - Real Path

Pin No.	Pin desc.	Net name	Connections
B10	VCC	VDD	+ 5V supply rail
A6	MSB		tie low
A5	MSS		tie high
B5	B15		tie low
C5	B14		tie low
A4	B13		tie low
B4	B12		tie low
A3	B11		tie low
A2	B10		tie low
B3	B9		tie low
A1	B8		tie low
B2	B7		tie low
C2	B6		tie low
B1	B5		tie low
C1	B4		tie low
D2	B3		tie low
D1	B2		tie low
E3	B1		tie low
E2	B0		tie low
E1	CEB		tie high
F2	CLK	CLK	external i/p - common to all ICs
F3	GND	GND	0V supply rail
G3	MSA0		tie high
G1	MSA1		tie low
G2	A15	AR15	external i/p ; IC1-H1
F1	A14	AR14	external i/p ; IC1-H2
H1	A13	AR13	external i/p ; IC1-G1
H2	A12	AR12	external i/p
J1	A11	AR11	external i/p
K1	A10	AR10	external i/p
J2	A9	AR9	external i/p
L1	A8	AR8	external i/p
K2	A7	AR7	external i/p
K3	A6	AR6	external i/p
L2	A5	AR5	external i/p
L3	A4	AR4	external i/p
K4	A3	AR3	external i/p
L4	A2	AR2	external i/p
J5	A1	AR1	external i/p
K5	A0	AR0	external i/p
L5	CEA		tie low
K6	MSC		tie high
K10	VCC	VDD	+ 5V supply rail
J6	IS0		tie low
J7	IS1		tie high
L7	IS2		tie low
K7	IS3		tie high
L6	SV0	SFTA0	IC1-C15
L8	SV1	SFTA1	IC1-D14
K8	SV2		tie low
L9	SV3		tie low
L10	SVOE		tie high
K9	RS0		tie high
L11	RS1		tie high
J10	RS2		tie high

IC 2 : PDSP1601 - Real Path (continued)

Pin No.	Pin desc.	Net name	Connections
K11	C0		N/C
J11	C1	DAR0	IC4-L11
H10	C2	DAR1	IC4-K10
H11	C3	DAR2	IC4-J10
F10	C4	DAR3	IC4-K11
G10	C5	DAR4	IC4-J11
G11	C6	DAR5	IC4-H10
G9	C7	DAR6	IC4-H11
F9	GND	GND	0V supply rail
F11	C8	DAR7	IC4-F10
E11	C9	DAR8	IC4-G10
E10	C10	DAR9	IC4-G11
E9	C11	DAR10	IC4-G9
D11	C12	DAR11	IC4-F9
D10	C13	DAR12	IC4-F11
C11	C14	DAR13	IC4-E11
B11	C15	DAR14:15	IC4-E9,E10
C10	OE		tie low
A11	BFP		N/C
B9	CO		N/C
A10	RA0		L on even cycles, H on odd cycles
A9	RA1		tie high
B8	RA2		tie low
A8	CI		tie low
B6	IA0		tie low
B7	IA1		tie high
A7	IA2		tie high
C7	IA3		tie low
C6	IA4		tie high

IC 3 : PDSP1601 - Imaginary Path

Pin No.	Pin desc.	Net name	Connections
B10	VCC	VDD	+ 5V supply rail
A6	MSB		tie low
A5	MSS		tie high
B5	B15		tie low
C5	B14		tie low
A4	B13		tie low
B4	B12		tie low
A3	B11		tie low
A2	B10		tie low
B3	B9		tie low
A1	B8		tie low
B2	B7		tie low
C2	B6		tie low
B1	B5		tie low
C1	B4		tie low
D2	B3		tie low
D1	B2		tie low
E3	B1		tie low
E2	B0		tie low
E1	CEB		tie high
F2	CLK	CLK	external i/p - common to all ICs
F3	GND	GND	0V supply rail
G3	MSA0		tie high
G1	MSA1		tie low
G2	A15	AI15	external i/p ; IC1-F1
F1	A14	AI14	external i/p ; IC1-G3
H1	A13	AI13	external i/p ; IC1-G2
H2	A12	AI12	external i/p
J1	A11	AI11	external i/p
K1	A10	AI10	external i/p
J2	A9	AI9	external i/p
L1	A8	AI8	external i/p
K2	A7	AI7	external i/p
K3	A6	AI6	external i/p
L2	A5	AI5	external i/p
L3	A4	AI4	external i/p
K4	A3	AI3	external i/p
L4	A2	AI2	external i/p
J5	A1	AI1	external i/p
K5	A0	AI0	external i/p
L5	CEA		tie low
K6	MSC		tie high
K10	VCC	VDD	+ 5V supply rail
J6	IS0		tie low
J7	IS1		tie high
L7	IS2		tie low
K7	IS3		tie high
L6	SV0	SFTA0	IC1-C15
L8	SV1	SFTA1	IC1-D14
K8	SV2		tie low
L9	SV3		tie low
L10	SVOE		tie high
K9	RS0		tie high
L11	RS1		tie high
J10	RS2		tie high

IC 3 : PDSP1601 - Imaginary Path (continued)

Pin No.	Pin desc.	Net name	Connections
K11	C0		N/C
J11	C1	DAI0	IC5-L11
H10	C2	DAI1	IC5-K10
H11	C3	DAI2	IC5-J10
F10	C4	DAI3	IC5-K11
G10	C5	DAI4	IC5-J11
G11	C6	DAI5	IC5-H10
G9	C7	DAI6	IC5-H11
F9	GND	GND	0V supply rail
F11	C8	DAI7	IC5-F10
E11	C9	DAI8	IC5-G10
E10	C10	DAI9	IC5-G11
E9	C11	DAI10	IC5-G9
D11	C12	DAI11	IC5-F9
D10	C13	DAI12	IC5-F11
C11	C14	DAI13	IC5-E11
B11	C15	DAI14:15	IC5-E9,E10
C10	OE		tie low
A11	BFP		N/C
B9	CO		N/C
A10	RA0		L on even cycles, H on odd cycles
A9	RA1		tie high
B8	RA2		tie low
A8	CI		tie low
B6	IA0		tie low
B7	IA1		tie high
A7	IA2		tie high
C7	IA3		tie low
C6	IA4		tie high

IC 4 : PDSP16318 - Real Path

Pin No.	Pin desc.	Net name	Connections
B2	D7	B'R7	external o/p
C2	D8	B'R8	external o/p
B1	D9	B'R9	external o/p
C1	D10	B'R10	external o/p
D2	GND	GND	0V supply rail
D1	VDD	VDD	+ 5V supply rail
E3	D11	B'R11	external o/p
E2	D12	B'R12	external o/p
E1	D13	B'R13	external o/p
F2	D14	B'R14	external o/p
F3	D15	B'R15	external o/p
G3	C15	A'R15	external o/p
G1	C14	A'R14	external o/p
G2	C13	A'R13	external o/p
F1	C12	A'R12	external o/p
H1	VDD	VDD	+ 5v supply rail
H2	GND	GND	0V supply rail
J1	C11	A'R11	external o/p
K1	C10	A'R10	external o/p
J2	C9	A'R9	external o/p
L1	C8	A'R8	external o/p
K2	C7	A'R7	external o/p
K3	C6	A'R6	external o/p
L2	C5	A'R5	external o/p
L3	C4	A'R4	external o/p
K4	C3	A'R3	external o/p
L4	C2	A'R2	external o/p
J5	C1	A'R1	external o/p
K5	C0	A'R0	external o/p
L5	OED		tie low
K6	OEC		tie low
J6	SD2	SFTR2	IC1-E2 ; IC5-J6
J7	SD1	SFTR1	IC1-C1 ; IC5-J7
L7	SD0	SFTR0	IC1-E3 ; IC5-L7
K7	MS		tie low
L6	ASI1		tie high
L8	ASI0		tie low
K8	DEL		tie low
L9	CLR		tie low
L10	ASR1		tie low
K9	ASR0		tie low
L11	A0	DAR0	IC2-J11
K10	A1	DAR1	IC2-H10
J10	A2	DAR2	IC2-H11
K11	A3	DAR3	IC2-F10
J11	A4	DAR4	IC2-G10
H10	A5	DAR5	IC2-G11
H11	A6	DAR6	IC2-G9
F10	A7	DAR7	IC2-F11
G10	A8	DAR8	IC2-E11
G11	A9	DAR9	IC2-E10
G9	A10	DAR10	IC2-E9
F9	A11	DAR11	IC2-D11
F11	A12	DAR12	IC2-D10
E11	A13	DAR13	IC2-C11
E10	A14	DAR14	IC2-B11
E9	A15	DAR15	IC2-B11

IC 4 : PDSP16316 - Real Path (continued)

Pin No.	Pin desc.	Net name	Connections
D11	CEA		tie low
D10	B15	PR15	IC1-B14
C11	B14	PR14	IC1-A15
B11	B13	PR13	IC1-A14
C10	B12	PR12	IC1-B12
A11	B11	PR11	IC1-C11
B10	B10	PR10	IC1-A13
B9	B9	PR9	IC1-B11
A10	B8	PR8	IC1-A12
A9	B7	PR7	IC1-C10
B8	B6	PR6	IC1-B10
A8	B5	PR5	IC1-A11
B6	B4	PR4	IC1-A10
B7	B3	PR3	IC1-A9
A7	B2	PR2	IC1-B8
C7	B1	PR1	IC1-A8
C6	B0	PR0	IC1-C8
A6	CLK	CLK	external i/p - common to all ICs
A5	CEB		tie low
B5	OVR		N/C
C5	D0	B'R0	external o/p
A4	D1	B'R1	external o/p
B4	D2	B'R2	external o/p
A3	D3	B'R3	external o/p
A2	D4	B'R4	external o/p
B3	D5	B'R5	external o/p
A1	D6	B'R6	external o/p

IC 5 : PDSP16318 - Imaginary Path

Pin No.	Pin desc.	Net name	Connections
B2	D7	B'I7	external o/p
C2	D8	B'I8	external o/p
B1	D9	B'I9	external o/p
C1	D10	B'I10	external o/p
D2	GND	GND	0V supply rail
D1	VDD	VDD	+ 5V supply rail
E3	D11	B'I11	external o/p
E2	D12	B'I12	external o/p
E1	D13	B'I13	external o/p
F2	D14	B'I14	external o/p
F3	D15	B'I15	external o/p
G3	C15	A'I15	external o/p
G1	C14	A'I14	external o/p
G2	C13	A'I13	external o/p
F1	C12	A'I12	external o/p
H1	VDD	VDD	+ 5v supply rail
H2	GND	GND	0V supply rail
J1	C11	A'I11	external o/p
K1	C10	A'I10	external o/p
J2	C9	A'I9	external o/p
L1	C8	A'I8	external o/p
K2	C7	A'I7	external o/p
K3	C6	A'I6	external o/p
L2	C5	A'I5	external o/p
L3	C4	A'I4	external o/p
K4	C3	A'I3	external o/p
L4	C2	A'I2	external o/p
J5	C1	A'I1	external o/p
K5	C0	A'I0	external o/p
L5	OED		tie low
K6	OEC		tie low
J6	SD2	SFTR2	IC1-E2 ; IC4-J6
J7	SD1	SFTR1	IC1-C1 ; IC4-J7
L7	SD0	SFTR0	IC1-E3 ; IC4-L7
K7	MS		tie low
L6	ASI1		tie high
L8	ASI0		tie low
K8	DEL		tie low
L9	CLR		tie low
L10	ASR1		tie low
K9	ASR0		tie low
L11	A0	DAI0	IC3-J11
K10	A1	DAI1	IC3-H10
J10	A2	DAI2	IC3-H11
K11	A3	DAI3	IC3-F10
J11	A4	DAI4	IC3-G10
H10	A5	DAI5	IC3-G11
H11	A6	DAI6	IC3-G9
F10	A7	DAI7	IC3-F11
G10	A8	DAI8	IC3-E11
G11	A9	DAI9	IC3-E10
G9	A10	DAI10	IC3-E9
F9	A11	DAI11	IC3-D11
F11	A12	DAI12	IC3-D10
E11	A13	DAI13	IC3-C11
E10	A14	DAI14	IC3-B11
E9	A15	DAI15	IC3-B11

IC 5 : PDSP16318 - Imaginary Path (continued)

Pin No.	Pin desc.	Net name	Connections
D11	CEA		tie low
D10	B15	PI15	IC1-C2
C11	B14	PI14	IC1-D3
B11	B13	PI13	IC1-C3
C10	B12	PI12	IC1-C4
A11	B11	PI11	IC1-A2
B10	B10	PI10	IC1-B3
B9	B9	PI9	IC1-C5
A10	B8	PI8	IC1-B4
A9	B7	PI7	IC1-A3
B8	B6	PI6	IC1-A4
A8	B5	PI5	IC1-A5
B6	B4	PI4	IC1-B6
B7	B3	PI3	IC1-B7
A7	B2	PI2	IC1-A6
C7	B1	PI1	IC1-A7
C6	B0	PI0	IC1-C7
A6	CLK	CLK	external i/p - common to all ICs
A5	CEB		tie low
B5	OVR		N/C
C5	D0	B'I0	external o/p
A4	D1	B'I1	external o/p
B4	D2	B'I2	external o/p
A3	D3	B'I3	external o/p
A2	D4	B'I4	external o/p
B3	D5	B'I5	external o/p
A1	D6	B'I6	external o/p

References

For a general introduction to FFT's the following texts are recommended :

1. Rabiner and Gold, 'Theory and Application of Digital Signal Processing', Prentice - Hall, 1975.

2. Oppenhiem and Shafer, 'Digital Signal Processing', Prentice - Hall, 1975

Other GEC Plessey Semiconductors application notes and briefs of interest include:

AN47 'A Radix 2 Butterfly Processor'

AN49 'Complex Signal Processing with the PDSP16000 Family'

AN50 'A Fast FFT Processor Using the PDSP16000 Family'

AB01 'A 50ns Butterfly Processor'

AB10 'FIR Filtering with the PDSP16112 and PDSP16318'

In addition, many PDSP devices and applications are modelled on the 'PDSP Demonstrator' software, intended to be run on IBM - PC or compatibles.

CONFIGURING THE PDSP16256

1. INTRODUCTION

The PDSP16256 is a digital finite impulse response (FIR) filter capable of providing between 16 and 128 digital filtering stages at sampling rates from 3.125MHz for a 128 stage filter up to 25MHz for 16 stages. The device contains 16 multiplier-accumulator circuits which are multi-cycled to produce filters of greater than 16 stages. Input data and coefficents are both represented by 16 bit two's complement numbers with coefficients converted internally to 12 bits.

The PDSP16256 is designed such that devices may be easily cascaded. This allows filter lengths to be increased above the nominal 128 stages offered by a single device. Cascading may also be used to raise the filter sampling rate in those cases where a single device is unable to offer the required performance. For example, a single PDSP16256, when configured as a 64 stage filter, offers a maximum sampling rate of 6.25MHz. Using two devices in cascade, each configured as a 32 stage filter, offers a sampling rate of 12.5MHz. Expansion buses ensure that no loss of accuracy occurs when devices are connected in this configuration. If a system contains a number of PDSP16256 devices which each comprise separate unrelated filters, the inherent cascading features of the device may still be used in order to allow all of the devices to be configured from a single EPROM, thereby simplifying the system design.

The PDSP16256 contains a single 16 bit control register which determines the mode and speed of operation, along with other control functions. The filter coefficients are 12 bits wide, with one coefficient being required for each stage of the filter.

The data presented to the 16 bit coefficient input bus is automatically rounded to yield a 12 bit coefficient during device configuration. The rounding scheme employed is described in detail in the PDSP16256 data sheet.

For filters with 64 (single filter mode only), 32 and 16 stages the PDSP16256 offers the coefficient bank swap feature. If enabled, bank swapping allows one set of coefficients to be replaced by an alternative set without the need to load these from an external source.

Before filtering operations may commence, it is necessary to load the device with both the filter coefficients and the control register word. This may be accomplished in one of two ways;

- by loading the data from a local memory device (e.g. EPROM)
- via a system data bus under the control of an intelligent master device (e.g. a microprocessor)

The former has the advantage of providing a stand-alone system which is able to load the required data on system boot-up whilst using the minimum number of external components. The latter allows much greater flexibility in the types of filtering systems which may be implemented. Remote master mode could be used, for example, to implement an adaptive filtering system in which the filter coefficients are changed, depending on prevailing conditions, thus maximizing the performance of the filter at all times.

This application brief gives detailed information which will allow users to determine how to interface to the PDSP16256 in order to load coefficient and control information. This brief covers the use of the device in both automatic EPROM load and remote master modes. Information on cascading devices in both modes is also included.

2. TIMING

The PDSP16256 uses registers on all inputs and outputs to ensure that signals are synchronous with the system clock. This scheme ensures that all setup and hold times for incoming signals are the same. Similarly, all clock to data valid delays for output signals are also identical.

The GEC Plessey Semiconductors Digital Signal Processing IC Handbook (publication number PS 2252) lists these delays for the PDSP16256 as :

Parameter	Min	Max
Input signal setup	8	-
Input signal hold	4	-
SCLK to output valid	5	26

(All times in nanoseconds)

There is an exception to this rule; the output circuitry. Here, the delay from activating the output enable to the data being available on the pins is not the same as the value given above for other output signals. Similarly, the delays in changing from a high- to a low-impedance output state, and vice-versa, are also different. The values of these delays may be found by referring to the DSP IC Handbook.

3. AUTOMATIC EPROM LOAD MODE

The PDSP16256 generates all the required signals to interface directly to an EPROM and load the required data automatically without the need for external glue logic in this mode. Tying the EPROM input to ground selects this mode of operation. It is assumed that all data loaded consists of 8 bit bytes. The data load sequence is initiated by the RESET input to the PDSP16256. RESET, which is active low, should be held low for 16 cycles of the system clock, SCLK. After the sixth SCLK cycle, the BUSY signal will go high, signalling to other devices in the system that the PDSP16256 is performing a data load function. All output data is invalid while the BUSY signal is active. BUSY returns to its inactive state once all the required coefficient and control register data has been loaded.

The address bus used to access the EPROM is 8 bits wide. This allows a maximum of 256 data bytes to be read. For a digital filter containing the maximum 128 stages, 256 bytes will be required to store the coefficients as each one requires two bytes of storage. The byte held in the lower address of each byte-pair should be the least significant byte. If a smaller filter is to be implemented, 32 taps for example, then some of the address lines are redundant and will remain low at all times. These redundant address lines should not be connected to the EPROM when a number of PDSP16256 devices are to be cascaded. A further two bytes are required to store the 16 bit control register word. The control word always resides above the filter coefficients in the EPROM memory map.

The coefficient control pin, CCS, is used to determine whether the control register or filter coefficients are to be loaded. When CCS is high, the control register is loaded and when low, the filter coefficients are loaded. CCS is effectively used as the most significant address line in addition to the address bus A7:0. Hence, when CCS is high, the EPROM address that will be referenced is 256+A7:0, assuming that all eight address lines are in use. In order to load the control register, A7:0 takes the values 0 and 1 whilst CCS is held high, thereby accessing addresses 256 and 257 in 128 stage filter systems.

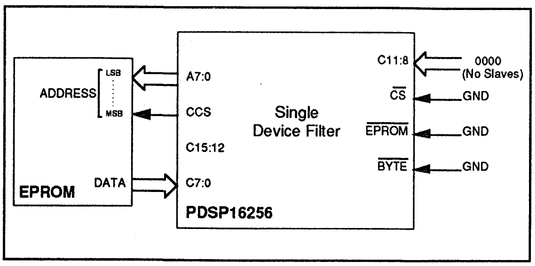

Fig. 1 Diagram of standalone EPROM system

Once the control register has been loaded, A7:0 starts counting upwards from address 0. As CCS has now returned low, the EPROM addresses will be in the range 0 to 255. Bits 12, 13 and 14 of the control register define the number of digital stages in the filter and whether the decimate option is enabled. This defines the number of coefficients required and is used to regulate the number of bytes loaded from the EPROM. Every 4 SCLK cycles a new address is generated and a new byte loaded using the 8 least significant bits of the coefficient input bus, C15:0. If only one PDSP16256 device is contained in a filtering system (i.e. devices are not being cascading), then bits C11:8 should all be tied to ground and C15:12 should be left unconnected. The reasons for this will become apparent after reading section 3.2.

3.1 Timing Information

Each address is held on the A7:0 outputs for four SCLK cycles. On the fourth cycle, the data presented to the lower 8 bits of the C15:0 input bus is loaded into the device. Using this information the type of EPROM required to correctly interface to a PDSP16256 may be specified.

The EPROM access time is calculated by subtracting the output delay between SCLK and A7:0 and the setup time required to correctly load data into the PDSP16256 from four SCLK periods. Hence,

EPROM Access Time = 4.T - Output Delay - Setup Time

(where T is the SCLK period)

We know that the output delay has a maximum value of 26ns and that the setup time is 8ns. Therefore,

EPROM Access Time = 4.T - 26 - 8 = 4T - 34

For a typical application where SCLK is running at 25MHz (T=40ns), it can be seen that the EPROM access time will be 126ns. This means that to correctly load data with SCLK running at 25MHz, an EPROM with an access time of 126ns, or less, is required.

3.2 Cascading devices in EPROM mode

An arbitrary number of PDSP16256 devices may be cascaded to form a given filtering system. However, only a maximum of 16 devices may access the same EPROM to load configuration data. Hence, if more than 16 devices are to be cascaded, these should be subdivided into blocks of 16 for configuration purposes, with each 16 device block having access to a separate EPROM.

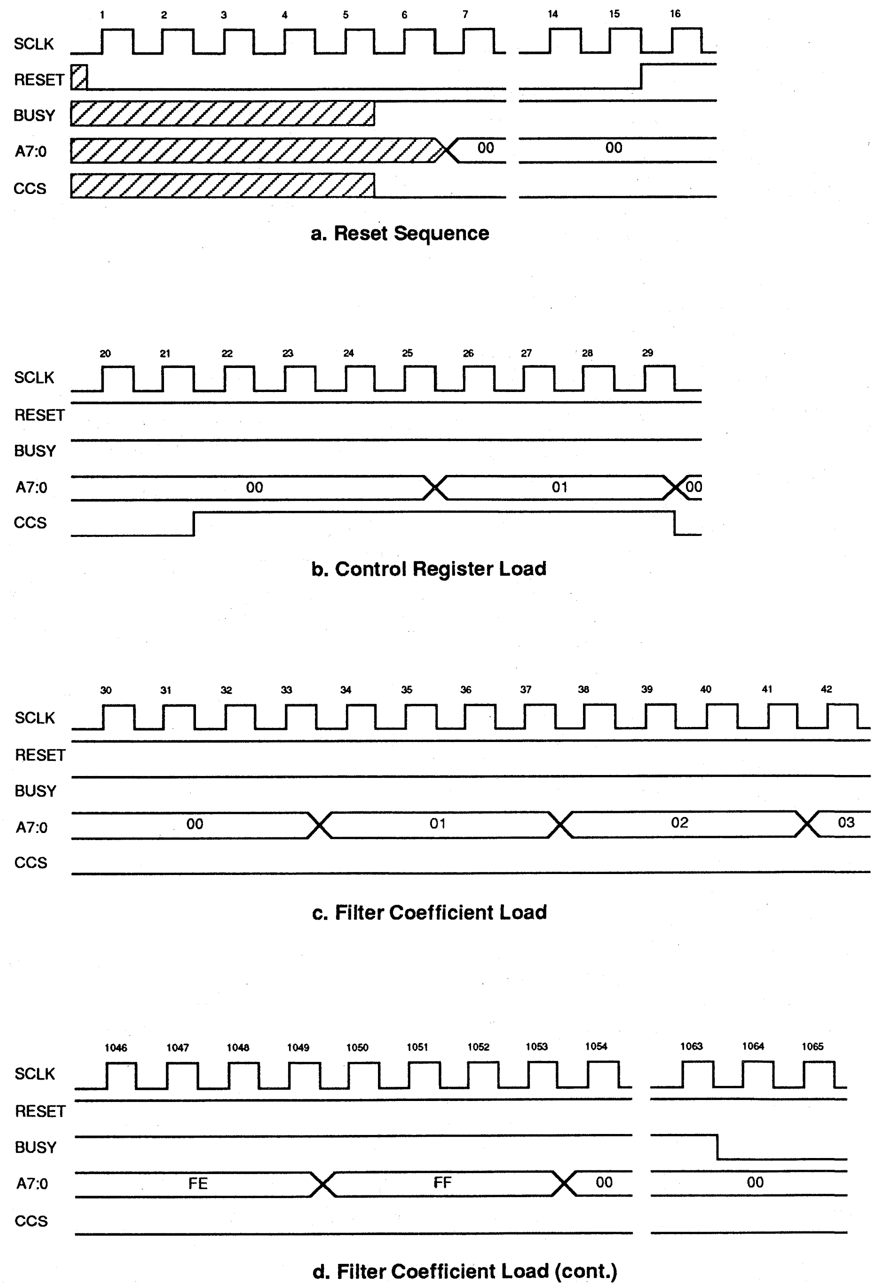

Fig.2 Reset & coefficient load sequences.

Configuring the PDSP16256

When several devices are connected in cascade, one of these devices is nominated the master (this should not be confused with a remote master discussed elsewhere in this application brief). The master device is responsible for controlling the loading of data for all devices to which it is connected. The master is identified by the fact that its EPROM input is grounded; all other cascaded devices, termed slaves, have their EPROM inputs connected to +5V. The master provides address, Write Enable, CCS and device decoding signals for all slave devices as well as for itself.

Only 8 of the 16 bits of the coefficient input bus, C15:0, are used in automatic EPROM load mode. The remaining eight lines are used to provide device decoding information. These signals help pass the required coefficient and control data to the correct device under the control of the master. The four bits C11:8, which are programmed as inputs on both slave and master devices, are used to assign a unique code to each slave. The first slave should have C11:8=0001, the second, C11:8=0010 and so on. The master device is therefore implicitly labelled 0000. The master uses C11:8 inputs to identify the total number of slaves contained in the cascaded system. If four devices are to be cascaded, for example, these will comprise one master and three slaves. Therefore, the C11:8 inputs on the master would be set to 0011.

The four bits C15:12 are programmed as outputs on the master and inputs on all slaves. Only when the value indicated on C15:12 is equal to the slave device code defined by C11:8 will the slave in question load the data passed to it via the coefficient input bus. This arrangement effectively provides a unique decode signal for each slave device in turn. The Chip Select signal, CS, should be tied to ground on all devices; slaves and master. C15:12 are used in conjunction wth A7:0 and CCS to access different pages of the EPROM, as shown by the timing diagram, fig. 4.

3.3 EPROM Memory Requirement

When the PDSP16256 loads coefficient data, if the filter length is such that bank swapping is possible, the device will always attempt to load a second bank of coefficients, irrespective of whether the bank swapping option has been enabled in the control word or not. Hence, in single filter mode, the amount of EPROM space required for each of the filter lengths, in bytes, is:

Filter Length	EPROM Space
128	512
64	512
32	256
16	128

In dual filter mode, the EPROM memory requirements are as follows:

Filter Length	EPROM Space
64	512
32	512
16	256
8	128

If n devices are connected in cascade, with each device implementing the same size of filter, then the EPROM space required is simply the product of the space required by a single device and the number of devices.

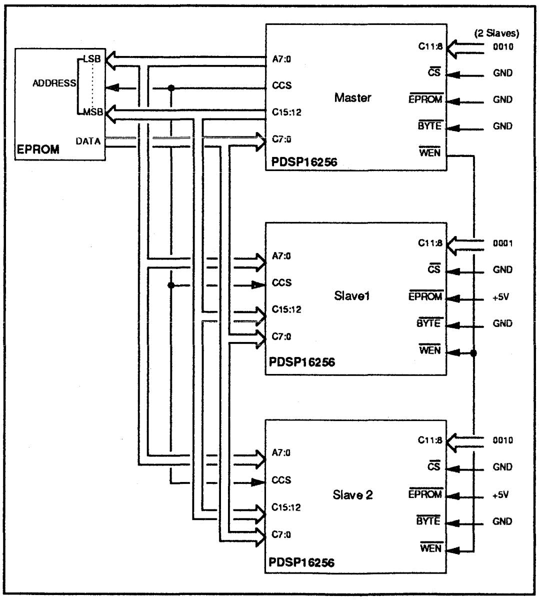

Fig.3 Cascaded system

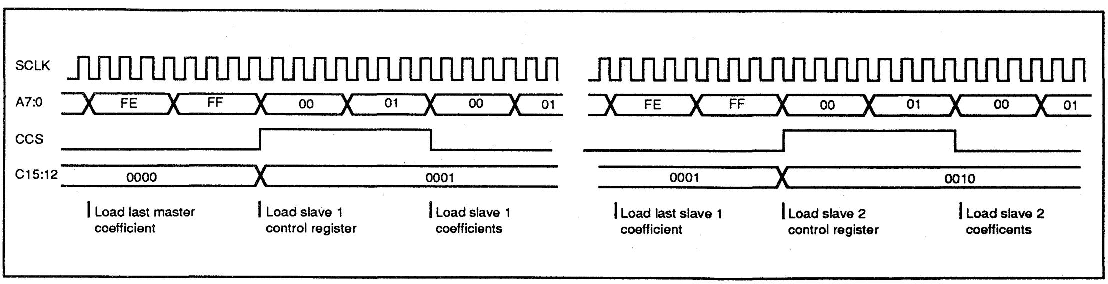

Fig.4 Load sequence for a cascaded system

Consider a system which contains four devices each of which contains a single 128 stage filter. In order to configure all four devices, an EPROM containing a total of

4 x 512 = 2048 bytes

will be required. The memory map for this EPROM is shown in Fig.5. The diagram shows the values of C15:12, CCS and A7:0 required to access each different area of the EPROM.

4. REMOTE MASTER MODE

In this mode, the remote master, is charged with supplying the address and data buses to the PDSP16256 via a synchronous peripheral interface. This mode allows all or selected coefficients to be updated under the control of the remote master whenever conditions dictate. Remote master mode is selected by tying the EPROM input pin of the PDSP16256 to +5V.

A filter initialisation sequence, where both the control register and filter coefficients are loaded, is initiated by a RESET sequence. RESET should be held low for 16 cycles of SCLK. After the sixth SCLK cycle the BUSY signal goes high, indicating that the device is performing internal initialisation operations. BUSY will remain active for 31 SCLK cycles. Only when it goes low can data be written to the device. A RESET sequence, as described above, is not needed when filter coefficients alone are to be updated.

When valid data appears on the address and data bus, this data is written to the PDSP16256 by the application of synchronous Chip Select (CS) and Write Enable (WEN) signals. Data may be written to the PDSP16256 either in 8 bit bytes or 16 bit words, the data width being selected by the BYTE input pin.

If BYTE is tied low, data is loaded as 8 bit bytes, and if high, as 16 bit words. The comparison that is made in automatic EPROM load mode between C15:12 and C11:8 is also made in byte mode, irrespective of the fact that remote master mode is selected. Hence, to satisfy this comparison, pins C15:8 should all be tied to ground when using 8 bit data transfers. In byte mode, the least significant byte (which is accessed when A0=0) should be written first, followed by the most significant (accessed when A0=1). If the device is operating in byte mode, it is necessary to hold CS low for an extra two SCLK cycles once all the coefficients have been loaded. This can be clearly seen in fig. 8c.

When 16 bit words are used, a maximum of 128 transfers are required to load all the required filter coefficients. Hence, only 7 address lines are needed and A7 is therefore redundant in word mode. Unlike automatic EPROM load mode, new data may be written to the device on every SCLK cycle, if required, as long as the timing constraints are honoured.

