


Digital Signal Processing

IC Handbook



GEC PLESSEY
SEMICONDUCTORS

DIGITAL SIGNAL PROCESSING

IC Handbook

G E C P L E S S E Y

S E M I C O N D U C T O R S

Foreword

In 1988, the first family of Plessey Digital Signal Processing building block components was launched, offering a significant increase in integration over the traditional Multipliers, ALUs and Address Generators.

Despite the fact that the data in almost all DSP applications is complex-valued (that is, of the form $A + jB$), the hardware of current DSP processors is only capable of operation on real data.

The GEC Plessey Semiconductors PDSP building block family is the only chip set to offer Complex Arithmetic as standard. This provides the system designer with four-fold improvements in speed and substantial reduction in board area and power consumption.

More recently, GEC Plessey Semiconductors has 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.

The application areas addressed by the PDSP family (building block and application specific) include:

- Digital filtering
- Pulse compression
- Digital modulation/demodulation
- Correlation
- Convolution
- Image processing
- Digital waveform synthesis

Contents

	PAGE
Foreword	2
Product index	4
Product list – alphanumeric	5
Technical data	7
Application notes	
A 50ns Butterfly processor	163
A 50ns complex multiplier/accumulator	165
The Pythagoras processor	167
FIR filtering with the PDSP16112 and PDSP16318	169
Interfacing the PDSP family	171
Three dimensional co-ordinate transforms with the PDSP16330	173
A radix 2 Butterfly processor	175
Complex signal processing with the PDSP16000 family	187
A fast FFT processor using the PDSP16000 family	194
FFT address generation using the PDSP1640	208
2-D edge detector board AP16401	217
Sobel v. PDSP operators	225
A high resolution FFTprocessor using the PDSP16116	226
Optimising the accuracy of an FFT system	249
Digital filtering using the PDSP16256	253
Support tools	
PDSP demonstrator	261
PDSP16256/PDSP16350 evaluation system	263
PDSP16488 real-time digital image processor board	265
Data converters for digital signal processing	269
Package outlines	271
Locations	279

Product index – Building block DSP ICs

Arithmetic logic units

Type	Function	Speed	Page
PDSP1601A	ALU and barrel shifter	20MHz	9

Complex arithmetic

Type	Function	Speed	Page
PDSP16112A	Complex multiplier (16 × 12 bits)	20MHz	31
PDSP16116A	Complex multiplier (16 × 16 bits)	20MHz	37
PDSP16318A	Complex accumulator	20MHz	51

Support functions

Type	Function	Speed	Page
PDSP1640	Address generator	20MHz	23
PDSP16520	Quad port synchronous RAM	20MHz	146

Product index – Algorithm specific DSP ICs

Image processing

Type	Function	Speed	Page
PDSP16401A	2-D edge detector	22MHz	95
PDSP16488	Single-chip 2-D convolver	40MHz	100

Digital filtering

Type	Function	Speed	Page
PDSP16256A	Programmable FIR filter	25MHz	57

Co-ordinate conversion

Type	Function	Speed	Page
PDSP16330B	Pythagoras processor	25MHz	73
PDSP16340	Polar to Cartesian converter	20MHz	79

Waveform generation and modulation

Type	Function	Speed	Page
PDSP16350	I/Q splitter and NCO	20MHz	86

Frequency domain processing

Type	Function	Speed	Page
PDSP16510	Stand alone FFT processor	40MHz	127
PDSP16540	32K bucket buffer	40MHz	155

Product list – alpha numeric

TYPE	DESCRIPTION	PAGE
PDSP1601/A	ALU and barrel shifter	9
PDSP1640	Address generator	23
PDSP16112/A	Complex multiplier (16 × 12)	31
PDSP16116/A	Complex multiplier (16 × 16)	37
PDSP16256/A	Programmable FIR filter	57
PDSP16318/A	Complex accumulator	51
PDSP16330/A/B	Pythagoras processor	73
PDSP16340	Polar to Cartesian converter	79
PDSP16350	I/Q splitter and NCO	86
PDSP16401/A	2-D edge detector	95
PDSP16488	2-D convolver	100
PDSP16510	Stand alone FFT processor	127
PDSP16520	Quad port synchronous RAM	146
PDSP16540	32K bucket buffer	155

Technical data



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 be used to select a 16-bit field from a 32-bit input in 100ns.




Fig.1 Pin connections - bottom view

PIN DESCRIPTIONS

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	SVOE	72	B11	C15
10	B3	B9	31	J2	A9	52	K9	RS0	73	C10	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	C1
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	CEB	41	L5	CEA	62	G11	C6	83	A7	IA2
21	F2	CLK	42	K6	MSC	63	G9	C7	84	C7	IA3

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

APPLICATIONS

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

ASSOCIATED PRODUCTS

- PDSP16112 Complex Multiplier
- PDSP1640 20MHz Address Generator
- PDSP16116 16 x 16 Complex Multiplier
- PDSP16318 Complex Accumulator
- PDSP16330 Pythagoras Processor

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 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 high 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 a 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.
V _{cc}	54 & 75	+5V supply. Both V _{cc} 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 this manner, multicyle, 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 one 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 arithmetic 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 the 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-IA0	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-IA0	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 = Value

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 that 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 IS0-3 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, these are:

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

Inst	IS3-IS0	Mnemonic	Function	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	BLSLV	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

- LSXYX Logical Shift, X = Direction YY = Source of Shift Value
- BSXYX Barrel Shift, X = Direction YY = Source of Shift Value
- ASXYX 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, NRRM1, 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 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.

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 may be forced to be input only by tying the 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 (NRRM1 and NRRM2) 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.

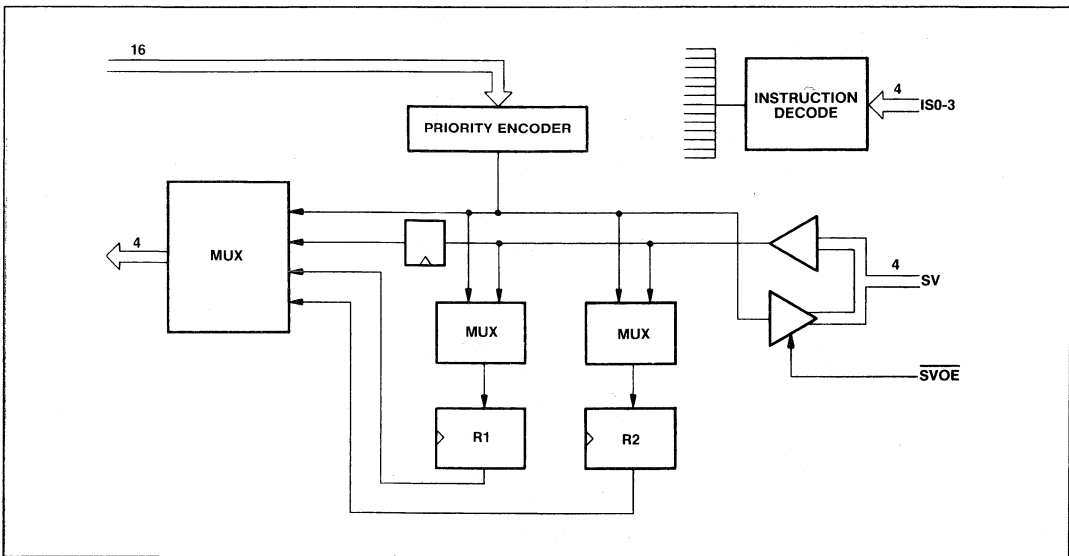


Fig.3 Shift control block

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	RS2-RS0	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 OUTPUT
	MAAPR	0	1	A-PORT INPUT
	MABPR	1	0	B-PORT INPUT
	MARSX	1	1	SHIFTER REGISTER FILE OUTPUT
B-MUX	MSB		Output	
	0	B-PORT INPUT		
	1	SHIFTER REGISTER FILE OUTPUT		
S-MUX	MSS		Output	
	0	B-PORT INPUT		
	1	ALU REGISTER FILE OUTPUT		
C-MUX	MSC		Output	
	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.
MIACI	<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 left 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 right 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.
BMACI	<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 of the 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 discarded, 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 MSBs are discarded, and the vacant LSBs 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 MSBs that exit the 16 bit field to the left, reappear in the vacant LSBs 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 MSBs 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 MSBs 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		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 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 loaded into the R1 register at the end of the cycle, and is also output on the SV port (provided 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 loaded into the R2 register at the end of the cycle, and also output on the SV port (provided 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, 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, the data on the register inputs will be loaded into both Left and Right Registers.
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 V_{cc}	-0.5 to 7.0V
Input Voltage V_{in}	-0.9 to $V_{cc} + 0.9V$
Output Voltage V_{out}	-0.9 to $V_{cc} + 0.9V$
Clamp diode current per pin I_k (See Note 2)	$\pm 18mA$
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

- Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
- 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
LC	12	35
AC	12	36

PDSP1640

20MHz ADDRESS GENERATOR

The PDSP1640 is an 8-bit address generator which will cascade efficiently to wider address fields at very high speed, without the need for external carry logic.

Three PDSP1640s allow 24-bit addresses to be generated at up to 10MHz; four allow 32-bit addressing at up to 7MHz. A single PDSP1640 will address an 8-bit field at 20MHz.

FEATURES

- 20MHz 8-Bit Address Generator
- Fast Cascade Logic gives 10MHz Operation at 24Bits
- Five On-chip User-Programmable Registers
- Output Mask Logic
- 300mW Maximum Power Dissipation
- 28-Pin DIL or LCC Package

APPLICATIONS

- DSP Address Generation
- Database Addressing
- DMA Controllers
- Modulo Counting




Fig.1 Pin connections - top view

ASSOCIATED PRODUCTS

- PDSP16318 Complex Accumulator
- PDSP16112 Complex Number Multiplier
- PDSP1601 ALU and Barrel Shifter




Fig.2 PDSP1640 simplified block diagram




Fig.3 PDSP1640 block diagram

PIN DESCRIPTIONS

Symbol	Pin No.	Pin name and description
CLK	1	Common clock to all registered internal elements. All registers are loaded, and outputs change on the rising edge of CLK.
I0-3	3,2, 27,26	Instruction inputs. The 16 instructions executable by the PDSP1640 are encoded onto these four lines. The instruction for any cycle must be valid at the inputs prior to the rising edge of CLK defining the start of the cycle in which the instruction is to be executed. The I0-3 inputs are internally latched by the rising edge of CLK.
CCEN	4	Conditional Instruction Enable. The Conditional Instructions during the current cycle are enabled if CCEN goes high before the end of the cycle. CCEN may be controlled directly by microcode or, where multiple 1640's are used, this input is used for expansion. See Figs.6 and 7.
COMP	25	Comparator Flag Output. This indicates that the comparator has detected an 'equal to' condition, COMP changes when CLK goes HIGH.
CI	13	Carry In. Carry in to least significant bit of the 8-bit adder.
CO	16	Carry Out. Carry out from the MSB of the adder.
DIO-7	5-12	Data Inputs. 8-bit data input to PDSP1640. The data on this port is loaded into the on-board registers on the rising edge of CLK.
DOO-7	24-17	Data Outputs. The 8-bit output from the counter. The output changes on the rising edge of CLK.
OE	15	3-State Output Control. When high, this signal forces the DOO-7 and COMP outputs into a high-impedance state.
GND	14	0V supply.
Vcc	28	+5V supply.

FUNCTIONAL DESCRIPTION

The PDSP1640 contains six main blocks; the five user programmable registers, an 8-bit Adder, the Mask Logic, a Comparator, the Control Decoder and the Next Address MUX and Counter Register.

The Registers

There are five user programmable registers; MASK, START1, START2, INC and COMP.

MASK Data loaded into the MASK register operates on the data fed to the Mask Logic from the Counter Register. Loading new data into the MASK register automatically enables the Mask Logic. The Mask Logic is disabled either by loading zeros into the MASK register or by executing OP CODE <7> (Clear Counter Register/Mask disable). The Mask Logic will remain disabled until new data is loaded into the MASK register.

N.B. The MASK register can only be loaded from DI0-7.

START1 The START1 register can be loaded from either the DI0-7 inputs, or from the Counter Register. The contents of the START1 register may be forced into the Counter Register (OP CODE <6>), or may be used as a jump address in a conditional instruction.

START2 The START2 register can be loaded from either the DI0-7 inputs, or from the Counter Register. The contents of the START2 Register may be used as the jump address in a conditional instruction.

INC The INC register contains the value by which the counter will increment. This may be a positive or negative number, represented in 2's complement.

The INC register may be loaded either from the DI0-7 inputs, or from the Counter Register.

COMP The COMP register contains the value used by the comparator. It may be programmed from either the DI0-7 inputs or from the Counter Register.

8-Bit Adder

The 8-BIT ADDER adds the contents of the Inc and Counter Registers and loads the result into the Counter Register conditional on the current instruction.

The ADDER has a fast carry system which eliminates the

need for external carry look-ahead circuitry when cascaded. Cascading is achieved by chaining CO to CI of the next most significant stage (see Fig.6).

Mask Logic

The MASK LOGIC is controlled by the contents of the Mask Register. 1's in the Mask Register will cause the corresponding outputs from the PDSP1640 to be frozen, even if the data in the Counter Register changes. In this manner 'windows' can be created within the counter's address field.

The MASK LOGIC is enabled whenever new data is loaded into the Mask Register. A zero word in the Mask Register will disable the MASK LOGIC as will executing OP CODE <7>.

Comparator

The COMPARATOR compares the value in the Comp Register with the output from the NEXT ADDRESS MUX. If the values are the same, then a signal is sent to the Control Decoder and onto the COMP output via the output register.

Control Decoder

The CONTROL DECODER has six inputs; the four instruction lines I0-3, CCEN, and an internal Comp Flag.

The CONTROL DECODER latches all except the CCEN input and the internal Comp Flag, on the rising edge of CLK. The I0-3 inputs are decoded to implement the operations shown in Table 2. CCEN and the Comp Flag change the instruction executed in the current cycle, where appropriate.

Next Address Mux and Counter Register

The contents of the COUNTER REGISTER at the start of a new cycle are determined by the NEXT ADDRESS MUX under the control of the Control Decoder.

The NEXT ADDRESS MUX selects between the DI0-7 inputs, the contents of START1 and 2, the contents of the Inc Register, and the output of the 8-bit Adder. The COUNTER REGISTER may be cleared by OP CODE <7>. The COUNTER REGISTER clock is inhibited for all instructions except those marked * in Table 2, this is to prevent the counter incrementing during register loads.




Fig.4 Register delays

**PDSP1640
INSTRUCTION SET**

Mnemonic	Op Code	Function
CCJDI	<0>	On the next cycle the contents of the Counter Register will be added to the contents of the Increment Register and the result loaded into the Counter Register. If the COMP flag and CCEN become active, then on the following cycle, the Counter Register will be loaded with the data on DI0-7.
CCJS1	<1>	On the next cycle the contents of the Counter Register will be added to the contents of the Increment Register and the result loaded into the Counter Register. If the COMP flag and CCEN become active, then on the following cycle, the Counter Register will be loaded with the contents of the Start1 Register.
CCJS2	<2>	On the next cycle the contents of the Counter Register will be added to the contents of the Increment Register and the result loaded into the Counter Register. If the COMP flag and CCEN become active, then on the following cycle, the Counter Register will be loaded with the contents of the Start2 Register.
LMRDI	<3>	Data present on DI0-7 will be loaded into the Mask Register by the rising edge of CLK at the end of the cycle in which this instruction is executed.
LCRDI	<4>	The data present on DI0-7 will be loaded into the Counter Register by the rising edge of CLK at the end of the cycle in which this instruction is executed.
LCRIR	<5>	The Counter Register will be loaded with the contents of the Inc Register by the rising edge of CLK at the end of the cycle in which this instruction is executed.
L CRS1	<6>	The Counter Register will be loaded with the contents of the Start1 Register by the rising edge of CLK at the end of the cycle in which this instruction is executed.
CLRCR	<7>	The Counter Register will be cleared and the Mask Logic disabled by the rising edge of CLK at the end of this instruction cycle.
LS1DI	<8>	The data present on DI0-7 will be loaded into the Start1 Register by the rising edge of CLK at the end of this instruction cycle.
LS1CR	<9>	The Start1 Register will be loaded with the contents of the Counter Register by the rising edge of CLK at the end of this instruction cycle.
LS2DI	<A>	The data present on DI0-7 will be loaded into the Start2 Register by the rising edge of CLK at the end of this instruction cycle.
LS2CR		The Start2 Register will be loaded with the contents of the Counter Register by the rising edge of CLK at the end of this instruction cycle.
LIRDI	<C>	The data present on the DI0-7 input will be loaded into the Inc Register by the rising edge of CLK at the end of this instruction cycle.
LIRCR	<D>	The Inc Register will be loaded with the contents of the Counter Register by the rising edge of CLK at the end of this instruction cycle.
LCPDI	<E>	The data present on the DI0-7 inputs will be loaded into the Comp Register by the rising edge of CLK at the end of this instruction cycle.
LCPCR	<F>	The Comp Register will be loaded with the contents of the Counter Register by the rising edge of CLK at the end of this instruction cycle.

Table 1 Instruction descriptions




Fig.5 Comparison flag timing

INSTRUCTION SET

Mnemonic	Code	I ₃	I ₂	I ₁	I ₀	Operation	Jump To
* CCJDI	0	0	0	0	0	Count by IR	DI0-7
* CCJS1	1	0	0	0	1	Count by IR	START1
* CCJS2	2	0	0	1	0	Count by IR	START2
LMRDI	3	0	0	1	1	Ld MR from DI0-7	
* LCRDI	4	0	1	0	0	Ld CR from DI0-7	
* LCRIR	5	0	1	0	1	Ld CR from IR	
* LCRS1	6	0	1	1	0	Ld CR from S1	
CLRCR	7	0	1	1	1	Clear CR/MR	
LS1DI	8	1	0	0	0	Ld S1 from DI0-7	
LS1CR	9	1	0	0	1	Ld S1 from CR	
LS2DI	A	1	0	1	0	Ld S2 from DI0-7	
LS2CR	B	1	0	1	1	Ld S2 from CR	
LIRDI	C	1	1	0	0	Ld IR from DI0-7	
LIRCR	D	1	1	0	1	Ld IR from CR	
LCPDI	E	1	1	1	0	Ld CP from DI0-7	
LCPCR	F	1	1	1	1	Ld CP from CR	

All instructions executed on the next rising edge of CLK. * indicates instructions which do not inhibit the counter register clock.

Table 2 Instruction set codes

Key

- IR = Increment Register
- MR = Mask Register
- CR = Counter Register
- S1 = Start1 Register
- S2 = Start2 Register
- CP = Comparator Register

Mnemonics

- CCJXX = Conditional Count, Jump to XX
- LXXYY = Load Destination XX from Source YY
- CLRCR = Clear Counter Register/Reset Mask Logic

COUNTER CONFIGURATIONS

Fig.6 illustrates chaining of PDSP1640s to 16 bits and Fig.7 the configuration for a 24-bit address generator. The cascaded devices have exactly the same functions as a single device.




Fig.6 Chaining to 16 bits

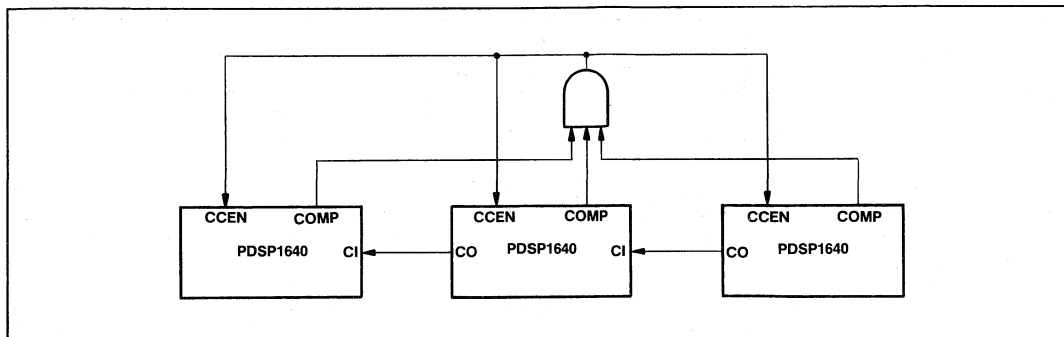


Fig.7 Chaining to 24 bits

TYPICAL APPLICATION

In the application shown in Fig.8 two PDSP1640s are used as an address generator in a digital waveform generator capable of producing complex waveforms at very high speeds. The programmable registers in the PDSP1640 allow

the host microprocessor to control both frequency (by altering Step Size) and waveshape (by selecting different wavetables by Start Address).




Fig.8 Arbitrary waveform generation

TYPICAL APPLICATION

Mnemonic	Op Code	Operation	Data
CLRCR	<7>	Clear CR/MR	X
LCRDI	<4>	Load CR	X
LIRDI	<C>	Load Inc Register	Start addr
LS1DI	<8>	Ld SR1 with branch addr	Step size
LCPDI	<E>	Ld COMPR with stop addr	Branch addr
CCJS1	<1>	Count by INC/goto SR1	Stop addr

Table 3 Typical instruction sequence for Fig.8

ELECTRICAL CHARACTERISTICS

Test conditions (unless otherwise stated):

T_{amb} (Industrial) = -40°C to +85°C, V_{CC} = 5.0V ± 10%, GND = 0V
 T_{amb} (Military) = -55°C to +125°C, V_{CC} = 5.0V ± 10%, GND = 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}			0.6	V	I _{OL} = -8mA
Input high voltage	V _{IH}	2.2			V	
Input low voltage	V _{IL}			0.8	V	
Input leakage current	I _{IL}	-10		10	μA	GND ≤ V _{in} ≤ V _{CC}
Output leakage current	I _{OZ}	-50		50	μA	GND ≤ V _{out} ≤ V _{CC} = V _{CC} max.
Output short cct current (Note 2)	I _{OS}	40		250	mA	V _{CC} = max.
Input capacitance	C _{IN}		9		pF	LC package
			12		pF	DG package

Switching Characteristics

Characteristic	Value				Units	Conditions
	Industrial		Military			
	PDSP1640 B0		PDSP1640 A0			
	Min.	Max.	Min.	Max.		
CLK frequency		20		20	MHz	
CLK high period	20		20		ns	
CLK low period	15		15		ns	
CLK to CO		44		44	ns	1 LSTTL + 5pF load
CLK to DO		34		34	ns	2 LSTTL + 20pF load Opcode 3
CLK to DO		28		28	ns	2 LSTTL + 20pF load, remaining Opcodes
CLK to COMP		35		35	ns	50pF load (Opcodes 0, 1, 2)
CI to CO		20		20	ns	1 LSTTL + 5pF load
Setup DI to CLK	10		10		ns	
Hold DI to CLK	3		3		ns	
Setup CI to CLK	20		20		ns	
Hold CI to CLK	3		3		ns	
Setup I to CLK	15		15		ns	
Hold I to CLK	3		3		ns	
Setup CCEN to CLK	30		30		ns	
Hold CCEN to CLK	0		0		ns	
\overline{OE} high to DO high Z		30		30	ns	See \overline{OE} test diagrams, Fig. 9
\overline{OE} low to DO/COMP valid		22		22	ns	
V_{CC} current		20		20	mA	$V_{CC} = \text{Max.}$, outputs unloaded, CLK freq = Max.