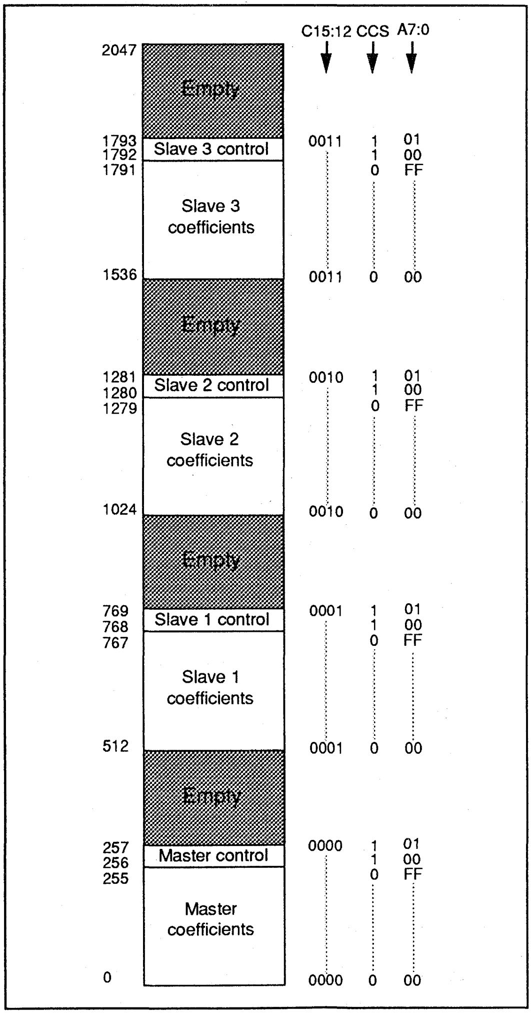

Fig.5 EPROM memory map for a four device system

As with automatic EPROM load mode, the control register should be loaded first, followed by the filter coefficients. A delay of at least two clock cycles needs to be inserted after the control register has been loaded and before any filter coefficents are written to the device. CS needs to be held low for a total of three clock cycles when loading the control register, as shown in fig. 8a. CCS becomes an input to the PDSP16256 in this mode and, when high, indicates that a load to the control register is to take place. Hence, as with automatic EPROM load mode, CCS is used as the most significant address line in conjunction with A7:0. The filter coefficients may be addressed at random, i.e. they need not be loaded in address order, and an arbitrary number may be modified under the control of the remote master at any time. It must be borne in mind that one of the effects of loading new coefficients will be that mathematically correct results will not

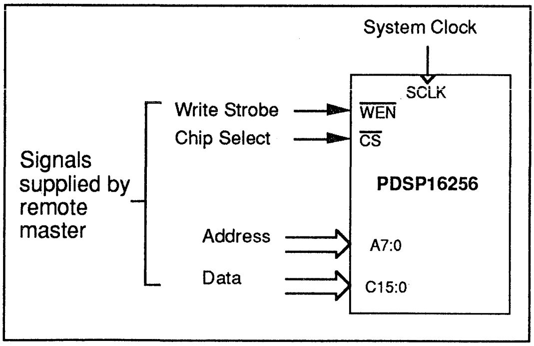

Fig.6 Remote Master standalone system

be obtained for a number of cycles, the exact number depending upon which coefficients were modified and the length of the filter. This is due to the effect of partial results, calculated using the old coefficients, being propagated through the device. It is not recommended that the control register be modified whilst the device is operating. The control register should only be loaded after a reset sequence, as described above.

The PDSP16256 is a synchronous device. This means that all external signals are assumed to be synchronous to the system clock. In general, the remote master's Write Enable and Chip Select signals will be asynchronous to the PDSP16256 system clock. Hence, the signals supplied to the PDSP16256 should be synchronised to SCLK and care taken to ensure that all setup and hold times are honoured. Double buffering the signals generated by the remote master, as shown in fig. 7a, will ensure that these constraints are met, whilst at the same time minimizing metastability problems.

The diagram below shows one possible implementation of a system that will ensure that all timing constraints are met when operating in remote master mode. The coefficient load state machine is used to generate the WEN and CS signals. As can be seen in fig. 7b, WEN is low only for the required one clock edge period. In more complicated systems this state machine could be used to perform other functions such as device decodes etc.

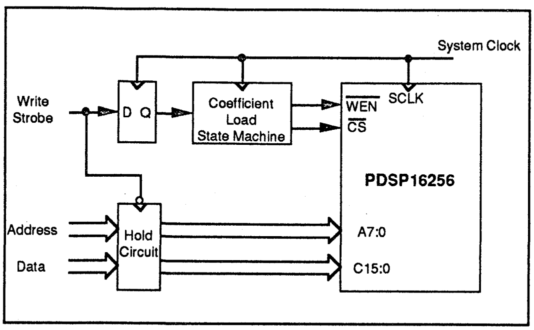

Fig.7a Remote master interface

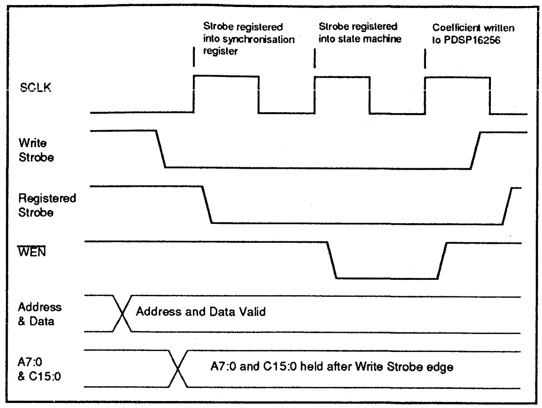

Fig.7b Remote master signal synchronisation

4.1 Timing

Figs. 8a and 8b show the setup and hold constraints which must be met if data is to be loaded correctly. The setup and hold time values are as listed in section 2 above. Fig. 9 shows examples of reset and load sequences for both byte and word modes.

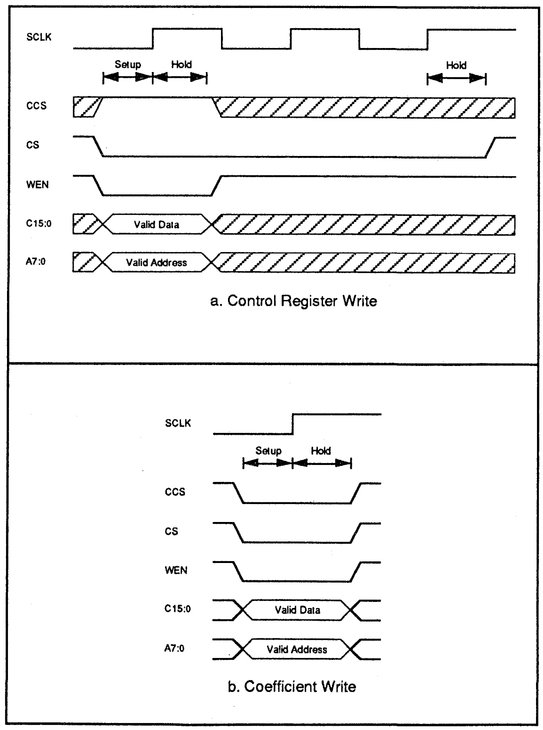

Fig.8a,b Remote master timing constraints

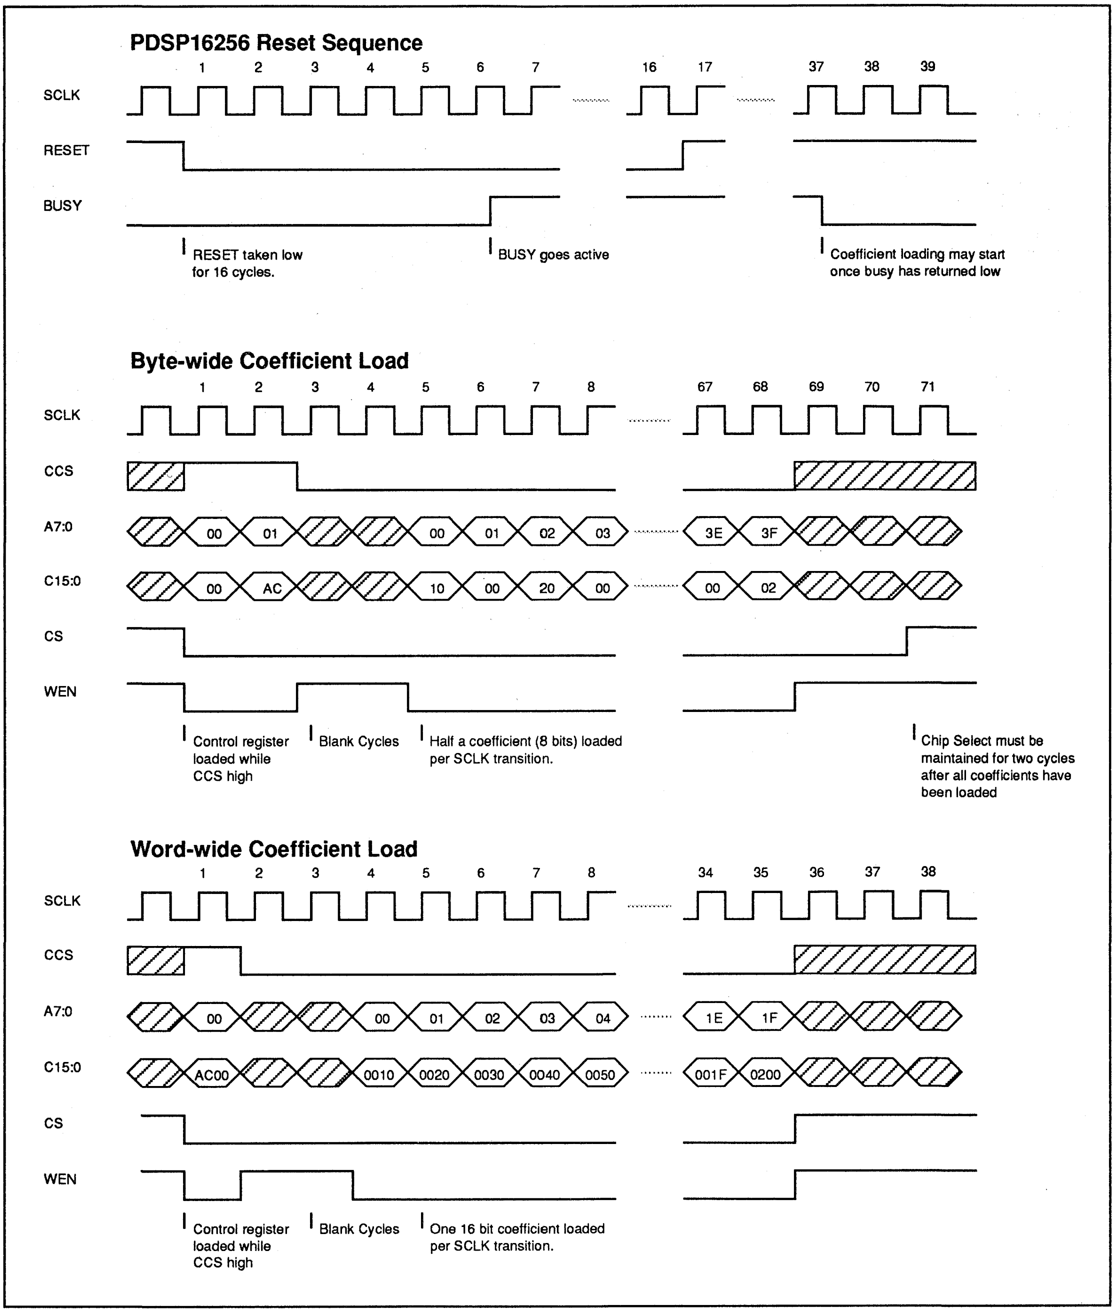

Fig.9 Reset & coefficient load sequences

4.2 Cascading in Remote Master mode

When devices are cascaded in this mode, the master-slave distinction that is evident in automatic EPROM load mode is not necessary as all devices are treated in the same way. Chip Select and Write Enable signals are used to selectively enable the required device and to write to it the relevant data. In many ways, the method used to write data to a series of PDSP16256 devices is analogous to the process of writing to a bank of RAM devices. Fig. 10 shows a typical cascade configuration for this mode.

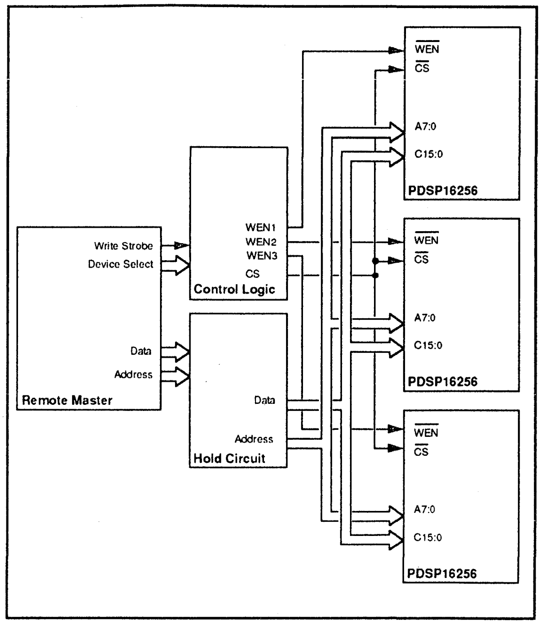

Fig. 10 Cascaded remote master system

DIGITAL FILTERING USING THE PDSP16256

INTRODUCTION

In the field of high performance filtering, engineering solutions are making increasing use of digital techniques. Digital filters are known to typically offer improved accuracy, complete predictability, flexibility and performance improvements. They are also highly suitable for integration with modern CAD tools and techniques, thus reducing development times and simplifying the design process.

General purpose DSP processors can implement digital filters with sampling rates upto approximately 250kHz, but until recently sampling rates beyond this threshold required complex custom design. However the latest CMOS design techniques now enable dedicated standard parts of the necessary speed and complexity to be fabricated, rendering custom designs obsolete.

WHAT ARE DIGITAL FILTERS?

'When a signal that is sampled in time and quantized in amplitude is processed such that the spectral characteristics of the signal are altered in a controlled manner then the resulting operation is termed digital filtering'.

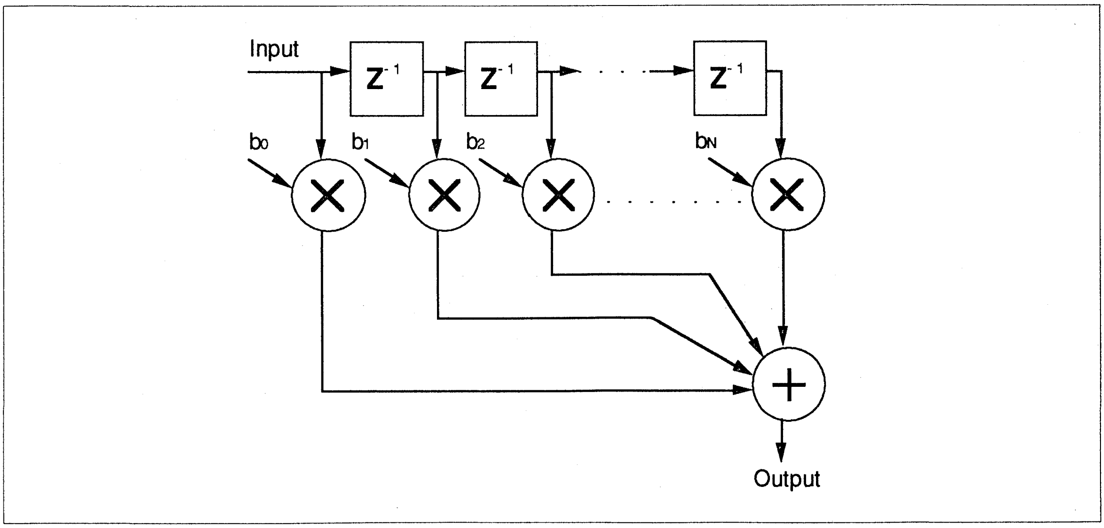

Fig 1 FIR Filter Structure.

Digital filters fall into two groups, those with infinite impulse response (IIR) and those with finite impulse response (FIR). The main difference between these two types is that the output from an FIR filter may be calculated from only current and previous inputs, whereas the output of an IIR filter depends on previous output states as well. Although IIR filters may be designed to be more efficient than an FIR for a given filter order, consideration must always be given to the stability of any design. FIR filters on the other hand, are inherently stable, are generally easier to design and implement in hardware and have the additional advantage that they may be designed such that they are free of phase distortion (i.e. constant group delay).

The output, y_n of an FIR may be calculated as the convolution of the input samples with the filter impulse response and can be represented by a difference equation such as:

$$y_n = b_0x_n+b_1x_{n-1}+.....+b_{N-1}x_{n-N+1}$$

or more generally:

$$y(n) = \sum_{k=0}^{N-1} h(k).x(n-k)$$

where coefficients b_k represent the N samples of the impulse response, h(k), of the desired filter.

WHAT IS THE PDSP16256?.

The PDSP16256 is a single chip FIR filter solution that is capable of operation at sample rates upto 25MHz. Internally it is arranged as two banks of eight multiplier/accumulators that are configurable in a number of ways. Each bank can be configured as a filter of 8, 16, 32, or 64 taps each doubling in length resulting in a halving of the maximum sample rate. The banks can be internally arranged as one single long filter, 2 independent filters, or 2 filters in connected in series. In addition a decimate option allows the output sample rate to be half the input sample rate, thus doubling the filter length. This mode ideal for low pass filter implementations since the high frequencys present in the input can be removed so the output still satisfies Nyquist's sampling criterion.

If the realization of the desired filter is beyond the capabilities of a single device then a number of devices in single filter mode can be cascaded to produce a filter with more taps, due to the provision at external pins of the full 32bit intermediate results.

DEVELOPMENT SYSTEM

A complete development system for the PDSP16256 is available from ERA Technology Ltd, consisting of a software package for filter design specifically tailored to the operating modes of the device and an IBM PC compatible board and control software.

The design package uses special procedures to quantize the filter coefficients in such a manner as to ensure an optimal filter response, given the internal bit accuracies. Low pass, high pass, Hilbert, delay, bandpass or bandstop filters are all supported. The user is given the option to leave any one of the filter design parameters free, and the software then determines this free parameter using the remaining specified parameters. Thus, for example, when designing a low pass filter the user can fix the number of taps to suit the maximum provided by the PDSP16256 at the required sampling rate. Either the transition band, pass band ripple, or stopband attenuation can then be left free, and the software will determine the best that can be achieved for that parameter, given the parameters which are fixed.

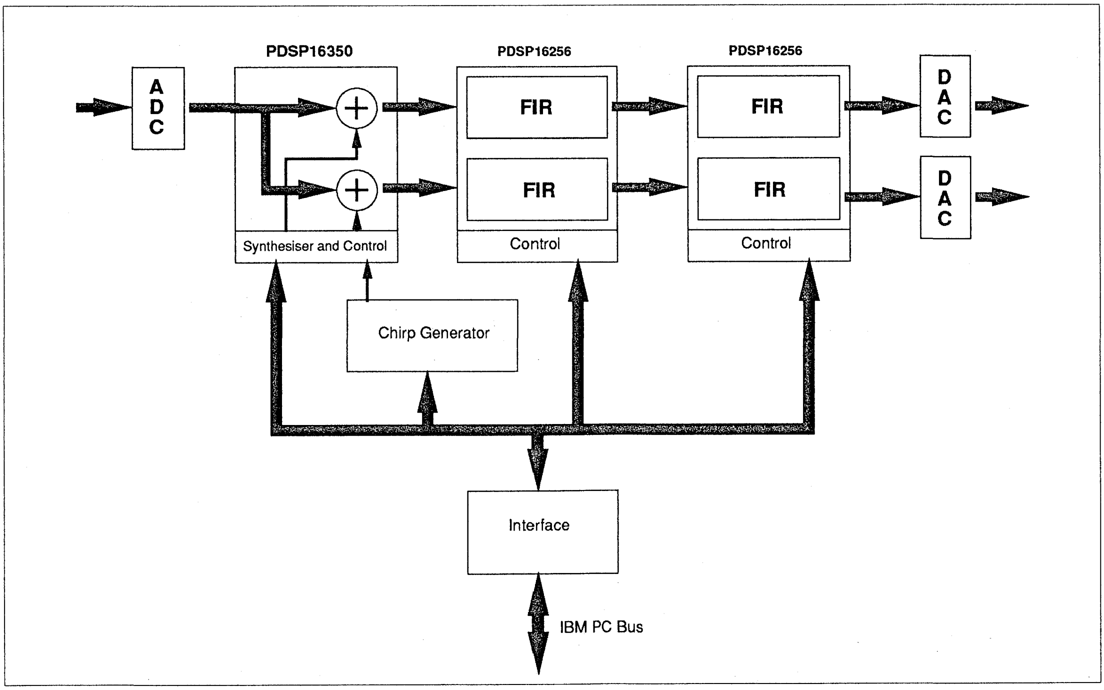

Fig 2 Schematic of Development Board.

The development board is arranged as shown in diagram 2. Digitization is undertaken using either a 20MHz eight bit ADC or a 1MHz twelve bit ADC. A quadrature mixing operation may be applied prior to filtering by means of a PDSP16350 I/Q splitter and numerically controlled oscillator. The digital filtering itself is undertaken using either one or two PDSP16256's. This configuration enables filters of upto 256 coefficients to be implemented using 16-bit data. It also provides the capability for cascaded filtering stages, or for two completely separate filters. The latter would be needed if the complex mixing option is in use. The output signal is available in digital form and in analogue form via dual 12-bit DAC's. The software supplied for the board controls configuration, enables loading of coefficients and can synthesize various waveforms.

EXAMPLES OF FILTERS IMPLEMENTED ON THE PDSP16256

The following filters are designed on the IBM PC design software and show in detail some of the performance characteristics achievable with the PDSP16256.

Figure 3 shows the frequency response of a 128 tap low pass filter designed for a cutoff frequency of 0.1 of the sampling frequency. As implemented on a single PDSP16256, configured in single filter, decimating mode, it exhibits a stop band rejection characteristic approaching -50dB.

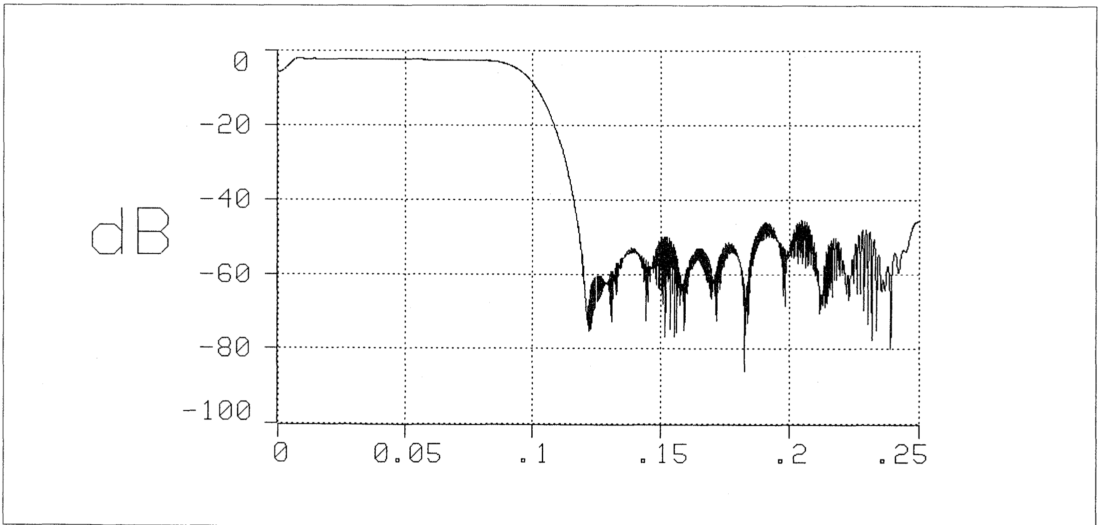

Fig 3 128 Tap Lowpass Filter on Single PDSP16256.

Figure 4 shows the frequency response of a filter designed to the same specification as the one shown in figure 3, but implemented as two 64 tap filters in series, again using a single PDSP16256.

It is clear that the series solution offers a much greater stop band rejection in practical applications but only as a trade off against the width of the transition band and at the expense of greater passband ripple.

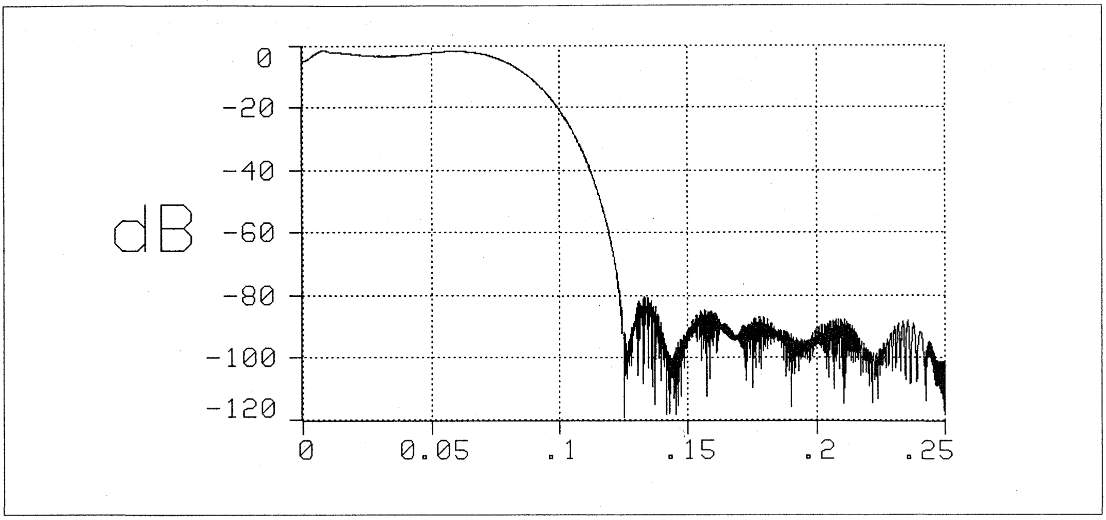

Fig 4 128 Tap Lowpass Filter Composed of two 64 Tap Filters in Series.

Figure 5 Illustrates the implementation of a narrow notch filter on a PDSP16256.

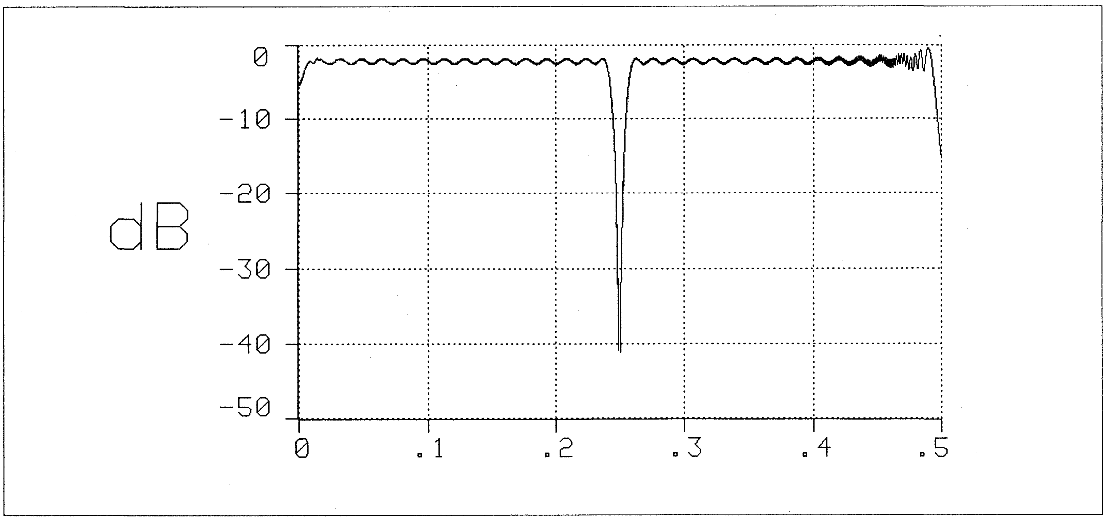

Fig 5 128 tap Bandstop Filter on a PDSP16256.

PRACTICAL SYSTEM CONFIGURATIONS

The PDSP16256 is designed with flexible interfacing characteristics to enable its use in a wide variety of system configurations. At it's simplest it can be configured to auto load the filter coefficients directly from EPROM on power up and be directly connected to ADC's and DAC's (figure 6). Alternatively it could be configured as a dedicated co-processor for a general purpose programmable DSP processor with the DSP device controlling the PDSP16256 configuration, an architecture ideal for adaptive filtering applications for instance (figure 7).

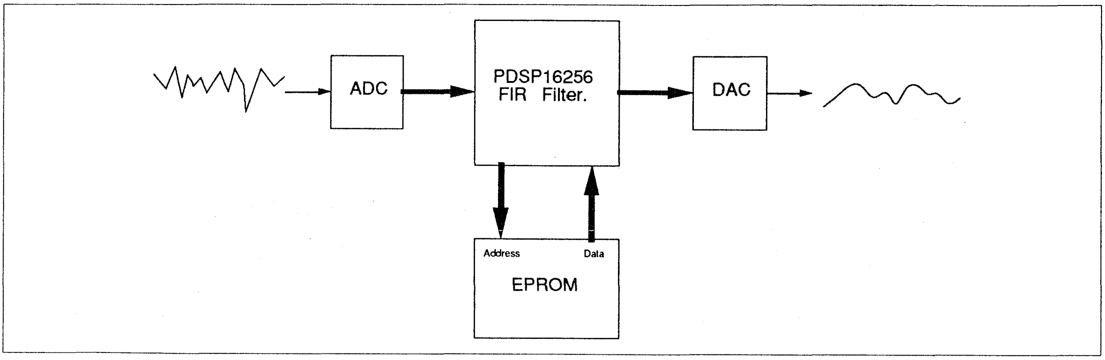

Figure 6 Simple Auto Load Configuration.

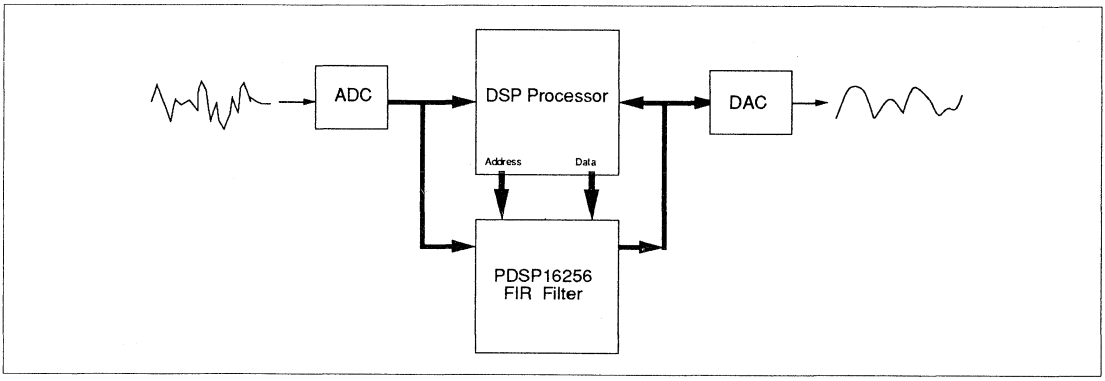

Figure 7 Slave Processor Configuration.

THE PYTHAGORAS PROCESSOR

In a signal processing system it is frequently necessary to calculate the modulus and argument of Complex numbers. This operation is particularly common after Fast Fourier Transforms or in coherent receiver systems. The evaluation of $\sqrt{(x^2 + y^2)}$ and arctan (y/x) are far from easy, so approximations are often used. A common technique for estimating the magnitude of x + jy is to take the larger value of x or y and add to it half the smaller value. The PDSP16330 Pythagoras Processor is a dedicated DSP engine capable of accurate calculation of both magnitude (modulus) and phase (argument) of Complex data at a rate of 100ns per sample

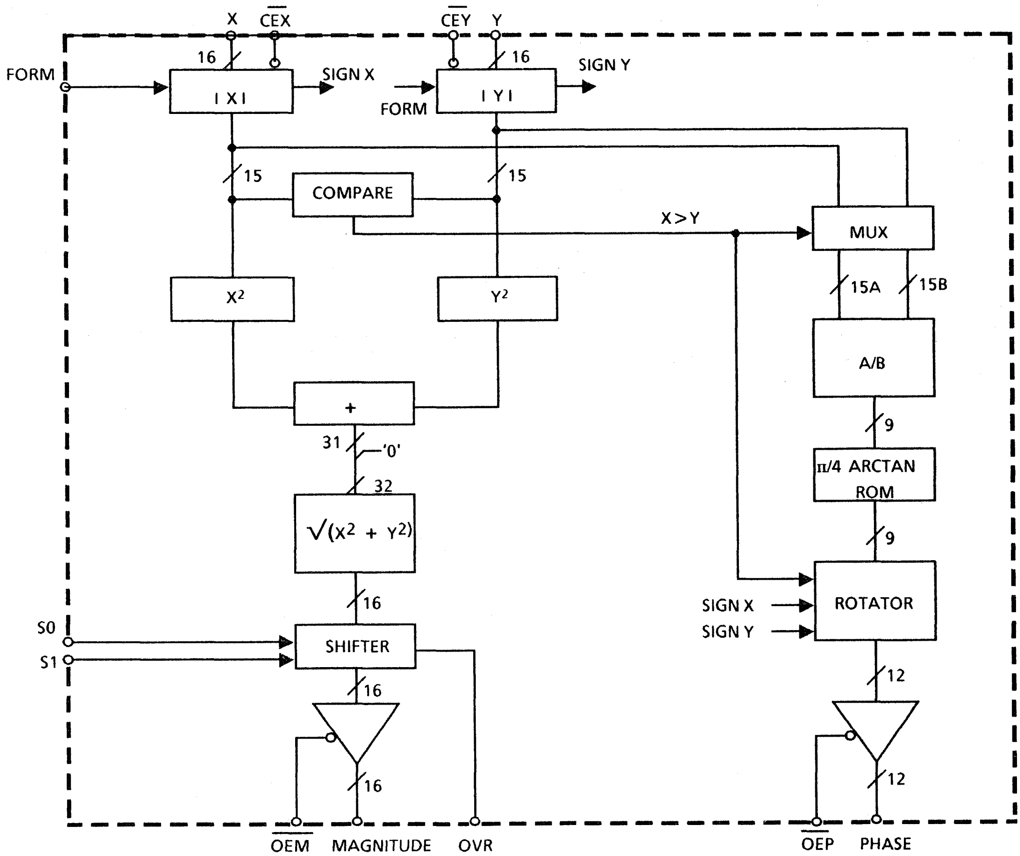

Fig.1 PDSP 16330 Block diagram

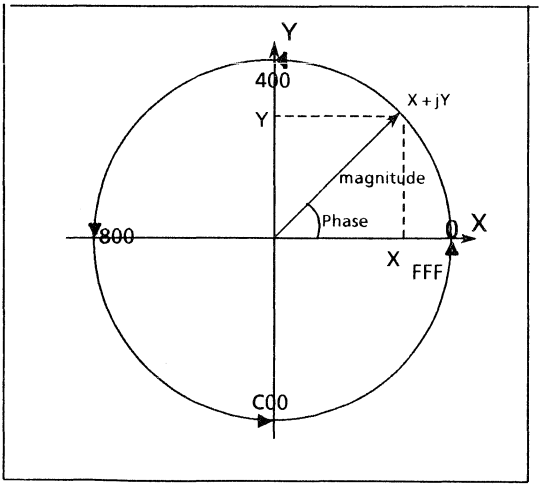

Fig 2

Fig 1 is the block diagram of the device, showing the separate paths for the root sum of squares and arctan (y/x). Fig 2 shows the relationship between the complex input x + jy and the magnitude and phase outputs. Input data can be either 2s complement or sign/magnitude format, depending on the state of the FORM control line.

The magnitude output has a range from 0 to FFFF, four degrees of magnitude output scaling are available via the shift control lines S0 and S1. If the MSB is shifted out of the 0 to FFFF range the OVR flag becomes active, indicating an invalid output. The range of the phase output is 0 to FFF representing a full 2π radians.

APPLICATIONS

FFT.
After an FFT has been carried out the resulting data is complex. This complex data contains information on the magnitude and phase of individual spectral components, but a Cartesian to Polar co-ordinate transformation is required to extract the desired information.

DEMODULATION
In coherent receiver systems the output from the IF stage will have two orthogonal components, I and Q. The carrier may be amplitude or phase modulated, or both. The Pythagoras Processor is used to extract the modulations from the I/Q data.

ROBOTICS
There are many requirements in robotics, position control and position monitoring where conversion from Cartesian space (X,Y, co-ordinates) to polar space (range and angular position) is needed. The Pythagoras Processor is capable of these transformations at very high speeds making it suitable for use even in fast moving machines.

THREE DIMENSIONAL COORDINATE TRANSFORMS WITH THE PDSP16330

The PDSP16330 is designed to carry out the coordinate transform $x,y, \rightarrow r,\theta$. Many applications in robotics and target positioning require three dimensional transformations of the form $x,y,z \rightarrow r,\theta,\phi$. This application brief shows how the Pythagoras Processor PDSP16330 can be used for three dimensional transforms.

Fig. 1 shows a point x,y,z in a three dimensional space. If we move down the z-axis to the point x,y,0, we are at a point whose distance from the origin is $h = \sqrt{(x^2 + y^2)}$ whose bearing is arctan (y/x). The distance from the origin to the point x,y,z is therefore given by $r = \sqrt{(h^2 + z^2)}$ and the elevation of that point given by arctan (z/h). In this way the three dimensional transform $x,y,z, \rightarrow r,\theta,\phi$ has been decomposed into two 2- dimensional transforms which can be carried out by the Pythagoras Processor.

Fig. 2 shows the most obvious implementation of a 3-D transform, using two Pythagoras Processors. The first processor is given x,y as its input, providing the bearing and the distance to the point x,y,0. The second processor has z (suitably delayed to match the pipeline delay through the first processor) and h as its inputs giving r and ϕ as its outputs. Output θ from the first processor is delayed so that all three outputs suffer the same pipeline delay.

Fig. 3 shows an alternative realisation employing a single Pythagoras Processor. In this case x and y data are input on every other cycle, the alternate cycle inputs being z and h. The z input has a pipeline delay to compensate for the delay on h relative to x and y. This configuration will achieve a throughput of 5MHz, half that of the previous circuit.

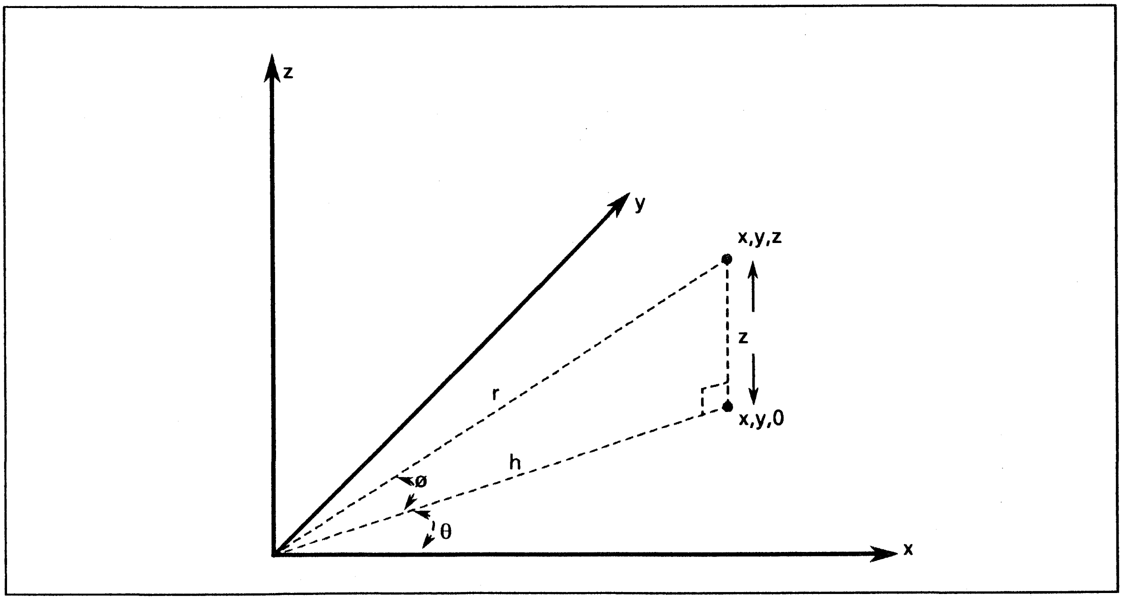

Fig.1

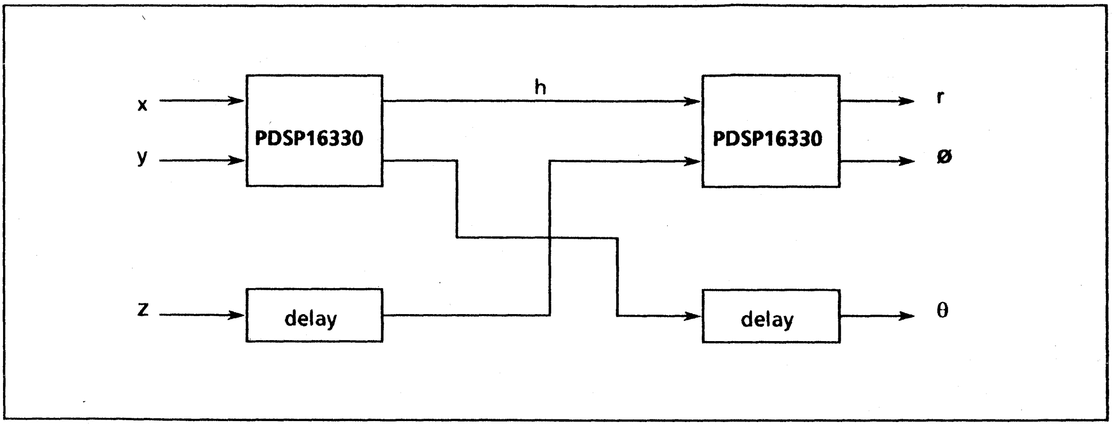
x
y
z
PDSP16330
delay
h
PDSP16330
delay
r
ø
θ

Fig 2

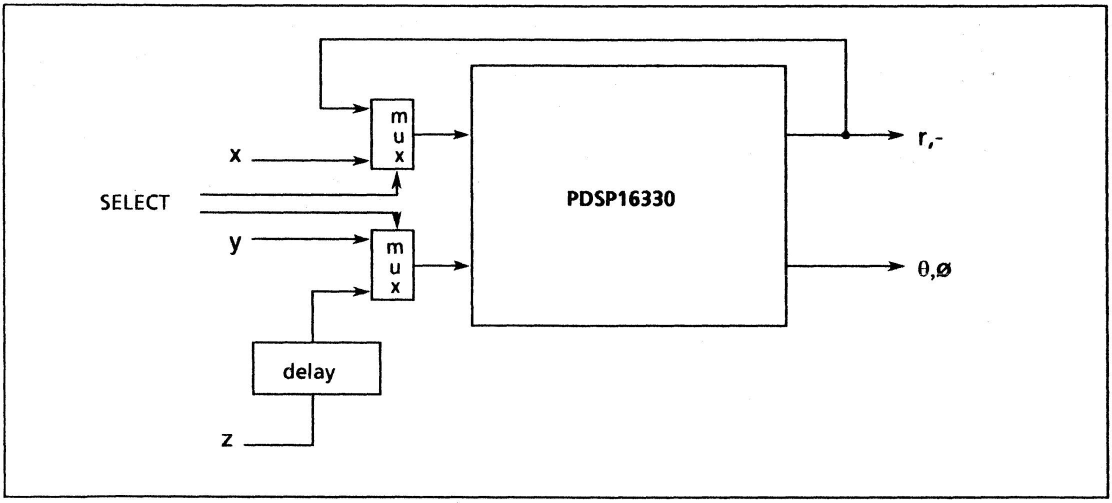
x
SELECT
y
z
m u x
m u x
delay
PDSP16330
r,-
θ,ø

Fig 3

THE PCS PDSP16510 SIMULATOR (Version V1.7)

PCS is a standalone program, running under DOS on PC platforms, which will perform a functional simulation of GEC Plessey Semiconductor's PDSP16510 FFT Processor. The model in the simulator is bit-accurate which means that it is accurate down to the least significant bit in replicating the function of the FFT processor.

The C language device model, which forms the heart of the PCS simulator, is supplied in source code format. This allows a system incorporating the PDSP16510 to be simulated, in software, such that its performance may be characterised and optimised prior to the prototyping stage.

The PDSP16510 FFT Processor

The PDSP16510 calculates Fast Fourier Transforms of up to 1024 points at sampling rates of up to 40MHz. It uses a decimation in time radix four algorithm to calculate both forward and inverse transforms. The device contains its own RAM, which allows new data to be loaded whilst transforming the last data block. The PDSP16510 is also able to dump transformed data whilst simultaneously loading and transforming other data blocks. These three operations: data load, data transform and data dump may all be performed concurrently, if required.

The PDSP16510 also provides built-in data windowing functions: both Blackman-Harris and Hamming windows are supported without the need for any external components, and without adversely affecting the data throughput rate of the device.

The PCS Simulator

For a device as complex as the PDSP16510, the need to be able to predict whether a given level of performance can be met, prior to system design and implementation, is great. It may prove to be difficult to realise this need if published performance data is the sole source of information. The principle of the PCS simulator is that it allows users to predict the in-system performance of this device with user-defined input data. The simulator also allows the user to ask "what if" questions by varying transform size and window operators to gauge their effect without the disadvantages of committing any design to hardware.

PCS offers five options from the master menu:

1. Edit Input Data File
2. Edit Control File
3. Run Single Pass Simulator
4. Run Batch Simulator
5. Return To DOS

1. EDIT INPUT DATA FILE

PCS provides a simple file editor which allows the entry of hexadecimal complex data points in a form suitable for entry to the PDSP16510 simulator. The editor offers block copy and block delete functions to ease the entry and correction of large data sets. The editor also incorporates a generate option which may be used to generate signals composed by the addition of a number of complex or real sinusoids of varying frequency and amplitude.

Upon selecting option 1 from the master menu, a two digit filename should be entered. All PCS filenames take the form

PCS##.EXT

##	=	a user-defined two figure decimal number
EXT	=	one of a series of extensions used to denote the file type: input data file, control file, output data file etc.

The menu, displayed at the bottom of the screen, offers the following options:

Fetch
Scroll
Copy
Insert
Delete
Generate
ESCape

An option is chosen by pressing the key shown in parentheses in the menu. This is normally the first letter of the option, e.g. F for Fetch.

The upper part of the screen is dedicated to displaying the input data entered in the selected file. The column headings in this area are Cyc, RIN and IIN. These are cycle number, real data component and imaginary data component respectively.

1. Fetch

Allows the user to move the display to a given cycle number.

2. Scroll

Scrolls the display upwards by one cycle each time it is invoked.

3. Copy

Allows a block of data to be copied from one part of the file and appended to the end of the file.

4. Insert

Used to insert data at a given point in a file. To insert the first data point in a new file, press carriage return when prompted for the cycle number.

5. Delete

Deletes all input data between two given points in the file.

6. Generate

This option is used to generate a signal consisting of a number of sinusoids of varying frequency and amplitude. For complex transforms, a maximum of four different sinusoids may be combined. For real transforms, each of the two real input signals may be composed of two different sinusoidal components.

Generate also offers the ability to load an externally generated tabular data file as an input data file. The file should have two columns, separated by at least one space character. The name should be of the standard PCS## format with the extension .DAT. Each column should hold 16 bit hexadecimal data; the first column representing real values and the second column imaginary values. An extract from such a file is shown below.

```
00FA 65B2
ABA3 562E
0000 34A2
12D2 8F02
```

For both internally and externally generated input data files, it should be noted that for an n point transform, the file should contain at least 3n data points. This is required such that the simulator has sufficient data to perform one complete load-transform-dump sequence. If a sufficient number of data points are not supplied, the data will be padded with zeros when loaded into the simulator. When Generate is used to create sinusoidal inputs, a file of the correct length is created automatically by the software.

Generate also has the added benefit of automatically generating a control file of the correct length. This is true for all three generate options: sinusoidal inputs, inverse transforms and externally supplied tabular data.

When Generate is selected, the following prompts will be displayed:

Forward or Inverse transform *	Transform type
Size	Size of transform
Real or Complex	Type R or C to select real or complex transform
Overlap	Percentage by which data is overlapped. Valid responses are 0%, 50% and 75%
External data supplied	Load data held in an external file ?

If external data is not supplied, then the subsequent prompts will request the frequencies and amplitudes of the sinusoids which comprise the input data stream.

Window Operator	You may choose a Rectangular, Hamming or Blackman-Harris window operator.
Overflow Detection (Y/N)	If the response to this question is Y, then the simulator will assume that bit 3 of the mode control word is set. See data sheet for further information.

* See example 3 for further details on inverse transforms

Once all the parameters have been entered, the samples will be generated and stored in the required file.

Pressing Escape returns control to the master menu.

2. EDIT CONTROL FILE

The control file defines the state of the INEN (Input Enable) and DEF (Reset) signals for each cycle of the input data file. If the Generate option is used to create the input data file, the control file will be generated automatically. The control file should contain three cycles for every data point in the data input file.

If the control file is longer than the input data file (.DIN) then the input file will be recycled until the end of the control file is reached. Conversely, if the control file is shorter than the input data file, the simulation will terminate at the last cycle defined in the control file.

If the control file is to be generated manually, the editing options available are similar to those for input data file creation; namely Fetch, Scroll, Copy, Insert and Delete.

Automatically generated control files contain nine preamble cycles, which are used to reset the device and, thereafter, the three cycles per data point required by the simulator. This means that for a 256 point transform, for example, the input data file will contain 768 points and the control file 2313 cycles.

3. RUN SINGLE PASS SIMULATOR

The simulator takes the control and data files specified by the user as the basis of the simulation run. The options offered by the simulator are described below:

1. Halt

This is used to temporarily or permanently halt a simulation run. A record is kept of how many cycles had been simulated and the simulator may be restarted from that point by using the Run option.

2. Run

Starts the simulator from the cycle shown in the status display.

Once the simulation run has been completed, interrupted by reaching a breakpoint or paused by the user, a new menu is displayed which offers the following options:

Scroll
Fetch
Extract
Escape

Scroll and Fetch are used to view the data in the same way as when editing an input data file. Extract generates a file, called PCS##.RES, which contains two columns of data; the first being the real component and the second the imaginary component of each frequency bin of the transform.

3. Breakpoint

Allows the user to define the cycle at which the simulator will pause prior to starting the simulation run, thereby allowing the results generated up to the breakpoint to be viewed.

4. Plot

Presents a graphical view of either the input or the output waveform. It also displays the magnitude of the complex output.

5. Initialise

Resets the cycle number to 1.

6. Variables

This option is used to define the variable watchlist, i.e. which variables will be displayed after each cycle when the simulation is completed or paused. Placing an 'X' next to a variable name adds that variable to the watchlist, pressing carriage return advances to the next variable. The variables that may be added to the watchlist are:

Input Variables:	RIN	Real input data
	IIN	Imaginary input data
Control Variables:	INEN	Active low signal which starts load process
	DEF	Active low signal which resets model
Output Variables:	ROUT	Real output data
	IOUT	Imaginary output data
	S3_0	Accumulated block floating point (BFP) shift value
	DAV	Data available signal
State Variables:	DUMP	Output data memory pointer
	LOAD	Input data memory pointer
	STAGE	Number of FFT pass
	GROUP	Number of butterfly group
	SAMPL	Number of butterfly
	WKSP	Beginning of work space address
	OPSP	Beginning of output space address
	MODE *	FFT mode word (Auxilliary Data Input)
	IPSHF	Butterfly input data BFP shift value
	OPSHF	Butterfly output data BFP shift value
	DLSHF	Delayed output data BFP shift value
	OVFLW	BFP overflow flag
Memory Variables:	MEM0-7	Memory word 0-7 x (offset + 1)

* Bit 15 of the MODE variable indicates the end of a valid load process and the start of the butterfly processing. This is not used in the real device. Refer to the data sheet for mode word format.

The display allows a total of twelve variables to be displayed simultaneously.

7. Escape

Returns to the master menu.

4. RUN BATCH SIMULATOR

This option has been included to speed up those simulation runs where a trace file, listing the values of a number of variables for each cycle, is unnecessary. In batch mode, up to six simulations may be performed serially, with each taking only two thirds of the time required in single pass mode, on average. The output of the simulator is a PCS##.RES file for each input file listing real and imaginary frequency bin data. The options offered by the batch simulator are:

1. Queue

The input data filenames are specified using this option. Up to six filenames may be entered, each separated by a carriage return.

2. Run

Initiates a simulation run for all the files selected via the Queue option.

3. Abort

This option will abort the simulation of the current input file and start the simulation of the next input file in the queue.

4. Halt

Aborts the current simulation and returns control to the user. Selecting Run at this stage will start the simulation process once more from the first file in the queue.

5. Escape

Returns to the master menu.

SIIMULATION EXAMPLES

Example 1 - Square wave using external data file

This example demonstrates how an external data file may be used to supply the data points for a simulation run.

The first task is to generate a text file, using a text editor, describing the square wave. The filename should be of the form PCS##.DAT, where ## is a two figure number between 00 and 99. The contents of the file will describe a square wave where the real component varies between zero and positive full scale values with the imaginary component always set to zero. Hence, the file should look like this:

```
0000 0000
(30 lines as above)
0000 0000
7FFF 0000
(30 lines as above)
7FFF 0000
```

This basic 64 line block should be repeated 12 times to yield a file 768 lines long. This is three times the required transform size, namely 256 points, in order to feed the simulator enough data points to perform a complete load-transform-dump cycle.

Now, invoke the PCS program and follow the steps listed below:

`<CR>`	Displays the master menu (note: <CR> = Carriage Return)
`1`	Selects Edit Input Data File option
`## <CR>`	Where ## is the name of the PCS##.DAT file containing the square wave data
`G`	Selects Generate option
`F`	Selects forward transform
`256 <CR>`	Selects 256 point transform
`C`	Complex transform
`0 <CR>`	0% overlap
`Y`	Signifies that an external .DAT file is the source of the input data
`R`	Use a rectangular window function
`N`	No overflow detection

The software will now start generating the .DIN file and displays the current sample number near the bottom of the screen. In total, 767 samples will be generated. When sample generation is complete, the prompt "View Data Y/N" will be displayed. Press N at this point to continue to the next phase.

Now type the following:

`3`	Selects single pass simulator option
`## <CR>`	Define input data file name
`P`	Select Plot
`I`	Select input signal
`R`	Selects real part of input signal
`1`	Zoom factor 1

The amplitude-time plot of the input signal is displayed. It can seen that it varies between zero and positive full scale with three cycles in total. To run the simulation type the following:

`ESCAPE`	Returns to the simulator menu
`V`	Select variables option. Select the following variables only: RIN IIN INEN DEF ROUT IOUT S3_0 DAV For the remainder, press <CR> when prompted
`0 <CR>`	To specify zero memory offset.
`R`	Starts the simulation run.

The simulator will run till cycle number 2313. When the simulation is complete, the trace file will be displayed. It lists the cycle number in the first column and the values of each of the watchlist variables after each cycle. Fetch and Scroll may be used to inspect this file. Extract creates an output file, named PCS##.RES, which lists the real and complex components of each of the frequency bins in tabular hexadecimal form. Pressing ESCAPE returns to the simulator menu.

When the simulator menu is displayed, type the following:

`P`	Select Plot
`M`	Magnitude of output data
`2`	Zoom factor 2

This will display the magnitude of the forward fast Fourier transform of the square wave. As expected, harmonics of the fundamental frequency can be seen with their amplitude decreasing as frequency increases.

Example 2 - 1024 Point With Sinusoidal Data

In this example, the Generate function will be used to create a waveform composed of two sine waves of different frequency and amplitude. To do this, follow the instructions below:

Invoke the PCS program and type the following:

<CR>	Displays master menu
1	Select Edit Input Data File
## <CR>	Where ## is the name to be given to the new data file
G	Invokes Generate function
F	Forward transform
1024 <CR>	Size of transform
C	Complex
0 <CR>	0% overlap
N	No external data file
10 <CR>	Frequency of first signal
0.8 <CR>	Amplitude of first signal
Y	Signifies that another signal is to be added to the input signal
20 <CR>	Frequency of second signal
0.4 <CR>	Amplitude of second signal
N	Signifies that no more sinusoids are to be added to the input signal
H	Hamming window
N	No overflow detection

The software now generates the required input waveform, which is 3071 cycles in length. The hexadecimal input data may be viewed at this stage. To continue to the next stage press N at the "View Data Y/N" prompt.

To simulate the device with the input defined above:

3	Selects single pass simulator
## <CR>	Specify input file name
P	Plot
I	Input waveform
R	Real component
1	Zoom factor

This displays the form of the input wave.

ESCAPE	Returns control to the simulator menu
V	Define watchlist. Select the first eight variables (as in example 1)
0 <CR>	Memory offset
R	Runs the simulation

The simulation runs to 6150 cycles. When complete, the trace file may be viewed. Pressing ESCAPE returns to the simulator menu. Now type:

P	Plot
M	Magnitude
3	Zoom factor

This will display the magnitude of the fast Fourier transform calculated by the simulator. The two peaks generated by the two sinusoids are clearly visible.

Note that higher frequency components peak is 6dB lower than that of the lower frequency peak. This is as predicted by theory.

Example 3 - Generating Inverse Transforms

Inverse transforms may be generated using output data files produced by PCS or by employing a user generated data file containing frequency data. The procedure employed to generate an inverse transform from both file formats is essentially the same.

1. User Generated Data Files

As with forward transforms, the PCS simulator accepts data files containing two columns of data, seperated by at least one space character, where the first column represents the real component of the data and the second column the imaginary component. For inverse transforms, each row of the data file represents a particular frequency bin; the first row being the D.C. bin and so on.

The inverse transformation of a pure sinusoid will be performed. This will illustrate clearly the link between the input data file containing frequency information and the output signal calculated by the simulator. The first step is to create a file 256 lines long, called PCS##.DAT. The contents of this file should be as shown below. This file may be easily generated using a line editor with block cut and paste facilities.

```
0000 0000
0000 0000
0000 0000
0000 0000
0000 0000
7FFF 0000
0000 0000
(248 lines of zero data)
0000 0000
```

Invoke PCS and type the following when the master menu is displayed:

1	Edit Input Data File
## <CR>	Enter name of .DAT file
G	Select Generate option
I	Inverse transform
256 <CR>	Selects 256 point transform
C	Complex data
0 <CR>	0% overlap
Y	Signifies that an external file holds the data
N	No overflow detection

The software will now start generating data samples. After the 255th sample, the warning "Insufficient data: padding with zero" will be displayed. This is due to the fact that the simulator always expects sufficient data to carry out a complete load transform dump cycle, whereas the data file contains only one third of the data required for such a cycle. However, padding the data with zeros gives the same end result. When the prompt "View Data Y/N" is displayed, press N to continue to the next phase:

3	Selects single pass simulator
## <CR>	Define input data filename
R	Run simulation

The simulation will run for 2313 cycles. Now type:

<ESCAPE>	Returns to simulator menu
P	Select plot option
O	Output signal
R	Real component
2	Zoom factor

The display will show a sinusoid whose first peak is at the left hand edge of the graph. Now type the following:

<CR>	To select a further plot
O	Output signal
I	Imaginary component
2	Zoom factor

The display shows another sinusoid of equal amplitude and frequency to the first, but lagging in phase by 90°.

2. PCS Generated Data Files

All PCS simulation runs allow the user to produce a .RES file listing the contents of each of the frequency bins in tabular format (see section 3, paragraph 2 for information on the Extract option). The data contained in this file is in the correct format to be fed directly into PCS as a .DAT file.

In order to use the .RES file in this manner it must be renamed from PCS##.RES to PCS##.DAT via the DOS RENAME command. Ideally a new filename should also be chosen, e.g. rename PCS01.RES to PCS02.DAT. This reduces filename confusion as other files with the name PCS01 will already exist in addition to PCS01.RES. Once the file is renamed, the procedure for calculating an inverse transform is exactly the same as that described above for user generated .DAT files.

Additional Example Files

The PCS distribution disk contains three external data files comprised of test data taken from the silicon compiler used to design the PDSP16510. Multiple 64 point transform data is provided in the file FFT64.DAT which includes four different signal blocks of 64 points each. Thus the 256 output points show four different responses in sequence. Normal 256 point transform data is provided in the file FFT256.DAT and 1024 point transform data is provided in the file FFT1024.DAT. All files are run with a rectangular window function applied to the data. To use these files they must first be copied to files which PCS recognises as external data files, i.e. they should be copied to filenames of the form PCS##.DAT. See example 1 for further information on external data files.

APPENDIX 1

PDSP16510 FFT Processor Model

Although the internal algorithms have been developed to mimic those in the device, the internal structure and method of processing are merely functional and not gate level equivalents. Pipeline effects on data values have been accounted for but the model does not provide an exact representation of device operation in real time.

The interface for this model has been rationalised for clarity and does not provide all the control signals used by the real device. The model incorporates complex data I/O paths which are controlled by INEN and DEF signals. These are considered sufficient to provide an accurate functional simulation of the device. Load and dump clocks are simulated internally and although load and dump rates are varied according to the active FFT process they are obviously tied to the system clock, represented by the simulator call.

The real device performs four butterflies in 12 system clock cycles. This model incorporates only one data path structure and is set to process one butterfly every three cycles, thus completing four butterflies in 12 cycles or calls from the simulator. This infers that for transform sizes with fewer or an equal number of butterflies (16, 64 or 256 point), each sample in the data file must be entered over three cycles. Output data will also be spread over three cycles. In order to fit 1024 point transform simulations into the maximum cycle count of 9999, the butterfly process rate has been increased to one per cycle. The data file need only repeat sample data over two cycles to allow one complete transform process (1280 cycles) to complete within the loading period (2048 cycles). This mechanism was developed to permit the model to handle the extended I/O periods required by overlapped transforms and does not represent any structure in the real device.

The GENERATE function in the input data editor allows for this pseudo clock structure. If you are using the INSERT function, this clock structure must be accounted for.

Internal state variables are used merely to monitor the state of the FFT process and do not represent any registers, storage or control structures in the real device. The memory area is defined as an 8x256 word array and the simulator will trace any column of 8 words by entering a memory offset value from 0 to 255 in the memory variable list. The memory is partitioned by each FFT process according to its particular needs. All FFT processes use a load space in which to deposit input data, a work space in which the butterfly processes are executed and an output space from which completed transforms are clocked out. This arrangement permits loading, processing and dumping of data to be simultaneous processes. The load space is normally twice the size of the selected transform to permit overlapped sampling modes. The work space and output space are equal to the size of the transform. Thus, for a 64 point transform the load space is 128 words long (offsets 0 to 15), the work space is 64 words long (offsets 16 to 23) and the output space is also 64 words long (offsets 24 to 31).

Model Limitations

Although the model was designed to be bit accurate, tests indicate that bit 0 of the result is not always the same as that produced by the PDSP16510 when processing identical data.

IMPLEMENTING LARGE AND NON STANDARD TRANSFORMS WITH THE PDSP16510 FFT PROCESSOR

BACKGROUND

The PDSP16510 is a stand-alone FFT Processor which performs 16, 64, 256, or 1024 point FFT's with input sampling rates of up to 40MHz - typically an order of magnitude faster than programmable DSP parts. A single device can window and transform up to 1024 complex points without the need to access external memory during the computation. The device internally uses 16 bit block floating point arithmetic, which provides sufficient precision to allow transform sizes of 4096 points, or even 16384 points in some circumstances, to be implemented with adequate dynamic range.

The purpose of this application note is to describe how the PDSP16510 can be used in systems requiring transform sizes of up to 16384 complex points. In such applications it is necessary to support the PDSP16510 with other arithmetic devices from the PDSP family. Several alternatives are presented but the optimum architecture has a regular structure, is easy to control, and can be repeated as required to accommodate high input sampling rates.

DOING 512 POINT TRANSFORMS

Before going on to consider transform sizes of greater than 1024 points, the special case of the 512 point transform will first be considered. The radix 4 algorithm used by the PDSP16510 requires that the transform length be a power of four, so how can this be modified to handle 512 points?

The simplest way is to zero pad the input data, perform a transform of twice the required length and then discard half the spectral results. There are, however, two possible ways of inserting the zero's; either an equal number of zeros as data points can be appended to the end of the data block or a zero can be inserted between each data point. Appendices A and B present the mathematics of both these techniques, but in order to exploit the window and overlap facilities of the PDSP16510, it is essential to use the interleaving zero approach.

Rember that the PDSP16510 will apply its internal window operator to the complete 1024 point block, and not just to the 512 samples of actual data. Thus, if all the zero's are appended at the end, the data will be incorrectly modified since the actual data will only use half the window operator. In fact if block overlapping is used there will be a phase shift in the results, even for the case of a rectangular window. By interleaving with zero's, the actual data will be modified by every other value in a window containing 1024 operators. This is the correct value for 512 samples using a window with 512 operators.

In the case of interleaving zero's Appendix B proves that the spectrum is repeated in the second half of the outputs. Thus, even though a 1024 point transform is calculated, only the first 512 results need be outputed. The PDSP16510 provides an option to only output half the results, and hence improve the efficiency of the calculation. When zero's are appended at the end, the required spectrum is contained in the even results. The PDSP16510 does not provide an option to output only even results.