ABSOLUTE MAXIMUM RATINGS (Note 1)

Supply voltage V_{CC}	-0.5 to 7.0V
Input voltage V_{IN}	-0.9 to V_{CC} +0.9V
Output voltage V_{OUT}	-0.9 to V_{CC} +0.9V
Clamp diode current per pin I_K (see Note 2)	± 18 mA
Static discharge voltage (HMB)	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

NOTES

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

THERMAL CHARACTERISTICS

Package Type	θ_{JC} °C/W	θ_{JA} °C/W
DG	12	40
LC	13	56




Fig.9 Three state delay measurement load

ORDERING INFORMATION

Industrial (-40°C to +85°C)

PDSP1640 B0 DG (Ceramic DIL package)

PDSP1640 B0 LC (Leadless chip carrier)

Military (-55°C to +125°C)

PDSP1640 A0 DG (Ceramic DIL package)

PDSP1640 A0 LC (Leadless chip carrier)

PDSP16112/PDSP16112A

16 x 12 BIT COMPLEX MULTIPLIER

The PDSP16112/PDSP16112A will multiply a complex (16 + 16) bit data 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
- Pipelined Architecture
- Power Dissipation only 500mW
- TTL Compatible Inputs

APPLICATIONS

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

ASSOCIATED PRODUCTS

PDSP1601 ALU and Barrel Shifter
PDSP16318 Complex Accumulator
PDSP16330 Pythagoras Processor

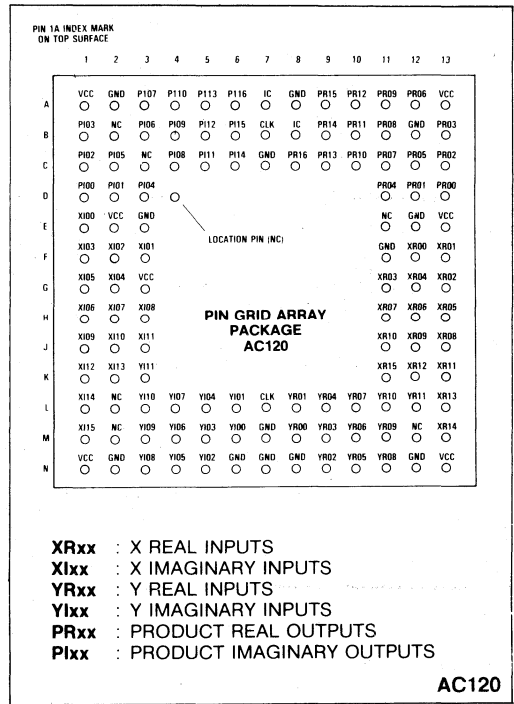


Fig.1 Pin connections - top view




Fig.2 Multiplier block diagram

PIN OUT - FUNCTION TO PIN

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

	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 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 Outputs : 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	PI mag 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	Common Clock to all on chip registers, both 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. $-1 \leq YI < 1$	VCC GND	All VCC and GND pins must be connected
		IC	Internally connected - do not use

FUNCTIONAL DESCRIPTION

The PDSP16112 Complex Multiplier contains four pipelined 16×12 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 eighth clock cycle.

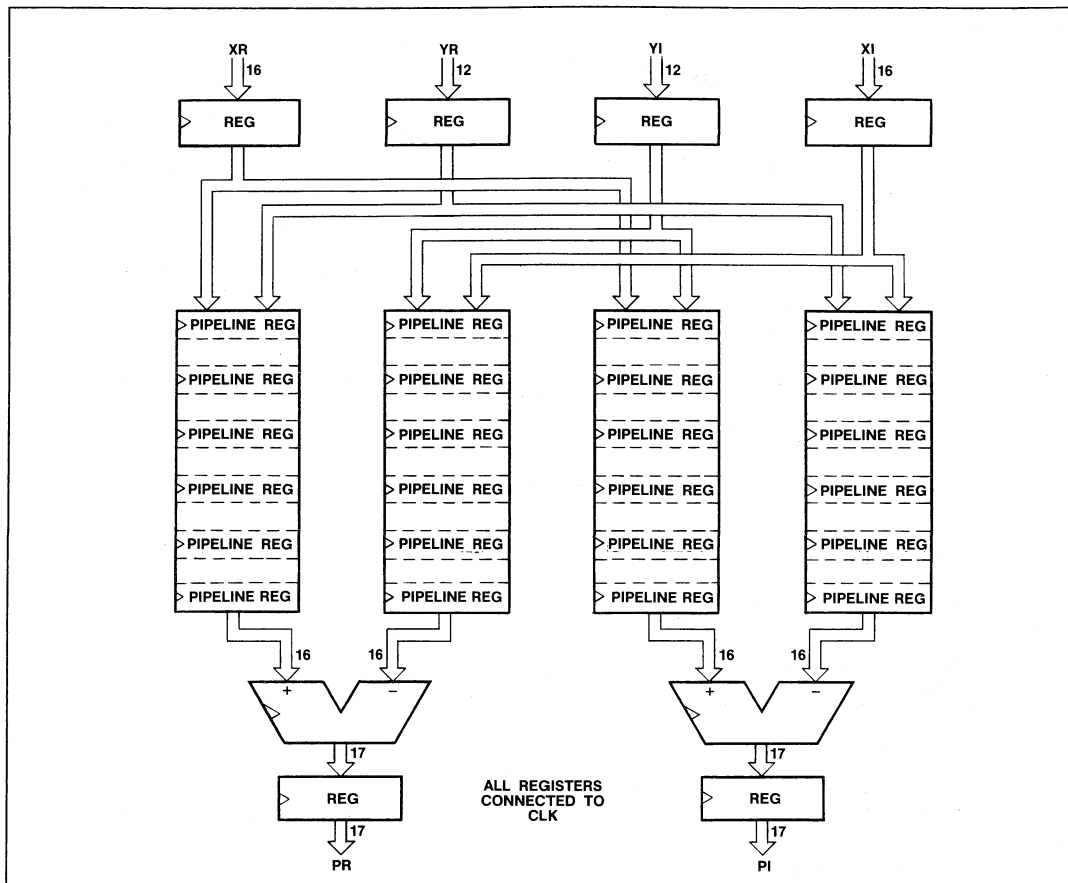


Fig.3 Pipelined multiplier structure


TYPICAL APPLICATION

The PDSP16112A may be configured as the main arithmetic element in a 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×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×W) and A-(B×W). One of the PDSP16318As calculates the real parts and the other calculates the imaginary parts.

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

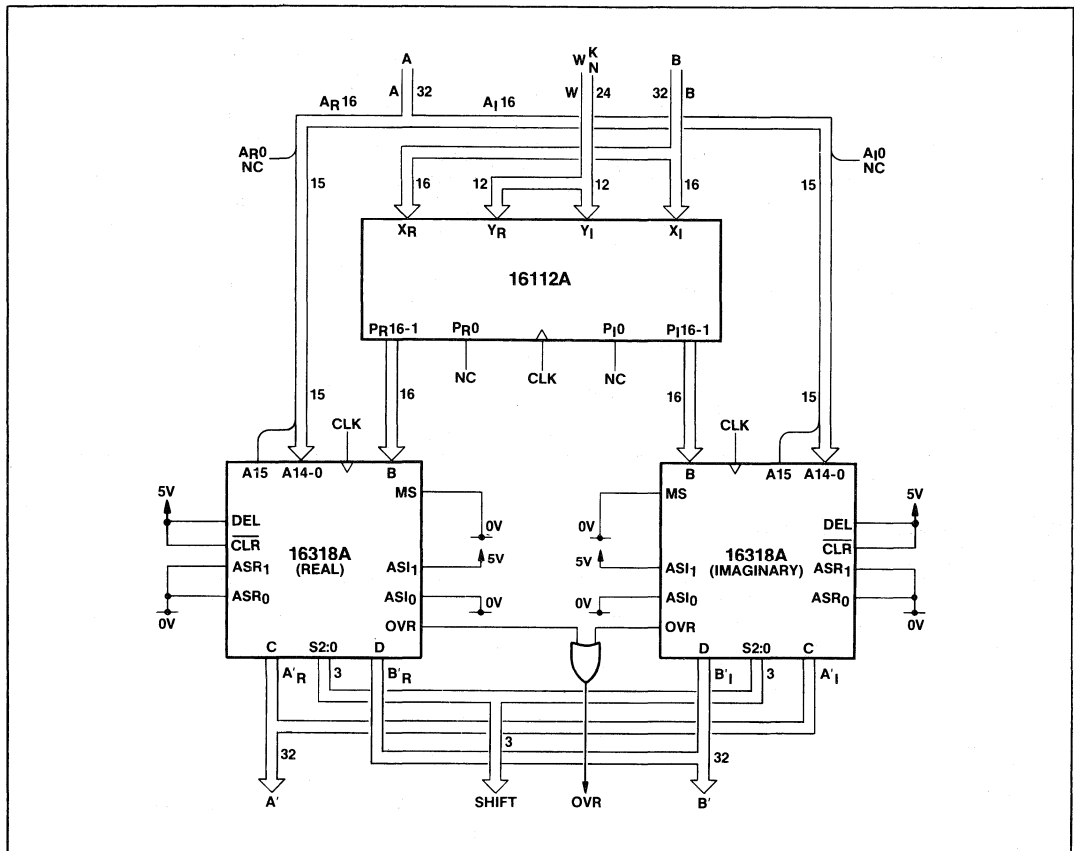


Fig.4 Radix 2 DIT butterfly processor

ELECTRICAL CHARACTERISTICS

Test conditions (unless otherwise stated):

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

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

Static Characteristics

Characteristic	Symbol	Value						Units	Conditions
		PDSP16112			PDSP16112A				
		Min.	Typ.	Max.	Min.	Typ.	Max.		
Output high voltage	V _{OH}	2.4		0.6	2.4		0.6	V	I _{OH} = 4mA
Output low voltage	V _{OL}							V	I _{OL} = 4mA
Input high voltage	V _{IH}	2.2			2.2			V	
Input low voltage	V _{IL}			0.8			0.8	V	
Input leakage current *	I _L	-1.2		+0.01	-1.2		+0.01	mA	GND ≤ V _{IN} ≤ 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-up resistor to V_{cc}.

PDSP16112/A

AC Characteristics

Characteristic	Symbol	Value Industrial						Value Military		Units	Conditions
		PDSP16112			PDSP16112A			Min.	Max.		
		Min.	Typ.	Max.	Min.	Typ.	Max.				
V _{CC} current	I _{CC}			90			170		90	mA	V _{CC} = 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 V _{CC} +0.5V
Output voltage V _{OUT}	-0.5V to V _{CC} +0.5V
Clamp diode current per pin 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}	
Industrial	-40°C to +85°C
Military	-55°C to +125°C
Package power dissipation P _{TOT}	1000mW

NOTES

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

ORDERING INFORMATION

Industrial (-40°C to +85°C)

PDSP16112 B0 AC (10MHz - LCC package)

PDSP16112A B0 AC (20MHz - LCC package)

Military (-55°C to +125°C)

PDSP16112 A0 AC (10MHz - PGA package)

THERMAL CHARACTERISTICS

Package Type	θ _{JC} °C/W	θ _{JA} °C/W
AC	12	36

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

PDSP16116 / A

16 BY 16 BIT COMPLEX MULTIPLIER

(SUPERSEDES JANUARY 1990 EDITION)

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 16x16 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 PDSP16318A, the PDSP16116A forms a two chip 20MHz Complex Multiplier-Accumulator with 20 bit accumulator registers and output shifters. The PDSP16116A in combination with two PDSP16318As and two PDSP1601As forms a complete 20MHz 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).

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

APPLICATIONS

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




Fig. 1 Simplified Block Diagram

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/A 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 streams. This operation has tradi-

tionally required an additional ALU to multiply the imaginary component by -1. The PDSP16116/A 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/A 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 real 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 imag p data	
CLK	INPUT	Clock, new data is loaded on rising edge of CLK	
CEX	INPUT	Clock enable X-port input register	
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
SOBFP	INPUT	Start of BFP operations **	Tie Low
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 *	Tie Low
WTOU1:0	OUTPUT	Word tag output**	
SFTA1:0	OUTPUT	Shift control for A-word / overflow flag *	
SFTR2:0	OUTPUT	Shift control for accumulator result **	
GWR4:0	OUTPUT	Global weighting register contents **	
OSEL1:0	INPUT	Selects the desired output configuration	
OER, 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

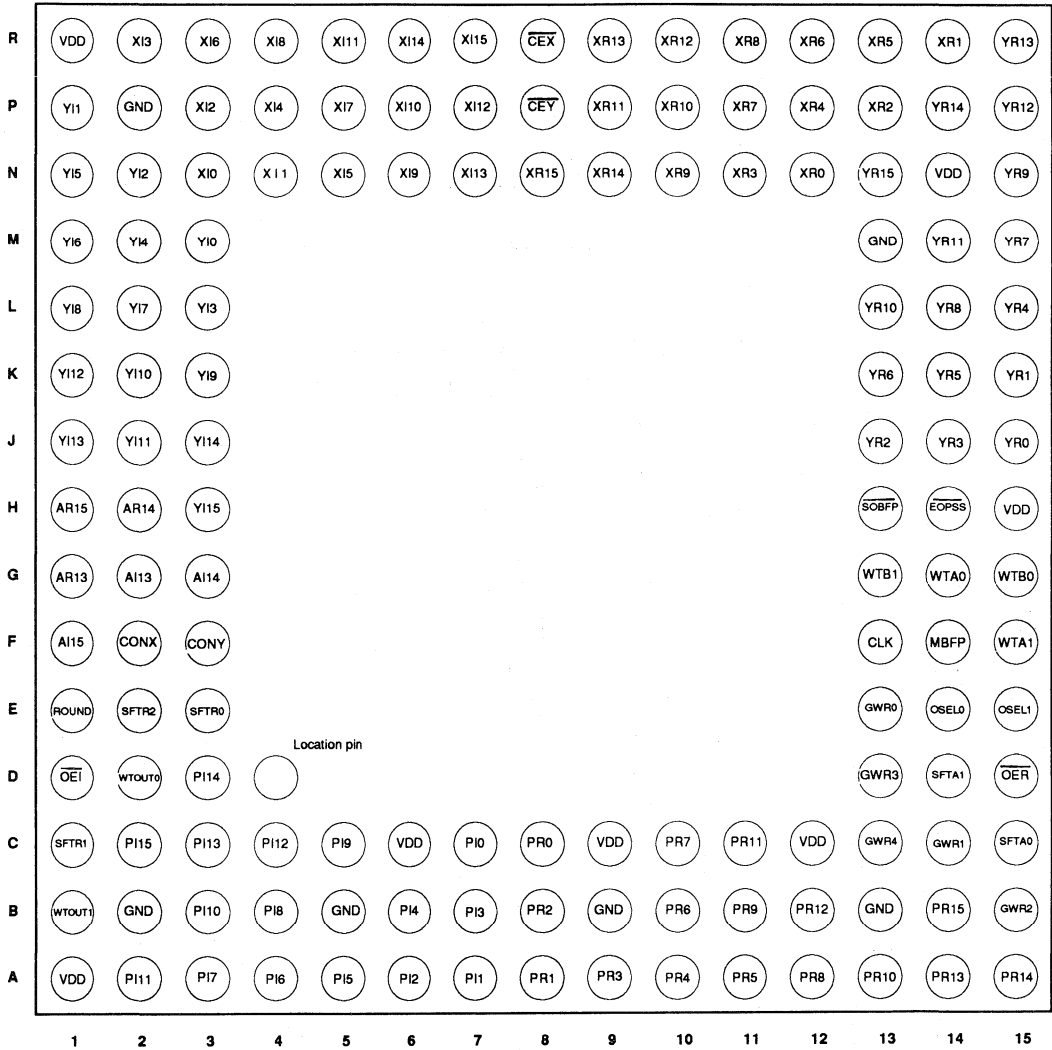


Fig. 3 Pin Allocation Diagram (Bottom View)

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 Stage

Complex twos 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

$$\text{Value} = (-1 \times S) + (\text{bit}14 \times 2^{-1}) + (\text{bit}13 \times 2^{-2}) + (\text{bit}12 \times 2^{-3}) \dots$$

The X & Y port registers are individually enabled by the $\overline{\text{CE}}_X$ & $\overline{\text{CE}}_Y$ signals respectively. 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 ($X_r \times Y_r$), ($X_r \times Y_i$), ($X_i \times Y_r$), ($X_i \times Y_i$), all of which are calculated in parallel by the four 16x16 Multipliers. Only one clock cycle is required to complete the multiply stage before the Multiplier results are loaded into the Multiplier output registers for passing on to the Adder/Subtractors in the next cycle. Each multiplier produces a 31bit 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 fractional 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/A includes a trap to ensure that the most positive number (value = $1 - 2^{-30}$), (hex = 7FFFFFFF) is substituted 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	$(X_R+X_I)x(Y_R+Y_I)$	low	low
X x Conj Y	$(X_R+X_I)x(Y_R-Y_I)$	low	high
Conj X x Y	$(X_R-X_I)x(Y_R+Y_I)$	high	low
Invalid	Invalid	high	high

Table 3 Conjugate Functions

Adder / Subtractor Stage

The 31bit Real and Imaginary results from the Multipliers are passed to two 32 bit Adder / Subtractors. The Adder calculates the imaginary result $((X_r \times Y_i) + (X_i \times Y_r))$ and the Subtractor calculates the Real result $((X_r \times Y_r) - (X_i \times Y_i))$. 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/Subtractor result. (The least significant 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 process.

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 bit 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 sign 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	25	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 / Subtractors 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 result 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 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 4 Output Selection

(Where MSR and LSR are the most and least significant 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/A. 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 DESCRIPTION

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 -2^0 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 registers.

CEX, 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 significant 16 bits of the Adder/Subtractor result prior to being passed to the output registers. The rounding operation takes place one cycle after the ROUND input is taken high providing rounded outputs on the second cycle after ROUND 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 applications. 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.

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.

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 determine the magnitude of the real part of the A- word and, hence, to determine if there will be any chance of word growth in the PDSP16318/A 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 the S2:0 pins on the PDSP16318/A 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 ALUs. When operating in normal mode, these outputs 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 Accumulator 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 occurring during that pass. (i.e. If WTOUT = 11, the the 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 itself is not processed in the PDSP16116/A, 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 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-word 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/A ALU. (The two words must have the same weighting if they are to be added).

In normal mode, SFTA0 performs 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.

SFTA1:0	FUNCTION
0 0	Shift A-word 1 place to the right
0 1	Shift A-word 2 place to the right
1 0	Shift A-word 3 place to the right
1 1	Shift A-word 4 place 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). There 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 PDSP16116A may be used as the main arithmetic unit of a butterfly processor which will allow the following FFT benchmarks :

- 1024 point complex radix-2 transform in 259µs
- 512 point complex radix-2 transform in 118µs
- 256 point complex radix-2 transform in 53µs

In addition, with pin MBFP tied high, the BFP circuitry within the PDSP16116/A 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/A 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 $\log_2(N)$ 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/A may be combined with two PDSP1601/As and two PDSP16318/As to form a complete BFP butterfly processor. The PDSP16318/As are used to perform the complex addition and subtraction of the butterfly operation, while the PDSP16116/As are used to match the data path of the A word to the pipelining and shifting operations within the PDSP16116/A.

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/A is intended for use in the FFT application described above. i.e. it is intended to prevent data degradation during the course of an FFT calculation. The operation of the PDSP16116/A 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 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/A, 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.XXXXXXXXXXXXXX word tag = 01
- XXXX.XXXXXXXXXXXXXX word tag = 10
- XXXXX.XXXXXXXXXXXXXX word tag = 11

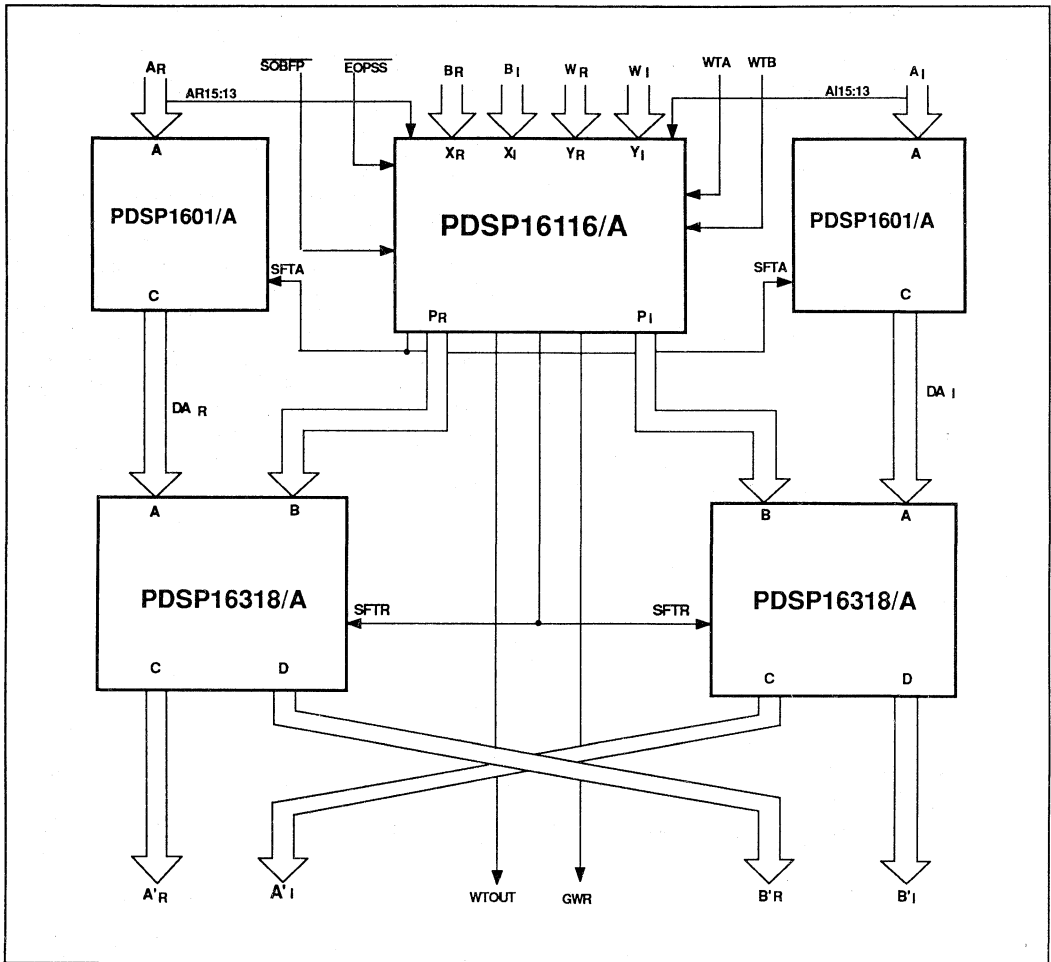


Fig. 4 FFT Butterfly Processor

At the end of each constituent pass of the FFT, the positions of the binary points 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.XXXXXXXXXXXXXX word tag = 00
XXXXX.XXXXXXXXXXXXXX word tag = 01
XXXXXX.XXXXXXXXXXXXXX word tag = 10
XXXXXXX.XXXXXXXXXXXXXX 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 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/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.

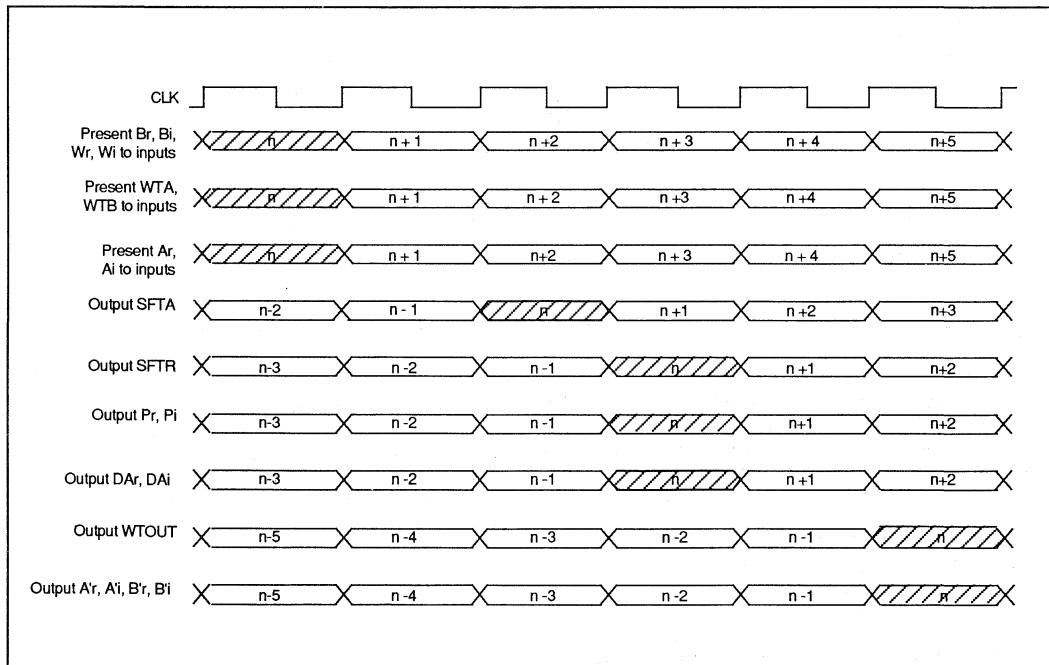


Fig. 6 Butterfly Data and Control Signals

Control of the FFT

To enable the block floating point hardware to keep track of the data, the following signals are provided :

- SOBFP - start of the FFT
- 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 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. 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.

* Should a longer pause be required between passes - to arrange the data for the next pass, for example, then EOPSS may be kept low as long as necessary - the next pass cannot commence until it is brought high again.