A suitable hardware configuration for performing 512 point transforms is shown in Figure 1. The PDSP16540 Bucket Buffer is needed to support continuous transforms, and also to implement block overlapping. Data is written to the buffer at twice the

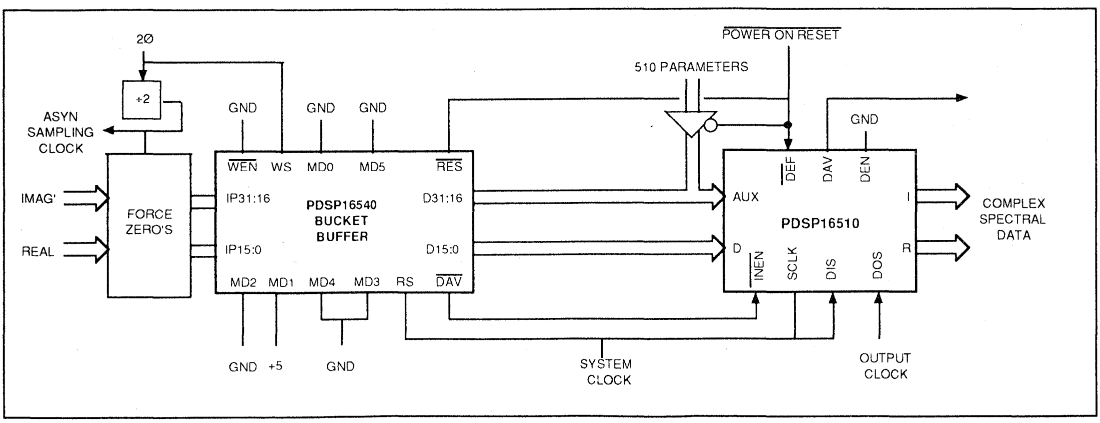

Figure 1. System for Performing 512 Point Complex Transforms

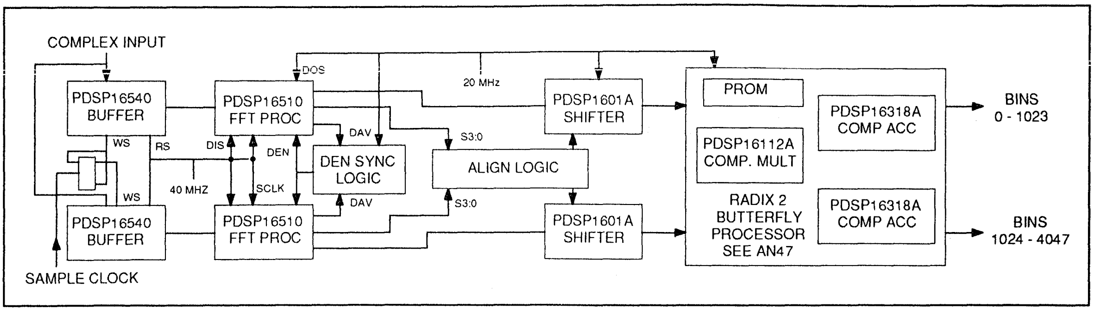

Figure 2. The Direct way to Implement 2048 Point Complex Transforms

original sampling rate, with external logic inserting zero's during the odd clock periods. Such a configuration will support a maximum data sampling rate of 3.4 MHz, but this can be increased to 20 MHz by connecting six devices in parallel. This multiple device arrangement is described in the PDSP16510 data sheet, and it should be noted that the Bucket Buffer is then not needed. The sampling rates achievable with a given number of devices will always be half those obtainable when doing true 1024 point transforms.

This highlights the disadvantages of the zero's insertion technique; the continuous sampling rates achievable are only half those obtainable when doing straightforward transforms of the same length. This stems from the fact that in the time taken to complete the whole transform operation, only half as many actual samples must have been written into the input buffer. Other system constraints must thus have dictated the need to only do 512 point transforms, rather than the transform time itself.

The techniques described in the following sections indicate how large transforms can be implemented, and are also applicable to 512 point transforms. In essence they trade the simple parallel approach to increasing the sample rate with additional multipliers, PROMS, and accumulators.

THE DIRECT WAY TO IMPLEMENT LARGE TRANSFORMS

Equation 1 in Appendix A indicates that an N point DFT can be expressed as the summation of the even and odd, N/2 point, DFT's; in fact a summation of the even points and the odd points twiddled by the otherwise missing sine and cosine values. Appendix A goes on to show that Equation 1 only represents the first half of the N point DFT. The second half is shown to be obtained by a twiddle and difference operation.

The direct way to implement a 2048 point transform would thus be to use two PDSP16510's; one performing a 1024 point transform on the odd inputs, and the other performing a 1024 point transform on the even inputs. As shown in Figure 2, the addition of a PDSP16112 complex multiplier and two PDSP16318 complex accumulators will combine the results to simultaneously produce both halves of the 2048 point transform. The use of the PDSP16540 Bucket Buffer allows the incoming data to be continuous, and any amount of block overlapping can be selected.

In this system,both PDSP16510's must produce their results simultaneously. This can only be gauranteed if the outputs are controlled by the DEN input. When both DAV outputs have gone valid, a DEN signal should be produced which is synchronized to the DOS strobe. The even sequence will then be exactly concurrent with the odd sequence, once the output circuit has been primed (see the DSP Handbook page 134). If the derived DEN signal is delayed by 14 DOS strobes (4 PDSP16510 priming delays plus 8 PDSP16112 delays plus 2 PDSP16318 delays) it can be used to provide a Data Valid signal to the rest of the system.

One of the complex accumulators can be avoided if the results produced by the PDSP16510's are read twice, firstly to produce bins 1-1023 with an add operation, and then bins 1024 - 2047 with a subtract operation.

This configuration has two shortcomings; the internal window operator cannot be used and shifters are required to align the

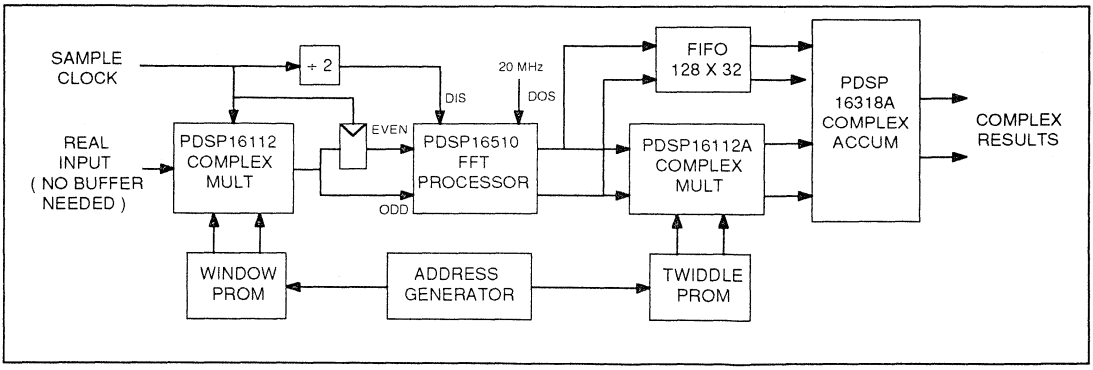

Figure 3. A System to do 512 Point Real Transforms

outputs from the PDSP16510's. The system thus needs an additional complex multiplier to window the data, which cannot be combined with the post twiddle operation needed in the backend butterfly processor. This is not shown in Figure 2. The shifters are needed since it cannot be guaranteed that both devices will produce the same block exponent from the internal variable shift operation, and turning this feature off would produce very poor dynamic range. PDSP1601's can be used for the shift operation, but unfortunately they have to be positioned before the combination of the results, and four devices are needed to shift two sets of 32 bit outputs. More efficient configurations are discussed in the next section.

This direct approach to computing larger transforms can, however, be used to produce an efficient system for 512 point real transforms. Such a system is shown in Figure 3, which illustrates a PDSP16510 configured to perform two simultaneous 256 point real transforms. This system will support incoming sampling rates of up to 19.5 MHz when doing 512 point real transforms (see Table 5 in the datasheet). With 50 % overlapping the rate reduces to 9.7 MHz, and with 75% overlapping it reduces to 4.8 MHz. To support these input rates the output rate should be at least equal to 19.5 MHz sampling rate. Since the PDSP16510 would need a 40 MHz system clock to achieve these throughputs, a convenient solution would be to divide this clock by two and then to use it as the output strobe. Both the PDSP16318A and PDSP16112A will support output clock rates of 20 MHz.

In this system the even samples are applied to the real input pins of the FFT Processor, and the odd samples to the auxiliary pins. No output shifters are needed since both sets of outputs will be internally shifted by the same amount. Since the PDSP16510 will output all the even results followed by all the odd results, it is necessary to buffer the even results in a FIFO. As the odd results are outputed they are twiddled and added to the even results coming from the FIFO.

In the case of real transforms the second half of the spectrum is a repeat of the first half. Thus the PDSP16510 will only output 128 even bins followed by 128 odd bins, even though two sets of 256 samples were applied. Only one PDSP16318 is needed at the output doing a complex add operation; the second half of the results do not exist and the twiddle and subtract is not needed.

This system can use the internal block overlapping features of the FFT Processor. Unfortunately the internal window operator cannot be used, and to perform this function a complex multiplier is needed at the input as shown in Figure 3.

MORE EFFICIENT WAYS OF IMPLEMENTING LARGE TRANSFORMS

Appendix C shows how an N point DFT can be performed by a combination of L and M point DFT's where N = L x M. After the L point transform the results are twiddled by exp(-j2πms/N) where m = 0 to M-1, and s is the row index. So how should L and M be chosen to make optimum use of the PDSP16510 and its supporting devices?

The most obvious approach is to use two FFT Processors, with an intermediate memory to store the results from the column transforms, and a complex multiplier to apply the intermediate twiddle factor. But this solution has three drawbacks; the memory cannot be a simple FIFO, since an address translation from columns to rows is needed, the internal window operators cannot be used, and results from the column transform will have different scaling factors. The last problem can be avoided by restricting the column transform to 16 points and turning off the internal block floating point option. With such a small transform size the dynamic range should not be compromised, and with the second FFT Processor doing 256 point transforms the system would handle 4096 points. In such a system, however, the window operator can only be applied to the data before it is stored in the input buffer. Thus the system needs two additional complex multipliers, one for the window operator and one for the intermediate twiddle.

This requirement leads to the possibility of making better use of the complex multipliers. Rather than using two PDSP16510's why not make the columns very small and use a complex multiplier and accumulator to perform a straightforward DFT, rather than an FFT? Thus a 2048 point transform can be done by calculating a 2 point DFT followed by a 1024 point FFT. The calculation of a 2 point DFT requires no multiplications, and the window and intermediate twiddles can be combined into one multiply operation.

We have thus chosen L in Appendix C to be 2, and M to be 1024. The row index, s, is thus 0 or 1. The data is thus arranged with samples 0 - 1023 in one row and 1024 - 2047 in the other row. **One** point of each of the two point DFT's is then calculated using samples 0 and 1024, followed by 1 and 1025, up to 1023 and 2047. The even point DFT calculation is just the sum the two inputs, and, since s = 0 in the first row, then the intermediate twiddle is unity . The resulting data is loaded into the PDSP16510,

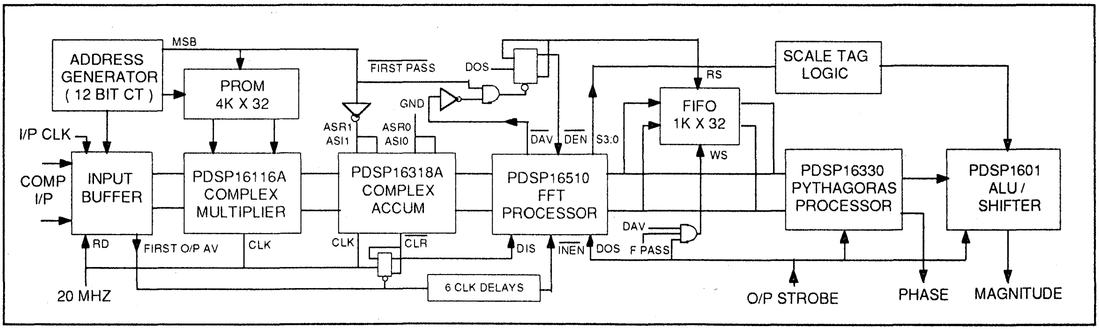

Figure 4. An efficeint system for performing 2048 Point Transforms which can adapted for up to 16384 Points

which does a 1024 point transform to calculate the even results of the 2048 point transform that is required. Since the twiddle values are unity, the PROM should contain the unmodified window operators which are read in the same sequence as the data from the input buffer i.e. 0, 1024, 1, 1025 etc.

The second point of the two point DFT is now calculated, and is simply the difference of the same sequence of values used previously. These must then be twiddled by $\exp(-j2\pi m/2048)$,before being transformed by the PDSP16510 to produce the odd values of the 2048 point transform. This can be done prior to the DFT by storing a second set of window values in the PROM which have been modified by these twiddle factors.

A system which will perform the necessary calculations is shown in Figure 4. The input buffer must handle any necessary block overlapping, and must allow sample 0 to be read out first, followed by sample 1024, then sample 1 and so on. The PDSP16116A performs the necessary complex multiplication using values from the PROM. The PROM must contain a set of 2048 window values plus a second set of modified window values as previously explained. A simple 12 bit counter can be used to provide the address sequences, with the most significant bit providing a page address bit, and the least significant bit used to provide the most significant address bit in each page.

The DEN and FIFO read strobe logic, shown in Figure 4, assumes that the odd results are calculated first. These are then dumped as quickly as possible into a FIFO. The outputs from this FIFO can then be mixed with the even results as they are produced by the FFT Processor during the next pass. The DEN pin on the PDSP16510 can be used to cause its outputs to go high impedance on alternate DOS output strobes. The real and imaginary complex outputs for the PDSP16510 can be converted to magnitude and phase by the PDSP16330 Pythagoras Processor.

The PDSP16318 does a complex subtraction whilst the first set of intermediate values are read into the PDSP16510, and then does an addition for the second set of inputs. Since the arithmetic could generate a 17 bit result, it is necessary to set the shifter within the PDSP16318 such that the 16 output pins for each complex component discard the least significant bit from the accumulator (S2:0 = 011).

The sequence of events within the PDSP16318 is as follows, and produces an input for the PDSP16510 on every other clock cycle (after an initial delay of 3 clock cycles).

1) Load complex point A0. Clear the accumulator (after the next clock edge)
2) Load complex point B0. Transfer A0 to the accumulator (ALU function is Accumulator + A0)
3) Load point A1, Clear the accumulator (after the next clock edge). Do Accumulator + B0 and load the output register.
4) Load point B1. Transfer A1 to the accumulator. A0 + B0 available at the output pins.
5) Load point A2. Clear the accumulator etc

The two sets of outputs produced by the PDSP16510 will have different bit significance, as indicated by the Scale Tag values. This tag indicates the number of left shifts which have occurred in order to compensate for too many right shifts introduced to prevent possible overflow in the internal data path. Thus a smaller scale tag indicates a larger output number. To preserve dynamic range it is necessary to keep the larger set of results, thus the results with the larger scale tag value must be shifted right (divided by 2) by the difference in the scale tag values. This can be done using the barrel shifter within a PDSP1601 ALU. This normalization is best done on the 16 bit magnitude output from the PDSP16330, since the phase outputs will always be correct without any shifts.

The same arrangement can be used to calculate transforms of 4096 points if the PDSP16116 and PDSP16318 are used to generate **each point in turn** of a 4 point DFT. Similarly 8192 points can be transformed if each point of an 8 point DFT is calculated,

FIRST PASS	SECOND PASS	THIRD PASS	FOURTH PASS
Wd' = Wd d = 3072 - 4095	Wd' = Wd. $\exp(-j3\pi/2)$. $\exp(-j2\pi m/4096)$ d = 3072 - 4095, m = 0 - 1023	Wd' = Wd. $\exp(-j3\pi)$. $\exp(-j4\pi m/4096)$ d = 3072 - 4095, m = 0 - 1023	Wd' = Wd. $\exp(-j9\pi/2)$. $\exp(-j6\pi m/4096)$ d = 3072 - 4095, m = 0 - 1023
Wc' = Wc c = 2048 - 3071	Wc' = Wc. $\exp(-j\pi)$. $\exp(-j2\pi m/4096)$ c = 2048 - 3071, m = 0 - 1023	Wc' = Wc. $\exp(-j2\pi)$. $\exp(-j4\pi m/4096)$ c = 2048 - 3071, m = 0 - 1023	Wc' = Wc. $\exp(-j3\pi)$. $\exp(-j6\pi m/4096)$ c =2048 - 3072, m = 0 - 1023
Wb' = Wb b = 1024 - 2047	Wb' = Wb. $\exp(-j\pi/2)$. $\exp(-j2\pi m/4096)$ b = 1024 - 2047, m = 0 - 1023	Wb' = Wb. $\exp(-j\pi)$. $\exp(-j4\pi m/4096)$ b = 1024 - 2047, m = 0 - 1023	Wb' = Wb. $\exp(-j3\pi/2)$. $\exp(-j6\pi m/4096)$ b = 1024 - 2047, m = 0 - 1023
W'a = Wa a = 0 - 1023	Wa' = Wa. $\exp(-j0)$.$\exp(-j2\pi m/4096)$ a = 0 - 1023, m = 0 - 1023	Wa' = Wa. $\exp(-j0)$. exp(-j4πm/4096 a = 0 - 1023, m = 0 - 1023	Wa' = Wa. $\exp(-j0)$. $\exp(-j6\pi m/4096)$ a = 0 - 1023, m = 0 - 1023

Figure 5. Modified Window Operators needed to perform 4096 point Transforms with one external multiplier

and16384 points can be transformed if 16 point DFT's are done. This will be illustrated using a 4096 point transform.

The original 4096 samples must be arranged in four rows, each containing 1024 columns. Thus Row 0 contains samples 0 - 1023; Row 1 contains samples 1024 - 2047; Row 2 contains samples 2048 - 3071; and Row 3 contains samples 3072 - 4095. To begin the complex multiplier / accumulator calculates the first value in the 4 point DFT using samples 0, 1024, 2048, and 3072. It then calculates the first value using samples 1, 1025, 2049, and 3073 and so on until the first points of all 1024 four point DFT have been calculated.

From the DFT equation given in Appendix A, each value is calculated by a twiddle and summation of the four inputs. The twiddles are expressed as exp($-j2\pi kn/4$) where n is 0, 1, 2, or 3 for the four samples selected, and k = 0 for the calculation of each of the first points or X(0) values. Thus for the trivial case when k = 0;

$X(0) = x(0) + x(1024) + x(2046) + X(3072)$ where X(0) is the first intermediate DFT value.

This calculation is repeated using samples 1, 1025, 2049, and 3073 etc. and the 1024 results are applied to the PDSP16510. This produces results which represent bins 0,4, 8,12 etc. in the 4096 transform actually required.

X(1), X(2), and X(3) must then be calculated using the same groups of 4 inputs, using k = 1 for the calculation of X(1), k = 2 for the calculation of X(2), and k = 3 for the calculation of X(3). These produce results from the PDSP16510 representing bins 1, 5, 9, 13 etc.; followed by 2, 6, 10, 14 etc.; followed 3, 7, 11, 15 etc.

In order to sequentially output the final 4096 bins it is necessary to provide three FIFO's to store the various sets of results. These are then combined with the last set of results before being converted to phase and magnitude by the PDSP16330A.

The PDSP16318 performs the summation of four inputs, and a two bit word growth can occur. The two least significant bits are thus ignored by setting the shift control S2:0 to 010. The results from the four 1024 point transforms will have different scale tag values, and the results must be normalized to the largest set. This is done by detecting the minimum scale tag value and subtracting it from each of the other values. Each set of results is then shifted right by this amount using a PDSP1601.

These twiddles for the 4 point DFT can be combined with the window operator and the intermediate twiddle. This requires that the 4096 original window operators are arranged in four sets, each of which contains four groups of 1024 different values. Each operator in each group is then modified by exp($-j2\pi n/4$) where n is the group number form 0 - 3. The four groups in each set are further modified by exp($-j2\pi ms/4096$) where m = 0 -1023 and s = 0 in the first set, 1 in the second set, 2 in the third set, and 3 in the fourth set. This arrangement is illustrated in Figure 5.

The total PROM size must be 16384 words, addressed as four pages of 4096 words. The first page is used when the first points are being calculated using an address sequence of 0,1024, 2048, and 3072 (to calculate the first point in the first DFT); followed 1,1025,2049, 3073; and so on up to 1023, 2047, 3071, and 4095. The next page is then used to calculate all1024 second points of the DFT, then another page to calculate all the third points, and finally the last page is used to calculate all the fourth points.

Performance

The performance of the system is dictated by the time taken to do the following operations ; load the PDSP16510 with the results of the column DFT's, do a 1024 point transform, and then normalize and produce phase and magnitude outputs. These operations must be repeated for the calculation of every point in the column DFT. Since the maximum clock frequency of the PDSP support devices is 20 MHz, it will take 100 nanoseconds to calculate one point of a 2 point DFT, or 200 nanoseconds to calculate one point of a 4 point DFT, or 400 nanoseconds to calculate one point of an 8 point DFT, or 800 nanoseconds to do one point in a 16 point DFT.

Thus the time taken to load 1024 intermediate results into the PDSP16510 is 102.4 microseconds when 2 point DFT's are done. The transform time itself takes 97.7 microseconds with a 40 MHz system clock, and to this must be added the time to dump the results. It should be noted that the dump rate of the PDSP16510 can be solely dictated by the requirements of the system.

It will actually support dump rates of 40 MHz, and these rates can be sustained if each set of results is loaded into a 40 MHz FIFO (not shown in Figure 4). The dump time is then only 25.6 microseconds as far as the performance calculation is concerned, and the total time needed to produce 2048 results is 2 x (102.4 + 97.7 + 25.6) = 451.4 microseconds. This corresponds to an input sampling rate of 4.5 MHz. It should be noted that the results need only be read out of each FIFO before the next load, transform, and FIFO write operation is complete i.e. data can be read out at the input sampling rate of 4.5 MHz. The combined rate going into the PDSP16330 is then 9 MHz, and standard grade parts can be used. From the complete system point of view, it might be more convenient if the output clock is obtained by dividing down the 40 MHz system clock needed by the PDSP16510. Either 5 or 10 MHz read rates could then be used, but an A grade PDSP16330 would be needed in the latter case.

This 4.5 MHz input sampling rate is the maximum sampling rate possible with this arrangement, and is only achievable when two 40 MHz output FIFO's are provided. Figure 4 , in fact, only shows the need for one output FIFO. In this situation the performance of the PDSP16330 and PDSP1601 limit the input sampling rate, and A grade parts should be used to allow 20 MHz outputs.The results from the second 1024 point transform can then only be dumped at 10 MHz, and are combined with the FIFO output to give a 20 MHz stream into the Pythagoras Processor. A more conservative performance figure is thus obtained if only one 20 MHz FIFO is assumed to be present, which results in a dump time of 51.2 microseconds for one transform and 102.4 microseconds for the other. The total 2048 point calculation time is then 553.8 microseconds, which corresponds to an input sampling rate of 3.7 MHz.

The above calculations are repeated for 4096, 8192, and 16384 point transforms and the results are summarized in Table 1. Both the maximum performance figures using two FIFO's, and the more conservative figures using one slower speed FIFO are given.

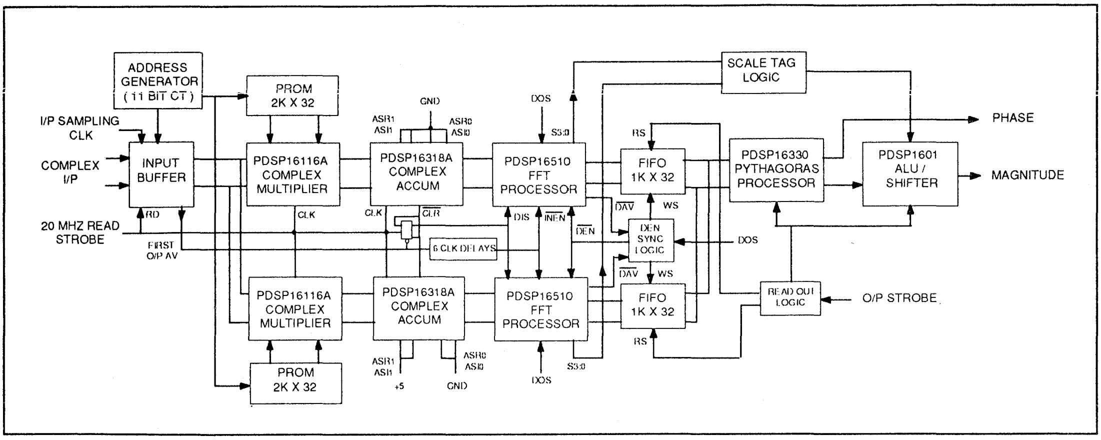

Figure 6. A 2048 point system using two PDSP16510 devices for increased performance

These performance figures can be improved by supplying a complete data path to compute each value in the DFT. Thus two data paths are needed to compute a 2048 point transform, four paths for a 4096 point transform, eight paths for a 8192 point transform, and 16 paths for a 16384 point transform. All the points in any given DFT are then produced simultaneously and applied to their own PDSP16510.

Figure 6 shows the use of two data paths in a 2048 point system, which will support sampling rates up to 9 MHz. These rates can be accomplished with only one PDSP16330A and one PDSP1601A. Because of internal synchronisation within the FFT Processor, one of the data paths may have finished its transform before the other. The paths can be pulled into synchronization by detecting when both DAV signals have gone active, and then generating a common DEN signal which has been synchronized to the DOS strobe. The slower read operations can commence as soon as the FIFO's are not empty. Alternatively both FIFO's can be allowed to fill in their own time , and data then read out when both are full. In this arrangement read and writes need not overlap, but sufficient space must be available in the FIFO's for the next set of results.

Figure 7 is a generic arrangement suitable for any number of data paths, which might be needed to sustain a particular sampling rate. It actually shows four data paths doing a 4096 point transform, with the results going into four FIFO's. These are then read out one after the other to give the sequential results of the 4096 point transform. The four sets of modified window operators are now split between four individual PROMS, but the total contents are the same. The results achieved are also summarized in Table1, and are given for FIFO's with both 40 MHz and 20 MHz writing rates. The reading rates needed are much lower, in fact the rate of combined data going into the PDSP16330 need only be the same as the input sampling rate. The reading rate of the individual FIFO's is proportionally less than this, and depends on the number of data paths in the system.

If the performance achieved with the full complement of additional data paths is too high, then the number of paths can be reduced to suite the sampling rates required.

This technique can be modified to perform 512 point transforms, with the PDSP16510 then doing 256 point complex transforms. In this mode load and dump operation can be concurrent with internal transform operations. With a 40 MHz system clock the transform time is 20.4 microseconds, but the load time is dictated by the 10 MHz maximum rate of producing DFT values from the PDSP16318. This results in a load time of 25.6 microseconds, which is greater than the transform time. The transform time is thus not the limiting factor, and it could be increased to 25.6 microseconds by using a slower system clock.The dump time can also be 25.6 microseconds without restraining the system level performance.

	MAXIMUM SAMPLING RATES			
	ONE DATA PATH		L COMPLETE DATA PATHS	
TRANSFORM SIZE	20 MHz FIFO	2 x 40 MHz FIFO	20 MHz FIFO'S	40MHz FIFO'S
512 (L = 2)	10MHz	- - - - - - - -	20 MHz	- - - - - - - -
2048 (L = 2)	3.7 MHz	4.7 MHz	8.1 MHz	9 MHz
4096 (L = 4)	2.6 MHz	3.1 MHz	11.5 MHz	12.4 MHz
8192 (L = 8)	1.7 MHz	1.9 MHz	14.6 MHz	15.3 MHz
16384 (L = 16)	1 MHz	1.08 MHz	16.9 MHz	17.3 MHz

Table 1. Maximum Sampling Rates

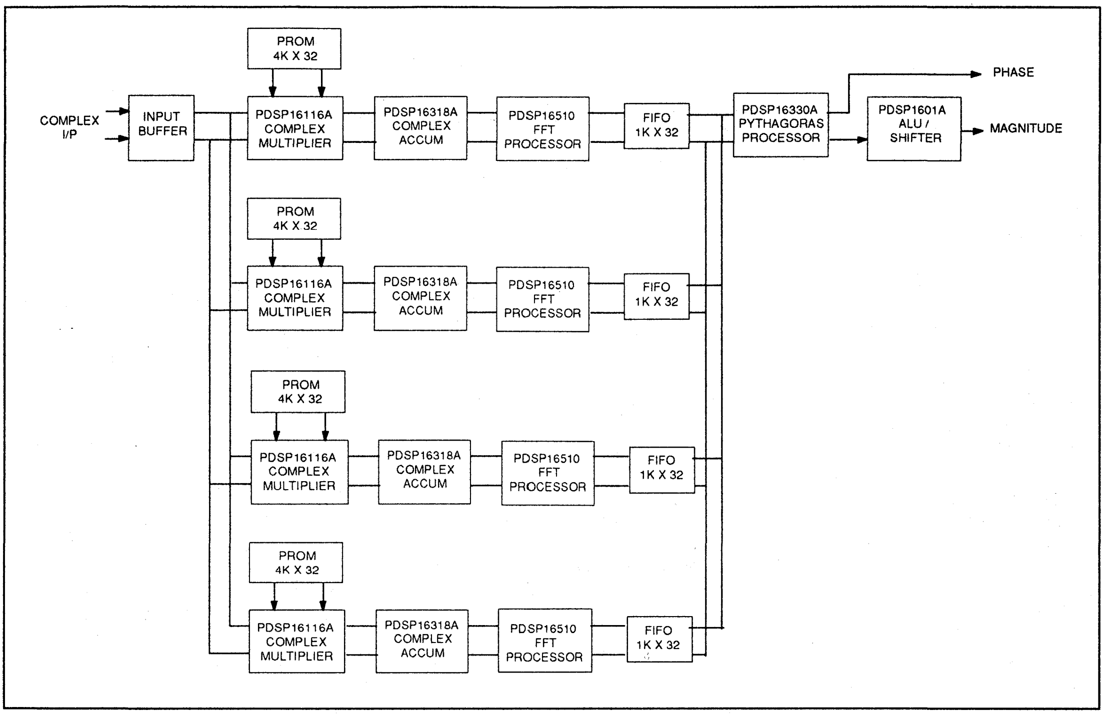

Figure 7. A high performance 4096 point system using four complete data paths

This level of perfromance can be sustained with only one FIFO to store the first set of results, and then dumping the second set of results at 10 MHz. The total time needed for the PDSP16510 to perform two 256 point transforms is thus 2 x 25.6 microseconds, which equates to an input sampling rate of 10 MHz. Two data paths would increase this sampling rate to 20 MHz.

INPUT BUFFERING NEEDS FOR THE PDSP16510 FFT PROCESSOR

BACKGROUND

The PDSP16510 FFT Processor contains 1K x 32 bits of internal RAM, enough to provide working memory for up 1024 point complex transforms. Once this memory is loaded with data no further external intervention is needed. On every pass of the transform, data is read from the RAM using the correct address sequence, and then written back after the butterfly operation. When the data has been completely transformed, it must be read out of the RAM and transferred to the next system element before new data is loaded.

All this, of course, takes many clock cycles, and in the meantime new data is being collected by the acquisition system. If this data is not to be lost it must be stored somewhere for future processing.

For transform sizes up to 256 points this presents no problem to the PDSP1510; it contains sufficient RAM to provide both working storage and input buffering for new data. In fact it contains sufficient RAM to also provide output buffering. This allows a single FFT Processor to handle higher sampling rates than might otherwise be expected.

The input and output buffers allow the time taken to transfer data in and out of the device to be effectively lost at the system level. Thus, whilst the working RAM is being used to transform a set of data, the output buffer can be dumping data previously transformed and the input buffer can be acquiring data to be next transformed. At the system level data is being continuously transformed, assuming, of course, that the time taken to do a transform is no greater than the time taken to load a new set of data.

When 1024 point transforms are to be undertaken no additional internal buffering is possible. Concurrent load, transform, and dump operations are thus not possible, and incoming data must be externally buffered if no information is to lost whilst a transform is in progress. For continuous transforms, the time taken to load this buffer must be greater than or equal to the sum of the time taken to read data from the buffer into the PDSP16510, then to transform it, and finally to transfer to the next device.

At first sight this buffer could be a simple FIFO, albeit a very wide 32 bit FIFO. It would also needs read rates of 40 MHz if maximum throughputs are to be possible. Once this FIFO were full additional logic would have to ensure that at least one location was transferred to the FFT Processor before the next word was written.

Many DSP applications, however, need to overlap data sets in the time domain before they are transformed to the frequency domain. This is the result of the need to apply a window operator to the data before it is transformed. Since only a finite segment of a signal can be observed at any time, discontinuities at the edge of the segment will introduce spectral errors. These are minimized by applying a window operator which weights the data more in the middle and less at the edges. There is thus a danger of missing some information at the edge of the segment, and this is avoided by overlapping the segments. Typically segments need to be overlapped by 50% or 75% to avoid loss of information, but the greater the overlap the less is the data sampling rate that can be achieved.

Overlapping data sets implies that old data must be re-read before new data is appended; an impossible task with a FIFO. For this reason the PDSP16540 Bucket Buffer has been introduced to support the FFT Processor. It allows data sets to be overlapped in 32 word increments, and requires no supporting logic. Although primarily designed to support 1024 point transforms, it can in fact help in smaller cases. The PDSP16510, itself, then supports 50% or 75% overlapping, but the PDSP16540 can be used when different amounts of overlapping are needed. This is discussed later.

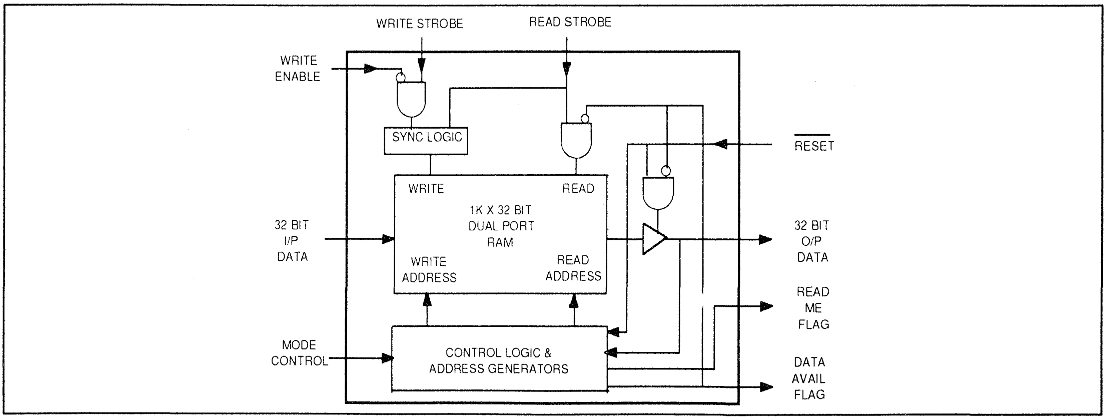

Figure 1. The PDSP16540 Bucket Buffer

THE PDSP16540 BUCKET BUFFER

This device is essentially a 1k x 32 bit synchronous RAM. Being synchronous it requires a continuously available clock, which normally would be the same as the PDSP16510 input and system clocks. It thus has the same 40 MHz maximum rate as the PDSP16510. Note that the data sheet for the PDSP16540 refers to this continuous clock as the read strobe - this does not imply that the strobe should only be present when a read operation is needed.

A write strobe with an enabling signal is needed to write data to the RAM. This write strobe can be asynchronous to the continuous read strobe, and is only needed when data is actually to be loaded. It would normally, however, be the data sampling clock used by the data acquisition circuitry, and thus is expected to be slower than the PDSP16540 read strobe (i.e. PDSP16510 input clock DIS). For example a single PDSP16510 will support continuous 1024 point transforms with sampling rates of some 6.7 MHz, when using 40 MHz clocks. The read strobe for the PDSP16540 is then 40 MHz, and the write strobe is 6.7MHz.

This biased read write ratio makes the use of a true dual port RAM unnecessary. Whenever a write operation is needed the read operation can be interrupted for one cycle, and the write operation actually internally performed with the continuously available read strobe (internal write strobe to read synchronization also takes place).

The device is designed to interface easily to the PDSP16510, and provides comprehensive data overlapping facilities. For correct operation both the block length of the data to be transferred to the PDSP16510, and also the amount of new data in that block must be defined. For commonly used set ups, these two parameters can be defined by tying mode pins high or low. For other alternatives tri-state buffers are needed, connected to up to 16 of the output pins. These are enabled during reset, when these outputs become inputs to an internal latch.

When the programmed amount of new data has been written to the RAM, a Data Available flag (DAV) goes active. This goes in-active for one cycle whenever more data is written to the buffer, and goes permanently in-active when the programmed block length has been transferred. When DAV is active the PDSP16540 will automatically produce new output data on every read strobe edge. The receiving device cannot halt the operation, and it must be dedicated to the transfer task. The DAV signal should be used to provide a clock enable signal for this receiving device.

DAV is the only signal needed to interface to the PDSP16510. For more general applications an additional Read Me Flag is provided. This can be programmed to go active before DAV, and thus warn the receiver that data is about to appear. This signal has no action internal to the PDSP16540.

CONNECTING THE BUCKET BUFFER TO THE FFT PROCESSOR

Figure 2 shows a typical 1024 point system with 50 % block overlapping. Grounding MD0 specifies that 1024 word blocks will be read from the RAM when DAV goes active. Forcing MD2:1 to logical 01 will ensure that DAV goes active when 512 new words have been written to the RAM. Thus the 1024 word block that is transferred to the FFT Processor consists of 512 previously used words and 512 new words. These new words are written to the RAM using the asynchronous Write Strobe, which is also the sampling clock used by the data acquisition circuit. Inputs MD4:3 are really don't care inputs defining when the unused Read Me Flag goes active, but are grounded for electrical reasons rather than leaving them open circuit.

MD5 should be grounded when complex words are to be processed. It should only be tied high if 1024 real transforms are to be performed with no block overlapping (i.e. MD2:0 must be tied low). In this particular case the Bucket Buffer will acquire two blocks of 1024 point real data, through inputs IP15:0, before DAV goes active. These two blocks are then transferred concurrently using all 32 outputs, and the PDSP16510 must be programmed to expect dual real blocks (Control Register Bits 8:6 = 101).

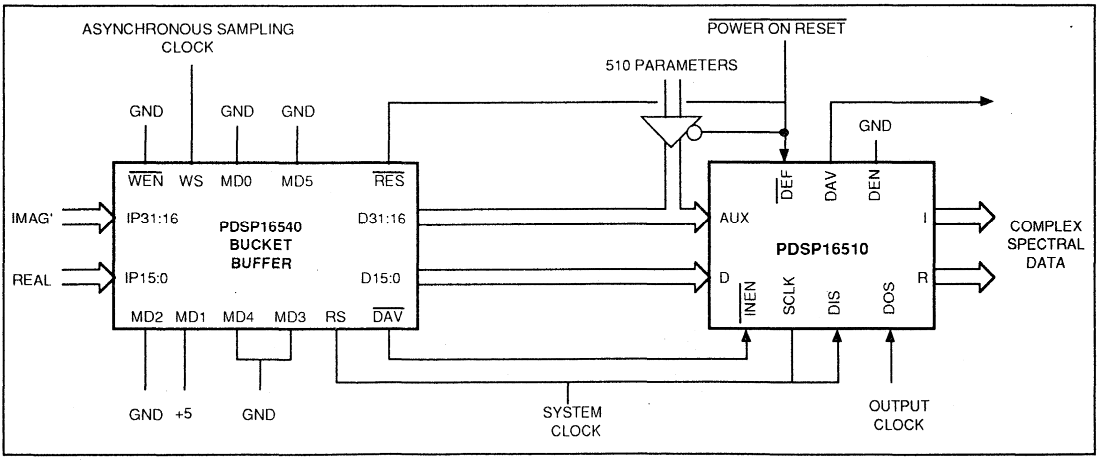

Figure 2. A Typical 1024 Point System with 50% Overlaps

The DAV output is directly connected to the INEN input on the PDSP16510. For correct operation the PDSP16510 must be programmed to use this input as a simple enable i.e. Control Register Bit 12 must be set. The following equation must be obeyed to prevent the loss of any incoming data when doing continuous transforms:

$$NS > 1024B(S/(S - B)) + T + D$$

where N is the amount of new data, S is the input sampling period, B is the read strobe period, T is the transform time which in the case of 1024 point transforms is 3907 system clock cycles, and D is the time to transfer data into the next device.

The factor S/(S - B) arises because the read sequence is interrupted for one B period every time new data is written to the buffer. It thus requires more than 1024 B periods to transfer 1024 words to the PDSP16510. For example if the read rate (B) is 4 times faster than the write rate (S), every 4th read cycle will be inhibited. Thus only 3 out of every 4 read cycles will actually result in data being transfered from the PDSP16540 to the PDSP16510. To achieve the maximum sampling rate possible (i.e. minimum S) data should be transferred in and out of the PDSP16510 at 40 MHz, and the system clock should also be 40 MHz.

Solving for these values gives;

$$NS \text{ must be} > \frac{25600 \times S}{S - 25} + 123375 \quad \text{where S is in nanoseconds.}$$

Rearranging the above equation into the standard quadratic form (i.e. S^2 + pS + q = 0) and solving for the routes gives the value for S.

For no overlapping N = 1024 and S must be greater than 150 nanoseconds. The maximum sampling rate is thus 6.66MHz.

For 50% overlapping N = 512, and the minimum S period is 296 nanoseconds. The maximum sampling rate is then 3.37 MHz

For 75% overlapping N = 256, and the minimum S period is 589 nanoseconds. The maximum sampling rate is then 1.69 MHz

Suppose system requirements, for example, dictate a sampling rate of 4 MHz and some overlapping is required. One solution, of course would be to use more than one FFT Processor as explained in the data sheet. The PDSP16540 bucket buffer would then not be needed. Since, in this particular example, the sampling rate achievable with 50 % overlapping is close to the 4 MHz requirement, it may be possible to compromise on the actual overlap used.

Solving the above equation for S = 250 results in the need to load at least 606 new samples before DAV goes active. By setting MD2:1 to 11 it is possible to define the required number of new words in multiples of 32. A 5 bit code is then inputted through a tri-state driver connected to the D9:5 outputs, which become inputs during RESET. This binary code specifies up to 31 additional blocks of 32 above the minimum of 32.

Rounding 606 up to 608 (32 x 19) results in the need to load the code 10010 through pins D9:5 to achieve the overlap possible. The actual percentage overlap is then 40.6% and the 4 MHz sampling rate will be possible.

In a similar manner the Bucket Buffer can be used to provide non standard overlapping when transform sizes smaller than 1024 points are required. Such a system is illustrated in Figure 3. By forcing pin MD0 high it is possible to define block sizes of less than 1024 words. A 5 bit code on pins D4:0, during RESET, defines up to thirty one additional 32 word blocks after the basic 32 word block. Thus to define a 256 word block it is necessary to input the code 00111 via a tri-state driver on pins D4:0.

The calculation needed to define the minimum sampling period with a given overlap is different when 1024 point transforms are not being performed. As explained earlier, the time to transfer data in and out of the PDSP16510 can then be effectively lost

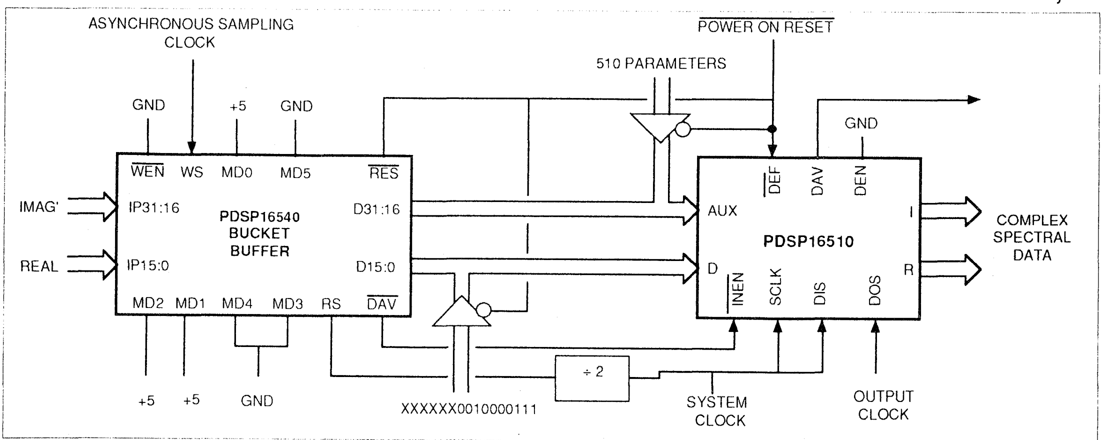

Figure 3. A 256 Point System with Non Standard Overlaps

at the system level. The required equation then simplifies to:

$NS > T$ if transfer times in and out of the PDSP16510 are less than T.
N, S, and T are as previously defined

When concurrent load, transform, and dump operations occur in the PDSP16510, it is not possible for the input clock rate (DIS) to be same as the system clock rate. The actual input rate must be reduced by the factor F from the system clock rate, where F is given by;

$$F = \frac{4}{6 + 0.001L}$$ where L is the system clock low time.

Thus if the system clock rate is to be 40 MHz, the clock low time would be say 12 nanoseconds, and the factor F would be 0.665. The maximum input rate is thus 26.6 Mhz.

In practise this is much greater than that needed to ensure that the time, to transfer 256 words from the Bucket Buffer to the FFT Processor, is less than the transform time. From the PDSP16510 data sheet the time taken to perform a 256 point transform, with a 40 MHz system clock, is 20.4 microseconds. Thus the input clock period needed to load 256 points in that time is 79 nanoseconds: or an input rate of 12.65MHz.

In this example the most convenient approach is to simply divide the system clock by two in order to provide the read strobe for the Bucket Buffer. There is, however, a relationship between the read strobe rate and the maximum write strobe rate. The write strobe period must be at least twice the read strobe period plus 10 nanoseconds. Thus with a 20 MHz read strobe the maximum write strobe rate is 9 MHz.

This writing rate is only achievable with read overlaps up to 50 %. Beyond this a second read rate requirement comes into effect. The write strobe period must also be greater than the read period multiplied by L/N, where L is the read block length and N is the amount of new data. This is another way of saying that the time taken to read the complete block must be no more than the time taken to load the required amount of new data.

In practise with a 20 MHz read strobe these considerations will not limit the writing rate in any way, and the maximum rates will be solely governed by the transform time in the FFT processor. Suppose, for example, we need to support a sampling rate of 7MHz (144 ns period) when doing 256 point transforms with some overlap. Then;

$$N \times 144 \text{ must be} > 20400 \text{ (the transform time)}$$

Thus N must be greater than 142 for 7MHz sampling rates. Since N must be rounded up to a multiple of 32 it is thus necessary to load 160 (32 x 5) new samples in the 256 word block. This requires the code 00100 to be present on pins D15:10 during RESET, and gives 37.5% block overlapping.

Note : In all the above equations any requirement for the input clock (DIS) to be asynchronous to the system clock (SCLK) of the PDSP16510, will have to be modified in a practical system. The PDSP16510 has a requirement that its input and system clocks must be synchronised to each other. It may be possible to burst into the PDSP16510 data at a higher rate than the equations specify, such that DIS is synchronous to SCLK, but that on average the required DIS rate is achieved. When using the PDSP16540 on the input, however, it is very easy to burst data into the PDSP16510.

INCREASE THE PERFORMANCE OF THE PDSP16510 FFT PROCESSOR BY USING MULTIPLE DEVICES

BACKGROUND

A single PDSP16510 FFT Processor ,using a 40 MHz system clock, will support sampling rates up to 12.3 MHz when doing 256 point complex transforms, and up to 6.8 MHz when doing 1024 point complex transforms. These rates can be increased to 40 MHz when several devices are connected in a ring arrangement.

In such a system one device is loaded with a complete block of data, and then starts a transform operation using its internal RAM. In the meantime incoming data is loaded into another FFT Processor which will then start its transform operation when all the data is available. Sufficient devices are needed to ensure that the first device has finished before its turn comes round again for new data.

The PDSP16510 actually supports two multiple device modes of operation. One mode always does separate load, transform, and then dump operations regardless of the actual transform size. The other mode does concurrent load, transform, and dump operations, but cannot be used to perform 1024 point transforms. It can in some circumstances allow less devices to be used in order to achieve a given sampling rate.

Note that an input buffer is not needed in either mode of operation (even when doing 1024 point transforms), and interdevice flags support block overlapping. With the standard 50% and 75% block overlapping no external logic is needed.

Note also that this arrangement is only intended to increase the sampling rates possible with the transform sizes supported by a single device. If larger transform sizes are needed see Application Note AB35.

GENERAL CONSIDERATIONS

Figure 1 illustrates the basic ring arrangement, using three devices for convenience. It can, of course, be expanded to any required number of devices. It shows that both inputs and outputs are commoned together. A block of data is loaded into the first device, then the next block is loaded into the second device, and so on. Sufficient devices are provide to ensure that continuous data can be supported without any loss.

The LFLG output and INEN input are used to co-ordinate the splitting of incoming data between the devices. This requires

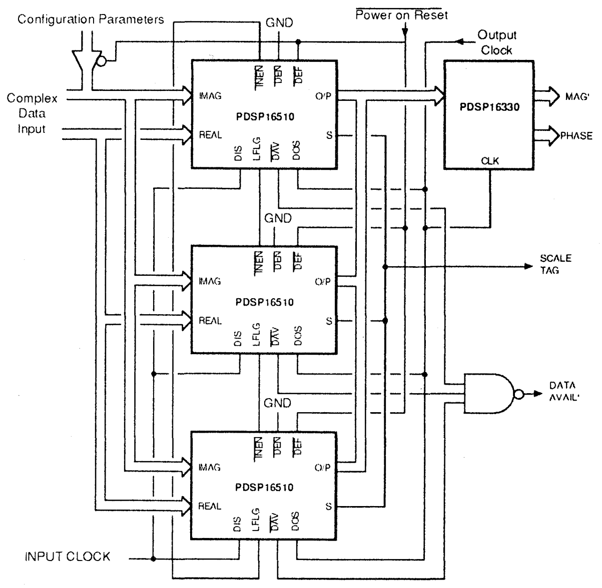

Figure 1. The Basic Multiple Device Arrangement

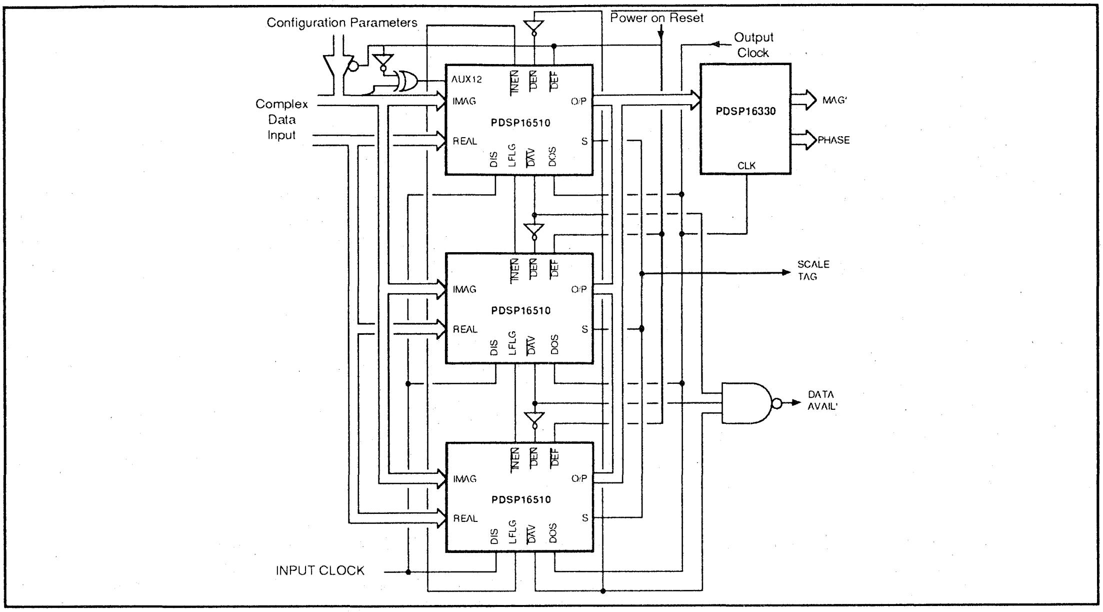

Figure 2. Adaptation of the Basic System

that the input data enabling signal (INEN) is no longer a simple DIS strobe enabling input. Whenever the Control Register specifies a multiple device mode of operation, it becomes a global ,edge activated, block enabling input. When it goes low a new block load operation commences, which will then continue until the programmed number of samples have been loaded. INEN can go back high at any time without causing inputs from the current block to be inhibited (early silicon requires that INEN stays low at least until all the block has been loaded).

The Load Flag (LFLG) goes active high after the first sample has been loaded, and goes low when the programmed number of samples have been loaded. This low going LFLG edge from a previous device provides the low going INEN edge to the next device. Thus once one device has received its full complement of data the next one will start loading data. As a further enhancement the LFLG will go in-active after half the block has been loaded if 50% overlapping is in use, or after a quarter the block has been loaded if 75% overlapping is in use. Thus, if for example, 50% overlapping is in use, the second device will also start loading data half way through the load operation into the first device. Two devices thus receive the same data in order to mechanize the overlap requirement.

By tying all the outputs together the system need only provide one output processor. The actual processing needed on the outputs is application dependent, and is beyond the scope of this article. It is assumed, however, that unless speed requirements make it impossible, the use of only one output processor will provide some economic advantage. It is also assumed that the reason for using multiple devices was to continuously process incoming data, without any loss of information. Thus once results are available they must be passed on to the rest of the system.

This can be achieved by tying all the Data Enable (DEN) inputs low, as illustrated in Figure 1. The Data Available Flag (DAV) will then be synchronized to the output strobe, and will go active when the first result is available on the output pins. The output bus goes low impedance with respect to the same output clock edge which causes DAV to go active. Figure 1 in fact shows all the DAV flags OR'ed together to give a common indicator for the rest of the system. In practice this only makes sense if the dump time is less than the load time. The combined output would then go in-active between individual devices for a period equivalent to the difference between the load and dump times.

For convenience we have so far indicated that all the DEN inputs are tied low, but there is some uncertainty in the time taken to complete a transform once all the data has been loaded. In fact the device uses an internal 12 cycle sequence, which will lead to a 12 system clock variation in the time needed to complete a transform. The number of system clock cycles needed to complete each transform, as given in Table 4 of the data sheet, and are worst case numbers. This uncertainty makes tying DEN low potentially dangerous.

If device two effectively completes its transform in less time than device one, then it could start dumping its results before device one had finished. If DEN is tied low this can only be prevented if the dump time (plus 4 DOS periods) is at least 12 system clock periods less than the load time. DIS and DOS can then not be tied together, and it is usually convenient to derive DOS by dividing down from the system clock.

An alternative approach has, however, been provided which still allows DIS and DOS still to be commoned together.

Whenever the control register specifies a multiple device node of operation, the operation of the DAV / DEN circuitry is modified. IF DEN is not active (i.e. high) when the device is internally about to generate DAV, then the output pin will not "asynchronously" go active as happens in the single device mode. Instead the DAV output pin will stay in-active until DEN goes active low. The output bus will also stay high impedance.

By connecting the inverted DAV output from one device to the DEN input of the next, this second device cannot commence its output sequence until the first one has finished, and its DAV has gone in-active. The DAV output from the second device will go active as soon as its DEN input goes low, but it will still be effectively synchronized to DOS since it is derived from the previous DAV signal going in-active. This arrangement is shown in Figure 2.

When DOS is physically connected to DIS, the combined Data Available Flag shown in Figure 1 will only glitch in-active as one device finishes and the next one starts. This glitch would safely occur after a DOS active edge, but even so the flag would only usefully indicate the initial delay from start up before valid results are obtained. When a reliable flag is needed to indicate the end of a set of results, then each PDSP16510 should be provided with a D type latch. This is set by the inactive going edge of DAV and reset when DOS again goes low.

The circuit shown in Figure 1 needs an edge after power on in order to initiate the load procedure into the first device. Rather than providing external logic to generate a start pulse which is OR'ed into the first INEN line, an alternative scheme is supported. By setting bit 12 in the Control Register contained within the first device, it is possible to cause the power on reset signal (DEF) to initiate a load procedure.

All Control Registers are loaded from the common AUX port whilst DEF is active, and we now require one bit in one of the registers to be different from the others! The easiest way to mechanize this is to include an EX-OR gate in the AUX12 input to the first device. The other input is driven from the DEF signal such that it causes a logical inversion as the Control Register is loaded, but not when imaginary data is loaded. This is also shown in Figure 2.

SAMPLING RATES POSSIBLE WITH SEPARATE LOAD TRANSFORM, AND THEN DUMP (MODE 2)

This mode can be used with all transform sizes and the maximum DIS and DOS rates can theoretically be equal to the system clock rates used. The DIS rate is the data sampling rate and can, of course, only be equal to the system clock rate if sufficient devices are provided in the ring to make this sensible at the system level. The DOS rate can be any rate greater than or equal to the system clock rate, and would normally be limited by the capabilities of the output processor.

The number of devices, N, needed to achieve a sampling period of S with a block size of n, is governed by;

$$NnS > nS + PK + D \qquad \text{where D is less than } nS$$

and K is the system clock period, P is the number of system clock periods needed to complete a transform as given in Table 4 of the datasheet, and D is the total dump time allowing for the 4 extra DOS periods needed for the internal output circuitry.

With 50% block overlapping the above equation becomes;

$$NnS > 2(nS + PK + D) \qquad \text{where D is less than } nS/2$$

With 75% overlapping it becomes;

$$Nns > 4(nS + PK + D) \qquad \text{where D is less than } nS/4$$

Table 1 gives the maximum sampling rates possible with 3, 4, 5, or 6 devices and output rates of 20 MHz and 40 MHz. It covers transform sizes of 256 and 1024 complex points.

	1024 POINT COMPLEX TRANSFORMS						256 POINT COMPLEX TRANSFORMS					
Number of	Dump at 20 MHz			Dump at 40 MHz			Dump at 20 MHz			Dump at 40 MHz		
Devices	0%	50%	75	0%	50%	75%	0%	50%	75%	0%	50%	75%
3	13.7	6.8	-	16.6	8.2	-	15.3	7.6	-	19	9.5	-
4	20.6	10*	-	24.8	16.6	-	22.9	10*	-	28.5	19	-
5	27.4	10*	5*	33.2	20*	8.2	30.6	10*	5*	38	20*	9.5
6	34.3	10*	5*	40*	20*	10*	38.4	10*	5*	40*	20*	10*

* indicates that the sampling rate is limited by the maximum dumping rate

Table 1. Maximum Sampling Rates with separate load transforms and dumps.
Where sampling rate is asynchronous to SCLK, a PDSP16540 (or similar) is assumed on the input.

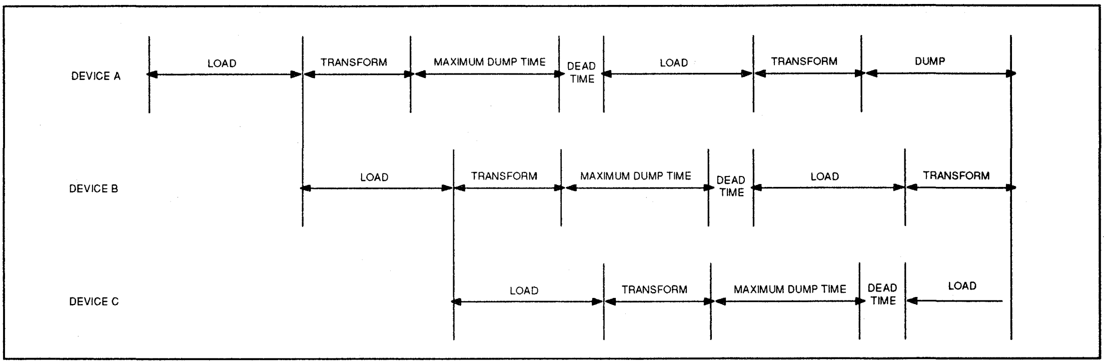

Figure 3. Sequence of events in a 3 device system doing separte load, transforms and then dumps

Figure 3 shows the sequence of events taking place in such a system without overlaps, and Figure 4 shows the sequence in a five device system doing 50% overlaps. It should be noted that the time taken to dump the results must be no more than the time taken to load the block of data (if you put more in than you take out, then something has to give). If 50% block overlapping is used the dump must be complete in half the load time, and one quarter the load time if 75% overlapping is needed. If these criterion are not met, then one device will not have finished dumping before the next one starts. Thus two sets of outputs would be enabled at the same time. Remember that the dump time can easily be made faster than the load time when required.

SAMPLING RATES POSSIBLE WITH CONCURRENT LOAD, TRANSFORM, AND DUMP (MODE 1)

In this mode transfers in and out of the PDSP16510 are concurrent with transform operations. Internal RAM restrictions do not allow this mode to be used with 1024 point transforms. For other sized transforms the sampling rates possible are theoretically much higher for a given number of devices. A limitation is, however, imposed on the maximum I/O rates, which can no longer be increased up to the system clock rate. Instead the rates are reduced to a factor, F of the system clock rate where F is given by;

$$F = \frac{4}{6 + 0.001 \varnothing L} \qquad \text{where } \varnothing L \text{ is the system clock low time}$$

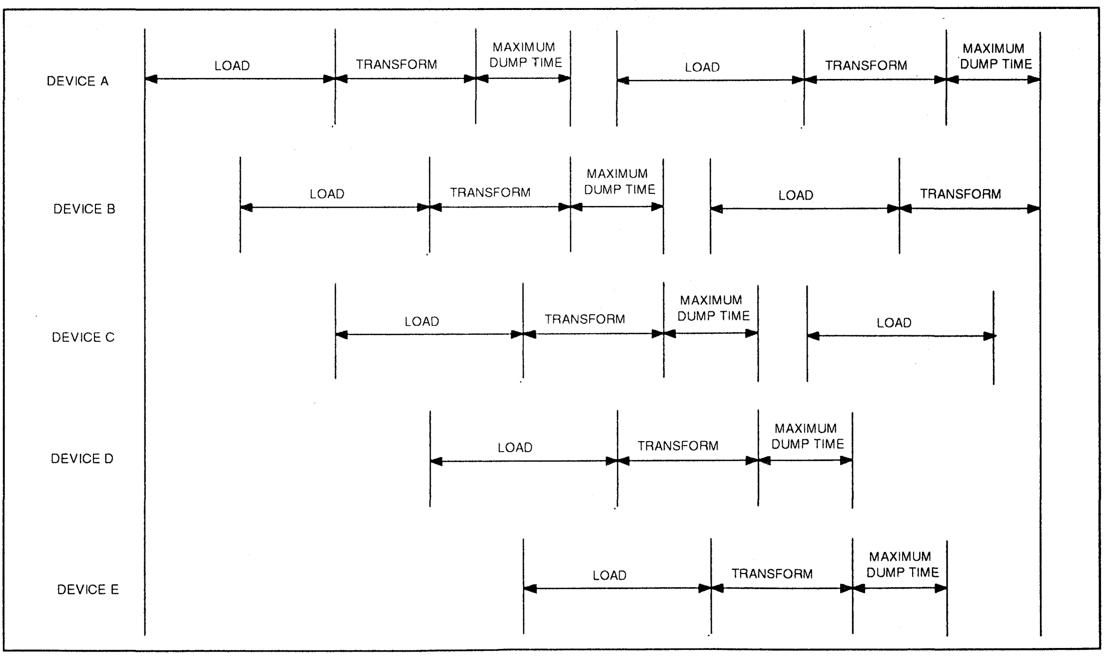

Figure 4. Sequence of events in a 5 device system with 50% overlaps

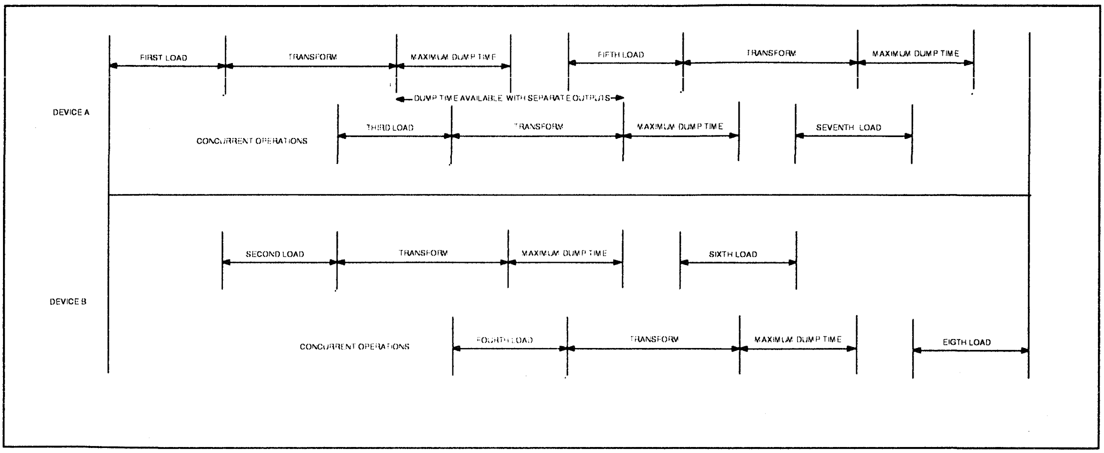

Figure 5. Sequence of events in a 2 device system with concurrent I/O Operations

With a 40 MHz system clock the low time might well be 15 nanoseconds. The factor, F, then becomes 0.66. The maximum theoretical I/O rate is thus 26.6 MHz. In this mode there are no benefits to be gained by increasing the output rate above the input rate (except when necessary because of overlaps). DIS and DOS are thus normally commoned together. In fact the device provides an internal divide DIS by or 2 or 4 feature. Thus the DIS and DOS pins can be externally connected to a source to match the DOS requirements, with the internal DIS strobe internally divided down to the correct frequency for 50% or 75% overlaps.

With N devices the theoretical sampling period, S, ignoring the F factor and with no overlapping is governed by;

$NnS > PK + 4$ DOS periods where P is the number of clock periods needed to complete the transform and K is the system clock period.

Figure 5 illustrates the sequence of events which occur with a two device system, with the outputs joined together. With a 40 MHz system clock, and common DIS and DOS, the theoretical maximum input sampling rate given by the above equation is 24.6 MHz when doing a 256 point complex transform. This sampling rate is just less than the limit which would be imposed by the factor F, but all other transform sizes supported by this mode of operation would be limited by the factor F, to sampling rates of 26.66 MHz. If each PDSP16510 has its own output processing circuit, then the outputs would not be joined together and each would have a dump time available equivalent to twice the load time.