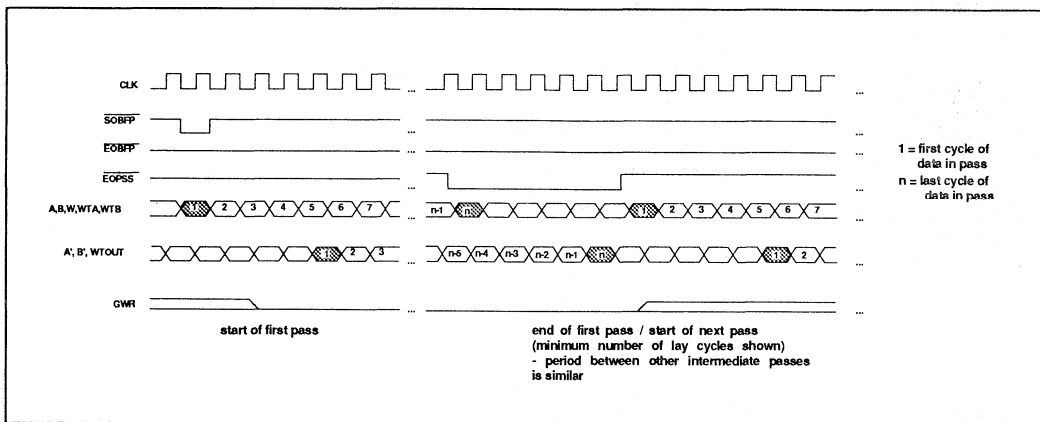


Fig. 7 Use of the BFP Control Signals

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 a pre-determined common weighting. This can be done by comparing the exponent of a given data word with the pre-determined universal 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 Universal 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 $\log_2(N)$ places to the right. Hence, if we choose the Universal Exponent to be $\log_2(N)$ 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 Universal 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 Universal Exponent to determine the shift required. This is done by subtracting it from the Universal 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 purposes of determining the shift required, instead of modifying it according to Table 6. To compensate for this, the Universal Exponent may be increased by one.

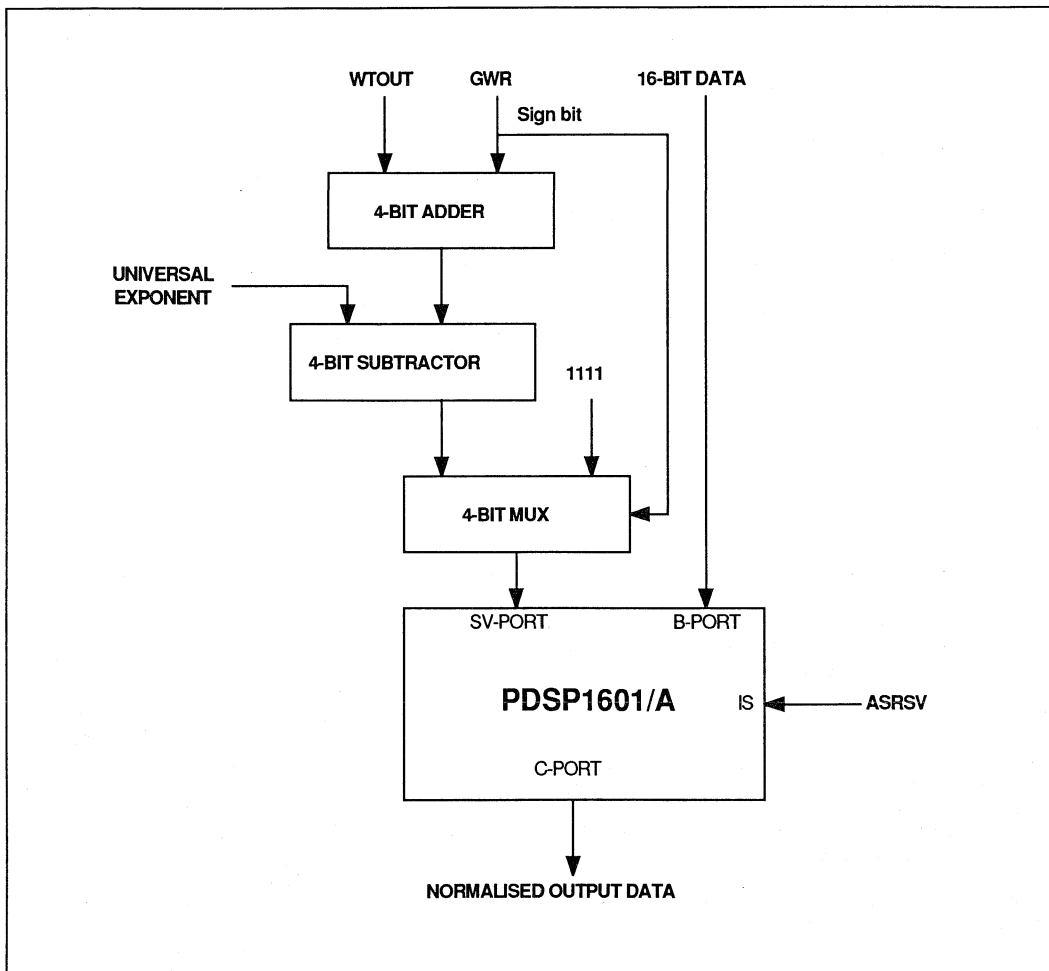


Fig. 8 Output Normalisation Circuitry

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 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}	
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 ambient 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 exceeded, 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: $T_{AMB} = -40^\circ\text{C}$ to $+85^\circ\text{C}$ $V_{CC} = 5.0V \pm 10\%$ Ground = 0V

Military: $T_{AMB} = -55^\circ\text{C}$ to $+125^\circ\text{C}$ $V_{CC} = 5.0V \pm 10\%$ Ground=0V

Static Characteristics

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	$I_{OH} = 8\text{mA}$
Output low voltage	V_{OL}	-		0.4	V	$I_{OL} = -8\text{mA}$
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	$GND < V_{OUT} < V_{CC}$
Output leakage current	I_{OZ}	-50		+50	μA	$V_{CC} = \text{Max}$
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 CEX or CEY to CLK rising edge	11	-	8	-	ns	
Hold CEX or 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 clock rising edge	-	2	-	0	ns	
Setup WTA1:0, WTB1:0, SOBFP or EOPSS inputs to CLK rising edge	14	-	8	-	ns	
Hold WTA1:0, WTB1:0, SOBFP or 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	-	8	-	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
OER or OEI rising PR-PORT or PI-PORT high to Z	-	35	-	25	ns	see Fig. 9
OER or OEI rising PR-PORT or PI-PORT low to Z	-	45	-	25	ns	see Fig. 9
OER or OEI falling PR-PORT or PI-PORT Z to high	-	22	-	18	ns	see Fig. 9
OER or 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 B0 AC	10MHz	Industrial
PDSP16116 A0 AC	10MHz	Military
PDSP16116A B0 AC	20MHz	Industrial
PDSP16116A A0 AC	20MHz	Military

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

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 throughput 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 FFT's to be executed in 256 μ s.

The PDSP16318/A is recommended for new designs instead of PDSP16316/A.

FEATURES

- Full 20MHz Throughput 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 Package

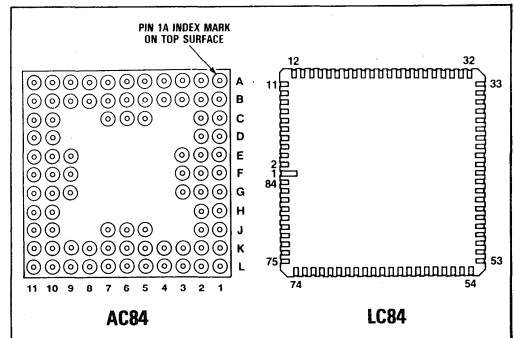


Fig.1 Pin connections - bottom view

APPLICATIONS

- High Speed Complex FFT or DFTs
- Complex Finite Impulse Response (FIR) Filtering
- Complex Conjugation
- Complex Correlation/Convolution

ASSOCIATED PRODUCTS

- PDSP16112 16 \times 12 Complex Multiplier
- PDSP16116 16 \times 16 Complex Multiplier
- PDSP1601 ALU and Barrel Shifter
- PDSP1640 20MHz Address Generator
- 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 Butterflies) 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 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.

PIN DESCRIPTIONS

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. D15 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.
AS11:0	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.

LC Pin	AC Pin	Function	LC Pin	AC Pin	Function	LC Pin	AC Pin	Function	LC Pin	AC Pin	Function
75	B2	D7	12	K2	C7	33	K10	A1	54	B10	B10
76	C2	D8	13	K3	C6	34	J10	A2	55	B9	B9
77	B1	D9	14	L2	C5	35	K11	A3	56	A10	B8
78	C1	D10	15	L3	C4	36	J11	A4	57	A9	B7
79	D2	GND	16	K4	C3	37	H10	A5	58	B8	B6
80	D1	VCC	17	L4	C2	38	H11	A6	59	A8	B5
81	E3	D11	18	J5	C1	39	F10	A7	60	B6	B4
82	E2	D12	19	K5	$\overline{\text{C0}}$	40	G10	A8	61	B7	B3
83	E1	D13	20	L5	$\overline{\text{OED}}$	41	G11	A9	62	A7	B2
84	F2	D14	21	K6	$\overline{\text{OEC}}$	42	G9	A10	63	C7	B1
1	F3	D15	22	J6	S2	43	F9	A11	64	C6	B0
2	G3	C15	23	J7	S1	44	F11	A12	65	A6	CLK
3	G1	C14	24	L7	S0	45	E11	A13	66	A5	CEB
4	G2	C13	25	K7	MS	46	E10	A14	67	B5	OVR
5	F1	C12	26	L6	AS11	47	E9	A15	68	C5	D0
6	H1	VCC	27	L8	AS10	48	D11	CEA	69	A4	D1
7	H2	GND	28	K8	$\overline{\text{DEL}}$	49	D10	B15	70	B4	D2
8	J1	C11	29	L9	$\overline{\text{CLR}}$	50	C11	B14	71	A3	D3
9	K1	C10	30	L10	ASR1	51	B11	B13	72	A2	D4
10	J2	C9	31	K9	ASR0	52	C10	B12	73	B3	D5
11	L1	C8	32	L11	A0	53	A11	B11	74	A1	D6

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 accumulator

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	0
1	1	0							15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
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)

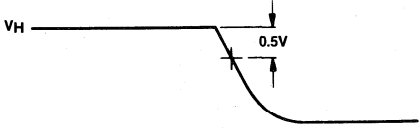
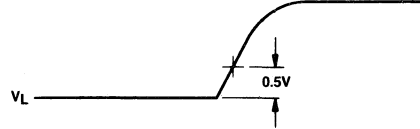
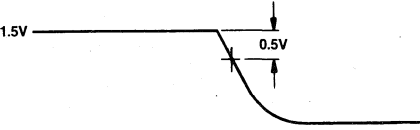
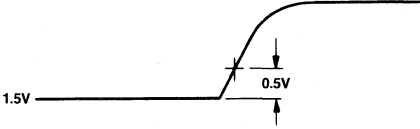
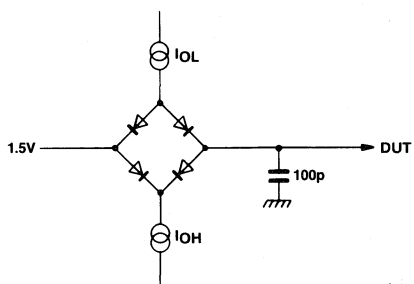
Supply voltage V_{CC}	-0.5V to 7.0V
Input voltage V_{IN}	-0.9V to $V_{CC} + 0.9V$
Output voltage V_{OUT}	-0.9V to $V_{CC} + 0.9V$
Clamp diode current per pin I_K (see Note 2)	18mA
Static discharge voltage (HMB) V_{STAT}	500V
Storage temperature range T_S	-65°C to +150°C
Ambient temperature with power 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

THERMAL CHARACTERISTICS

Package Type	θ_{JC} °C/W	θ_{JA} °C/W
LC	12	35
AC	12	36

NOTES

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

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	
<p>NOTES</p> <ol style="list-style-type: none"> 1. V_H - Voltage reached when output driven high. 2. V_L - Voltage reached when output driven low. <div style="text-align: center;">  </div>	

PDSP16318 /16318A

ELECTRICAL CHARACTERISTICS

Test conditions (unless otherwise stated):

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

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

Static Characteristics

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
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}	2.0		-	V	
Input low voltage	V _{IL}	-		0.8	V	
Input leakage current	I _{IL}	-10		+10	μA	GND ≤ V _{IN} ≤ V _{CC}
Output leakage current	I _{OZ}	-50	-	+50	μA	GND ≤ V _{OUT} ≤ V _{CC}
Output S/C current	I _{OS}	20	-	200	mA	V _{CC} = Max
Input capacitance	C _{IN}	-	9	-	pF	

Switching Characteristics

Characteristic	Value Industrial				Value Military		Units	Conditions
	PDSP16318		PDSP16318A		PDSP16318			
	Min.	Max.	Min.	Max.	Min.	Max.		
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, CLR setup to clock rising edge	8	-	5	-	8	-	ns	
DEL, ASR, CLR, MS, S2:0, ASI hold after clock rising edge	2	-	2	-	2	-	ns	
CEA, CEB setup to clock falling edge	2	-	2	-	2	-	ns	
CEA, 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
OEC/OED low to C15:0/D15:0 high data valid	-	40	-	30	-	40	ns	2 x LSTTL + 20pF
OEC/OED low to C15:0/D15:0 low data valid	-	40	-	30	-	40	ns	2 x LSTTL + 20pF
OEC/OED high to C15:0/D15:0 high impedance	-	40	-	30	-	40	ns	2 x LSTTL + 20pF
V _{CC} current	-	70	-	110	-	70	mA	V _{CC} = max, TTL input levels Outputs unloaded, f _{CLK} = max
V _{CC} current	-	30	-	60	-	30	mA	V _{CC} = 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 positive into the device.
3. CMOS input levels are defined as
V_{IL} = 0.5
V_{IH} = V_{DD} - 0.5

ORDERING INFORMATION

Industrial (-40°C to +85°C)

PDSP16318 B0 AC PDSP16318 B0 LC
 PDSP16318A B0 AC PDSP16318A B0 LC

Military (-55°C to +125°C)

PDSP16318 A0 AC PDSP16318 A0 LC

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



PDSP16256 / A

PROGRAMMABLE FIR FILTER

(SUPERSEDES MARCH1990 EDITION)

The PDSP16256/A 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 planes

APPLICATIONS

- High Performance Digital Filters
- Pulse Compression for Radar & Sonar
- Matrix Multiplication
- Correlation

ASSOCIATED PRODUCTS

- PDSP16350 I/Q Splitter / NCO
- PDSP16510 FFT Processor

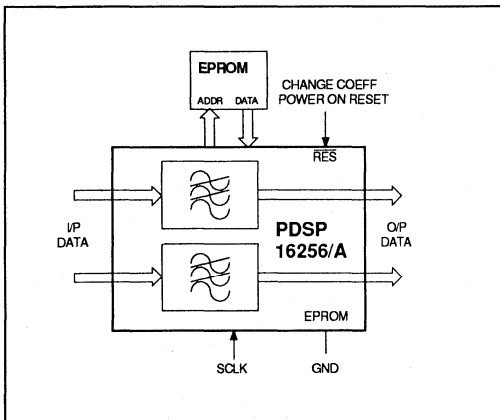


Fig. 1 Dual Filter

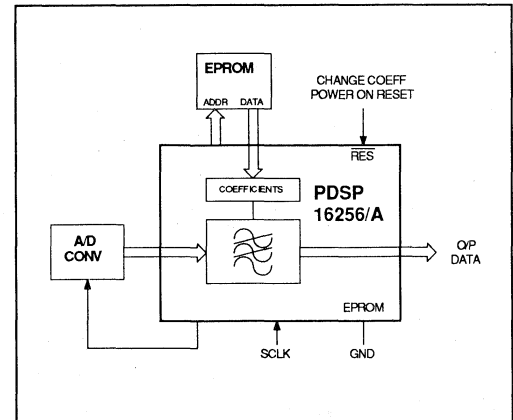


Fig. 2 Typical System Application

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. $\overline{\text{OEN}}$ 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 buses (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.




Fig. 3 Device Pinout - Bottom view




Fig. 4 Block Diagram

OPERATIONAL OVERVIEW

The PDSP16256/A is an application specific FIR filter for use in high performance digital signal processing systems. Sampling rates can be up to 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 fourth 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/A device. An EPROM file can be automatically generated in Motorola S-record format.

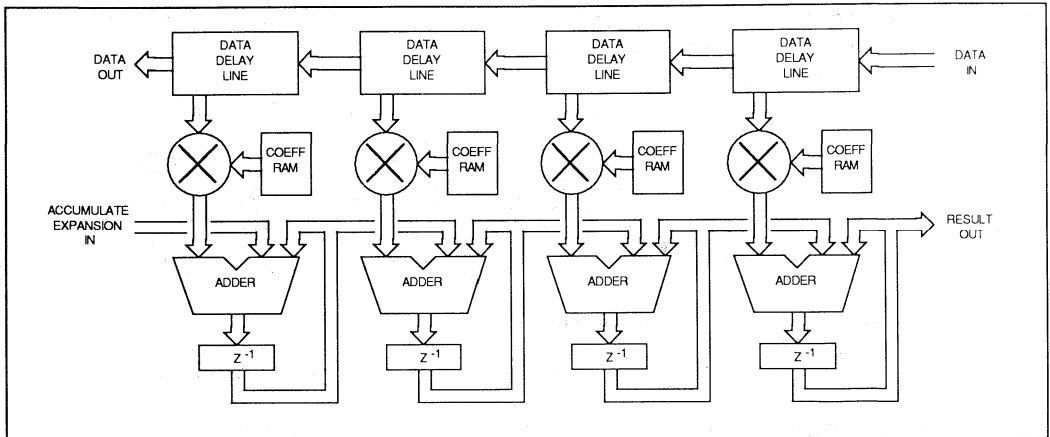


Fig. 5 Filter Network Diagram

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 buses 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


SPEED MODE 0 (Data input and output at the full SCLK rate) CR14:13 = 00, CR12 = 0




SPEED MODE 1 (Data input and output at half the SCLK rate) CR14:13 = 01, CR12 = 0



SPEED MODE 2 (Data input and output at a quarter of the SCLK rate) CR14:13 = 10, CR12 = 0



SPEED MODE 3 (Data input and output at an eighth of the SCLK rate) CR14:13 = 11, CR12 = 0



SPEED MODE 1 Decimating (Data input at half the SCLK rate and output at a quarter of the SCLK rate) CR14:13 = 01, CR12 = 1




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 re-organization. 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

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. 8 Dual Independent Filter Bus Utilisation

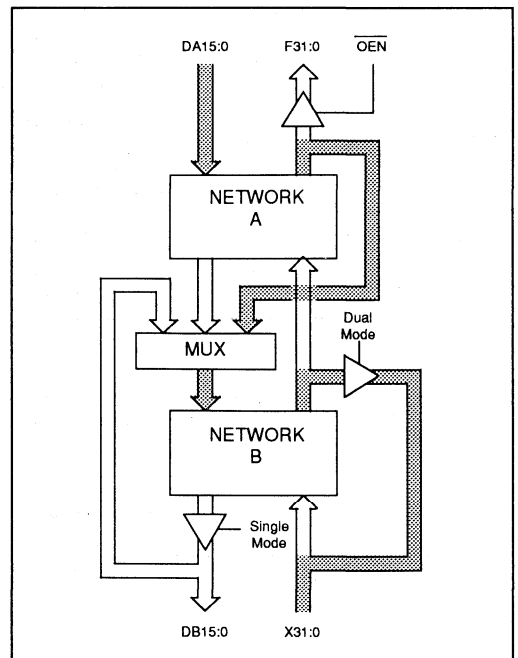


Fig. 9 Dual Cascaded Filter Bus Utilisation

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.

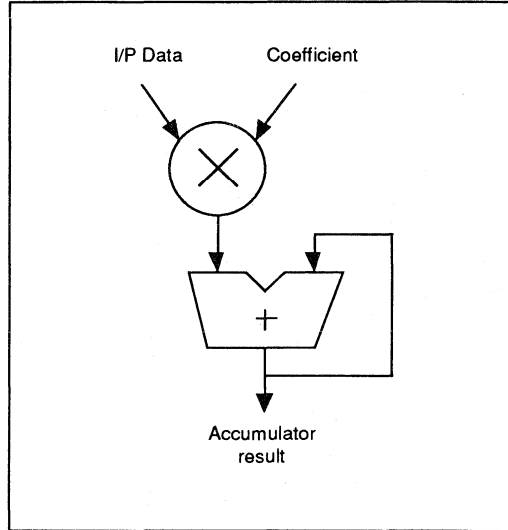


Fig. 10 Multiplier Accumulator

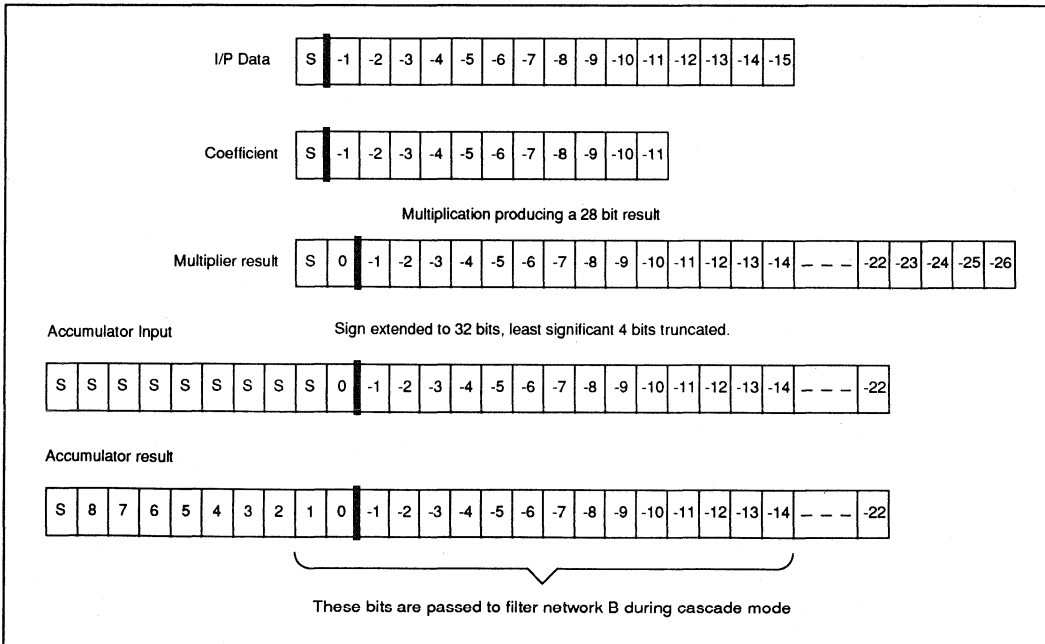


Fig. 11 Filter Accuracy

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 (FRUN) 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.