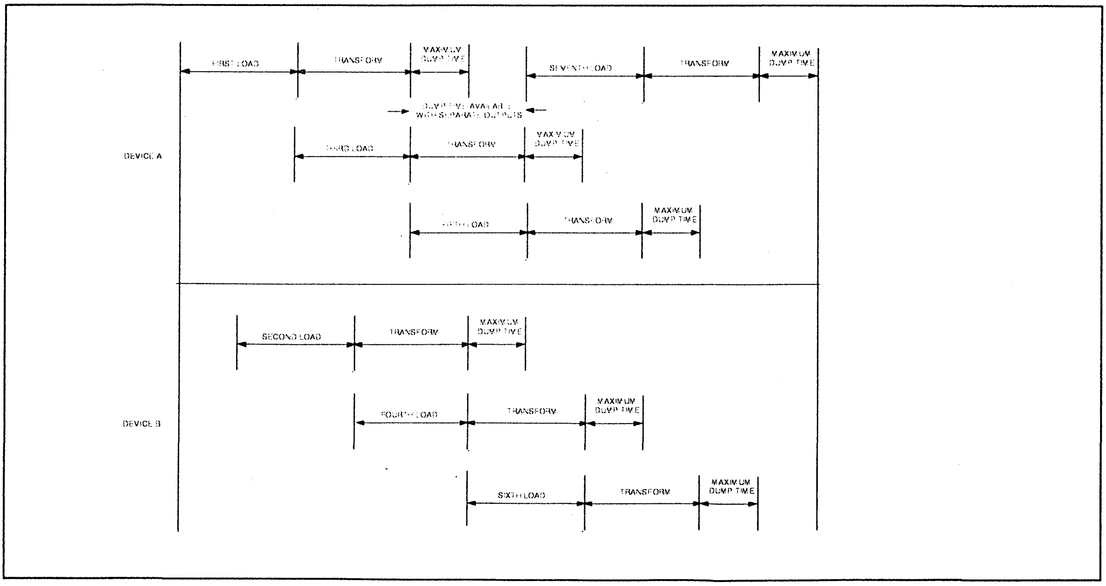

Figure 6. Sequence of events in a 2 device system with concurrent I/O and 50% overlaps

In this mode two devices can be used to gave the same performance when doing 50% overlaps as a single device with no overlaps. But the dump must be completed in half the load time if the outputs are commoned together. This does not limit the performance when doing 256 point complex transforms, and two devices will handle 50% overlaps at up to 12.3 MHz sampling rates. All other sizes are limited to up to 13.3 MHz sampling rates by the F factor, unless each device has its own output processing circuitry. Figure 6 illustrates the sequence of events in such a system.

As long as each PDSP16510 has its own output processor, then extra devices can be added as needed to increase the overlapped sampling rate up to that defined by the F factor i.e. 26.66 MHz. The minimum sampling period down to this limit is governed by;

$$NnS > 2(PK + 4DOS)$$

With 75% overlapping at least four devices are needed, and the minimum sampling period under the same conditions as above is given by;

$$NnS > 4(PK + 4DOS)$$

When all the outputs are commoned the dump must be completed in one quarter the load time. Since the maximum output rate is 26.66 MHz, then the maximum input rate is 6.66 MHz. For all transform sizes this is much less than suggested by the above equation. Thus separate output circuits must always be provided if the maximum performance is to be achieved.

Note that if the input or output rate chosen is asynchronous to SCLK, then a PDSP16540 (or similar) is assumed at the PDSP16510 interface.

MULTIPLE CONCURRENT TRANSFORMS

The PDSP16510 will support 4 concurrent 64 point complex transforms, or 8 concurrent 64 point real transforms. When the performance is increased by using multiple devices configured in MODE 1, a double LFLG transition is provided to support block overlapping. This is illustrated in Figure 7.

The LFLG output goes high after the first sample is loaded, and then low half way through the first sub block if 50% overlaps are in use. This transition is used to instigate the load procedure into the second device. LFLG then returns high and goes low again half way through the last sub block. This second low going transition is used to instigate a new low procedure in the first device.

Note that once a load procedure has started, the occurrence of a second edge will not effect the device in any way and thus each device only responds to an edge as shown in Figure 7. Note also that in this arrangement the DAV / DEN connection MUST be made, even if DIS and DOS are not commoned. This ensures that each device has an equal amount of time to dump its data.

This double LFLG transition also supports 75% overlapping with four devices. The first low going edge occurs 25% through the first sub block, and the second low going edge 25% through the last sub block.

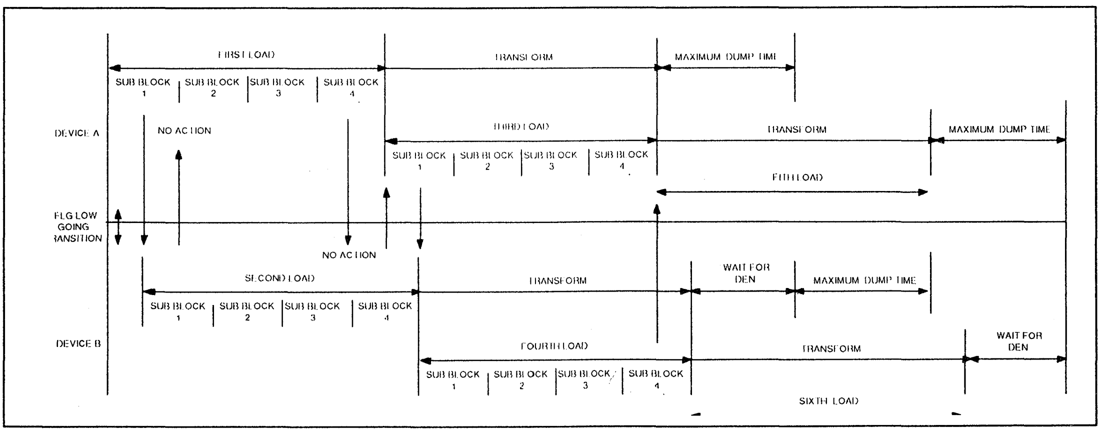

Figure 7. Sequence of events when performing multiple concurrent transforms.

Support Tools

GEC PLESSEY
SEMICONDUCTORS

GEC PLESSEY
SEMICONDUCTORS

DS3854 - 1.0

PDSP16256/PDSP16350 EVALUATION SYSTEM

AN IBM PC COMPATIBLE DEVELOPMENT SYSTEM FOR HIGH PERFORMANCE DIGITAL FILTERING AND SIGNAL GENERATION USING THE PDSP16256 AND PDSP16350

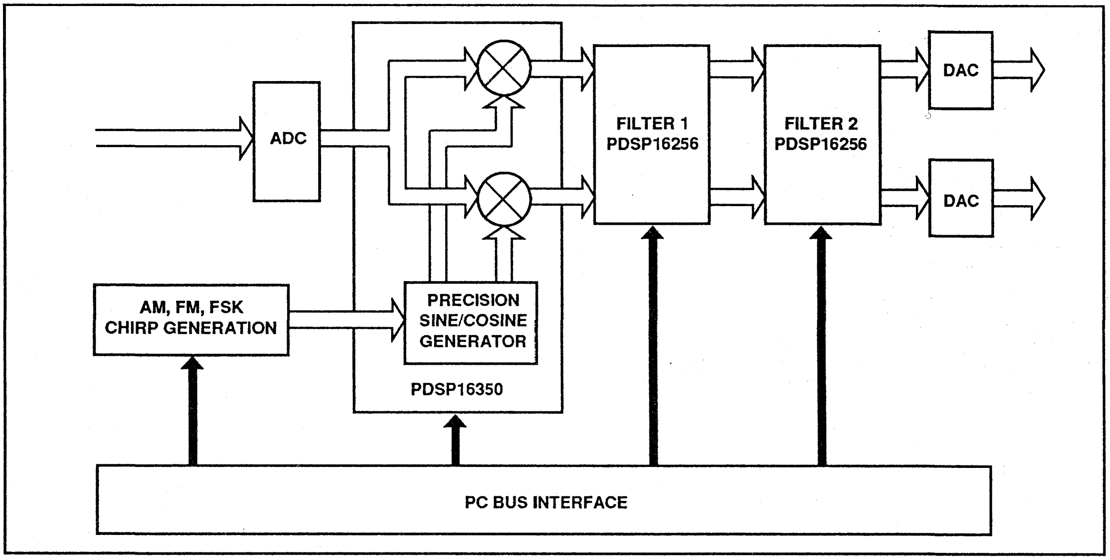

The evaulation system provides a complete environment for the development of high performance digital filtering and signal generation systems. The development system is based on the powerful GEC Plessey Semiconductors' PDSP16350 and PDSP16256 digital signal processing components and comprises the following elements:

- An IBM PC compatible board for high performance digital filtering and signal generation
- A powerful digital filter design software package optimised to the specific characteristics of the PDSP16256 programmable FIR filter
- A flexible digital filter development software package enabling the board to be configured and operated in a wide range of modes
- Comprehensive supporting documentation

These facilities constitute an easy-to-use development tool for digital filtering systems requiring sample rates of up to 20MHz and provide an ideal environment for rapidly evaluating the capabilities of these high performance GEC Plessey Semiconductors devices.

APPLICATIONS

The evaluation system is applicable to a wide range of areas, including:

- Radar
- Ultasonics
- Video processing
- Digital radio
- Sonar
- Data communications
- Instrumentation
- Satellite communications

DIGITAL FILTERING BOARD

Key features of the IBM PC compatible board are as follows:

General

- 8- or 12- bit ADC versions
- Dual 12-bit DACs
- 16-bit data
- Real or quadrature modes of operation
- Sample rates between 1kHz and 200MHz
- Digital output ports

Digital Filtering

- Configurable as separate I and Q channels or as a single real only channel
- Cascaded or single filter options
- Filtering lengths between 8 and 128 coefficents (dependent on sample rate and configuration)
- Decimate by 2 mode

Signal Generation

- High resolution sine/cosine generation
- Supports amplitude modulation (AM) and frequency modulation (FM)
- High precision quadrature chirp signal generation

DEVELOPMENT SOFTWARE

- Filtering designs optimised to the characteristics of the PDSP16256
- Cascaded, dual or single filter modes
- Frequency selective and Hilbert transform filters
- Bit-accurate frequency responses
- EPROM file generation
- Flexible board configuration options
- Filter coefficient loading
- Control of precision waveform generation
- Easy-to-use menu driven operation

ORDERING INFORMATION

The following options are available:

Part number	ADC sample rate	ADC resolution	No. of PDSP16256 ICs
PDSP DFDS-1	20MHz	8 bits	1
PDSP DFDS-2	20MHz	8 bits	2
PDSP DFDS-3	1MHz	12 bits	1
PDSP DFDS-4	1MHz	12 bits	2

All options include the following IBM PC compatible hardware and software:

- Digital filtering board
- Digital filter design package
- Digital filter development package
- Support documentation

PDSPFTDS

PDSP16510 FFT PROCESSOR EVALUATION BOARD

FEATURES

- Complete evaluation and prototyping system for performing high speed FFTs using the PDSP16510 FFT Processor IC from GEC Plessey Semiconductors
- IBM-PC/AT compatible plug-in board with menu driven software support
- Supports 64, 256 and 1024 point real and complex FFTs
- On board 10MHz, 8 bit ADC for real time signal acquisition and spectrum analysis
- Real time FFT magnitude & phase extraction with the PDSP16330 Pythagoras Processor IC
- Logarithmic LUT and 8 bit 10MHz DAC for real time dB scaled power spectrum output
- Complex (16 + 16 bits) or real (16 bits) digital input port
- High speed complex (16 + 16 bits) digital data outputs from PDSP16510
- High speed magnitude (16 bits) and phase (12 bits) digital outputs from the PDSP16330

EVALUATION BOARD OVERVIEW

The IBM-PC compatible board provides a convenient environment for rapidly evaluating the PDSP16510 FFT Processor IC. 64, 256 or 1024 point FFTs can be performed on either real only or complex input data. The FFTs can be calculated in real time on input data received either from the analog input port ,using the on board 8 bit, 10MHz ADC, or alternatively from a high speed (20MHz) complex (16 + 16 bit) digital data input port.

Magnitude and phase data is derived directly from the complex results generated by the PDSP16510 FFT Processor using the PDSP16330 Pythagoras Processor IC. An 8 bit, 10MHz DAC and logarithmic look up table are provided to allow dB scaled power spectrum of the transform results to be viewed in real time using a conventional oscilloscope trace. Alternatively, the high speed digital data output port may be used, giving access to the complex digital (16 + 16 bits) data output from the PDSP16510 or the magnitude (16 bit) and phase (12 bit) values output by the PDSP16330.

The hardware interface to the IBM-PC AT bus includes both an input buffer and a high speed output buffer. The input buffer can be used to allow the PC to display the captured input data or to provide data input from a PC file to the PDSP16510. A 32K word high speed buffer may be used to capture magnitude output data from the PDSP16330. Multiple FFT results can be captured in real time and displayed on the PC using the control software.

A comprehensive graphics-based software control program together with detailed technical documentation is provided with the hardware. The software is easy to use, with keyboard/mouse driven control, drop down menus and dialogue boxes. A library of low level C-code routines are also provided to allow the user to develop proprietary software control if required.

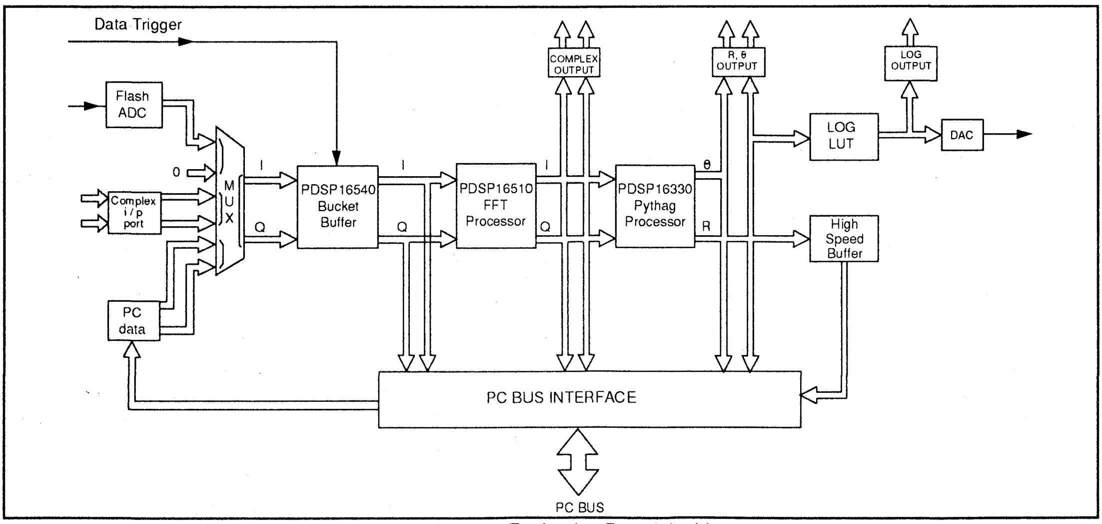

Fig.1 PDSP16510 Evaluation Board Architecture

HARDWARE FUNCTIONS

The architecture of the PDSP16510 Evaluation Board is shown in figure 1.

ADC

The board contains an 8 bit flash ADC (GEC Plessey Semiconductors part number SP973T8) capable of operating on board at rates up to 10MHz. The ADC is driven by an input buffer set to provide full scale inputs to the ADC from an analog input that is 2v peak to peak.

The ADC sample rate is software controllable from the PC and may be set at various divisions between 10MHz and 650KHz.

Digital Input Port

Complex or Real digital data can be input to the board via a ribbon cable connector. This input is fully compatible with the digital output connector from on the GEC Plessey Semiconductor PDSP16256/350 Digital Filter Evaluation Board.

Input Buffer

The PDSP16540 32K Bucket Buffer is used to buffer either the ADC or digital data input. The buffer and the PDSP16510 can be configured to operate with all the available modes of the PDSP16510 and with 0, 50% and 75% overlap when receiving data from the digital input.

Input Capture Circuitry

The PC has access to the input data to the PDSP16510 FFT Processor to allow the PC to display the captured input data set or alternatively to load data into the buffer from an external file.

FFT Processor

FFTs are calculated by the PDSP16510. Its input connects directly to the PDSP16540 Bucket Buffer, whilst the complex data output drives into the PDSP16330 Pythagoras Processor. The PDSP16330 calculates the corresponding magnitude and phase angle values from the complex PDSP16510 data output, to allow the power spectrum of the the input signal to be derived directly. Note that the PDSP16540 is only strictly required to support continuous 1024 point transforms with the PDSP16510. However, due to the board architecture it is also used in other modes.

High Speed Output Buffer

A 32K word deep high speed buffer is provided at the magnitude output of the PDSP16330. This may be used to capture the results of multiple FFTs in real time for subsequent analysis on the PC.

Trigger Input

A TTL level input trigger can be used to enable the Input Capture Circuity allowing the board to synchronise to external systems.

Digital Output Ports

High speed digital outputs are provided for either complex (16 bit real + 16 bit imaginary) data output from the PDSP16510 or the magnitude (16 bit) and phase angle (12 bit) data output from the PDSP16330.

Scope Output

An 8 bit DAC and log look up table are provided to give an output of log magnitude to a conventional oscilloscope input. The output is scaled such that 80dB is displayed with 10dB per division.

Test Source

An on-board test source is provided that will pass a modulated carrier through the FFT board. The test source sweeps approximately between 1.5MHz and 2.5MHz and can be enabled in any board mode in which the ADC input is selected.

CONTROL SOFTWARE

A comprehensive graphics-based software control program is provided with the hardware. The software is controlled via the keyboard/mouse, with drop down menus and dialogue boxes giving the user high level control over board operation.

The low level software modules used to drive the board are provided to the user as a set of object and C source libraries, allowing the user to easily create their own software control for the board or interface it to other packages.

The software control program allows the evaluation board to be operated in four modes: Free Run, Continuous Display, Single Shot and Transient Capture.

FREE RUN MODE

Free run mode allows the PDSP16510 to run at full speed, processing data from either the ADC or complex digital input ports. When the ADC is used, the digitised data is input to the real input of the PDSP16540 Bucket Buffer (the imaginary input is automatically zeroed). Results are available from both the digital output ports and also the oscilloscope DAC output. The oscilloscope display is scaled for an 80dB dynamic range with 10dB per scope division.

Free run mode configuration is shown in figure 2. The maximum data input rate depends on the source of the data (ADC or complex digital input). Tables 1 to 3 detail the possible data rates.

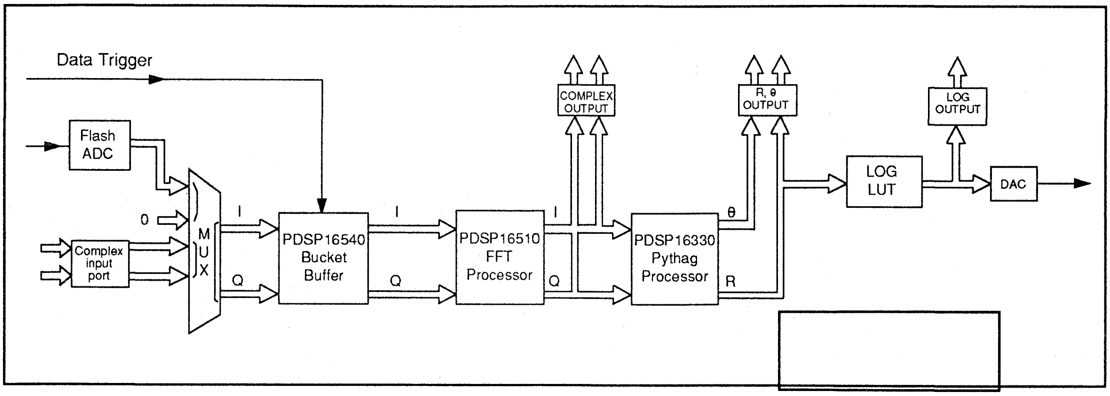

Fig.2 Free Run Mode

Table 1. Free Run Mode Complex Digital Input Data Rates - Complex FFTs

(Complex input data or real only input data with imaginary data zeroed by user connection)

FFT Size	Overlap	Max i/p rate (MHz)
1024 pt complex	0%	5.86
	50%	2.93
	75%	1.46
256 pt complex	0%	12.54
	50%	6.27
	75%	3.13
4x64 pt complex	0%	16.4
	50%	8.2
	75%	4.1

Table 2. Free Run Mode Complex Digital Input Data Rates - Real FFTs

(2 simultaneous real only FFT channels, one channel on real input plus one channel on imaginary input)

FFT Size	Overlap	Max i/p rate (MHz)
2x1024 pt real	0%	5.08
2x256 pt real	0%	9.9
	50%	4.95
	75%	2.47
8x64 pt real	0%	12.5
	50%	6.25
	75%	3.12

Table 3. Free Run Mode ADC input sample rates

(Data input via real port, imaginary data input automatically zeroed)

FFT Size	Overlap	Max i/p rate (MHz)
1024 pt complex	0%	5.86
	50%	2.93
	75%	1.46
256 pt complex	0%	10.0
	50%	5.0
	75%	2.5
4x64 pt complex	0%	10.0
	50%	5.0
	75%	2.5
2x1024 pt real	0%	10.0*

* NB. uses special mode of PDSP16540 to support 2x1024 point real FFTs

Continuous Display Mode

In continuous display mode the PC is used to control the FFT system to process one set of input data at a time (overlapped FFTs are not supported in this mode). The PC will take one FFT result from the board and display the spectrum on the PC screen. Whilst displaying this result, the PC accesses the board to read the next set of data, which is subsequently displayed. A continuous, "real time" display is thus generated.

Input data is captured in blocks by the PDSP16540 Bucket Buffer. When the buffer is full, the PC has access to the buffer output whilst strobing the data into the PDSP16510. The PDSP16510 calculates the FFT and when the output is valid, the PC strobes the PDSP16510 to output the results.

The output magnitude values are calculated by the PC and displayed on screen.

Input data can be taken taken from either the ADC, complex digital input port, or alternatively from a PC-based ASCII file.

The display can show a number of different windows and all of these may be enabled if required. The display options are :-

- Input data (Real or Complex, dependent on selected source)
- Output magnitude
- Infinite precision FFT
- Transient Response

The Infinite Precision FFT window is generated by processing the input data with a full floating point FFT on the PC. The performance of the 16 bit block floating point implementation of the PDSP16510 can then be compared directly.

The transient option displays successive magnitude data on the screen in a waterfall format.

Continuous display mode operation can be frozen at any time to allow detailed examinition of results. Keyboard driven cursors are provided to make measurements from the window display.

Output data may be optionally stored in a PC file. Board configuration for Continuous Display Mode is shown in figure 3.

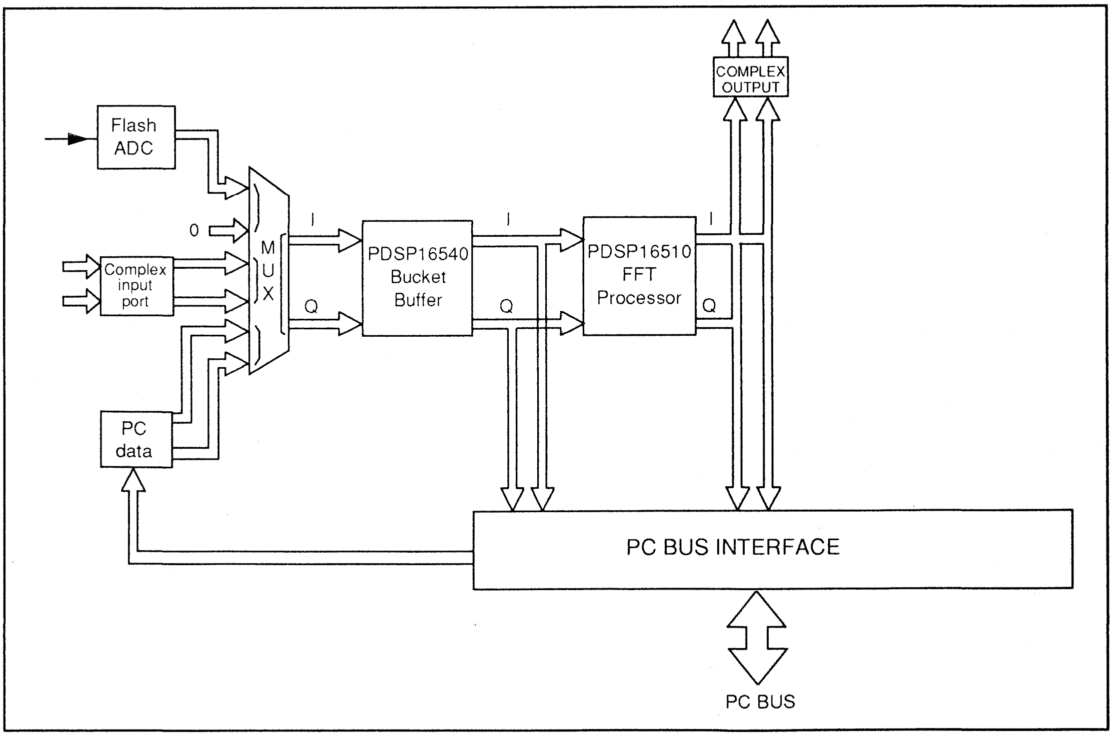

Fig.3 Continuous Display Mode

SINGLE SHOT MODE

Offers the same features as the continuous mode, but only takes one set of data.

TRANSIENT CAPTURE MODE

Transient Capture Mode (shown in figure 4) is similar in architecture to Free Run Mode, but operates as a single event FFT. The high speed output buffer is used to capture a block of spectra from the board in real time. The limit on the number of spectra is set by the RAM size of 32K words. i.e. it is possible to capture 16 x 1024 point spectra or 64 x 256 point spectra. The transient capture facility only stores magnitude data from the PDSP16330 output.

The results of transient capture are displayed on the PC as a 3D Spectrograph, and may also by stored to a data file.

TEST PATTERN UTILITY

A test pattern can be generated that creates an input waveform to the board comprising 3 sine waves with user defined relative amplitudes, phase and frequency separation.

ORDERING INFORMATION

PDSPFTDS/NA/NA - PDSP16510 Evaluation Board, software & manual

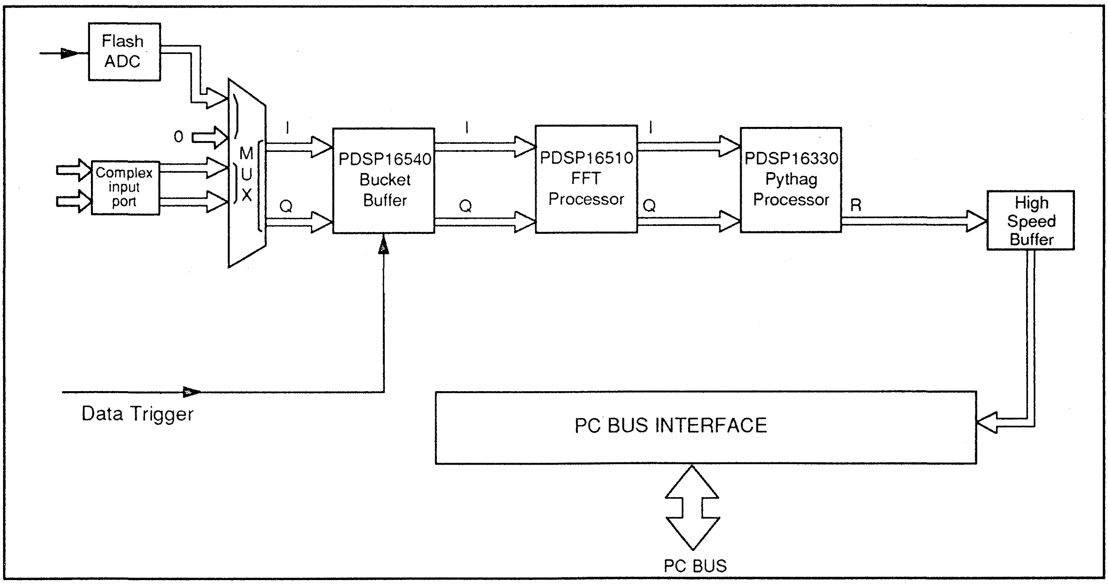

Fig.4 Transient Capture Mode

Package Outlines

Dimensions are shown thus: mm (in).
For further package information, please contact your local Customer Service Centre.

GEC PLESSEY
SEMICONDUCTORS

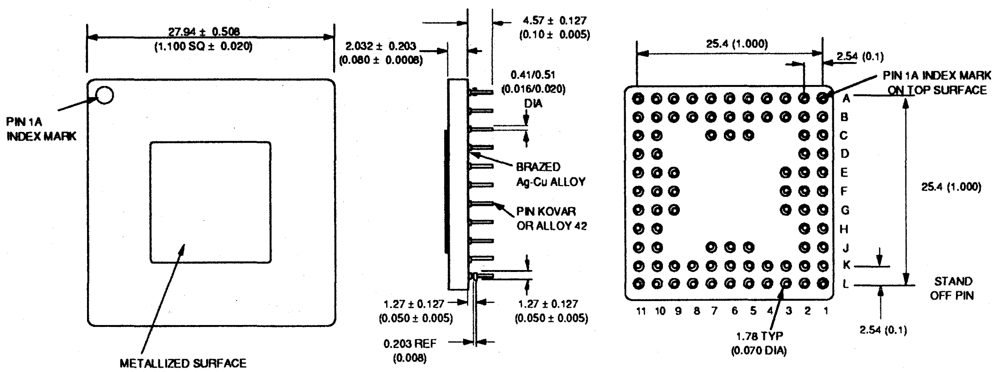

84-PIN GRID ARRAY PACKAGE - AC84

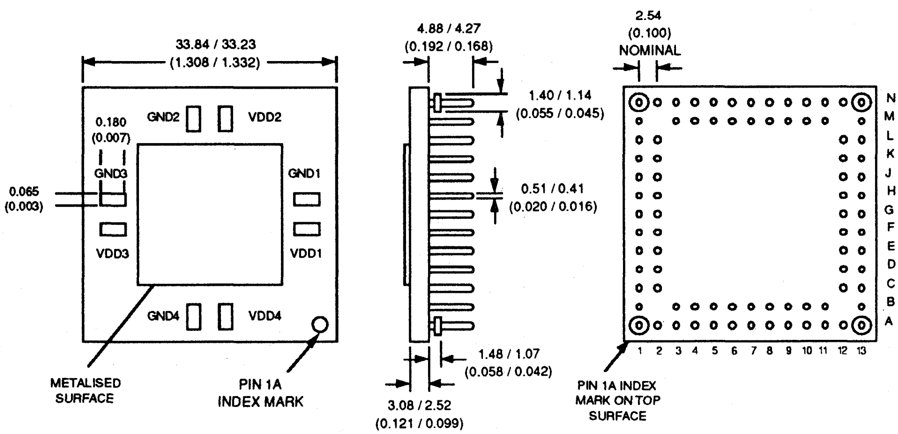

84-PIN GRID ARRAY PACKAGE - AC84

(Used for the PDSP16340, PDSP16350, PDSP16488 and PDSP16510)