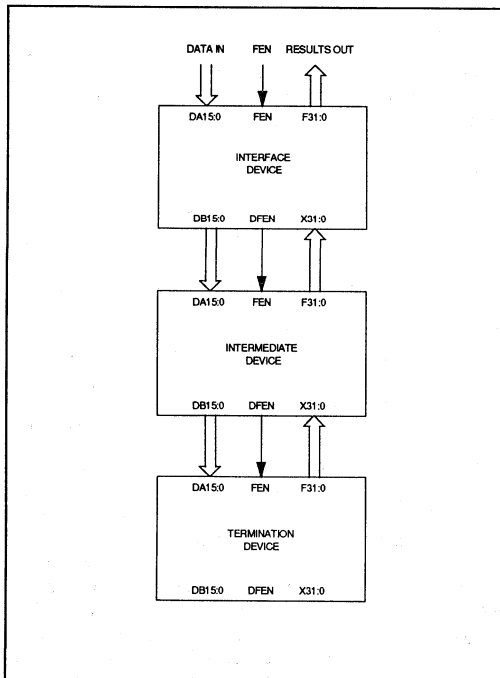


Fig. 12 Three Device Cascaded System

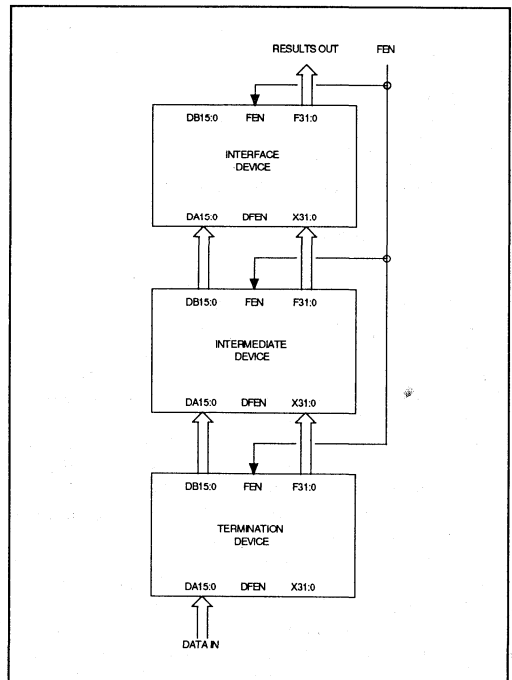


Fig. 13 Full Speed Cascaded System




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 its 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 latencies 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/A 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

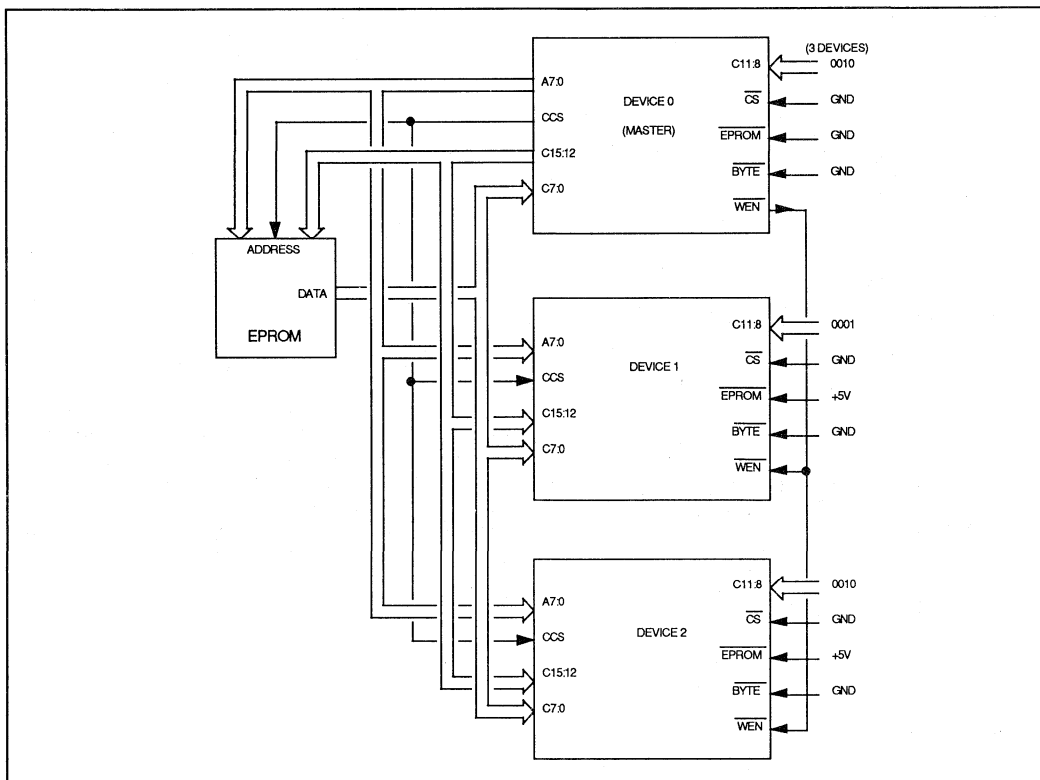


Fig. 15 Three device auto EPROM load

must also drive the address inputs on the slave devices.

When the filter length is less than the maximum, the PDSP16256/A 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

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/A 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.

Control Register			Number of Coefficients Loaded
14	13	12	
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.

The address and coefficient buses plus the 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/A. 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 RES 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

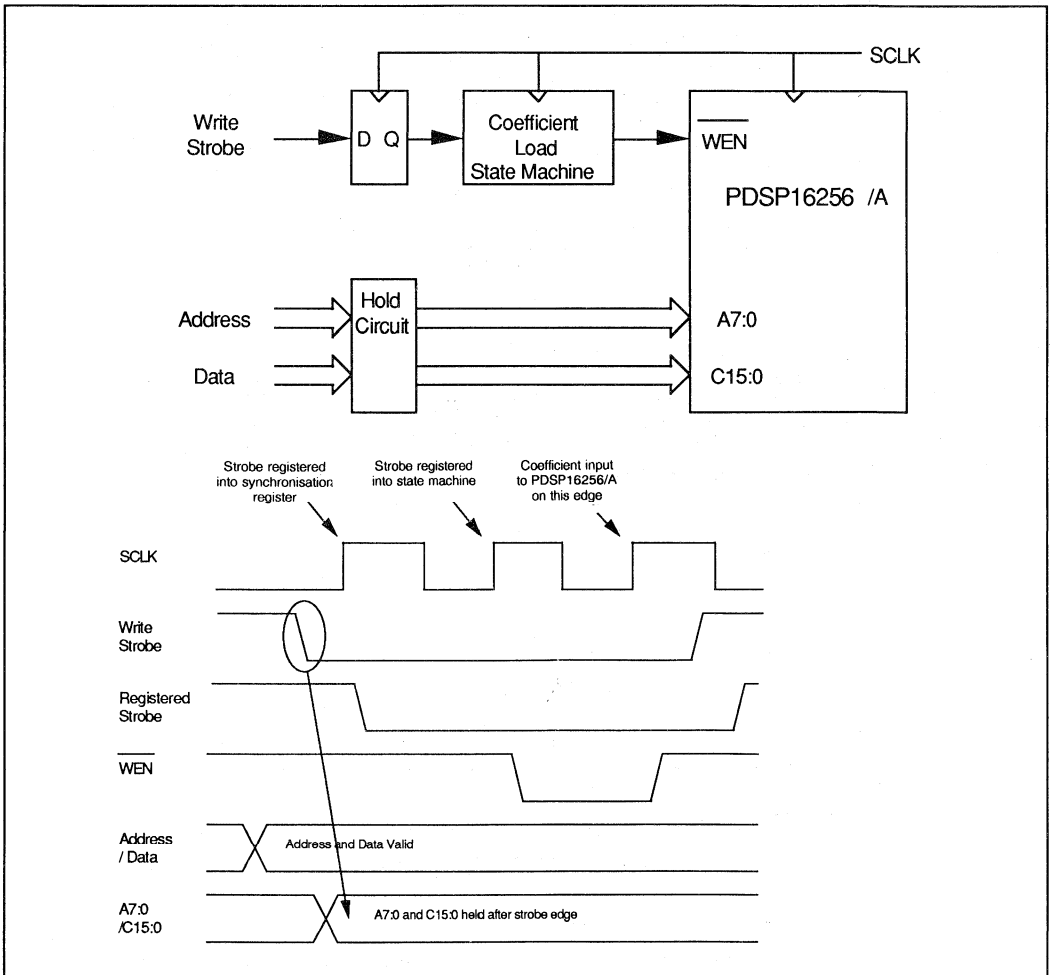


Fig. 18 Remote Master Synchronisation

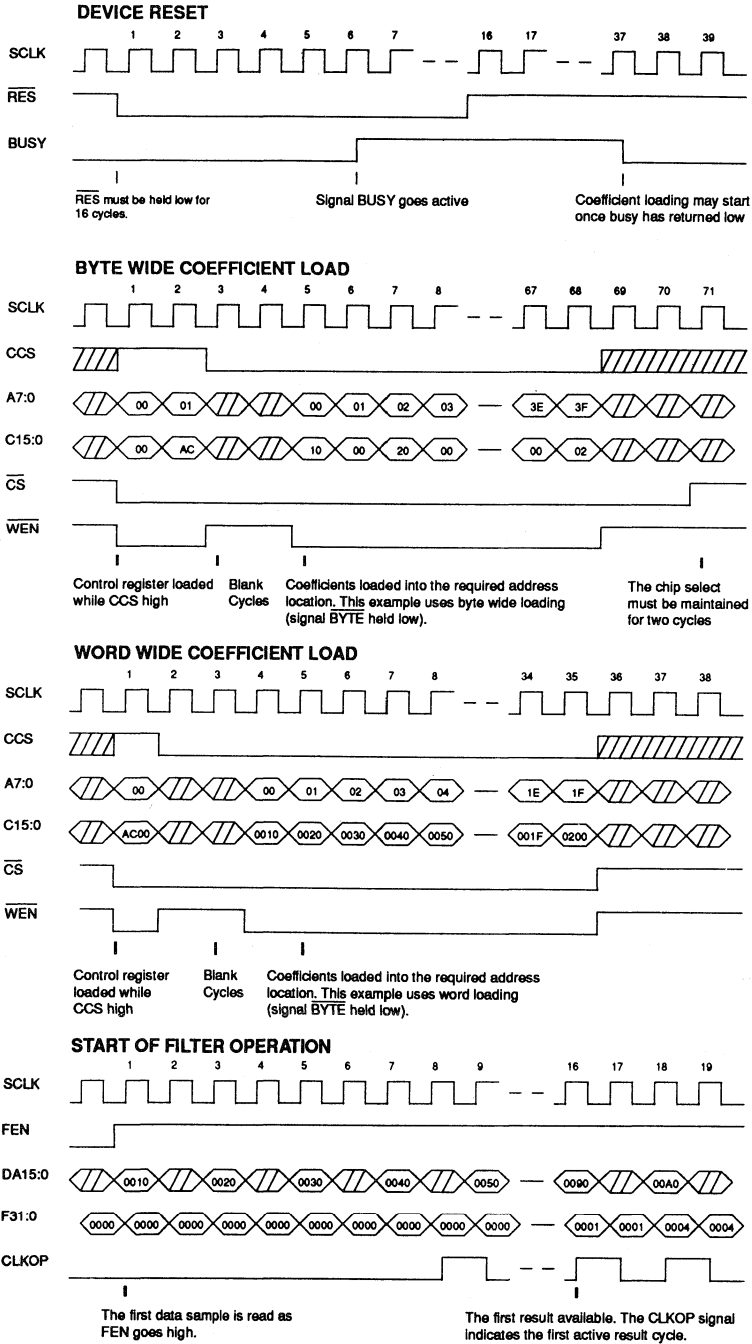


Fig. 19 Device Startup

CONTROL REGISTER

The internal operation of the PDSP16256/A 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 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 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 θ_{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 reliability.
4. Current is defined as positive into the device
5. $V_{cc} = \text{Max}$, Outputs Unloaded, Clock Freq = Max
6. The θ_{JC} data assumes that heat is extracted from the top face of the package.

ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)

Commercial: $T_{AMB} = 0^{\circ}\text{C}$ to $+70^{\circ}\text{C}$ $T_J(MAX) = 100^{\circ}\text{C}$ $V_{CC} = 5.0\text{V}\pm 5\%$ Ground = 0V
 Industrial: $T_{AMB} = -40^{\circ}\text{C}$ to $+85^{\circ}\text{C}$ $T_J(MAX) = 110^{\circ}\text{C}$ $V_{CC} = 5.0\text{V}\pm 10\%$ Ground = 0V
 Military: $T_{AMB} = -55^{\circ}\text{C}$ to $+125^{\circ}\text{C}$ $T_J(MAX) = 150^{\circ}\text{C}$ $V_{CC} = 5.0\text{V}\pm 10\%$ Ground = 0V

Static Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	$I_{OH} = 4\text{mA}$
Output low voltage	V_{OL}	-		0.4	V	$I_{OL} = -4\text{mA}$
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	$\text{GND} < V_{IN} < V_{CC}$
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	μA	$\text{GND} < V_{OUT} < V_{CC}$
Output S/C current	I_{OS}	10		300	mA	$V_{CC} = \text{Max}$

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

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

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 25MHz rate. Cartesian 16+16 bit twos' 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 $2 \times \pi$ field to eliminate phase ambiguities.

The PDSP16330 is offered in three speed grades: a basic 10MHz part (PDSP16330), a 20 MHz version (PDSP16330A) and a 25MHz version (PDSP16330B).

FEATURES

- 25MHz Cartesian to Polar Conversion
- 16-Bit Cartesian Inputs
- 16-Bit Magnitude Output
- 12-Bit Phase Output
- 2s' 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 LCC/PGA Package

APPLICATIONS

- Digital Signal Processing
- Digital Radio
- Radar Processing
- Sonar Processing
- Robotics

ASSOCIATED PRODUCTS

PDSP16112	16 × 12 Complex Multiplier
PDSP16116	16 × 16 Complex Multiplier
PDSP16318	Complex Accumulator
PDSP16340	Polar to Cartesian Converter
PDSP16350	I/Q Splitter and NCO
PDSP16510	Stand Alone FFT Processor
PDSP16520	Quad Port RAM

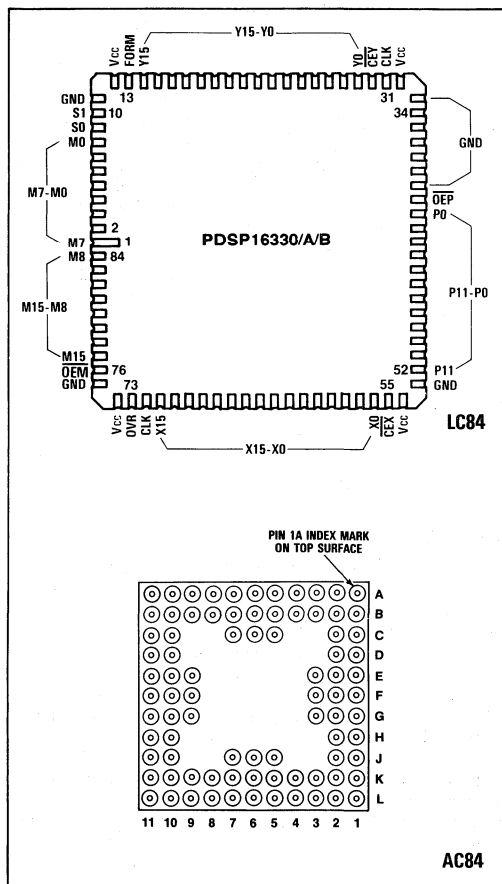


Fig.1 Pin connections - bottom view




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 Twos 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 $2 \times \pi / 4096$ 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 twos' 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 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 No.*	Pin Name and Description
CLK	31, 72	Clock: Common Clock to device Registers. Register contents change on the rising edge of clock. Both pins must be connected.
$\overline{\text{CE}}\text{X}$	55	Clock Enable: Clock Enable for X Port. The clock to the X port is enabled by a low level.
$\overline{\text{CE}}\text{Y}$	30	Clock Enable: Clock Enable for Y Port. The clock to the Y port is enabled by a low level.
X15-X0	71-56	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	14-29	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	77-84 1-8	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	52-41	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.
$\overline{\text{OEM}}$	76	Output Enable: Output Enable for M Port. The M Port is in a high impedance state when this input is high.
$\overline{\text{OEP}}$	40	Output Enable: Output Enable for P Port. The P Port is in a high impedance state when this input is high.
FORM	13	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 level indicates that two's complement data is applied, a high indicates Sign-Magnitude.
S1-S0	10, 9	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	73	Overflow: Overflow flag. This signal becomes active if the scaling currently selected causes an invalid value to be presented to the Magnitude output.
V _{cc}	12, 32, 54, 74	+ 5V supply. All V _{cc} pins must be connected.
GND	11, 33, 34, 35, 36, 37, 38, 39, 53, 75	0V supply. All GND pins must be connected.

* Pin numbers are for LC package. For AC package see Pin Function table.

INPUT DATA RANGE

Twos Complement	Sign Magnitude
7FFF	7FFF
.	.
.	.
0001	0001
0000	{0000}
	{8000}
FFFF	8001
.	.
.	.
8001	FFFF

PIN FUNCTION

Pin No.		Function	Pin No.		Function	Pin No.		Function
LC	AC		LC	AC		LC	AC	
1	F3	M7	29	L9	Y0	57	A9	X1
2	G3	M6	30	L10	\overline{CEY}	58	B8	X2
3	G1	M5	31	K9	CLK	59	A8	X3
4	G2	M4	32	L11	V _{CC}	60	B6	X4
5	F1	M3	33	K10	GND	61	B7	X5
6	H1	M2	34	J10	GND	62	A7	X6
7	H2	M1	35	K11	GND	63	C7	X7
8	J1	M0	36	J11	GND	64	C6	X8
9	K1	S0	37	H10	GND	65	A6	X9
10	J2	S1	38	H11	GND	66	A5	X10
11	L1	GND	39	F10	GND	67	B5	X11
12	K2	V _{CC}	40	G10	\overline{OEP}	68	C5	X12
13	K3	FORM	41	G11	P0	69	A4	X13
14	L2	Y15	42	G9	P1	70	B4	X14
15	L3	Y14	43	F9	P2	71	A3	X15
16	K4	Y13	44	F11	P3	72	A2	CLK
17	L4	Y12	45	E11	P4	73	B3	OVR
18	J5	Y11	46	E10	P5	74	A1	V _{CC}
19	K5	Y10	47	E9	P6	75	B2	GND
20	L5	Y9	48	D11	P7	76	C2	OEM
21	K6	Y8	49	D10	P8	77	B1	M15
22	J6	Y7	50	C11	P9	78	C1	M14
23	J7	Y6	51	B11	P10	79	D2	M13
24	L7	Y5	52	C10	P11	80	D1	M12
25	K7	Y4	53	A11	GND	81	E3	M11
26	L6	Y3	54	B10	V _{CC}	82	E2	M10
27	L8	Y2	55	B9	\overline{CEX}	83	E1	M9
28	K8	Y1	56	A10	X0	84	F2	M8

ELECTRICAL CHARACTERISTICS

Test conditions (unless otherwise stated):

T_{amb} (Commercial) = 0°C to +70°C, T_{amb} (Industrial) = -40°C to +85°C, T_{amb} (Military) = -55°C to +125°C

V_{CC} (Commercial) = 5.0V ± 5%, V_{CC} (Industrial and Military) = 5.0V ± 10%, GND = 0V

STATIC CHARACTERISTICS

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V _{OH}	2.4			V	I _{OH} = 3.2mA
Output low voltage	V _{OL}			0.6	V	I _{OL} = -3.2mA
Input high voltage (CMOS)	V _{IH}	3.0			V	Inputs \overline{CEX} , \overline{CEY} and CLK only
Input low voltage (CMOS)	V _{IL}			1.0	V	Inputs \overline{CEX} , \overline{CEY} and CLK only
Input high voltage (TTL)	V _{IH}	2.2			V	All other inputs
Input low voltage (TTL)	V _{IL}			0.8	V	All other inputs
Input leakage current (Note 1)	I _{IL}	-10		+120	μA	GND ≤ V _{IN} ≤ V _{CC}
Input capacitance	C _{IN}		10		pF	
Output leakage current	I _{OZ}	-50		+50	μA	GND ≤ V _{IN} ≤ V _{CC}
Output S/C current	I _{OS}	-50		230	mA	V _{CC} = Max

NOTES

1 All inputs except clock inputs have high value pull-down resistors

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	2 × LSTTL + 20pF	
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 after 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		
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 × LSTTL + 20pF
\overline{OEM} , \overline{OEP} low to data low data valid		30		25		25	ns		2 × LSTTL + 20pF
\overline{OEM} , \overline{OEP} high to data high impedance		30		25		25	ns		2 × LSTTL + 20pF
\overline{OEM} , \overline{OEP} low to data high impedance		30		25		25	ns		2 × LSTTL + 20pF
V _{CC} current (TTL input levels)		110		180		225	mA	V _{CC} = Max. Outputs unloaded Clock freq. = Max.	
V _{CC} 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 positive into the device
3. CMOS input levels are defined as: V_{IH} = V_{DD}·0.5V, V_{IL} = +0.5V

ABSOLUTE MAXIMUM RATINGS

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 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
LC	12	35
AC	12	36

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.

ORDERING INFORMATION

Commercial (0°C to +70°C)

- PDSP16330 C0 LC (10MHz - LCC package)
- PDSP16330 C0 AC (10MHz - PGA package)
- PDSP16330A C0 LC (20MHz - LCC package)
- PDSP16330A C0 AC (20MHz - PGA package)
- PDSP16330B C0 AC (25MHz - PGA package)

Industrial (-40°C to +85°C)

- PDSP16330 B0 LC (10MHz - LCC package)
- PDSP16330 B0 AC (10MHz - PGA package)
- PDSP16330A B0 LC (20MHz - LCC package)
- PDSP16330A B0 AC (20MHz - PGA package)
- PDSP16330B B0 AC (25MHz - PGA package)

Military (-55°C to +125°C)

- PDSP16330 A0 LC (10MHz - LCC package)
- PDSP16330 A0 AC (10MHz - PGA package)
- PDSP16330A A0 LC (20MHz - LCC package)
- PDSP16330A A0 AC (20MHz - PGA package)

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

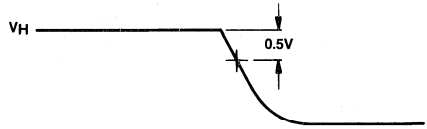
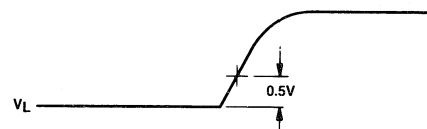
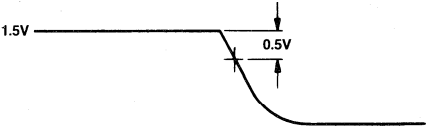
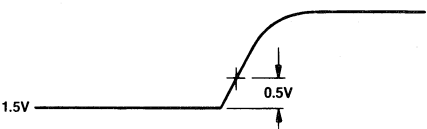
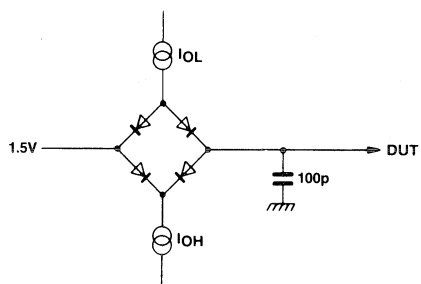
Test	Waveform - measurement level
Delay from output high to output high impedance	
Delay from output low to output high impedance	
Delay from output high impedance to output low	
Delay from output high impedance to output high	
<p>NOTES</p> <ol style="list-style-type: none"> 1. V_H - Voltage reached when output driven high. 2. V_L - Voltage reached when output driven low. <div style="text-align: center;">  </div>	

Fig.3 Three state delay measurement load

PDSP16340

POLAR TO CARTESIAN CONVERTER

(SUPERSEDES APRIL 1990 EDITION)

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 an coordinate conversion processor, the device converts data from 16 bit polar coordinates (R, θ) into 16 bit cartesian coordinates (Real, Imaginary). The translation is illustrated in Fig. 1, and uses the formula:-

$$X_r = R \cos(\theta)$$

$$X_i = R \sin(\theta)$$

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 $R \cos(\theta)$ and $R \sin(\theta)$ 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 Package

ASSOCIATED PRODUCTS

- PDSP16330 Pythagoras Processor
- PDSP16520 Quad Port Synchronous RAM
- 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




Fig. 3 Device Pinout - Bottom View

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 $\text{SQRT}(2)$ has been saturated to $\text{SQRT}(2)$. 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 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 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 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 MEN input allows the value in the input latch to be retained in a similar manner to the use of the 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.

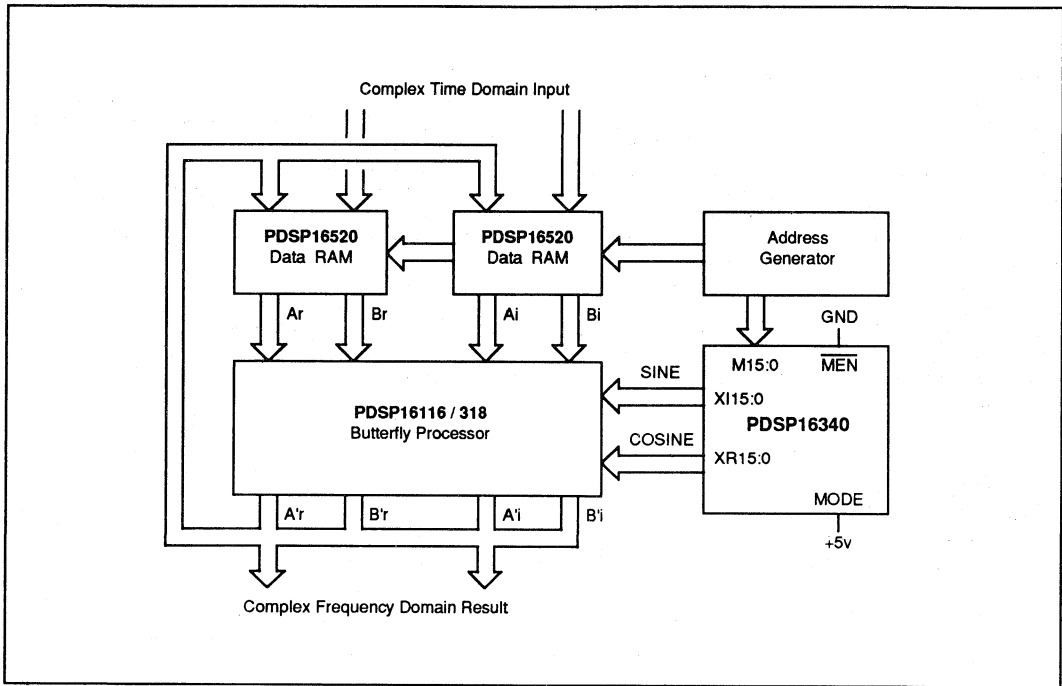


Fig. 5. Sin / Cos generator for 20 MHz FFT System

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	2	2	2	2	2	2	2	2	2	2	2	2	2	2
RANGE = 0	S	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2

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	0	-1	-2	-3	-4	-5	-6	-7	-8	-9	-10	-11	-12	-13	-14	-15
	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2

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	2	2	2	2	2	2	2	2	2	2	2	2	2	2
RANGE = 0 OR MODE = 1	S	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2

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 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 θ_{JC}	5°C/W

NOTES

- Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
- Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
- Exposure to absolute maximum ratings for extended periods may affect device reliability.
- $V_{CC} = \text{Max}$. Outputs Unloaded, Clock Freq = Max.
- CMOS levels are defined as
 $V_{FH} = V_{CC} - 0.5v$
 $V_{FL} = +0.5v$
- Current is defined as positive into the device.
- θ_{JC} data assumes that heat is extracted from the top face of the package.

ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)

Industrial: $T_{AMB} = -40^\circ\text{C}$ to $+85^\circ\text{C}$ $T_{J(MAX)} = 110^\circ\text{C}$ $V_{CC} = 5.0V \pm 10\%$ Ground = 0V

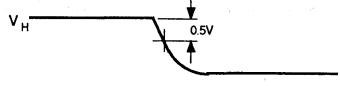
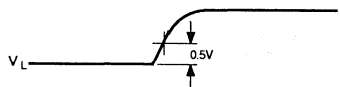
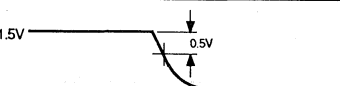
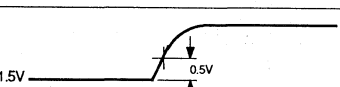
Military: $T_{AMB} = -55^\circ\text{C}$ to $+125^\circ\text{C}$ $T_{J(MAX)} = 150^\circ\text{C}$ $V_{CC} = 5.0V \pm 10\%$ Ground = 0V

Static Characteristics

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	$I_{OH} = 4mA$ $I_{OL} = -4mA$
Output low voltage	V_{OL}	-		0.4	V	
Input high voltage	V_{IH}	2.0		-	V	GND < V_{IN} < V_{CC}
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	μA	
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	μA	
Output S/C current	I_{SC}	10		250	mA	

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	30pF	
M15:0 or P15:0 hold after clock rising edge	4		-	4		-	ns		
\overline{MEN} or \overline{PEN} setup to clock rising edge	20		-	20		-	ns		
\overline{MEN} or \overline{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		
Clock freq	DC		20	DC		20	MHz		
Clock High Time	15		-	15		-	ns		
Clock Low Time	20		-	20		-	ns		
$\overline{OER}, \overline{OEI}$ low to data valid	-		20	-		20	ns		see Fig. 6
$\overline{OER}, \overline{OEI}$ high to data high impedance	-		20	-		20	ns		see Fig. 6
Pipeline delay V_{IN} to V_{OUT}	28		28	28		28	CLKs		
V_{CC} Current (CMOS inputs)	-		430	-		450	mA		see Note 4

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	




Fig. 6 Tri-state delay measurement load.

ORDERING INFORMATION

PDSP16340 B0 AC (Industrial - PGA package)

PDSP16340 A0 AC (Military - PGA package)

Call for availability of High Reliability parts and MIL-STD-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.

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 package

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




Fig. 1 Block Diagram

ASSOCIATED PRODUCTS

- PDSP16256 Programmable FIR Filter
- PDSP16510 FFT Processor
- PDSP16340 Polar to Cartesian Converter
- PDSP16488 2D Convolver

ORDERING INFORMATION

- PDSP16350 B0 AC (Industrial - PGA package)
 - PDSP16350 A0 AC (Military - PGA package)
- Call for availability of High Reliability parts and MIL-STD-883C screening.




Fig. 2 Pin Out Diagram Bottom View

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 lower 18 bits. The full 34 bit phase increment register can also be loaded using JUMP see below.
JUMP	<p>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.</p> <p>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.</p>
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 CEN low. If 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 :

$$\text{Resolution} = \frac{\text{Clock Frequency}}{2^{34}} \quad \text{Hz}$$

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 :

$$\text{Resolution} = \frac{\text{Clock Frequency}}{2^{18}} \quad \text{Hz}$$

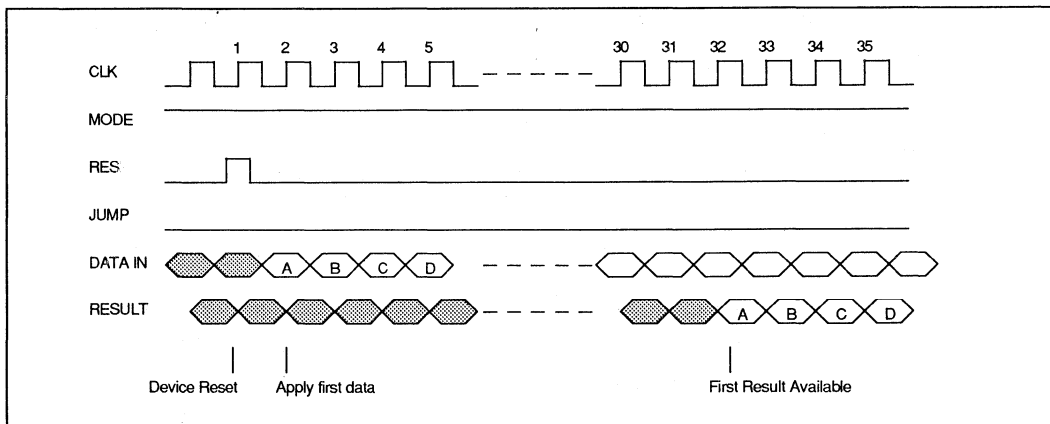


Fig. 4 Amplitude Modulation (18bit frequency accuracy)

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.

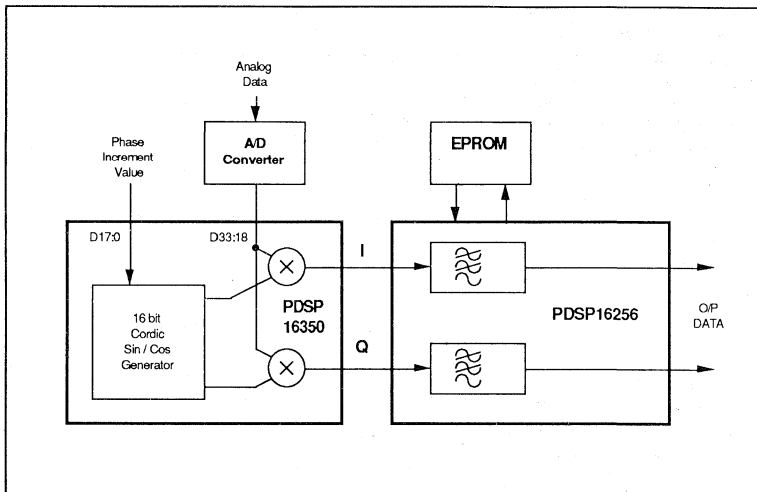


Fig. 5 IQ Split Function

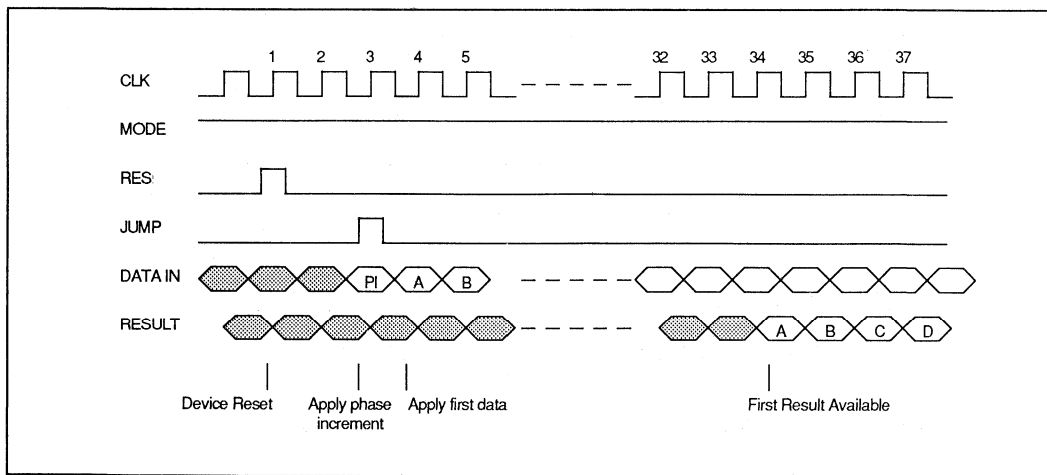


Fig. 6 Amplitude Modulation (34bit frequency accuracy)

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.

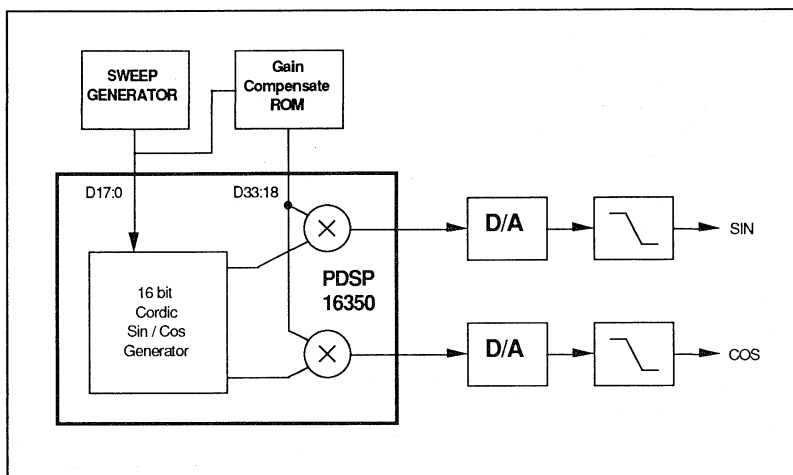


Fig. 7 Quadrature Chirp Generator




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

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

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 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 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 θ_{JC}	5°C/W

NOTES

- Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
- Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
- Exposure to absolute maximum ratings for extended periods may affect device reliability.
- V_{CC} = Max, Outputs Unloaded, Clock Freq = Max.
- CMOS levels are defined as
V_{IH} = V_{DD} - 0.5v
V_{IL} = +0.5v
- Current is defined as positive into the device.
- The θ_{JC} data assumes that heat is extracted from the top face of the package.

ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise stated)

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}	2.0		-	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	
OES, OEC low to data valid	-		20	-		20	ns	30pF
OES, OEC high to data high impedance	-		20	-		20	ns	30pF
Pipeline delay VIN to VOUT	31		31	31		31	CLKs	
V _{CC} Current (CMOS inputs)	-		430	-		450	mA	See Note 4
V _{CC} Current (TTL inputs)	-		460	-		500	mA	See Note 4

PDSP16401/PDSP16401A

2-DIMENSIONAL EDGE DETECTOR

The PDSP16401 is a single chip CMOS video signal processor which will determine the presence, direction and gradient magnitude of edges in a 3 x 3 pixel frame in a raster scanned image.

FEATURES

- 22 Megapixels/Sec Processing Rate
- 13 Bit Edge Magnitude Output
- 3 Bit Edge Direction Output
- Built-in Threshold Detector
- P_{tot} < 700mW at 22MHz

APPLICATIONS

- Machine Vision
- Image Enhancement
- Pattern Recognition
- Video Effects Generation

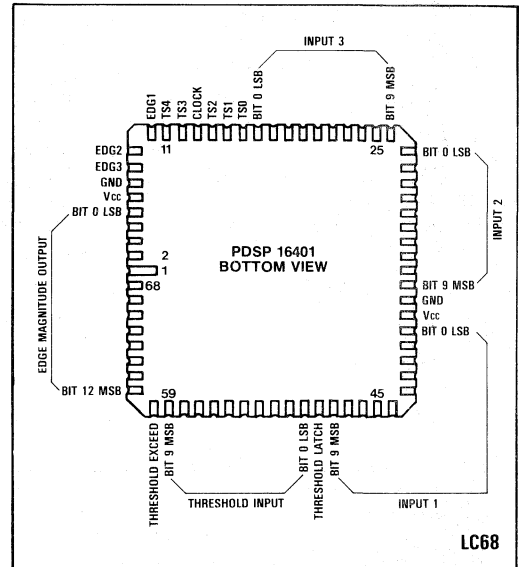


Fig.1 Pin connections - bottom view (not to scale)

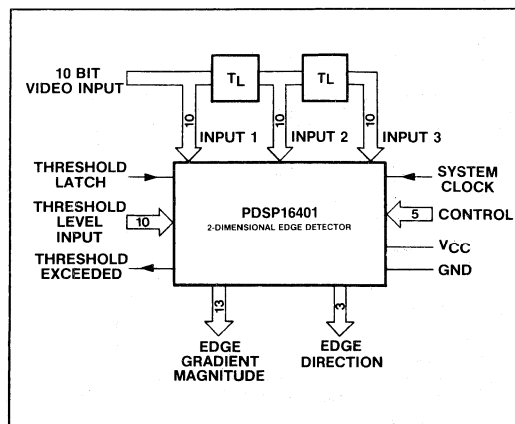


Fig.2 Typical system connection.
The two line stores should have a delay of one line period minus three clock periods.

PIN DESCRIPTION

Pin Nos.	Name	Function
39 to 48	INPUT 1	Input for 10 bits of digitised video corresponding to the current line of the image. Data is unipolar, black level is all zeros. Pin 39 is LSB.
27 to 36	INPUT 2	Input for 10 bits of digitised video corresponding to the previous line of the image (i.e. delayed by 1 line period - $3 \times t_{\text{clock}}$). Data is unipolar, black level is all zeros. Pin 27 is LSB.
17 to 26	INPUT 3	Input for 10 bits of digitised video delayed by (2 line periods) - $6 \times t_{\text{clock}}$. Data is unipolar, black level is all zeros. Pin 17 is LSB.
37, 7	GND	0V supply. Both GND pins must be connected.
38, 6	V _{cc}	+5V supply. Both V _{cc} pins must be connected.
49	THRESHOLD LATCH	Used in conjunction with the clock to latch data from THRESHOLD INPUT into DETECTOR input latch. Data is held when THRESHOLD LATCH is low.
50 to 59	THRESHOLD INPUT	10 bits input to THRESHOLD DETECTOR. Data is latched in by THRESHOLD LATCH. Data is compared with 10 MSBs of EDGE MAGNITUDE output. Pin 50 is LSB.
60	THRESHOLD EXCEEDED	Single bit output, logic high when threshold level is exceeded.
5 to 1 & 68 to 61	EDGE MAGNITUDE	13 bits of data representing the largest edge magnitude within the current 3×3 pixel image frame. Pin 5 is LSB. Data is in magnitude only format (see Table 2).
10, 9, 8	EDGE 1-3	3 bits representing the edge direction. See Fig.4 for encoding detail.
13	CLOCK	System clock, no minimum frequency.
11, 12, 14, 15, 16	T _{S4} to T _{S0}	Control pins, normally all held low, but can be set to give single filter operation (see Table 2).




Fig.3 Edge detector block diagram

FUNCTIONAL DESCRIPTION

The PDSP16401 requires three 10-bit wide digitised video inputs, corresponding to three lines of the input image (see Fig.2).

Edges are detected by concurrent convolution of a window consisting of three adjacent pixels in each of the three input lines, with four 3 x 3 masks (Table 1 gives the coefficients of the convolution masks). The four masks are implemented by four separate FIR filters operating in parallel (see Fig.3). Each filter passes video data associated with the orientation of each particular mask. Since within the 3 x 3 window the horizontal or vertical distance between pixels is less than the diagonal distance by a factor of $\sqrt{2}$, the horizontal and vertical mask functions are scaled by a factor of 1.5.

The outputs from the four filters are fed into the comparison network which compares the magnitude of the 13-bit outputs, producing a 3-bit word representing the

largest output plus its sign. The sign represents the direction of the edge ie black to white or white to black.

Fig.4 illustrates the coding, where the arrow represents a direction perpendicular to a black-white transition. A 13 bit 2's complement word, which is the magnitude of the output of the filter producing the maximum output is also produced.

The 10 MSBs of the 13-bit magnitude output are internally fed into a threshold detector where they are compared with the external threshold level input. If the output magnitude exceeds the threshold level, the TE output goes high.

Control Inputs

Table 2 gives the operations associated with the various control inputs. When a single convolution mask is selected, the edge magnitude comparator is disabled.

ORIENTATION	MASK
HORIZONTAL (Filter 1)	$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix} \times 1.5$
VERTICAL (Filter 2)	$\begin{bmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{bmatrix} \times 1.5$
45° Bottom left to top right (Filter 3)	$\begin{bmatrix} 2 & 1 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & -2 \end{bmatrix}$
45° Top left to bottom right (Filter 4)	$\begin{bmatrix} 0 & -1 & -2 \\ 1 & 0 & -1 \\ 2 & 1 & 0 \end{bmatrix}$

Table 1 Convolutional masks




Fig.4 Detail of 3-bit word representation of edge orientation

		Pin Nos. and Codes					Magnitude Output Format	
		16	15	14	12	11		
Package Pin Number		16	15	14	12	11		
Package Pin Name		TS0	TS1	TS2	TS3	TS4		
Logic Levels	Normal Edge Detection Mode	0	0	0	0	0	Magnitude	
	Single* Direction Sensing Mode	Filter 1**	1	1	1	1	0	2's Comp
		Filter 2**	0	0	0	1	0	2's Comp
		Filter 3**	1	0	1	1	0	2's Comp
		Filter 4	0	0	1	1	0	2's Comp

* 14th Bit (MSB) appears on EDG3 in Single Filter Mode.

** When using Filters 1, 2 and 3 - all output bits are inverted.

Table 2 Control pin codes and functions




Fig.5 Timing diagram

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)	$\pm 18mA$
Static discharge voltage (HMB)	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

- Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
- Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
- 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	7	36

ELECTRICAL CHARACTERISTICS

Test conditions (unless otherwise stated):

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

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

Static Characteristics

Characteristic	Symbol	Value Ind./Military			Value Ind. Only			Units	Conditions
		PDSP16401			PDSP16401A				
		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}	2.2			2.2			V	
Input low voltage	V _{IL}			0.8			0.8	V	
Input leakage current	I _L	-10		+10	-10		+10	μA	GND ≤ V _{IN} ≤ V _{CC}
Output short circuit current (Note 2)	I _{OS}	12		80	12		80	mA	V _{CC} = max.
Input capacitance	C _I		10			10		pF	

Switching Characteristics

Characteristic	Symbol	Value Industrial						Value Military		Units	Conditions
		PDSP16401			PDSP16401A			Min.	Max.		
		Min.	Typ.	Max.	Min.	Typ.	Max.				
V _{CC} current	I _{CC}			100			140		100	mA	V _{CC} = max f _{CLK} = max. No O/P loading I/Ps low
CLK frequency	f _{CLK}			15			22		15	MHz	
Min. CLK low		25			20			25		ns	
Min. CLK high		25			20			25		ns	
Input setup time (data)	t _{ds}	30			25			30		ns	
Input hold time (data)	t _{dh}	3			3			3		ns	
Input setup time (control)	t _{is}	50			40			50		ns	
Input hold time (control)	t _{is}	3			3			3		ns	
Delay, clock to output	t _{od}			50			35		50	ns	
Threshold latch to clock setup time	t _{TLS}	10			5			10		ns	
Threshold latch to clock hold time	t _{TLH}	3			3			3		ns	
Threshold input to clock setup time	t _{TIS}	15			10			15		ns	
Threshold input to clock hold time	t _{TIH}	3			3			3		ns	
Latency, input to edge magnitude output		20		20	20		20	20		cycles	
Latency, threshold input to threshold exceeded		3		3	3		3	3		cycles	

ORDERING INFORMATION

Industrial (-40°C to +85°C)

PDSP16401 B0 LC (Industrial - LCC package)

PDSP16401A B0 LC (Industrial - LCC package)

Military (-55°C to +125°C)

PDSP16401 A0 LC (Military - LCC package)

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

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.

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

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 package

ASSOCIATED PRODUCTS

- PDSP16401 2D Edge Detector
- PDSP16510 FFT Processor
- PDSP16256 Programmable FIR Filter

ORDERING INFORMATION

PDSP16488 C0 AC (Commercial - PGA package)

PDSP16340 B0 AC (Industrial - PGA package)

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

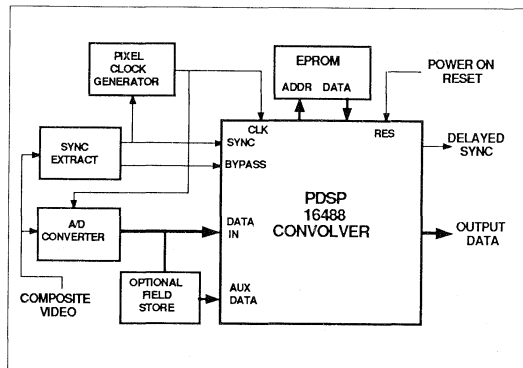


Fig. 1 Typical, Stand Alone, Real Time System




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

PDSP16488

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.
$\overline{PC1}$	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 $\overline{PC0}$ input on the next device.
$\overline{PC0}$	INPUT	This pin is used in conjunction with $\overline{PC1}$ 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{DS}	I/O	The data strobe from a host computer. Active low. This pin will be an output from an EPROM supported MASTER or SINGLE device, which provides strobes to any other devices.
\overline{CE}	INPUT	An active low enable which is internally gated with $\overline{R/W}$ and \overline{DS} to perform reads or writes to the internal registers. In a SINGLE or MASTER device, which is supported from an EPROM, \overline{CE} should be tied high. In such a system, however, \overline{CE} can be pulsed low to initiate a new load procedure after reset has gone inactive.
$\overline{R/W}$	INPUT	Read / not write line from the host CPU. When an EPROM is used this pin should be tied low.
\overline{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{RES}	INPUT	Active low power on reset signal.
\overline{SINGLE}	INPUT	Tied to ground to indicate a SINGLE device system. Internal pull up resistor.
\overline{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{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 $N \times N$ 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 8×8 window, it is possible for the accumulated sum to grow to 22 bits within a single device. With 16 bit pixels, and an 8×4 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 determine if the result was above or below a value contained within an internal register.

MULTIPLIER ARRAY

The PDSP16488 contains sixteen 8×8 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 if 16 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.

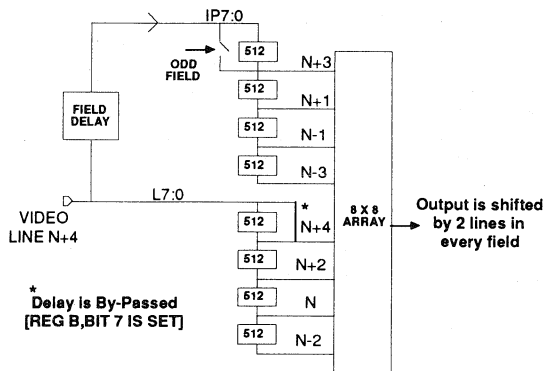
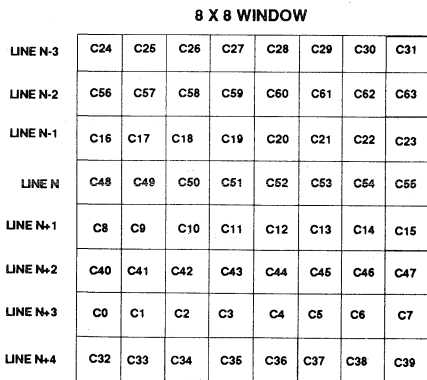
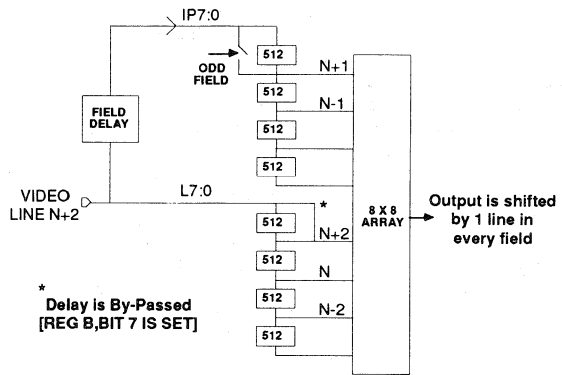
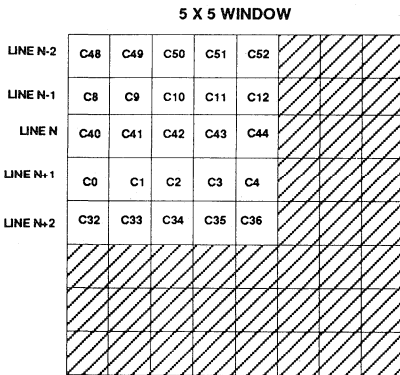




Figure 3. Line Delay Allocations in Single Device Interlaced Systems

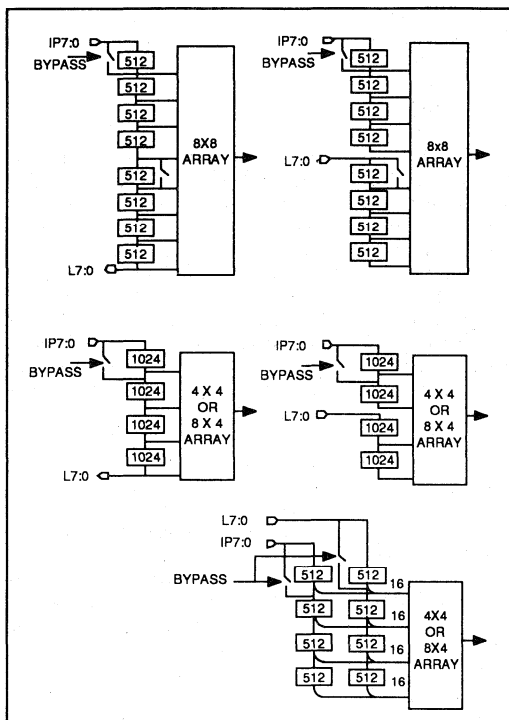


Fig. 4. Line Delay Configurations

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.

Figure 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 occurrences 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 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




Fig. 5. Gain Control Operation

values, but otherwise leaves them unchanged. Occasional overflows can be tolerated in some systems, and this option prevents any gross errors.

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.

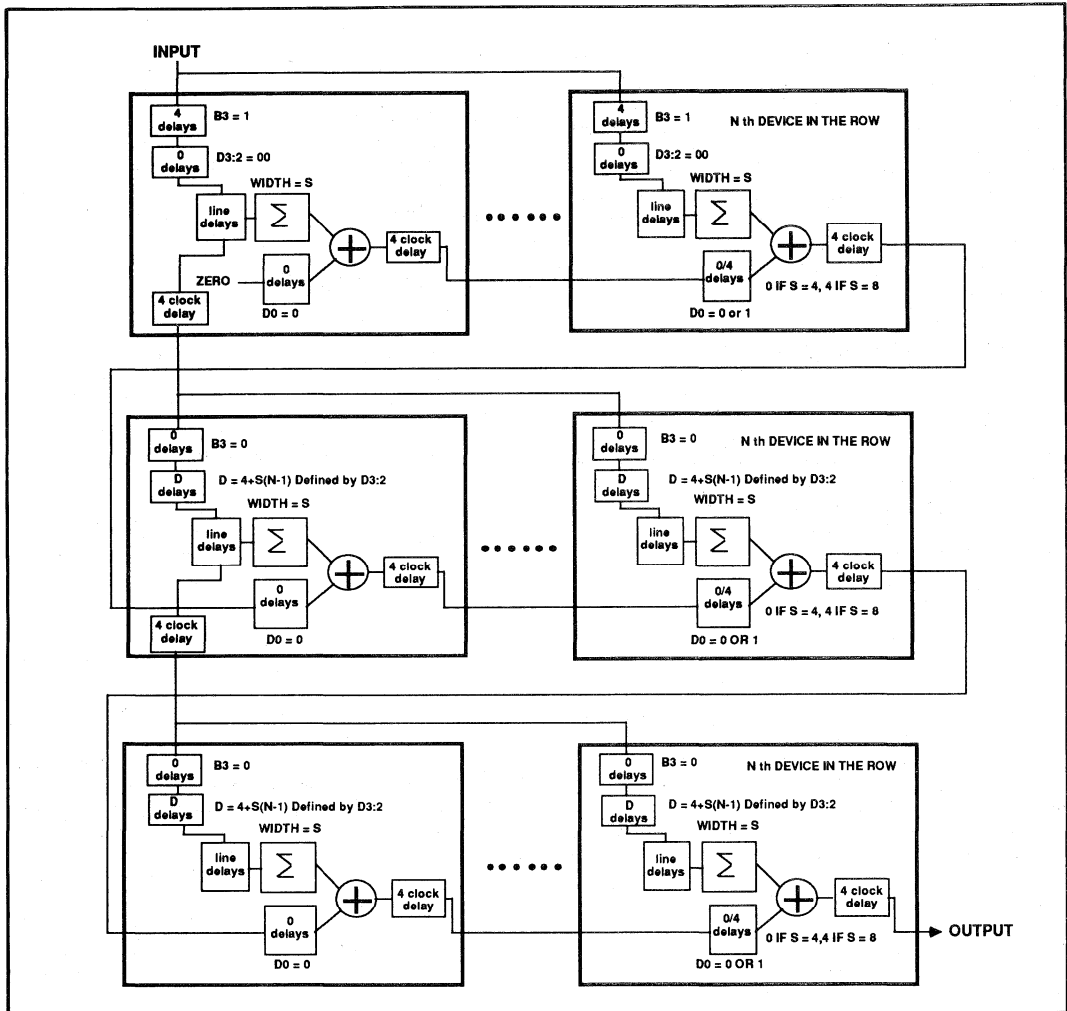


Fig. 6. Multi-Device Delay Paths

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;

$$DELAY = 4 + [N-1].S \text{ where } N \text{ is the number of devices in a row and } S \text{ is the partial window width, ie 4 or 8.}$$

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 coefficient 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

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	35
8	8x4	30
8	8x8	29
16	4x4	29
16	8x4	29

Table 4 Pipe line delays

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 PSDP16488 contains facilities for outputting 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.





Fig.7 Pixel Input Delays



NOTE
Two coefficients occurring 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{PROG} signal from the host.

Registers are individually addressed and can be loaded in any sequence once the global \overline{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 $\overline{PC1}/\overline{PC0}$ connections. The $\overline{PC0}$ output from the final device is used as a host \overline{REPLY} signal, and indicates that the last device has received data after the propagation delay of previous devices. The timing strobe is produced in the MASTER device from the host data strobe, and will appear on the $\overline{PC0}$ output. This feature allows the user to cascade any number of devices without knowing the propagation 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{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{REPLY} goes active, as shown in Figure 8.

If the \overline{REPLY} signal is not to be used, the $\overline{PC0}/\overline{PC1}$ connections are not necessary, and the host data strobe for a write operation must be wide enough to allow for the worst case propagation delay through all the devices (TDEL). If the data or address from the host does not meet the set up time given in Figure 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: $TACC + (N - 1) \cdot TDEL$ 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

- $\overline{PC0}$ An input from the previous $\overline{PC1}$ output in a multiple device chain. Not needed on a SINGLE device or if the self timed feature is not used.
- $\overline{PC1}$ 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 $\overline{PC0}$ input of the next device at intermediate points in the chain if the self timed feature is used.
- R/\overline{W} Read/Not Write line from the host CPU which is connected to all devices in the system.
- \overline{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{DS} An active low host data strobe which is connected to all devices in the system.
- \overline{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. \overline{PROG} and \overline{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 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{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/\overline{W} input should be tied low on all devices.

The width of the data strobe is determined by the feedback connection from the $\overline{PC1}$ output on the last device to the $\overline{PC0}$ input on the MASTER. The $\overline{PC0}/\overline{PC1}$ 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 relevant 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{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.
$\overline{R/W}$	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. The input should be tied high on a MASTER or SINGLE device, unless it is to be used to initiate a new load procedure, after \overline{RES} has gone inactive. The low transition initiates this load procedure.
\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 occurring, either after power on, or as the result of \overline{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 SINGLE and 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 from 1-15
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 inter-connected 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 multiplication 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 is 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 3:2.

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
4	1	2's complement pixel data
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 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 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 resistance, junction-to-case θ_{JC}	5°C/W

NOTES ON MAXIMUM RATINGS

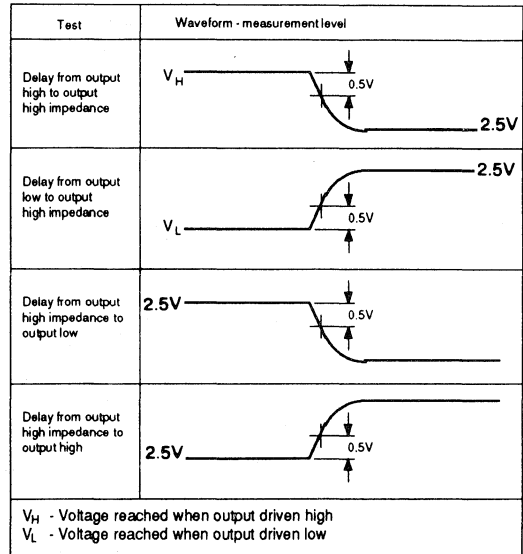
- Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
- Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
- Exposure to absolute maximum ratings for extended periods may affect device reliability.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

STATIC ELECTRICAL CHARACTERISTICS

Operating conditions (unless otherwise stated):

$T_{amb} = -40^\circ\text{C}$ to $+85^\circ\text{C}$, $V_{CC} = 5.0V \pm 10\%$

NOTE: Signal pins PROG, PC0, X15, MASTER, SINGLE, DS, BYPASS and OV 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		Units	Notes
		Min.	Max.		
\overline{DS} Hold Time after \overline{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 rising edge of strobe when set up must be $N \times T_{DEL}$, for N devices
Read Set Up Time 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 T_{DEL} 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}	5		ns	
Chip Enable Hold Time	T_{CH}	5		ns	
$\overline{PC1}$ 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	Single device	
Delay from $\overline{PC0}$ Input to Write in-active	T_{WH}	5		ns		
$\overline{PC1}$ In-Active Delay	T_{PCH}		50	ns		
Write from 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		
\overline{CE} Set Up Time	T_{CSU}	0		ns		
\overline{CE} Hold Time	T_{CH}	0		ns		
Available 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	Increases to 40ns for DELOP output
Line Store Output Delay	T_{LD}		20	ns	
HRES In-active Set Up Time	T_{RSU}	10		ns	
Output Enable Time	T_{DLZ}		15	ns	} Measured with a 15k Ω series resistor and 30pF load capacitance
Output Disable Time	T_{DHZ}		25	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 SINGLE and X15 pins must be tied tied low, and the PC0, PC1, and DS pins are redundant. The PROG pin

becomes an output and indicates that a register load sequence is occurring. 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 RES input. Alternatively the 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 systems 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/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 PROG and CE pins may be connected together, and then driven by a host address decode. The output on PC1, which provides a 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 irrelevant in the SINGLE mode, as is bit 7 if the gain control is used.

With 8-bit pixels, 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 D15:0 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 MASTER pin; the other then receives control signals

in the normal way and has its MASTER and 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 outputting 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 outputting 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 second 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.

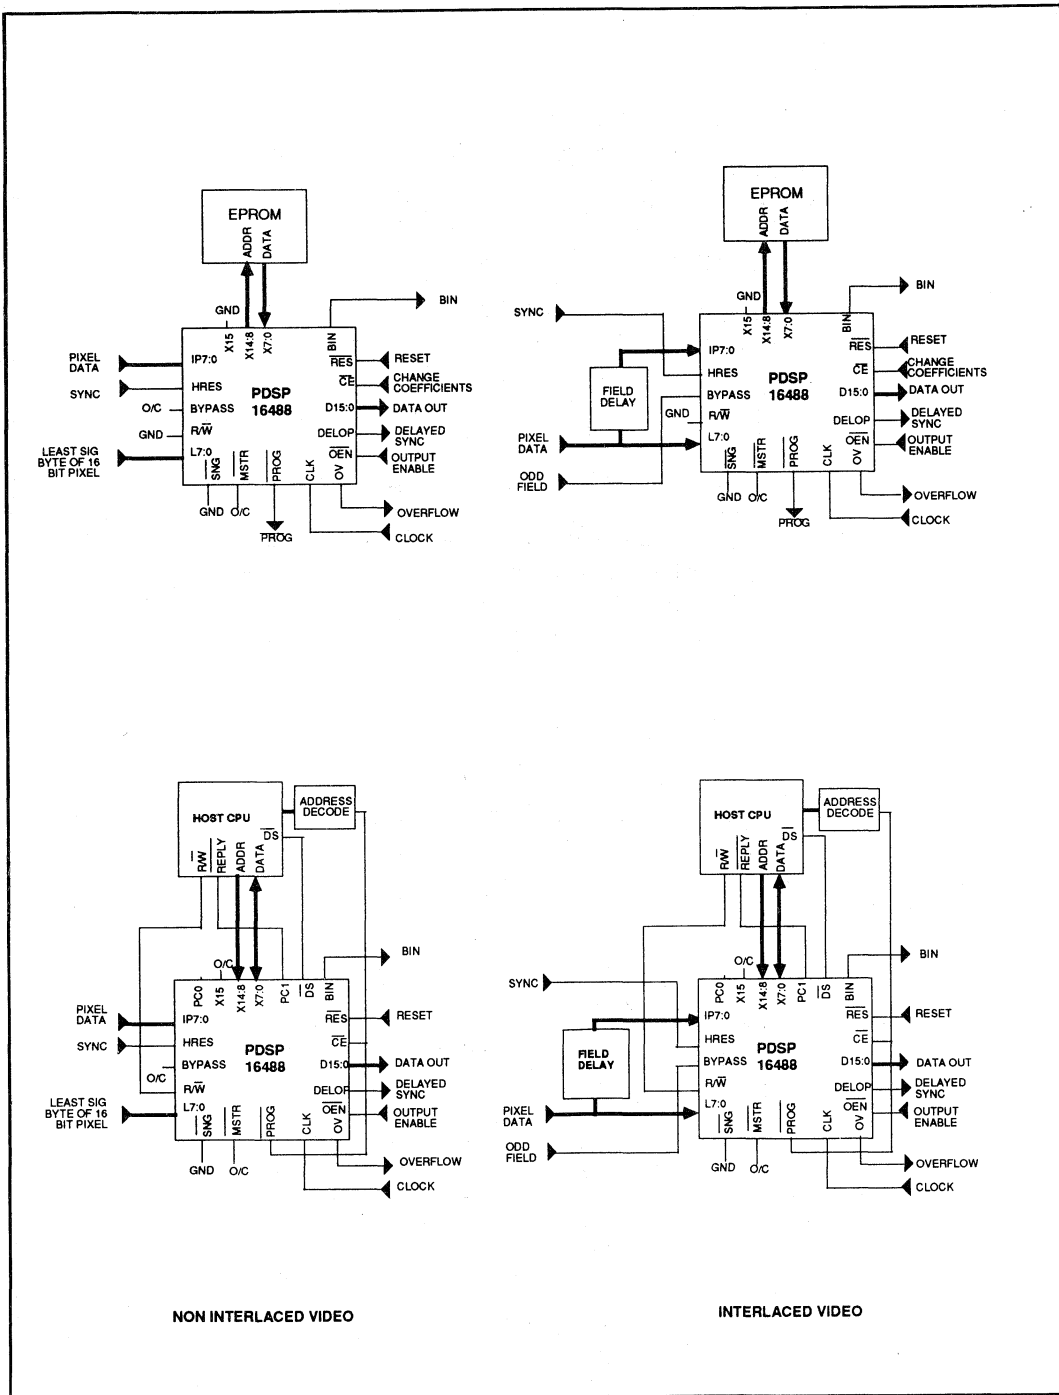
EIGHT DEVICE SYSTEMS

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.



NON INTERLACED VIDEO

INTERLACED VIDEO

Figure 11 Single Device Systems

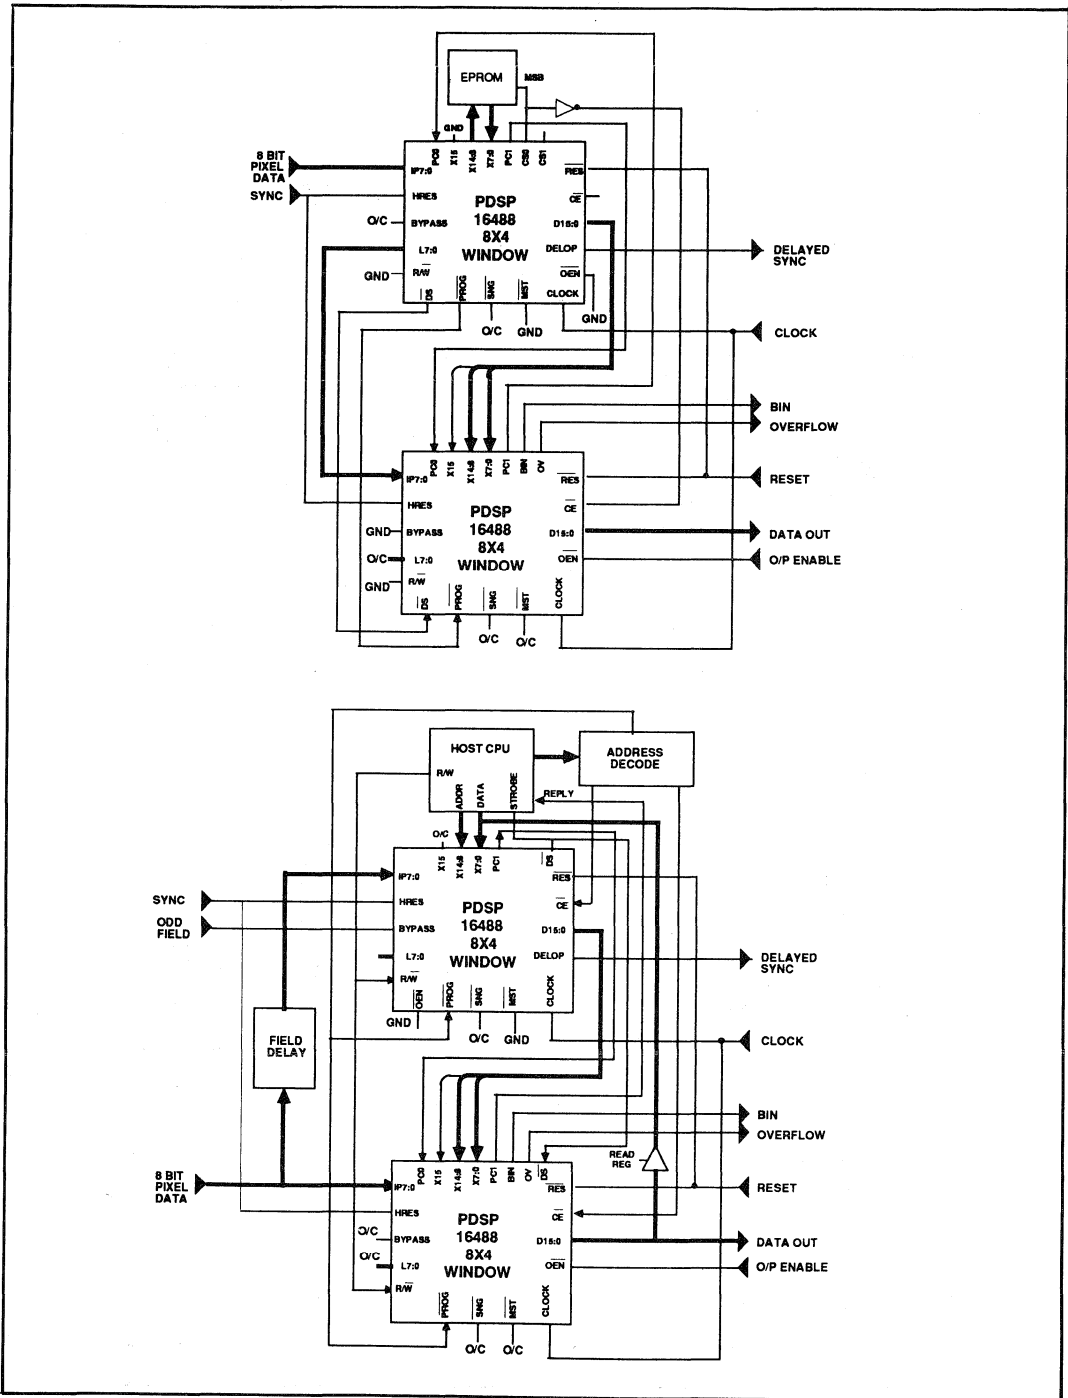


Figure 12. 8 Bit Dual Device Systems

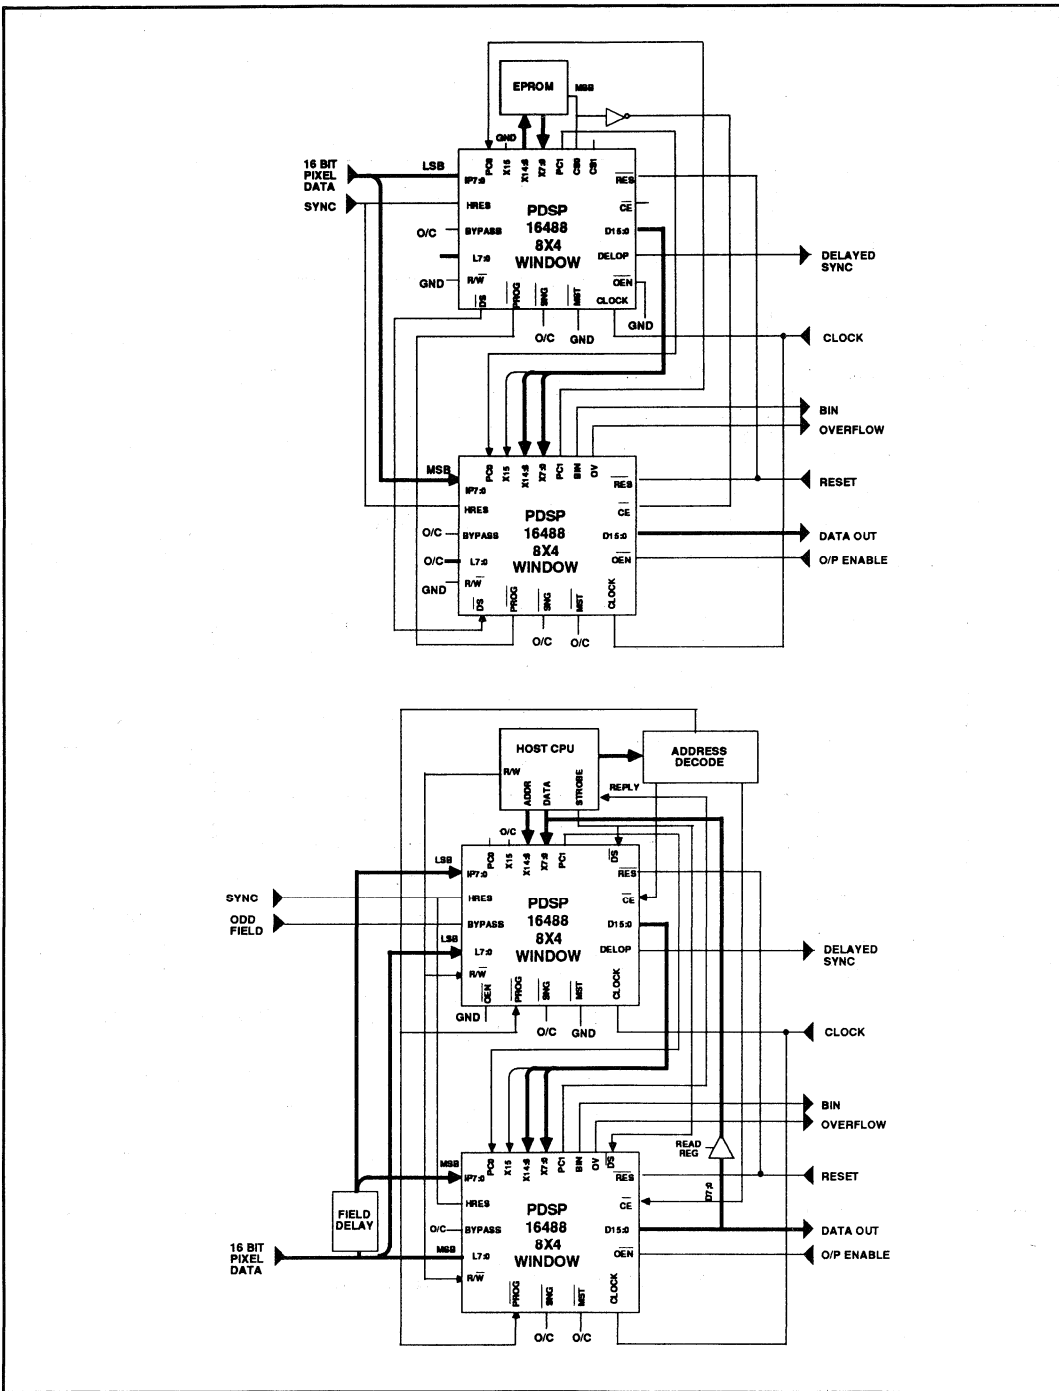


Figure 13. Dual Device 16 Bit Systems.