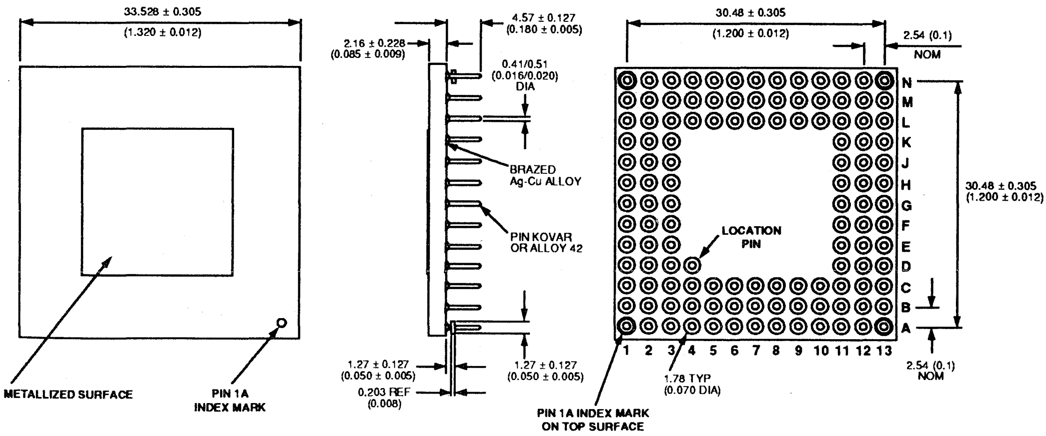

120-PIN GRID ARRAY PACKAGE - AC120

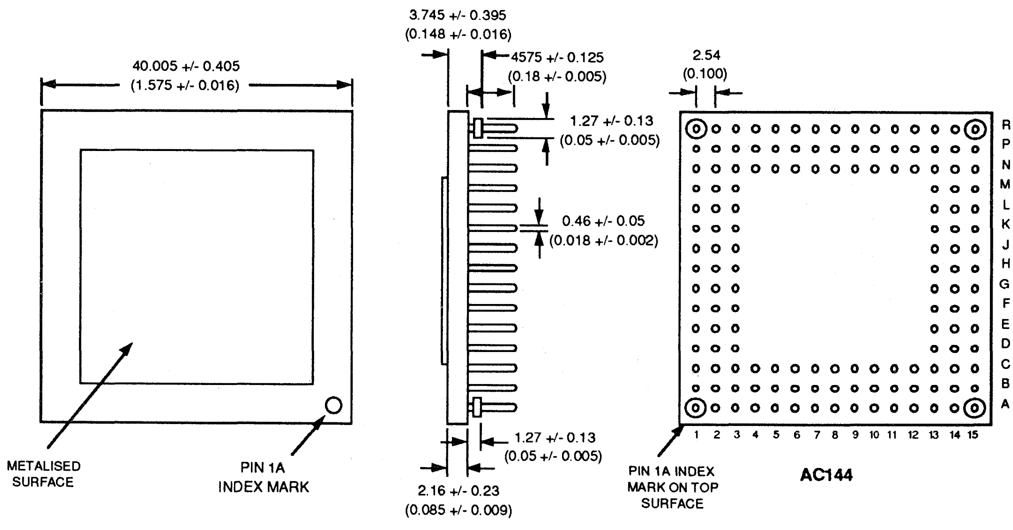

144-PIN GRID ARRAY PACKAGE - AC144

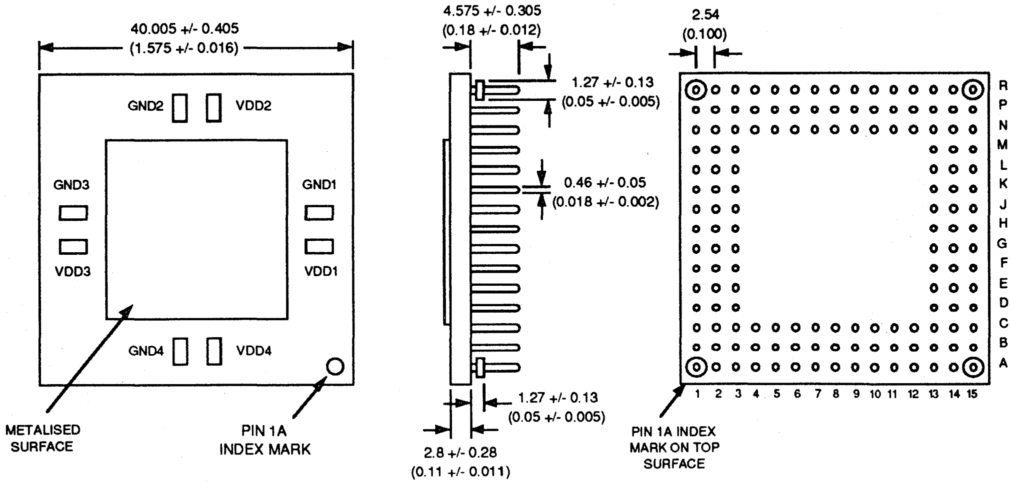

144-PIN GRID ARRAY PACKAGE - AC144

(Used for the PDSP16256)

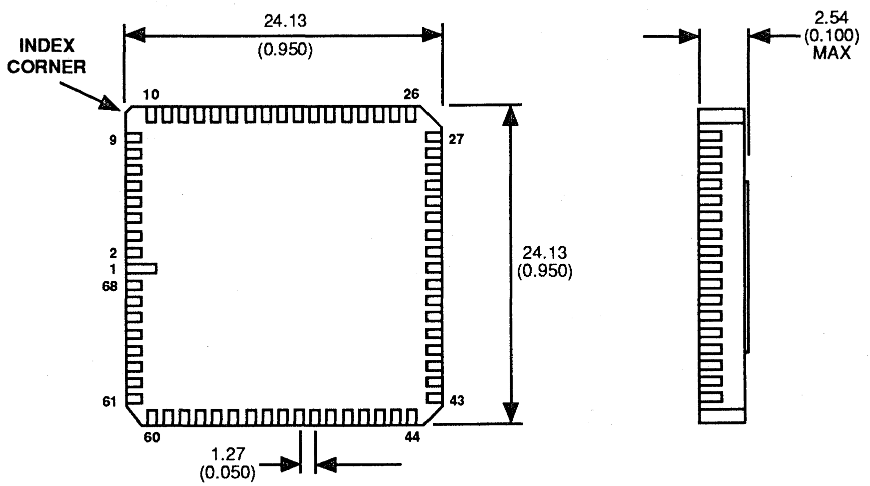

68 CONTACT LCC PACKAGE - LC68

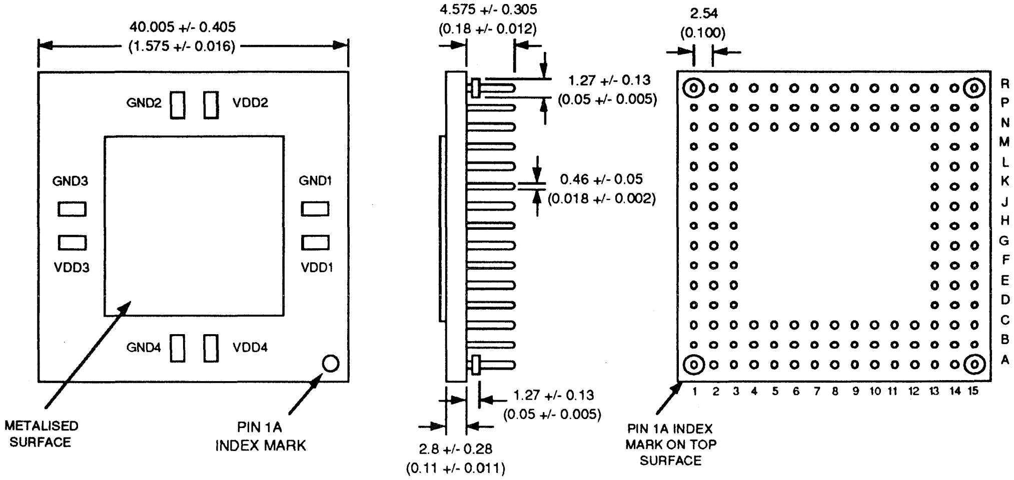

144-PIN GRID ARRAY PACKAGE - AC144

(Used for the PDSP16256)

28 LEAD P.L C.C. (HP) HP 28

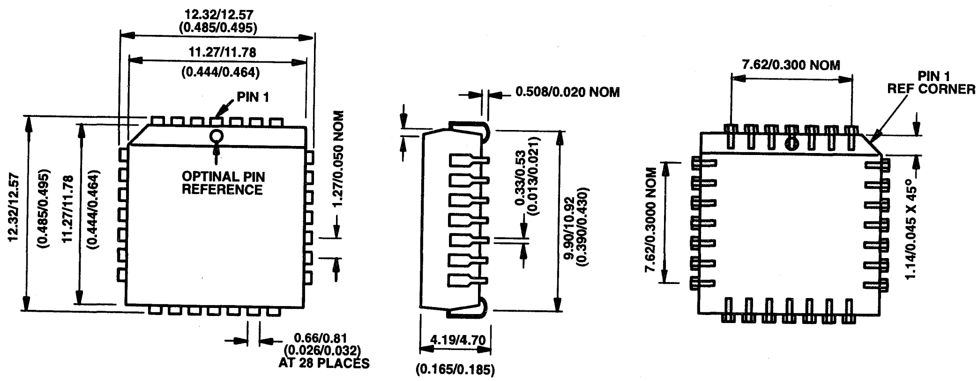

28-LEAD CERAMIC DIL - DG28

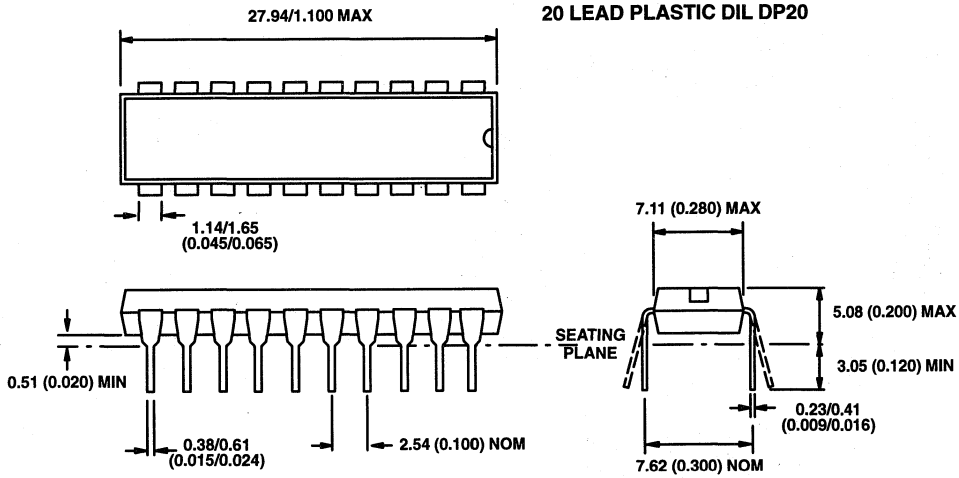

20-LEAD PLASTIC DIL - DP20

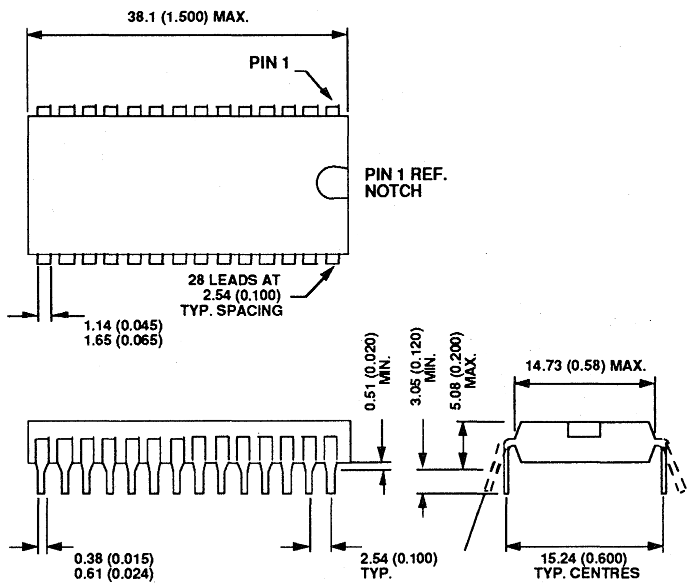

28-LEAD PLASTIC DIL - DP28

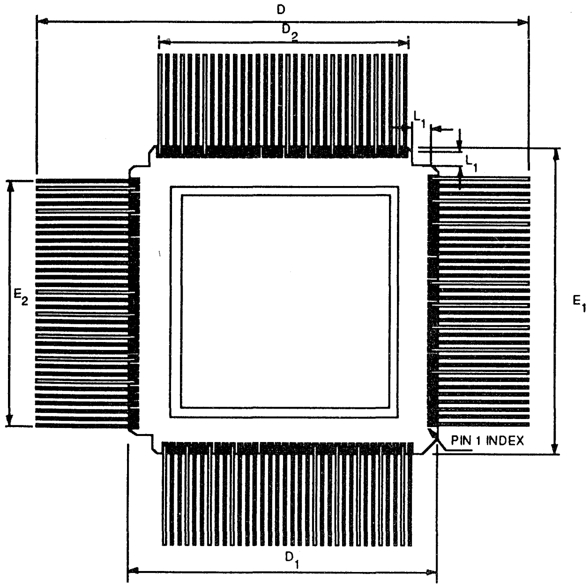

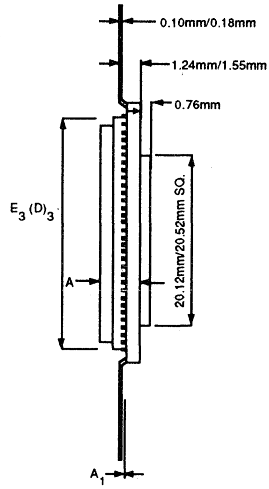

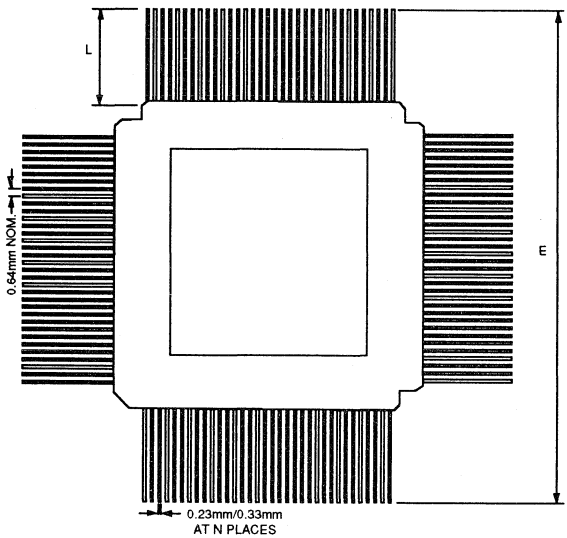

Sym-bol	MILLIMETRES			Note
	MIN	Nominal	MAX	
A	2.08		2.59	
A_1		0.20		
D	36.96		38.35	
D_1	24.00		24.38	
D_2		20.32		
D_3	20.62		21.30	
E	36.96		38.35	
E_1	24.00		24.38	
E_2		20.32		
E_3	20.62		21.03	
L	6.48		6.99	
L_1		1.27		
N		132		
NOTE				

Sym-bol	INCHES			Note
	MIN	Nominal	MAX	
A	0.082		0.102	
A_1		0.008		
D	1.455		1.510	
D_1	0.945		0.960	
D_2		0.800		
D_3	0.812		0.828	
E	1.455		1.510	
E_1	0.945		0.960	
E_2		0.800		
E_3	0.812		0.828	
L	0.255		0.275	
L_1		0.050		
N		132		
NOTE				

132-PIN POWER CERAMIC QFP - GC132

(Used for the PDSP16340, PDSP16350, PDSP16488, PDSP16510 and PDSP16540)

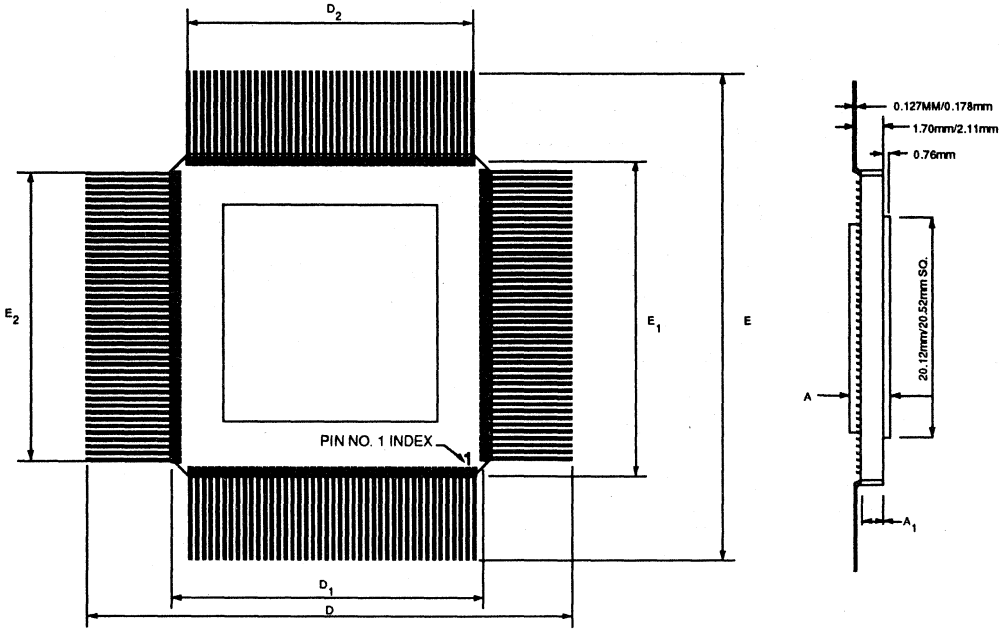

Sym-bol	MILLIMETRES MIN	Nominal	MAX	Note
A	2.95		3.66	
A_1			2.10	
D	45.21	45.97	46.36	
D_1	28.83	29.21	29.59	
D_2		26.67		
E	45.21	45.97	46.36	
E_1	28.83	29.21	29.59	
E_2		26.67		
L	7.87		8.38	
N		172		
NOTE				

Sym-bol	INCHES MIN	Nominal	MAX	Note
A	0.116		0.144	
A_1			0.083	
D	1.780	1.810	1.825	
D_1	1.135	1.150	1.165	
D_2		1.050		
E	1.780	1.810	1.825	
E_1	1.135	1.150	1.165	
E_2		1.050		
L	0.310		0.330	
N		172		
NOTE				

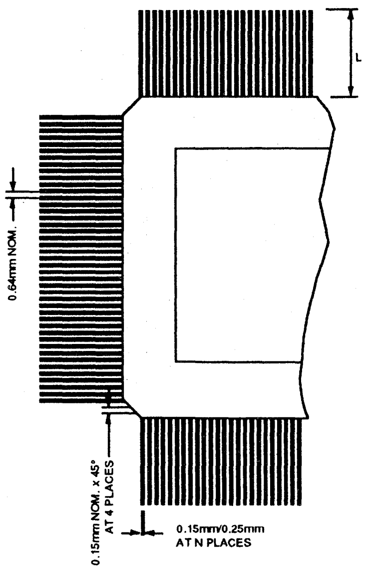

172-PIN POWER CERAMIC QFP - GC172

(Used for the PDSP16256)

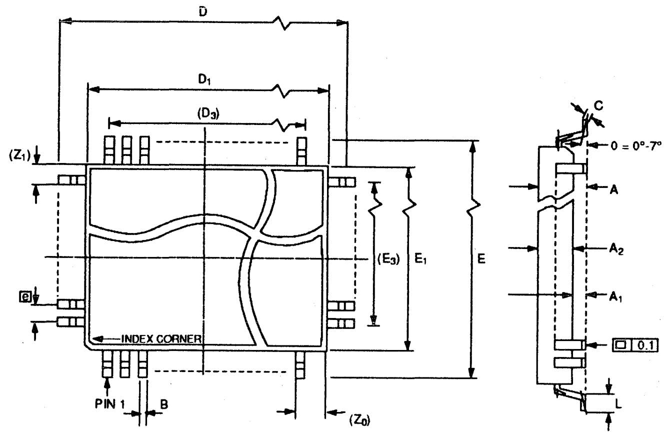

Sym -bol	MILLIMETRES			Note
	MIN	Nominal	MAX	
A			3.25	
A_1	0.15			
A_2			2.84	
D	23.65	24.20	24.50	
D_1		19.99		
D_3		18.85	REF	
E	17.65	18.20	18.50	
E_1		14.00		
E_3		12.60	REF	
L	0.51	0.81	1.11	
e		0.65	BSC	
B	0.30	0.36	0.46	
C	0.13	0.15	0.23	
N		100		
N_D		30		
N_E		20		
NOTE	RECTANGULAR			

Sym -bol	INCHES			Note
	MIN	Nominal	MAX	
A			0.128	
A_1	0.006			
A_2			0.112	
D	0.931	0.953	0.965	
D_1		0.787		
D_3		0.742	REF	
E	0.695	0.717	0.729	
E_1		0.551		
E_3		0.486	REF	
L	0.022	0.032	0.044	
e		0.026	BSC	
B	0.012	0.014	0.018	
C	0.005	0.006	0.009	
N		100		
N_D		30		
N_E		20		
NOTE	RECTANGULAR			

100-PIN CERAMIC QFP - GG100

(Used for the PDSP16318 and PDSP16330)

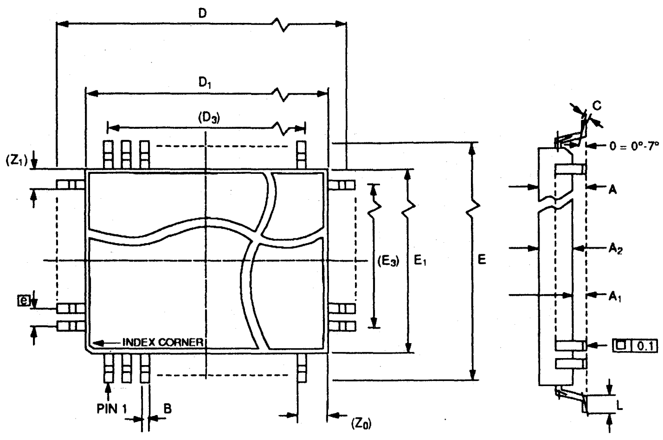

Sym -bol	MILLIMETRES			Note
	MIN	Nominal	MAX	
A			4.06	
A_1	0.25			
A_2			3.61	
D	31.65	31.90	32.15	
D_1		27.99		
D_3		23.20	REF	
E	31.65	31.90	32.15	
E_1		27.99		
E_3		23.20	REF	
L	0.65	0.80	0.95	
e		0.80	BSC	
B	0.30	0.36	0.45	
C	0.15	0.175	0.23	
N		120		
N_D		30		
N_E		30		
NOTE	SQUARE			

Sym -bol	INCHES			Note
	MIN	Nominal	MAX	
A			0.160	
A_1	0.010			
A_2			0.142	
D	1.246	1.256	1.266	
D_1		1.102		
D_3		0.913	REF	
E	1.246	1.256	1.266	
E_1		1.102		
E_3		0.913	REF	
L	0.026	0.031	0.037	
e		0.0315	BSC	
B	0.012	0.014	0.018	
C	0.006	0.007	0.009	
N		120		
N_D		30		
N_E		30		
NOTE	SQUARE			

120-PIN CERAMIC QFP - GG120

(Used for the PDSP16112)

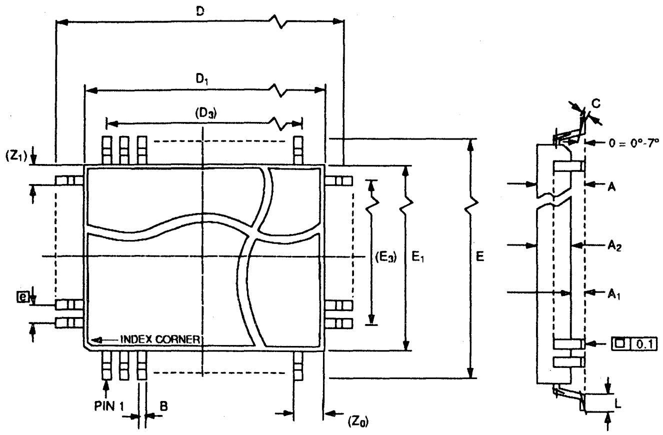

Sym-bol	MILLIMETRES			Note
	MIN	Nominal	MAX	
A			4.06	
A_1	0.25			
A_2			3.61	
D	31.65	31.90	32.15	
D_1		27.99		
D_3		22.75	REF	
E	31.65	31.90	32.15	
E_1		27.99		
E_3		22.75	REF	
L	0.65	0.80	0.95	
e		0.65	BSC	
B	0.22	0.30	0.35	
C	0.15	0.175	0.23	
N		144		
N_D		36		
N_E		36		
NOTE	SQUARE			

Sym-bol	INCHES			Note
	MIN	Nominal	MAX	
A			0.160	
A_1	0.010			
A_2			0.142	
D	1.246	1.256	1.266	
D_1		1.102		
D_3		0.896	REF	
E	1.246	1.256	1.266	
E_1		1.102		
E_3		0.896	REF	
L	0.026	0.031	0.037	
e		0.0256	BSC	
B	0.009	0.012	0.014	
C	0.006	0.007	0.009	
N		144		
N_D		36		
N_E		36		
NOTE	SQUARE			

144-PIN CERAMIC QFP - GG144

(Used for the PDSP16116)

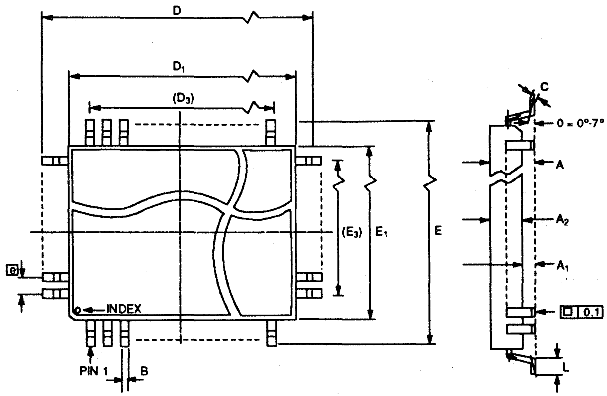

Sym-bol	INCHES MIN	INCHES Nominal	INCHES MAX	Note
A			0.128	
A_1	0.006		0.016	
A_2	0.100		0.112	
D	0.931	0.941	0.951	
D_1	0.783	0.787	0.791	
D_3		0.742	REF	
E	0.695	0.705	0.715	
E_1	0.547	0.551	0.555	
E_3		0.486	REF	
L	0.026	0.031	0.037	
e		0.026	BSC	
B	0.010		0.016	
C	0.005		0.008	
N		100		
N_D		30		
N_E		20		
NOTE	RECTANGULAR			

Sym-bol	MILLIMETRES MIN	MILLIMETRES Nominal	MILLIMETRES MAX	Note
A			3.25	
A_1	0.15		0.40	
A_2	2.55		2.85	
D	23.65	23.90	24.15	
D_1	19.90	20.00	20.10	
D_3		18.85	REF	
E	17.65	17.90	18.15	
E_1	13.90	14.00	14.10	
E_3		12.60	REF	
L	0.65	0.80	0.95	
e		0.65	BSC	
B	0.25		0.40	
C	0.13		0.20	
N		100		
N_D		30		
N_E		20		
NOTE	RECTANGULAR			

100-PIN PLASTIC QFP - GP100

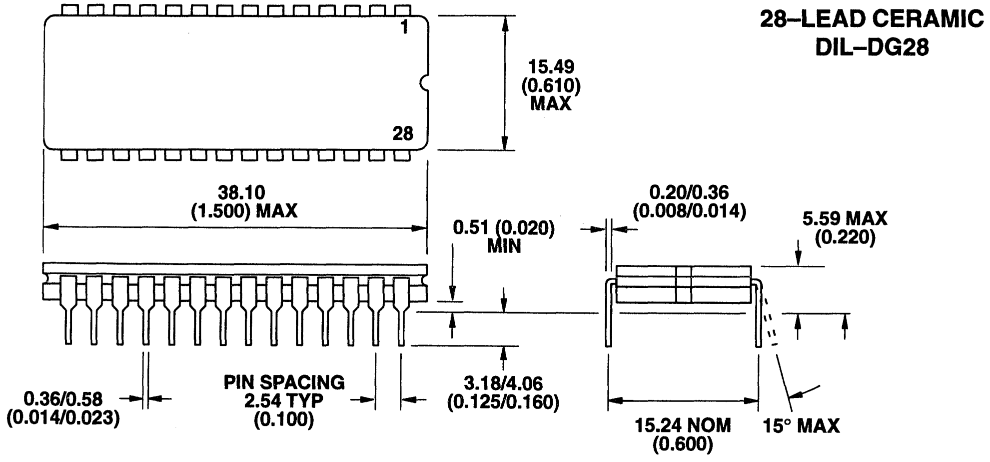

HP28

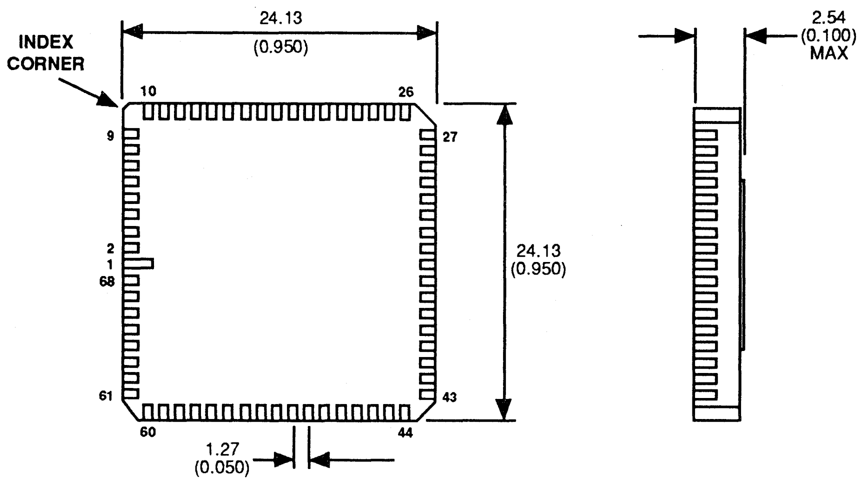

68 CONTACT LCC PACKAGE - LC68

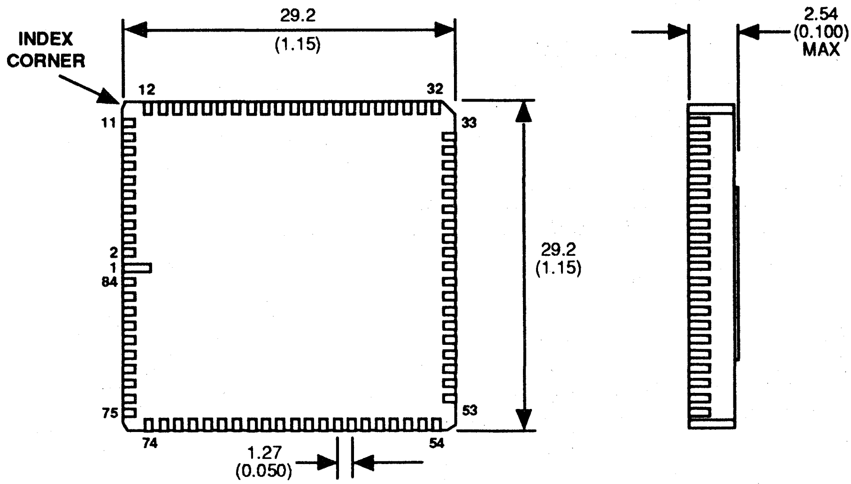

84-PIN LEADLESS CHIP CARRIER - LC84
(HERMETIC)

20 LEAD MINIATURE PLASTIC MP20

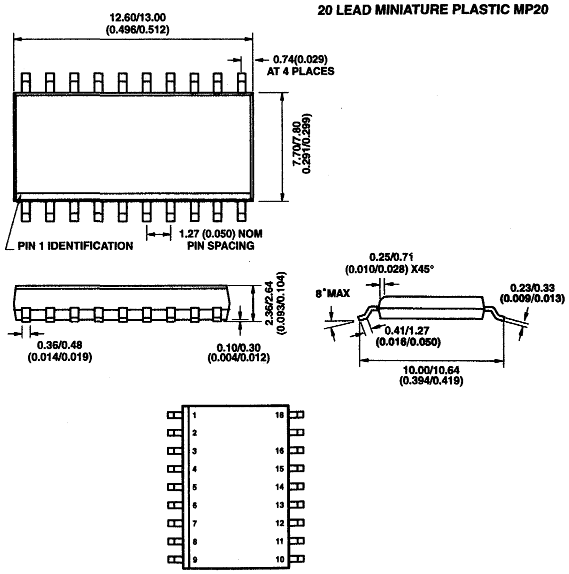

MP20

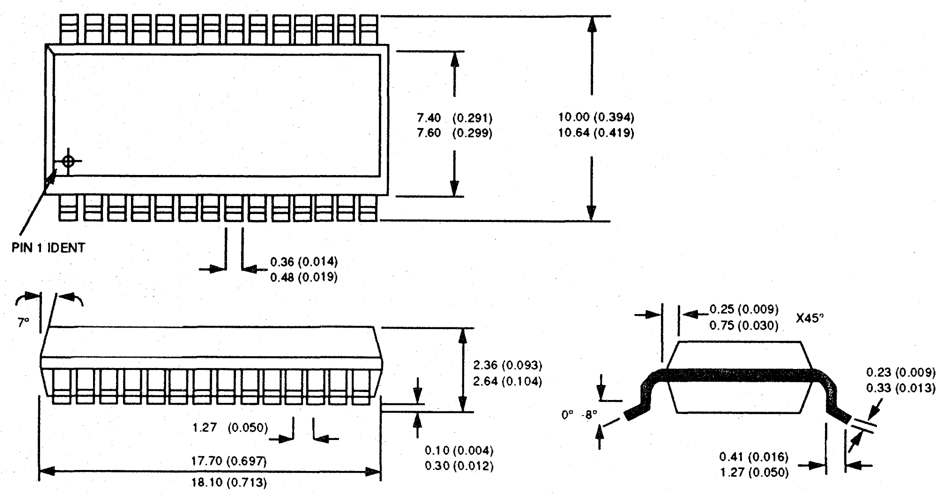
7.40 (0.291)
7.60 (0.299)
10.00 (0.394)
10.64 (0.419)
PIN 1 IDENT
0.36 (0.014)
0.48 (0.019)
7°
2.36 (0.093)
2.64 (0.104)
1.27 (0.050)
17.70 (0.697)
18.10 (0.713)
0.10 (0.004)
0.30 (0.012)
0.25 (0.009)
0.75 (0.030)
X45°
0.23 (0.009)
0.33 (0.013)
0° -8°
0.41 (0.016)
1.27 (0.050)

MP28

GEC Plessey Semiconductors Locations

GEC PLESSEY
SEMICONDUCTORS

Headquarters Operations

UNITED KINGDOM	Cheney Manor, Swindon, Wiltshire, United Kingdom, SN2 2QW. Tel: (0793) 518000 Fax: 0793 518411
NORTH AMERICA	P O Box 660017, 1500 Green Hills Road, Scotts Valley, California 95067-0017, USA. Tel:(408) 438 2900 Fax: (408) 438 5576

Customer Service Centres

FRANCE & BENELUX	Z.A. Courtaboeuf, Miniparc-6, Avenue des Andes, Bat. 2-BP 142, 91944, Les Ulis Cedex A. France.Tel: (1) 64 46 23 45. Fax: (1) 64 46 06 07
GERMANY, AUSTRIA and SWITZERLAND	Ungererstraße 129, 80805 Munchen 40, Germany. Tel: 089/36 0906-0. Fax: 089/36 0906-55
ITALY	Via Raffaello Sanzio, 4, 20092 Cinisello Balsamo (MILAN). Tel: (02) 66040867. Fax: (02) 66040993.
JAPAN	CTS Kojimachi Building (4th Floor), 2-12, Kojimachi, Chiyoda-ku, Tokyo 102. Tel: (03) 5276-5501. Fax: (03) 5276-5510.
NORTH AMERICA	**Integrated Circuits and Microwave Products** P O Box 660017, 1500 Green Hills Road, Scotts Valley, California 95067-0017, USA. Tel: (408) 438 2900. Fax: (408) 438 7023 **Hybrid Products** 100 Smith Street. Farmingdale, NY11735, USA. Tel: (516) 293 8686. Fax: (516) 293 0061.
SOUTH EAST ASIA	No. 3 Tai Seng Drive, GEC Building, Singapore 1953. Tel: (65) 3827708. Fax: (65) 3828872
SWEDEN	Katarina Bangata 79, 116 42 Stockholm. Tel: 46 8 7029770. Fax: 46 8 6404736
UK, EIRE, DENMARK, FINLAND and NORWAY	Unit 1,Crompton Road,Groundwell Industrial Estate, Swindon, Wilts, U.K., SN2 5AF. Tel: (0793) 518510. Fax: (0793) 518582. **Power Division** Carholme Road, Lincoln, U.K., LN1 1SG. Tel: (0522) 510500. Fax: (0522) 510550

North American Sales Offices

NATIONAL SALES	P O Box 660017, Scotts Valley, CA 95067-0017. Tel: (408) 438-2900. ITT Telex: 4940840. Fax: (408) 438-5576.
METRO NY/NJ	160 Smith Street, Farmingdale, NY11735. Tel: (516) 293-8686. Fax: (516) 293-0061.
EASTERN	Two Dedham Place, Suite 1, Allied Drive, Boston, MA 02026. Tel: (617) 320-9790. Fax: (617) 320-9383.
WESTERN	1735 Technology Drive, Suite 100, San Jose, California 95110. Tel: (408) 451-4700. Fax: (408) 451-4710.
ARIZONA/NEW MEXICO	4635 South Lakeshore Drive, Tempe, AZ 85282. Tel: (602) 491-0910. Fax: (602) 491-1219.
SOUTH CENTRAL	9330 LBJ Freeway, Ste. 665, Dallas, TX 75243. Tel: (214) 690-4930. Fax: (214) 680-9753.
NORTHWEST	7935 Datura Circle West, Littleton, CO 80120. Tel: (303) 798-0250. Fax: (303) 730-2460.
DIXIE and FLORIDA	668 N. Orlando Ave., Suite 1015 B, Maitland, FL 32751. Tel: (407) 539-1700. Fax: (407) 539-0055. 385 Commerce Way, Longwood, FL 32750. Tel: (407) 339-6660. Fax: (407) 339-9355.
SOUTHWEST	13900 Alton Parkway #123, Irvine, CA 92718. Tel: (714) 455-2950. Fax: (714) 455-9671.