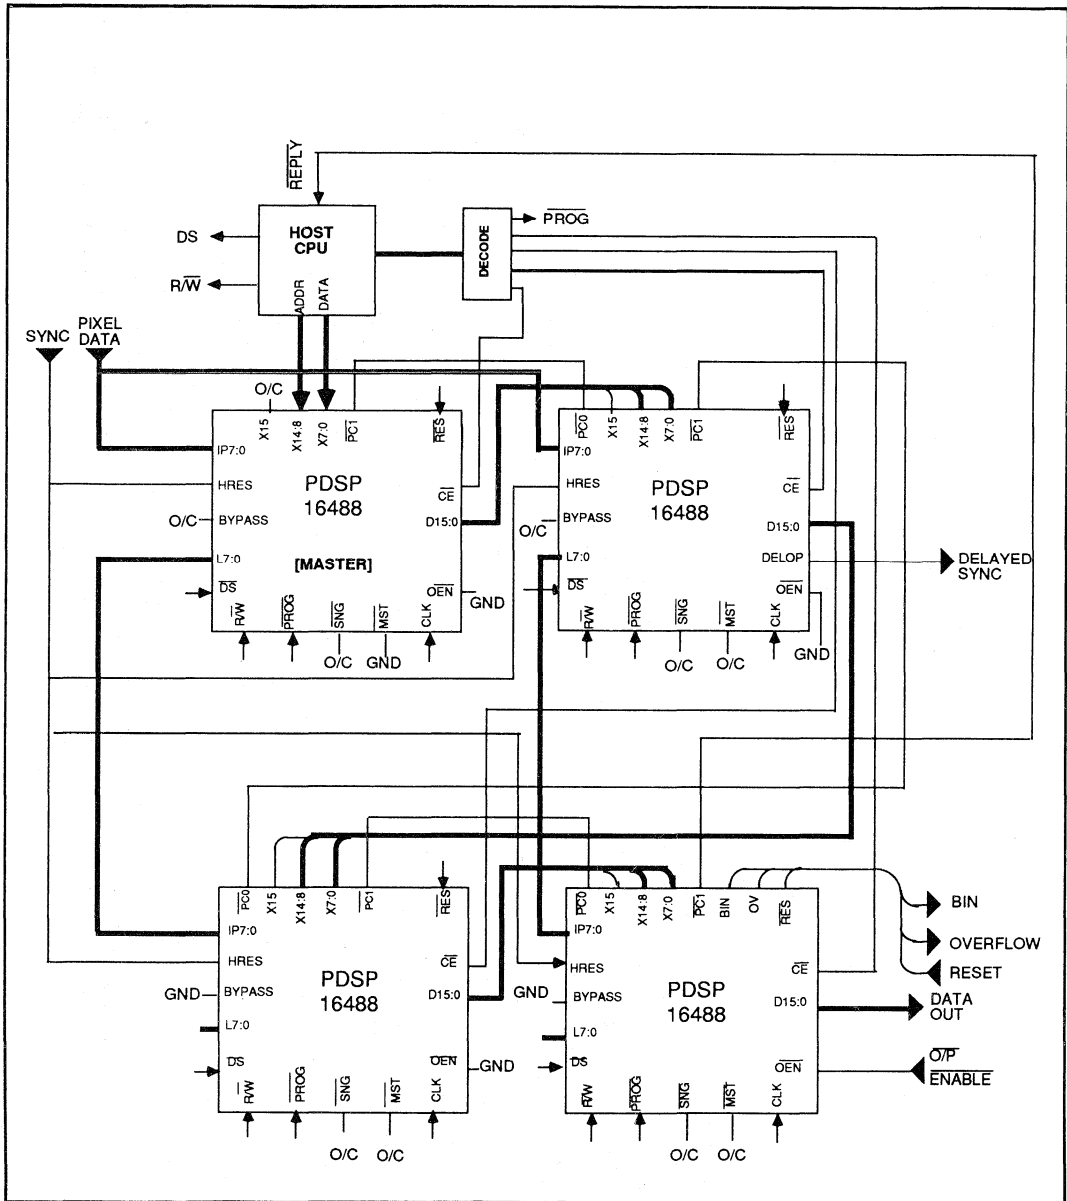


Figure 14. Four Device Non Interlaced System.

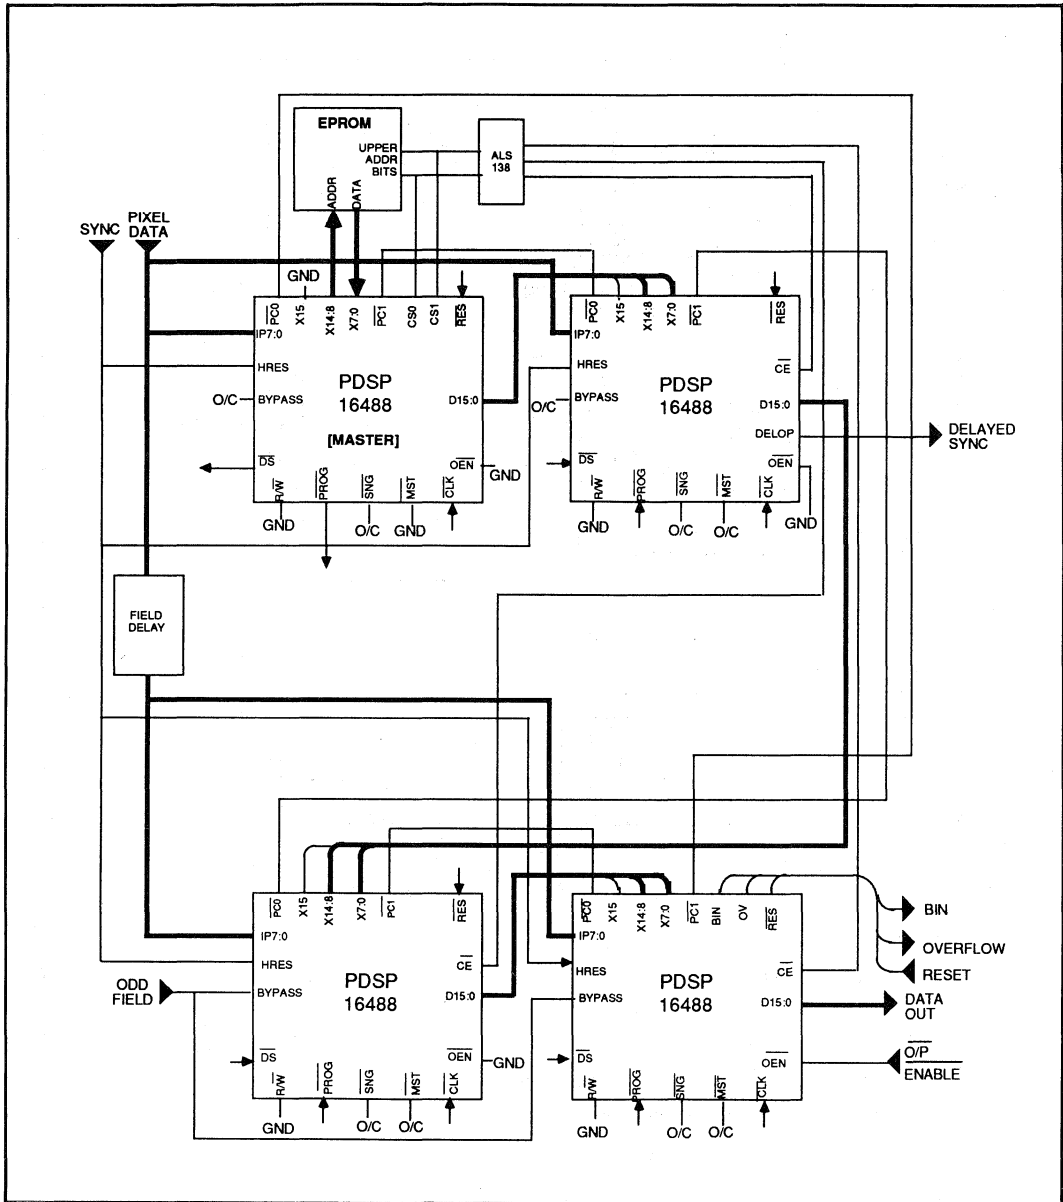


Figure 15. Four Device Interlaced System.

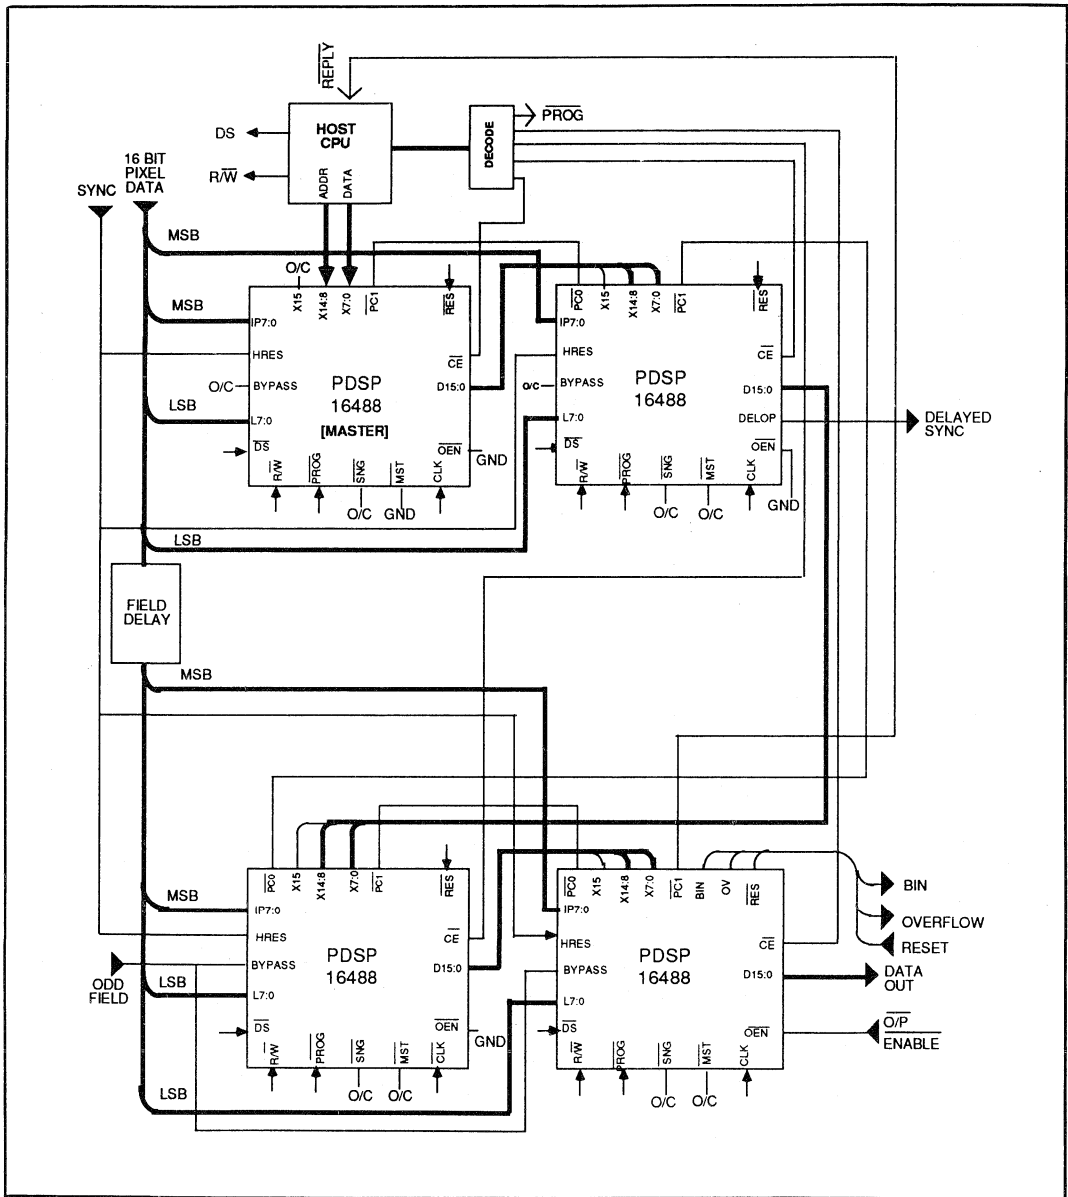


Figure 16. Four Device System with 16 Bit Pixels

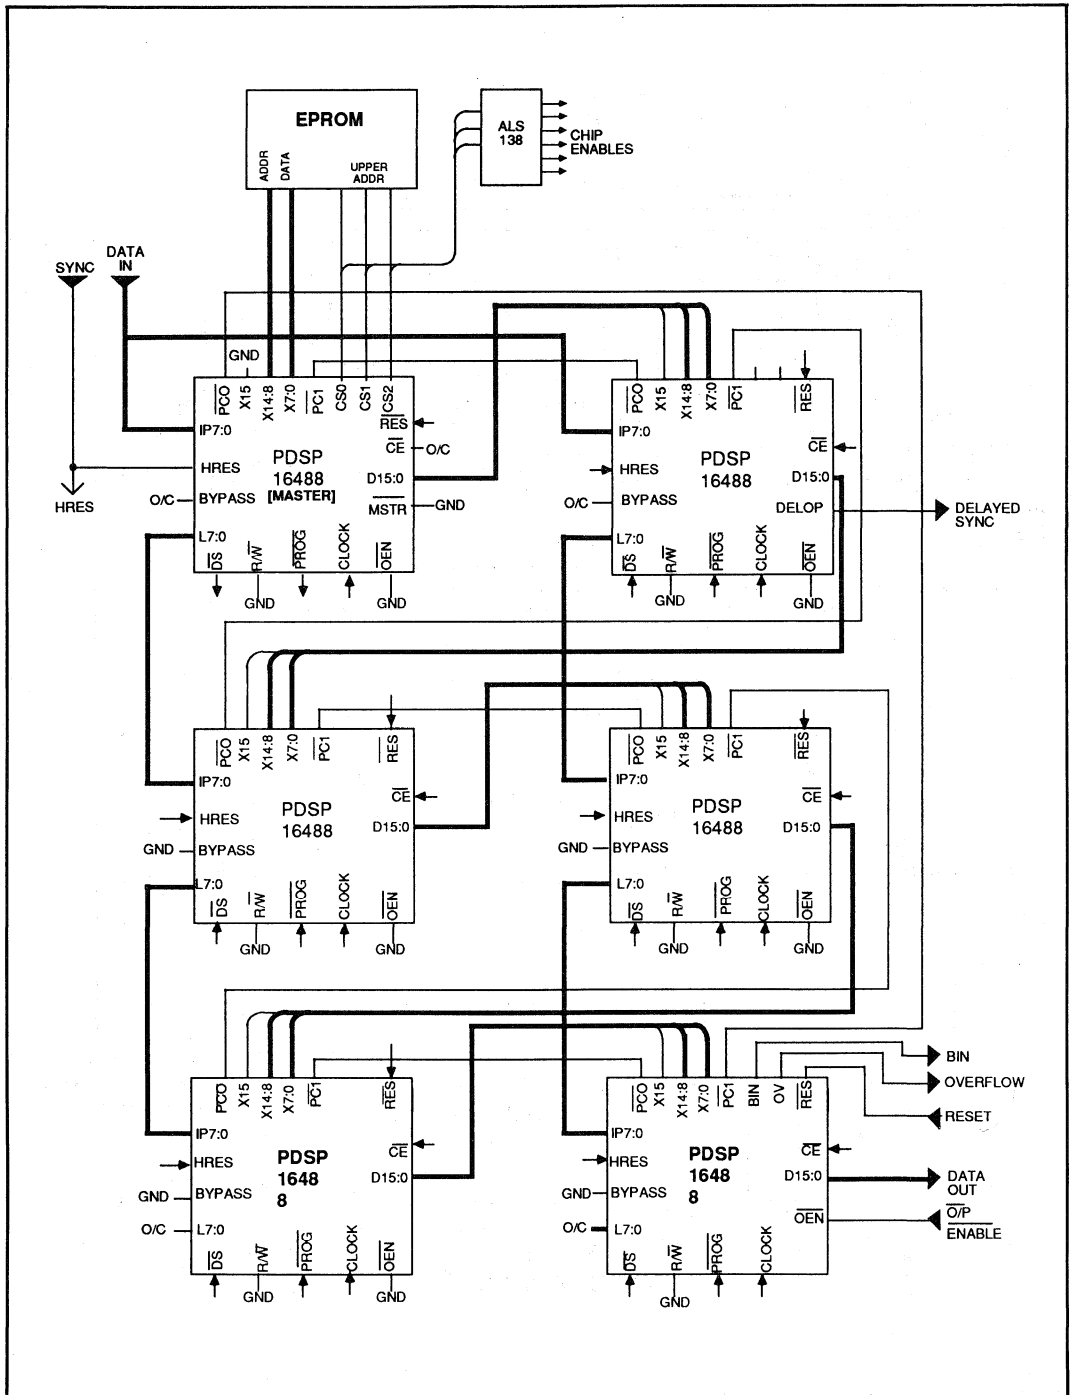


Figure 17. Six Device Non Interlaced System.

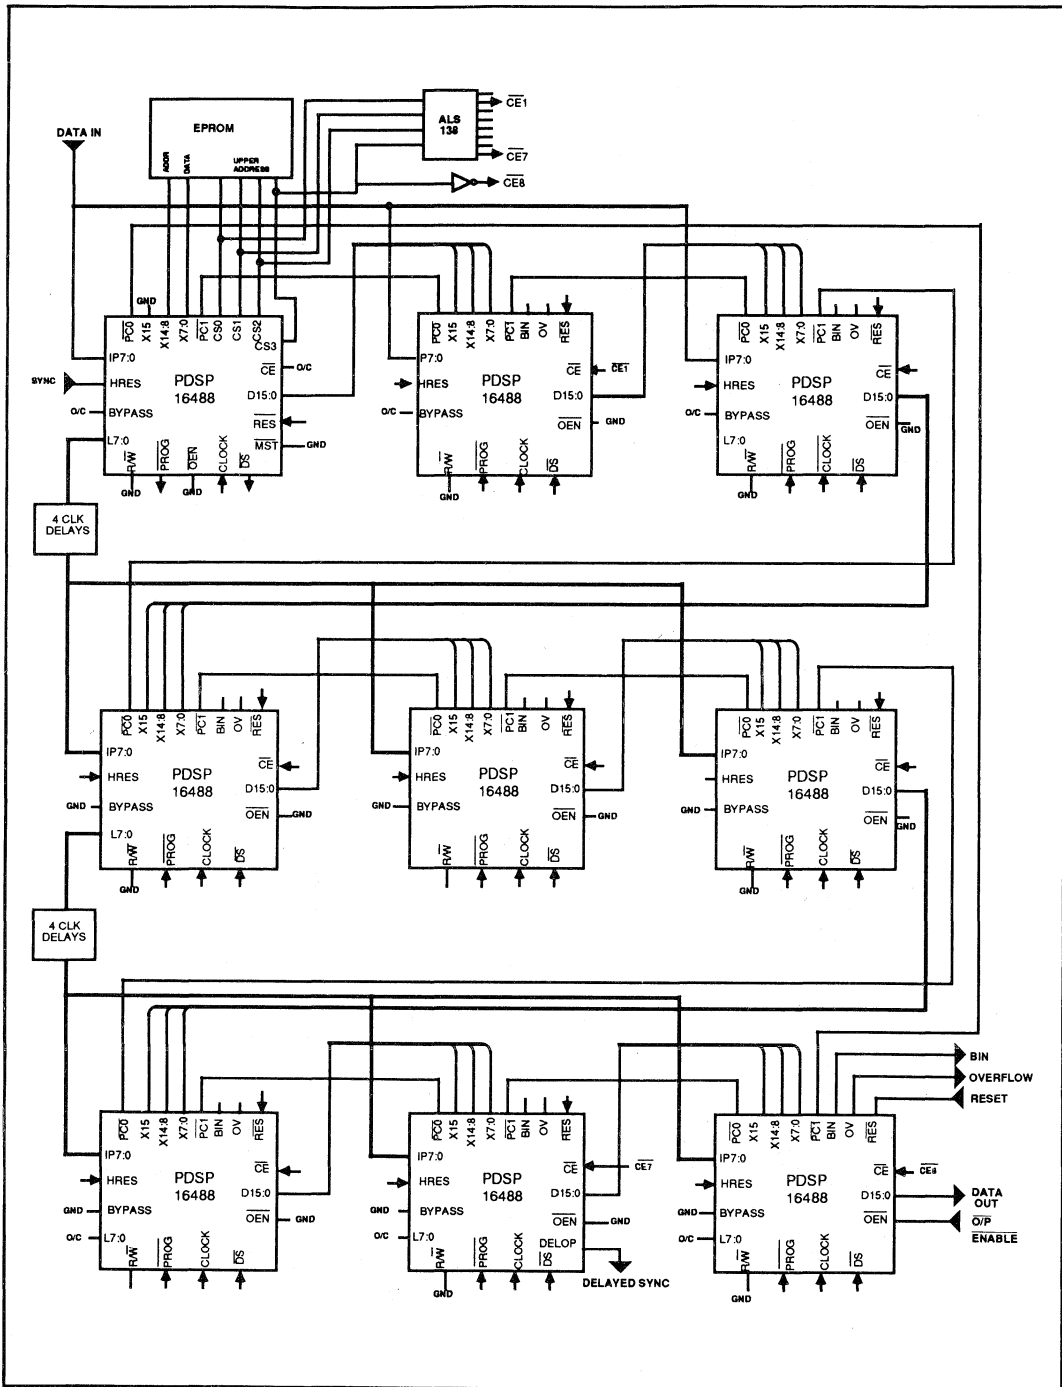


Figure 18. Nine Device Non Interlaced System.



PDSP16510

STAND ALONE FFT PROCESSOR

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 asynchronous to the 40MHz system clock used for internal operations.