DISTRIBUTION SALES 1500 Green Hills Road, Scotts Valley, CA 95066.
Tel: (408) 438-2900. ITT Telex: 4940840. Fax: (408) 438-7023.

CANADA 3608 Boul. St. Charles, Suite 9, Kirkland, Quebec, H9H 3C3.
Tel: (514) 697-0095. Fax: (514) 694-7006.
Shucho-Gu, CPO Box 7981, Seoul. Tel: 2 595 9101-6. Tx: K25981 KMLCORP.
Fax: 2 595 9107.

Semi-Custom Design Centres

AUSTRALIA Unit 1, 38 South Street, Rydalmere, NSW 2116, Australia.
Tel: 612 638 1888. Fax: 612 638 1798.

FRANCE & BENELUX Z.A. Courtaboeuf, Miniparc-6, Avenue des Andes, Bat.2-B.P. No.142 91944
Les Ulis Cedex A, France.Tel: (1) 64 46 23 45. Fax: (1) 64 46 06 07.

ITALY Via Raffaello Sanzio, 4, 20092 Cinisello Balsamo (MILAN).
Tel: (02) 66040867. Fax: (02) 66040993.

GERMANY Ungererstraße 129, D 80805 Munchen 40. Tel: (089) 3609 06 0.
Fax: (089) 3609 06 55.

JAPAN CTS Kojimachi Building (4th Floor), 2-12, Kojimachi, Chiyoda-ku, Tokyo 102.
Tel: (03) 5276-5501. Fax: (03) 5276-5510.

N AMERICA: Canada **Alberta Microelectronics Center,** 3553 31st St., N.W., Calgary, Alberta T2L2K7.
Tel: (403) 289-2043.
Microstar Technologies, 7050 Bramelea Rd., #27A Mississaugo, Ontario L5SITI
Canada. Tel: (416) 671-8111

USA P O Box 660017, 1500 Green Hills Road, Scotts Valley, California 95067-0017.
Tel: (408) 438-2900. Fax: (408) 438-5576.
Two Dedham Place, Suite 1, Allied Drive, Boston, Massachusetts 02026.
Tel: (617) 320-9369. Fax: (617) 320-9383.
13900 Alton Parkway #123, Irvine, California 92718.
Tel: (714) 455-2950 Fax: (714) 455-9671.

Colorado, USA **Analog Solutions,** 5484 White Place, Boulder, CO 80303. Tel: (303) 442-5083.

Illinois, USA **Frederikssen & Shu Laboratories, Inc.,** 531 West Golf Rd., Arlington Heights,
IL 60005. Tel: (312) 956-0710.

UNITED KINGDOM Cheney Manor, Swindon, Wiltshire SN2 2QW.
Tel: (0793) 518000. Fax: (0793) 518411.
Tweedale Way, Hollinwood, Oldham, Lancashire OL9 7LA.
Tel: 061 682 6844. Fax: 061 688 7898.
Doddington Road, Lincoln LN6 3LF. Tel: 0522 500500. Fax : 0522 500550.

North American Representatives

ALABAMA **Electramark, Inc.,** 500 Wynn Drive, Suite 521, Huntsville, AL 35816.
Tel: (205) 830-4400 Fax: (205) 830-4406.

ARIZONA **Fred Board Associates,** 7353 E, 6th Avenue, Scottsdale, AZ 85251.
Tel: (602) 994-9388. Fax: (602) 994-9477.

CALIFORNIA **Gary Chilcote & Associates,** P.O. Box 1795, 1902 Quite Ranch Road, Fallbrook,
CA 92028. Tel: (619) 728 7678.Fax: (619) 728 3738.
Jones & McGeoy, 5100 Campus Drive, Suite 300, Newport Beach, CA 92660.
Tel: (714) 724 8080. Fax: (714) 724 8090.

CONNECTICUT **Stone Components,** 123 Commerce St., Clinton, CT 06413.
Tel: (203) 669-4344. Fax: (203)669-9958.

FLORIDA **American Micro Sales,** 1325 N. Congress Ave., Ste 204, West Palm Beach,
FL 33401. Tel: (407) 689 3860.Fax: (407) 689 3168.
American Micro Sales, 274 Wilshire Blvd., Ste 241, Casselberry, FL 32707.
Tel: (407) 831 2505. Fax: (407) 831 1842.

	American Micro Sales, P.O. Box 399, 1033 Rosetree Lane, Tarpon Springs, FL 34688/9. Tel: (813) 938 3073. Fax: (407) 831 1842.
GEORGIA	**Electramark, Inc.,** 6030-H Unity Drive, Norcross, GA 30071. Tel: (404) 446-7915. Fax: (404) 263 6389
ILLINOIS	**Micro Sales, Inc.,** 901 West Hawthorn Drive, Itasca, IL 60043. Tel: (708) 285-1000. Fax: (708) 285-1008.
INDIANA	**Leslie M. DeVoe,** 4371 E. 82nd St., Suite D, Indianapolis, IN 46250. Tel: (317) 842-3245. Fax: (317) 845-8440.
IOWA	**Lorenz Sales,** 5270 N. Park Place N.E., Cedar Rapids, IA 52402. Tel: (319) 377-4666. Fax: (319) 377-2273.
KANSAS	**Lorenz Sales, Inc.,** 8645 College Blvd., Suite 220, Overland Park, KS 66210. Tel: (913) 469-1312. Fax: (913) 469-1238.
	Lorenz Sales, Inc., 1530 Maybelle, Wichita, KS 67212. Tel: (316) 721-0500. Fax: (316) 721-0566.
MARYLAND	**Walker Associates,** 1757 Gablehammer Road, Westminster, MD 21157. Tel: (410) 876-9399. Fax: (410) 876-9285.
	Walker Associates, 169 Queen Anne Bridge Road, Mitchellville, MD 20716. Tel: (410) 249-7145.
MASSACHUSETTS	**Stone Components,** 2 Pierce Street. Framingham, MA 01701. Tel: (508) 875-3266. Fax: (508) 875-0537.
	Stone Components, 10 Atwood Road, Newburyport, MA 01950. Tel: (508) 462-1079.
MICHIGAN	**Greiner Associates Inc.,** 15224 E. Jefferson Avenue, Grosse Point Park, MI 48230. Tel: (313) 499-1088. Fax: (313) 499-0665.
MINNESOTA	**High Technology Sales,** 4801 West 81st Street, Suite 115, Bloomington, MN 55437. Tel: (612) 844-9933. Fax: (612) 844-9930.
MISSOURI	**Lorenz Sales, Inc.,** 10176 Corporate Square Dr., Suite120, St. Louis, MS 63121. Tel: (314) 997-4558. Fax: (314) 997-5829.
NEBRASKA	**Lorenz Sales,** 2801 Garfield Street, Lincoln, NE 68502. Tel: (402) 475-4660. Fax: (402) 474-7094.
	HLM Assoc., 333 Littleton Raod, Parsippany, NJ 07054. Tel: (201) 263-1535. Fax: (201) 263-0914.
NEW HAMPSHIRE	**Stone Components,** 436 S. Baboosic Lake Rd., Merrimack, NH 03054. Tel: (603) 429-3462. Fax: (603) 429-3462.
NEW YORK	**HLM Assoc.,** 64 Mariners Lane, Box 328, Northport, NY 11768. Tel: (516) 757-1606. Fax: (516) 757-1636.
	Regan Compar, 3301 Country Club Road, Endwell, NY 13760. Tel: (607) 754-2171 Fax: (607) 754-4270.
	Regan Compar, 214 Dorchester Avenue, Syracuse, NY 13202. Tel: (315) 432-8232. Fax: (315) 432-8238.
	Regan Compar, 37A Brookhill Lane, Rochester, NY 14625. Tel: (716) 271-2230. Fax: (716) 381-2840.
OHIO	**Scott Electronics, Inc.,** 3131 S. Dixie Dr., Suite 200, Dayton, OH 45434. Tel: (513) 294-0539. Fax: (513) 294-4769.
	Scott Electronics, Inc., 30 Alpha Park Drive, Cleveland, OH 44143. Tel: (216) 473-5050. Fax: (216) 473-5055.
	Scott Electronics, Inc., 916 Eastwind Dr., Westerville, OH 43081. Tel: (614) 882-6100. Fax: (614) 882-0900.
	Scott Electronics, Inc., 10901 Reed Hartman Hwy., Suite 301, Cincinnati, OH 45242. Tel: (513) 791-2513. Fax: (513) 791-8059.
PENNSYLVANIA/ NEW JERSEY	**Metz-Jade Associates, Inc.,** 7 Wynnewood Rd, Suite 203, P.O.Box 276, Wynnewood, PA 19096. Tel: (215) 896-7300. Fax: (215) 642-6293.
TEXAS	**Oeler & Menelaides, Inc.,** 8430 Meadow Rd., Suite 224, Dallas, TX 75231. Tel: (214) 361-8876. Fax: (214) 692-0235.
	Oeler & Menelaides, Inc., 8705 Shoal Creek Rd., #103, Austin, TX 78758. Tel: (512) 453-0275. Fax: (512) 453-0088.

WISCONSIN **Micro Sales, Inc.,** 210 Regency Ct., Suite L101, Waukesha, WI 53186-0545.
Tel: (414) 786-1403. Fax: (414) 786-1813.

CANADA **GM Assoc. Inc.,** 7050 Bramalea Road, Mississauga, Ontario L5S 1T1 (Toronto).
Tel: (416) 671-8111. Fax: (416) 671-2422.
GM Assoc. Inc., 3860 Cote-Vertu, St. Laurent, Quebec H4R 1N4.
Tel: (514) 335-9572. Fax: (514) 335-9573.
GM Assoc. Inc., 301 Moodie Dr., Nepean, Ontario K2H 9CR (Ottawa).
Tel: (613) 820-3822. Fax: (613) 820-7633.
GM Assoc. Inc., 300-3665 Kingsway, Vancouvr, BC V5R 5W2.
Tel: (604) 439-3383. Fax: (604) 439-8479.

Distributors

AUSTRALIA and NEW ZEALAND **GEC Electronics Division,** Unit 1, 38 South Street, Rydalmere, NSW 2116, Australia.
Tel: 612 638 1888. Fax: 612 638 1798.

AUSTRIA **Eurodis Electronics GmbH,** Lamezanstrasse 10, A-1232 Wien.
Tel: 1 610620. Fax: 1 61062151.

BELGIUM **Tekelec Belgium NV,** Bergensesteenweg 501, B-1500 Halle.
Tel: 02 360 1288 Fax: 02 360 3807,

BULGARIA **Macro Group,** P.O. Box 140, Sofia 1618. Tel/Fax: 359 2 525537.

CANADA **Semad,** 1825 Woodward Dr., Ottawa, Ontario K2C 0R3.
Tel: (613) 727-8325. Fax: (613) 727-9489.
Semad (Corp.), 85 Spy Court, Markham, Ontario L3R 4Z4.
Tel: (416) 475-8500. Fax: (416) 475-4158.
Semad, 243 Place Frontenac, Pointe Claire, Que H9R 4Z7.
Tel: (514) 694-0860. Fax: (514) 694-0965.
Semad, 6815 8th St. N.E., Ste. 175, Calgary, Alberta T2E 7H7.
Tel: (403) 252-5664 Fax: (604) 420-0124.
Semad, 8563 Government St, Burnaby, BC V3N 4S9.
Tel: (604) 420-9889. Fax: (604) 420-0124.

CHINA **World Peace Industrial Co Ltd,** Room 709, Oriental Building, 39 Jianshe Road,
q Shenzen, China.Tel: 755-2284970. Fax: 755-2284972.

CZECH REPUBLIC **Macro Semiconductors,** Bechynova 3, 160 00 Prague 6.
Tel: 42 2 3112182. Fax: 42 2 3113454.

DENMARK **Scansupply A/S,** Gladsaxevej 356, DK-2860 Soeborg.
Tel: 45 39 66 50 90. Fax: 45 39 66 50 40.
Scansupply A/S, Marselisborg Havnevej 36, 8000 Arhus C.
Tel: 45 86 12 77 88. Fax: 45 86 12 77 18.

EASTERN EUROPE **FA Bernhart GmbH,** Melkstattweg 27, PO Box 1628, D 8170 Bad Toelz., Germany.
Tel: 80 41 41 676 Fax: 80 41 71 504 Tx: 526246 FABD.

FRANCE **3D:**
6-8 rue Ambroise Croisat, Z.I. des Glaises, 91120 Palaiseau. Tel: 1 64 47 29 29.
Fax: 1 64 47 00 84.
3 rue Berthelot, 69627 Villeurbanne CEDEX. Tel: 72 35 22 00. Fax: 72 34 67 72.
Z.I. du terroir, rue de l'industrie, 31140 Saint Alban.
Tel: 61 37 44 00. Fax: 61 37 44 29.
Parc Club du golf, Batiment 1, 13856 Aix-en-Provence CEDEX 3. Tel: 91 61 80 20.
Omnitech Sertronique:
C.A. de Montheard, 11 rue Edgar Brant, 72016 Le Mans CEDEX.
Tel: 43 86 74 74. Fax: 43 86 74 86.
165 boulevard de Valmy, Batiment Evolic 1, 92706 Colombes.
Tel: 1 46 13 07 80. Fax: 1 46 13 07 90.
ZAC des Petites Landes, 44470 Thouare-sur-Loire. Tel: 40 68 06 07.
Fax: 40 68 06 72.
99 boulevard de l'Artillerie, 69007 Lyon. Tel: 72 73 11 87. Fax: 72 73 18 00.

37 rue Saint-Eloi, 76000 Rouen. Tel: 35 88 00 38. Fax: 35 15 06 22.
18/20 rue Cabanis, 59007 Lille. Tel: 20 33 21 97. Fax: 20 56 00 49.
Parc Cadera, Batiment F, 33700 Bordeaux. Tel: 56 34 46 00. Fax: 56 34 47 13.

GERMANY **AS Electronic Vertriebs GmbH,** In den Garten 2, D-61352 Bad Homburg. Tel: 06172 458931. Fax: 06172 42000.

Astronic GmbH, Gruenwalder Weg 30, D-82041 Deisenhofen. Tel: 089 6130303. Fax: 089 6131668.

Micronetics GmbH, Dieselstrasse 12, D-71272 Renningen. Tel: 07159-925830. Fax: 07159-9258355.

Weisbauer Elektronik GmbH, Heiliger Weg 1, D-44135 Dortmund. Tel: 0231 579547. Fax: 0231 577514.

GREECE **Impel Ltd.,** 30 Rodon Str, Korydallos, Piraeus, Greece. Tel: 010 30 1 49 67815. Tlx: 213835. Fax: 01 49 54041.

HONG KONG **Gain Tune Ltd,** Room 809, 8th Floor, Hunghom Commercial Centre, Tower B, 37-39 Ma Tau Wai Road,Hunghom, Kowloon, Hong Kong. Tel: 852-3654860. Fax: 852-7641129.

YES Products Ltd., Block E, 15/F Golden Bear Industrial Centre, 66-82 Chaiwan Kolk Street,Tsuen Wan,N.T. Hong Kong. Tel: 4144241-6. Tx: 36590. Fax: 4136078.

HUNGARY **Macro Group,** Sza'mado u.15, FSZT 3, Budapest 1016. Tel: 36 1 1868033. Fax: 36 1 1852318.

INDIA **Mekaster Telecom PVT Ltd.,** 908 Ansal Bhawan, 16 Katuba Ghandi Marg, New Delhi, 100 001 India Tel: 11 3312110 Fax: 11 3712155.

ITALY **Adelsy - Divisione della Generalmusic SpA,** Via Novara 570, 20153 Milano. Tel: 02 3810310. Fax: 02 38002988.

Eurelettronica SpA, Via E. Fermi 8, 20090 Assago (MI). Tel: 02 457841. Fax: 02 4880275.

Fanton Srl, Milano - Torino - Bologna - Firenze - Roma - Padova. Tel: 02 48912963. Fax: 02 4890902.

JAPAN **Cornes & Company Ltd.,** Ryukakusan Building (5th Floor), 2-5-12 Higashi-Kanda, Chiyoda-ku. Tokyo 101. Tel: 3-5821-1628. Fax: 3-5821-1632.

Cornes & Company Ltd., Cornes House, 13-40 Nishi-honmachi 1-chome, Nishi-Ku, Osaka 550. Tel: 6 532 1012.Tx: 525-4496. Fax: 6 541-8850.

KOREA **KML Corporation,** 6th Floor, Dukmyung Building, (A-Dong) 113-3, Banpo-Dong, Shucho-Gu, CPO Box 7981, Seoul. Tel: 2 595 9101-6. Tx: K25981 KMLCORP. Fax: 2 595 9107.

MALAYSIA **Adequip Enterprise Sdn Bhd,** #6-01 6th Floor, Wisjma Stephens, 88 Jalan Raya Chulan, 50200 Kualar Lumpur, Malaysia. Tel: 2423522. Fax: 2423264.

NETHERLANDS **Tekelek-Airtronic B.V.,** PO Box 63, NL-2700 AB Zoetermeer. Tel: 079 310100. Fax: 079 417504

NORWAY **Skandinavisk Elektronikk A/S,** Ostre Aker Vei 99, Postboks 99 Veitvet, Oslo 5. Tel: 22 64 11 50. Tx: 71963 Fax: 22 64 34 43.

POLAND **Macropol Co. Ltd,** 03-301 Warsaw, Ul. Jagiellonska 80. Tel/Fax: 48 22 112933.

PORTUGAL **Anatronic SL,** Urbanisazao Do, Infantado, Lote 18, 2° ESQ, Loures. Tel: 1 79 33 2 99.

SLOVAKIA **Macro Zilinia,** Vysokoskolakov 6, 010-01 Zilinia. Tel: 42 89 34181. Fax: 42 89 34109.

SOUTH AFRICA **T.E.C. Tellumat Electronic Components,** 1 Jansen Road, Jet Park, Boksburg 1459, P.O. Box 8174, Elandsfontein 1406. Tel: 2711 823 2950. Fax: 2711 823 2668.

SPAIN **Anatronic SA,** Avda de Valladolid 27, 28008 Madrid. Tel: 91 542 4455/6. Fax: 91 2486975.

Anatronic SA, Bailen, 176, Estresuelo 1°, 08037 Barcelona. Tel: 93 258 1906/7. Fax: 93 258 7128.

SWEDEN **Pronesto AB (Microwave Products Only),** Finlandsgatan 18, 16475 Kista. Tel: (08) 752 9080. Fax: (08) 751 4111.

SWITZERLAND **Basix für Electronik AG,** Hardturmstr 181, CH-8010 Zuerich. Tel: 1 2761111. Fax: 1 2764199.

	Elbatex AG, Hardstr 72, CH-5430 Wettingen. Tel: 56 2751 11. Fax: 56 2755 32.
TAIWAN	**Prospect Technology Corporation,** 5, Lane 55, Long-Chiang Road, Taipei, Taiwan. Tel: (886-2) 721-9533. Fax: (886-2) 773-3756.
	World Peace Industrial Co Ltd, 8th Floor, 76 Cheng Kung Road, Sec. I., Nankang District, Taipei, Taiwan ROC. Tel: (886-2) 788-5200. Fax: (886-2) 788-3255.
THAILAND	**Westech Electronics Co. Ltd,** 77/113 Moo Ban Kitikorn, Ladprao Soi 3, Ladprao Road, Ladyao, Jatujak, Bangkok 10900. Thailand. Tel: 2 5125531. Fax: 2 2365949.
TURKEY	**EMPA,** 34630, Besyol Londra Asfalti Florya Is Merkezi Sefakoy, Istanbul. Tel: (1) 599 30 50 Tx: 21137 Etna tr Fax: (1) 598 53 53.
UNITED KINGDOM	**Celdis Ltd.,** 300 Kings Road, Reading, Berks RG1 4GA. Tel: 0734 666676. Tx: 818142. Fax: 0734 263100.
	Farnell Electronic Components Ltd., Canal Road, Leeds LS12 2TU. Tel: 0532 636311. Tx: 55147. Fax: 0532 633404.
	Gothic Crellon Ltd., 3 The Business Centre, Molly Millars Lane, Wokingham, Berkshire RG11 2EY. Tel: 0734 788878, 787848. Tx: 847571. Fax: 0734 776095.
	Gothic Crellon Ltd., P.O.Box 301, Elizabeth House, 22 Suffolk Street, Queensway, Birmingham B1 1LZ. Tel: 021 643 6365. Tx: 338731. Fax: 021 633 3207.
	The Macro Group, Burnham Lane, Slough SL1 6LN. Tel: 0628 604383. Tx: 847945 Fax: 0628 66873.
	Semiconductor Specialists (UK) Ltd., Carroll House, 159 High Street Yiewsley, West Drayton, Middlesex UB7 7XB Tel: (0895) 445522. Tx: 21958. Fax: (0895) 422044.
	ESD Distribution Ltd, Edinburgh Way, Harlow, Essex CM20 2DF. Tel: 0279 626777. Tlx: 818801. Fax: 0279 441687.
	Unitel Ltd., Unitel House. Fishers Green Road, Stevenage, Herts. Tel: (0438) 312393. Tx: 825637. Fax: (0438) 318711.
USA: Alabama	**Pioneer Technologies,** 4835 University Square #5, Huntsville, AL 35816. Tel: (205) 837-9300. Fax: (205) 837-9358.
	Hammond, 4411-B Evangel Circle NW, Huntsville, AL 35816. Tel: (205) 830-4764. Fax: (205) 830-4287.
Arizona	**Insight Electronics,** 1515 W. University Dr., Suite 103. Tempe, AZ 85281. Tel: (602) 829-1800. Fax: (602) 967-2658.
California	**Insight Electronics (Corp.),** 6885 Flanders Dr., Unit C, San Diego, CA 92121. Tel: (619) 587-0471. Fax: (619) 587-0903.
	Insight Electronics, 2 Venture Plaza, Ste. 340, Irvine, CA 92718. Tel: (714) 727-3291. Fax: (714) 727-1804.
	Insight Electronics, 4333 Park Terrace Dr., Ste. 101, Westlake Village, CA 91361. Tel: (818) 707-2101. Fax: (818) 707-0321.
	Insight Electronics, 1295 Oakmead Parkway, Sunnyvale, CA 94086. Tel: (408) 720-9222. Fax: (408) 720-8390.
	Pioneer Standard, 217 Technology Dr., Ste. 110, Irvine, CA 92718. Tel: (714) 753-5500. Fax: (714) 753-5074.
	Pioneer Standard, 5850 Canogo Ave., Ste. 400, Woodland Hills, CA 91367. Tel: (818) 883-4640. Fax: (818) 883-9721.
	Pioneer Technologies, 134 Rio Robles, San Jose, CA 95134. Tel: (408) 954-9100. Fax: (408) 954-9113.
North Carolina	**Hammond,** 2923 Pacific Ave., Greensboro, NC 27406. Tel: (919) 275-6391. Fax: (919) 272-6036.
	Pioneer Technologies, 2200 Gateway Centre Blvd., Ste. 215, Morrisville, NC 27560. Tel: (919) 460-1530. Fax: (919) 460-1540.
South Carolina	**Hammond,** 1035 Lowndes Hill Rd., Greenville, SC 29607. Tel: (803) 232-0872. Fax: (803) 232-0320.
Colorado	**Insight,** 384 Inverness Dr. S., Ste. 105, Engelwood, CO 80112. Tel: (303) 649-1800. Fax: (303) 690-4754.

Connecticut	**Pioneer Standard,** 2 Trapp Falls Rd., #100, Sheton, CT 06484. Tel: (203) 929-5600. Fax: (203) 929-9791.
Florida	**Hammond Ft. Lauterdale,** 6600 N.W. 21st Ave., Ft. Lauderdale, FL 33309. Tel: (305) 973-7103. Fax: (305) 973-7601.
	Hammond Orlando (Corp.), 1230 West Central Blvd., Orlando, FL 32805. Tel: (407) 849-6060. Fax:(407) 648-8584.
	Pioneer Technologies, 674 South Military Trail, Deerfield Beach, FL 33442. Tel: (305) 428-8877. Fax: (305) 481-2950.
	Pioneer Technologies, 337 South Northlake Blvd., #1000, Altamonte Springs. FL 32701 Tel: (407) 834-9090. Fax: (407) 834-0865.
Georgia	**Hammond,** 5680 Oakbrook Pkwy., Suite 160, Norcross, GA 30093.] Tel: (404) 449-1996. Fax: (404) 242-9834.
	Pioneer Technologies, 4250C Rivergreen Pkwy., Duluth, GA 30136. Tel: (404) 623-1003. Fax: (404) 623-0665
Illinois	**Pioneer Standard,** 2171 Executive Drive #104, Addison, IL 60101. Tel: (708) 495-9680. Fax: (708) 495-9831
Indiana	**Pioneer Standard,** 9350 N. Priority Way W. Drive, Indianapolis, IN 46240. Tel: (317) 573-0880. Fax: (317) 573-0979
Maryland	**Pioneer/Tech. Group. Inc.,** 9100 Gaither Rd., Gaithersburg, MD 20877. Tel: (301) 921-0660. Fax: (301) 921-4255
Massachusetts	**Pioneer Standard,** 44 Hartwell Avenue, Lexington, MA 02173. Tel: (617) 861-9200. Fax: (617) 863-1547
Michigan	**Pioneer Standard,** 13485 Stamford, Livonia, MI 48150. Tel: (313) 525-1800. Fax: (313) 427-3720.
	Pioneer Standard, 4595 Broadmoor Ave. S.E., Ste. 235, Grand Rapids, MI 49512. Tel: (616) 698-1800. Fax: (616) 698-1831.
Minneapolis	**Pioneer Standard,** 7625 Golden Triangle Dr., Eden Prairie, MN 55344. Tel: (612) 944-3355. Fax: (612) 944-3794.
Missouri	**Pioneer Standard,** 2029 Woodland Pkwy. #101, St. Louis, MO 63146. Tel: (314) 432-4350. Fax: (314) 432-4854.
New Jersey	**Pioneer Standard,** 14-A Madison Rd.. Fairfield, NJ 07006. Tel: (201) 575-3510. Fax: (201) 575-3454.
New Mexico	**Alliance Electronics,** 10510 Research Rd. SE, Albuquerque, NM 87123. Tel: (505) 292-3360 Fax: (505) 275-6392
New York	**Mast,** 710-2 Union Pkwy., Ronkonkoma, NY 11779. Tel: (516) 471-4422. Fax: (516) 471-2040.
	Pioneer Standard, 68 Corporate Drive, Binghamton, NY 13904. Tel: (607) 722-9300. Fax: (607) 722-9562.
	Pioneer Standard (Corp.), 60 Crossways Park West, Woodbury, NY 11797. Tel: (516) 921-8700 Fax: (516) 921-2143.
	Pioneer Standard, 840 Fairport Park, Fairport, NY 14450. Tel: (716) 381-7070. Fax: (716) 381-5955.
Ohio	**Pioneer Standard,** 4800 East 131st St., Cleveland, OH 44105. Tel: (216) 587-3600. Fax: (216) 587-3906.
	Pioneer, 4433 Interpoint Blvd., Dayton, OH 45424. Tel: (513) 236-9900. . Fax: (513) 236-8133.
Oregon	**Insight,** 8705 SW Nimbus Ave., #200, Beaverton, OR 97005. Tel: (503) 644-3300
Pennsylvania	**Pioneer Technologies,** Keith Valley Business Center, Horsham, PA 19044. Tel: (215) 674-4000. Fax: (215) 674-3107.
	Pioneer Standard, 259 Kappa Drive, Pittsburgh, PA 15238. Tel: (412) 782-2300. Fax: (412) 963-8255.
Texas	**Insight,** 12701 Research Bl.Ste 301, Austin, TX 78759.
	Insight, 1778 Plano Rd., Suite 320, Richardson, TX 75081. Tel: (214) 783-0800. Fax: (214) 680-2402.
	Insight,15437 McKaskle, Sugarland, TX 77478.

	Pioneer Standard, 1826-D Kramer Ln., Austin. TX 78758. Tel: (512) 835-4000. Fax: (512) 835-9829.
	Pioneer Standard, 13765 Beta Road, Dallas, TX 75244. Tel: (214) 386-7300. Fax: (214) 490-6419.
	Pioneer Standard, 10530 Rockley Rd. #100, Houston, TX 77099. Tel: (713) 495-4700. Fax: (713) 495-5642.
Washington	**Insight,** 12002 115th Ave. N.E., Kirkland, WA 98034. Tel: (206) 820-8100. Fax: (206) 821-2976.
Wisconsin	**Pioneer Standard,** 120 Bishop's Way #163, Brookfield, WI 53005. Tel: (414) 784-3480. Fax: (414) 784-8207.

UK Export

(To countries other than those listed)	**GEC Plessey Semiconductors ,** Unit 1, Crompton Road, Groundwell Industrial Estate, Swindon, Wilts, UK SN2 5AF. Tel: (0793) 518510. Fax: (0793) 518582.
	Whiteaway Laidlaw (Overseas) Ltd., PO Box 93, Ambassador House, Devonshire Street North, Manchester M60 6BU Tel: 061 273 3228 Fax: 061 274 3757

Publication No. HB3923-1 December 1993

Supersedes PS2252 November 1990 Digital Signal Processing IC Handbook

GEC PLESSEY
SEMICONDUCTORS