A 1024 point complex transform can be completed in 96µs, 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
- PDSP16488 2D Convolver (8 x 8)




Fig. 1. Block Diagram

FEATURES

- Completely self contained FFT Processor
- Internal RAM supports up to 1024 complex points
- 16 bit data and coefficients plus block floating point for increased dynamic range
- 450 MIP operation gives 96 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 package

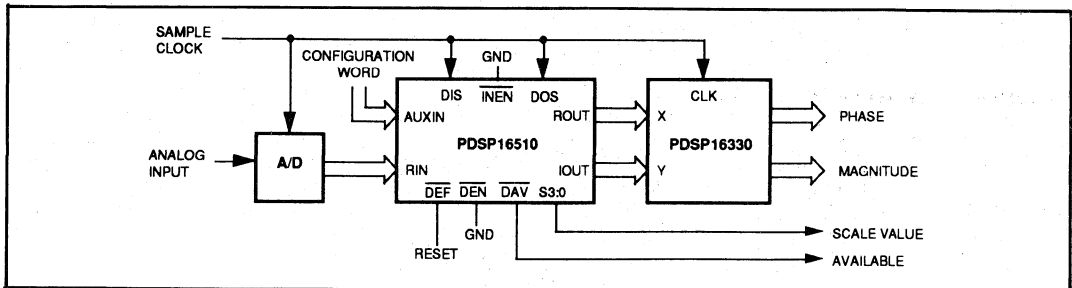



Fig. 2. Typical 256 Point Real Only System Performing Continuous Transforms



Pin Out Diagram - Bottom View

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

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 inactive 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 \overline{DAV} and \overline{DEN} are active. Otherwise they are high impedance.
I15:0	O	These pins output the imaginary component of the transformed data when \overline{DAV} and \overline{DEN} are active. Otherwise they are high impedance.
\overline{DEF}	I	The high going edge of \overline{DEF} is used to internally latch the contents of AUX15:0, which then define the operating mode. In the simplest system \overline{DEF} is a power on reset. When \overline{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 inactive after the complete, one quarter, or one half a data block has been loaded.
\overline{INEN}	I	The use of this input is mode dependent. It is either used as an 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 asynchronous input is used to load data into the device.
DOS	I	The rising edge of this asynchronous 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{DAV}	O	An active low signal that indicates that a transform is complete. Transformed data will then be outputted under the control of the DOS strobe, in normal sequential order. It may be optionally programmed to be delayed by 24 DOS strobes to match the delay through a PDSP16330.
\overline{DEN}	I	This input is used to enable the data dump operation when \overline{DAV} has gone active. If it is tied low the device will automatically dump data when \overline{DAV} goes active. Otherwise the device will wait for the enabling signal before the dump operation commences.
VDD	P	Four +5V pins
GND	P	Four ground pins

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

PDSP16510

when a control signal \overline{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 \overline{DEF} is active. \overline{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 comprises 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 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 will NEVER occur, however, if unit input data is purely real or imaginary, or if a unit amplitude complex input has only positive components. 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

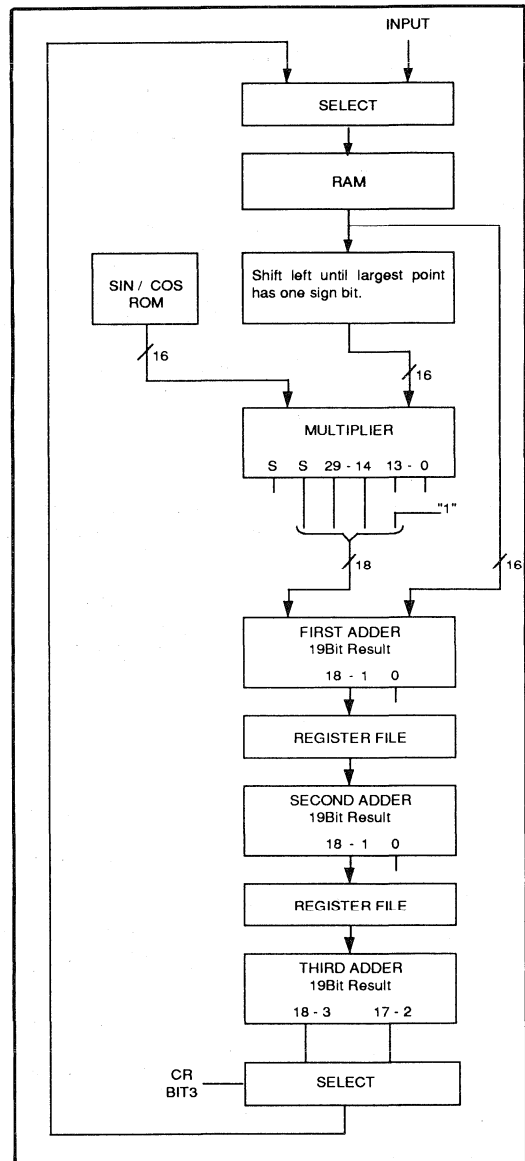


Fig. 3 One of Four Data Paths

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 RAM has been designed for use in a wide variety of applications. The provision of an asynchronous input strobe (DIS), allows data to be loaded without the need for additional external buffering. An asynchronous output strobe (DOS) also allows transformed data to be dumped with the sampling clock, this being particularly useful when the device is performing the inverse transform back to the time domain. 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 asynchronous 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. The internal RAM organi-




Fig. 4. RAM Organization with 256 Data Points




Fig. 5. RAM Organization with 1024 Point Transforms

sation is shown in Fig. 4. It should be 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 maximum transfer rate, which can be supported by the input and output circuits, ultimately limits the data sampling rate. When load and dump operations are separate

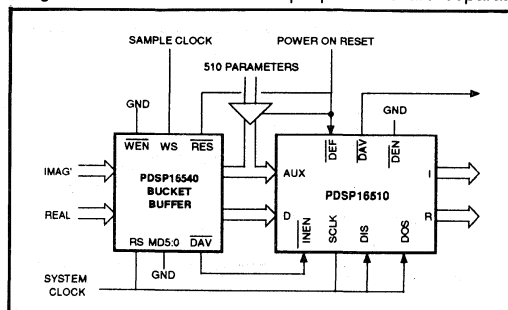


Fig. 6. 1024 Point Transforms with I/P Buffer

operations, which 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, \emptyset . When all three operations are concurrent the sampling rate, S , is reduced by a factor F , where F is defined as below if \emptyset is in MHz and L is the system clock low time in nanoseconds;

$$S = F\emptyset, \text{ where } F = \frac{4}{6 + 0.001\emptyset L}$$

LOADING DATA

Data loading is controlled by three signals; DIS an input strobe, \overline{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.

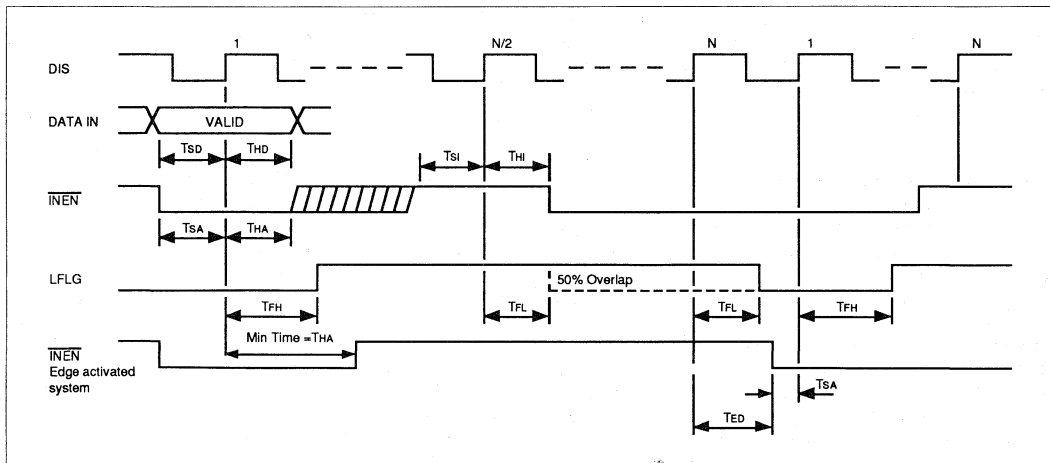
The data strobe is used to load data into the internal

workspace RAM, and data must meet the specified set up and hold times with respect to its rising edge. This strobe can be asynchronous to the system clock used internally, and the device will perform the necessary internal synchronisation. 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 \overline{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 Bit 12 is set in a SINGLE device system the \overline{INEN} signal is simply used as an enable for the DIS strobes. When \overline{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, \overline{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



Characteristic	Symbol	Min	Max	Units	Conditions
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} inactive Hold Time to ensure no load	T_{HI}	2		ns	
\overline{INEN} inactive going set up for no load operation	T_{SI}	8		ns	
Delay to LFLG going active	T_{FH}		10	ns	30pF
Delay to LFLG going in-active	T_{FL}		10	ns	30pF
Min time to \overline{INEN} low in a single device system	T_{ED}	15		ns	

Table 1. Timing Information with Continuous Inputs.

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 $\overline{\text{INEN}}$ 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 $\overline{\text{INEN}}$ going low. Loading will cease when a complete block (or group of blocks with multiple concurrent transforms) of data has been loaded, even if $\overline{\text{INEN}}$ remains low. $\overline{\text{INEN}}$ 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 $\overline{\text{INEN}}$ 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, $\overline{\text{INEN}}$ can be taken high to inhibit the load operation when gaps occur in the data stream. In the $\overline{\text{INEN}}$ edge activated mode gaps in the data stream can only be accommodated if the DIS clock is externally inhibited. Taking $\overline{\text{INEN}}$ 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 previously.

When Control Register Bit 12 is set in any multiple device mode, the $\overline{\text{DEF}}$ 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 $\overline{\text{DEF}}$ goes in-active. The other devices need the $\overline{\text{INEN}}$ 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 $\overline{\text{INEN}}$ has gone low. In the simple $\overline{\text{INEN}}$ 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 $\overline{\text{DEF}}$ is active and will go high in response to the first $\overline{\text{INEN}}$

enabled DIS edge after $\overline{\text{DEF}}$ has gone in- active.

The active going LFLG edge does not normally have any system significance, but in the block overlapping modes the inactive going edge will occur when 50% or 75% of the data has been loaded. By driving the $\overline{\text{INEN}}$ 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 asynchronous output strobe [DOS], a dump enable signal [$\overline{\text{DEN}}$], and a Data Available signal [$\overline{\text{DAV}}$]. The $\overline{\text{DAV}}$ signal is used to indicate that the internal output buffer contains transformed data, and the $\overline{\text{DEN}}$ 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 four DOS strobes once a transform is complete in order to transfer data to the output pins. These DOS strobes must be enabled by the $\overline{\text{DEN}}$ input, unless the device has been configured in one of the multiple device modes (See section on Multiple Device Systems).

The state of the $\overline{\text{DEN}}$ input at the end of a transform is used to control the transition of the active going edge of the $\overline{\text{DAV}}$ 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 $\overline{\text{DEN}}$ input is tied low in a single device system, the active going $\overline{\text{DAV}}$ transition will be internally synchronised to the rising edge of a DOS clock. If $\overline{\text{DEN}}$ 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 $\overline{\text{DEN}}$ to go high whilst $\overline{\text{DAV}}$ is active, if this $\overline{\text{DAV}}$ synchronous mode is needed.

In this $\overline{\text{DAV}}$ synchronised mode the first rising edge of the DOS clock, after $\overline{\text{DAV}}$ 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 $\overline{\text{DAV}}$ transition, since DOS must be a continuous clock, and there is then no delay before the first output becomes valid. The $\overline{\text{DAV}}$ 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. $\overline{\text{DAV}}$ 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 $\overline{\text{DAV}}$ goes active. If this is not the case the $\overline{\text{DAV}}$ output will not go active at the correct time, and the internal output circuitry will not be primed. Once $\overline{\text{DAV}}$ is active, however, it is possible for DOS to be irregular, and $\overline{\text{DEN}}$ 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 $\overline{\text{DEN}}$ remains low once $\overline{\text{DAV}}$ goes in-active.

If $\overline{\text{DEN}}$ is not active in a single device when the transform

is complete, then the device will wait for $\overline{\text{DEN}}$ 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 $\overline{\text{DAV}}$ 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 $\overline{\text{DEN}}$ input after $\overline{\text{DAV}}$ has gone active. $\overline{\text{DEN}}$ then functions as an enable for the host produced data strobes on the DOS pin. $\overline{\text{DEN}}$ may either stay active for the complete transfer, or may be used to enable each DOS input. When $\overline{\text{DEN}}$ and DOS are both active an internal read operation occurs, and an address generator is incremented. $\overline{\text{DAV}}$ goes inactive in response to the DOS edge needed to read the last output, unless Bit 15 in the Control Register is set. In this case $\overline{\text{DAV}}$ goes inactive when the next $\overline{\text{INEN}}$ edge is received for reasons given later.

Results are transferred from the device with the rising edge of the DOS strobe when $\overline{\text{DEN}}$ 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 a further three 'dummy' strobes are needed each time $\overline{\text{DAV}}$ goes active in order to prime the output

circuitry. The actual output sequence is given in Table 3, and illustrates that four $\overline{\text{DEN}}$ enabled DOS strobes are needed before the first frequency bin appears on the output pins. This is then read by the host with the fifth DOS strobe. $\overline{\text{DAV}}$ does not go inactive until the DOS edge after the last bin appeared on the output pins.

In addition to the above requirements it is necessary to provide at least four DOS strobes after $\overline{\text{DEF}}$ has gone inactive, but before $\overline{\text{DAV}}$ goes active. These initialize the internal address counters and do not rely on $\overline{\text{DEN}}$ also being active. They are needed every time $\overline{\text{DEF}}$ has been used to change the operating mode.

The tri-state drivers on the output busses are only enabled when both $\overline{\text{DAV}}$ and $\overline{\text{DEN}}$ are active. When $\overline{\text{DEN}}$ is tied permanently, low the output bus will start to become valid from the DOS edge which also generates the $\overline{\text{DAV}}$ output. The next DOS edge can then be used to transfer the first output to the next device. When $\overline{\text{DEN}}$ is driven low in response to the $\overline{\text{DAV}}$ output, the outputs start to become valid when $\overline{\text{DEN}}$ 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 $\overline{\text{DEN}}$ goes inactive, then both the 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 $\overline{\text{DAV}}$ going active until $\overline{\text{INEN}}$ again goes active, and a new set of data is loaded. The output tristate drivers, however, normally go high impedance when $\overline{\text{DAV}}$ goes inactive once a dump operation has been completed. In order to

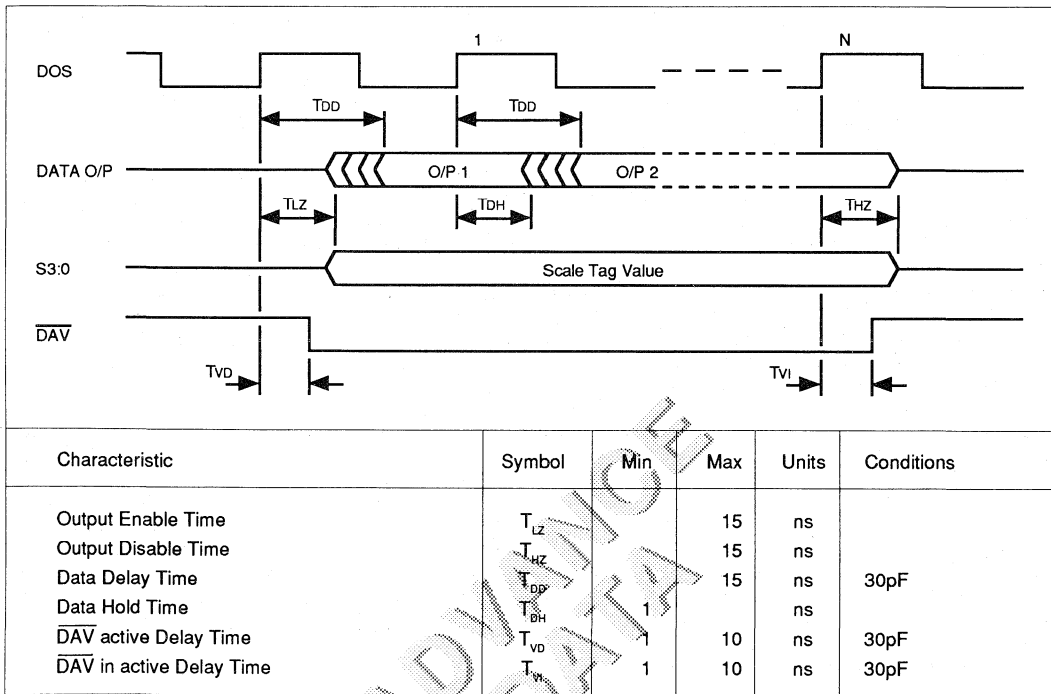


Table 2. Output Timing with $\overline{\text{DEN}}$ tied low.

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 \overline{DAV} output active until a further \overline{INEN} edge is received. The user must then keep track of how many outputs have been dumped before \overline{INEN} is generated to start a new load operation.

The \overline{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 \overline{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 \overline{DEN} is active, and not when \overline{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.

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 \overline{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 \overline{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 \overline{INEN} edge activated mode (Bit 12 reset) is used, and the inverted \overline{DAV} edge is used to drive the \overline{INEN} input. This then initializes a new load operation only when the previous dump has been completed. In such a system the \overline{INEN} edge will be asynchronous to the DIS strobe, and the set up time given in Table 1 may not be obeyed. This will simply cause an extra input sample to be possibly ignored, but will not cause data corruption.

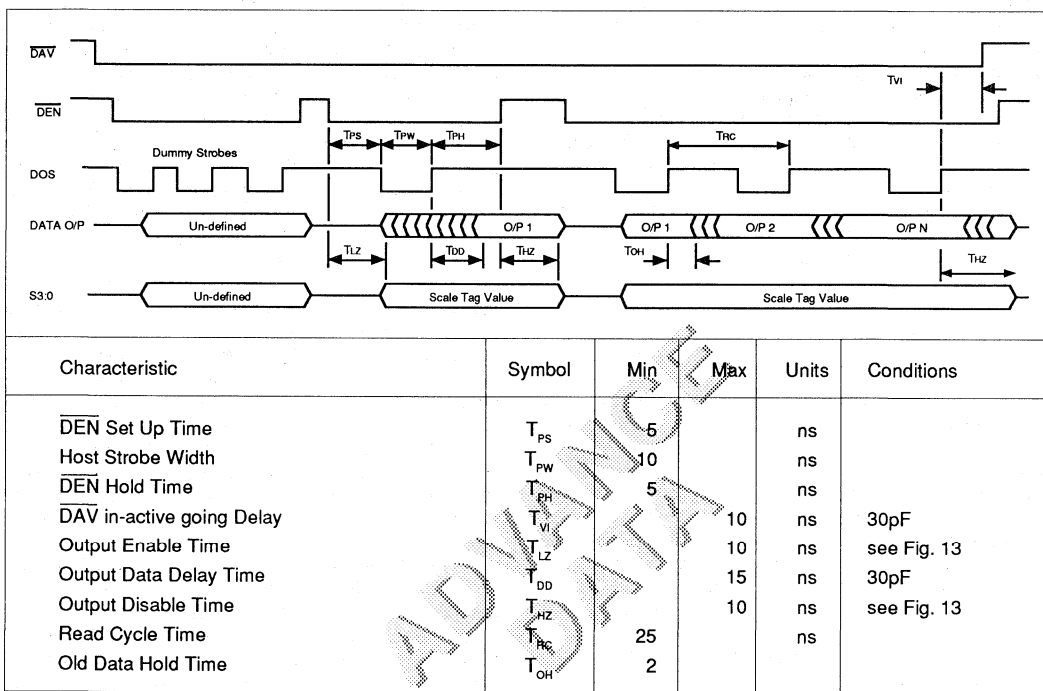


Table 3. Host Controlled Output Timing

PDSP16510




Fig. 7. Host Controlled System

FULL CO - PROCESSOR OPERATION

A single device to 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 \overline{DEN} is a host provided enable for host read operations, and \overline{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 \overline{INEN} , which is then automatically transformed. The \overline{DAV} output provides a flag indicating that the transform is complete, and results are then read by the host using DOS enabled by \overline{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 pro-

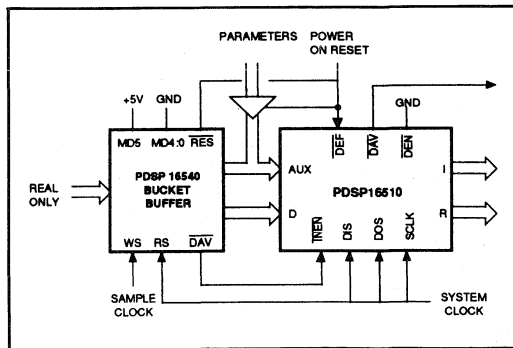


Figure 8. 1024 Point Real Transforms

duces 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.

In a single device system, performing non overlapped transforms on data from a SINGLE source, only the Real input pins are used, and the Imaginary inputs are redundant except when configuring the device. If block overlapping is needed it will be necessary to load pairs of data blocks simultaneously, using both the real and imaginary inputs. An external FIFO is then needed to provide a simple delay for a block of data (or 4 blocks in the 8 x 64 mode), the output of which must provide data for the real inputs. Continuous inputs can still be accepted, and each block (or 4 blocks) will initially occur on the imaginary inputs, and then occur again on the real inputs as an output from the FIFO. The data output sequence will consist of the results from a pair of inputs, followed by the results obtained after the required overlap. Thus with 50% overlapping the sequence is 1 & 2 followed by 1.5 & 2.5 followed by 3 & 4 followed by 3.5 & 4.5 etc, where 1 2 3 4 are the sequential inputs to the external FIFO, 1.5 is the overlap between 1 and 2, and 2.5 is the overlap between 2 and 3.

When any block overlapping mode is selected the device will be automatically configured to expect pairs of inputs, and explicit Control Register bits are not needed. If data is not from a single source, any block overlapping must be handled by an external buffer, and not by the device itself.

With no overlapping the device will normally expect data to only occur on the real inputs, as previously stated. By setting Control Register Bits 8:6 to 101, however, it is possible for a single device to accept data from two independent sources using the real and imaginary inputs. Maximum sampling rates will then only be half those possible when a single source is used, if no incoming data is to remain un-processed. With two sources a transform must be completed in the time to load parallel blocks, otherwise incoming data will be lost. With one source a transform need not be finished until two data blocks have been acquired. In this dual input mode results from data on the real inputs always precede those from the imaginary inputs.

When eight simultaneous 64 point transforms are performed, the sampling rates given in Table 5 assume that data is from a common source. The data output sequence will be 1 5 2 6 3 7 4 8 corresponding to inputs 1 through 8 in normal order from a single source. When data is from two sources the sampling rates will be halved, and the output sequence will be 1A 1B 2A 2B 3A 3B 4A 4B, where A and B are the dual simultaneous sources on the real and imaginary inputs respectively. If data block overlapping is used in either of the above cases, the eight outputs will be followed by results from the same basic eight blocks but time displaced to give the required overlap. If more than two sources are to be handled the user must provide appropriate buffering and multiplexing, and the sampling rates must be proportionally reduced.

When two 1024 point transforms are performed with a single device, on data from a single source, the input buffer must be arranged to acquire two blocks before initialising a transfer to the device. In order to improve the maximum sampling rates possible, data should be read simultaneously from each half of the buffer, and loaded into the real and imaginary inputs. This halves the transfer time from the buffer to the device, but requires the device to expect dual inputs. Thus if block overlapping is not needed Control Register Bits 8:6 should be set to 101.

This fast transfer mode is supported by a special option on the PDSP16540 Bucket Buffer. It will acquire two 1024 point non overlapping blocks using the sampling clock, and then transfer the results to the FFT processor at the full system clock rate. Figure 8 shows the system arrangement. It does not support block overlapping.

With 1024 point transforms all block overlaps are handled by the buffer logic, and not by the internal RAM, but the device must still be programmed to expect the required overlap if the external buffer makes use of the inactive LFLG edge to mark the overlap point. To achieve the performance given in Table 5 with 50% overlaps, the buffer must provide sufficient storage for at least 2.5 data blocks. With 75% overlaps it must provide storage for 2.75 blocks. This extra storage allows transfers between devices to be only needed when a complete new block has been acquired for 50% overlaps, and when half a new block has been acquired for 75% overlaps. If storage is restricted to two data blocks, only half the sampling rates given will be possible. Transfers between devices must then occur when a half or a quarter of as new block has been acquired. Since the minimum time between transfers must be no less than the transform time itself, the sampling rates must be proportionally reduced to prevent loss of data.

SINGLE DEVICE SAMPLING RATES

In a single device system the maximum sampling rate is dependent on the transform size, the data overlap, and whether real or complex data is applied. Table 4 gives the times taken to complete the transforms for the various block

Configuration	Clock Periods
16 X 16PT COMP	420
4 X 64PT COMP	624
256PT COMP	816
1024PT COMP	3907
8 X 64PT REAL	816
2 X 256PT REAL	1032
2 X 1024PT REAL	4699

Table 4. Computation Times in Clock Periods

sizes, which include an allowance for synchronisation between the DIS strobe and the system clock. If continuous data is to be transformed, the time to acquire a new block of data (or partial block with overlapping) must be at least equal to these transform times. Load and dump times must also be added in the 1024 modes. For non continuous transforms the sampling rate is only limited by the maximum input / output rates as discussed 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;

$$nS \text{ or } \frac{nS}{2} \text{ or } \frac{nS}{4} \text{ must be greater than or equal to } PK + 4W$$

where n is the transform size, S is the input sampling rate, 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.

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;

$$(n - 4)S \text{ must be greater than or equal to } PK$$

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. 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

16 X 16 COMPLEX	4 X 64 COMPLEX	256 COMPLEX	1024 COMPLEX	8 X 64 REAL	2 X 256 REAL	2 X 1024 REAL
0% 50% 75%	0% 50% 75%	0% 50% 75%	0% 50% 75%	0% 50% 75%	0% 50% 75%	0% 50% 75%
23.9 - -	16.1 8.0 4.0	12.3 6.1 3.0	6.8 3.4 1.7	24.6 12.3 6.1	19.5 9.7 4.3	12.1 6.0 3.0

Table 5 Sampling Rates in MHz Obtained with a Single Device and various overlaps.

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:

$$1024S \text{ or } 512S \text{ or } 256S > 1024B + PK + D$$

for 0%, 50%, or 75% overlapping respectively. S, P, and K were defined previously. 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

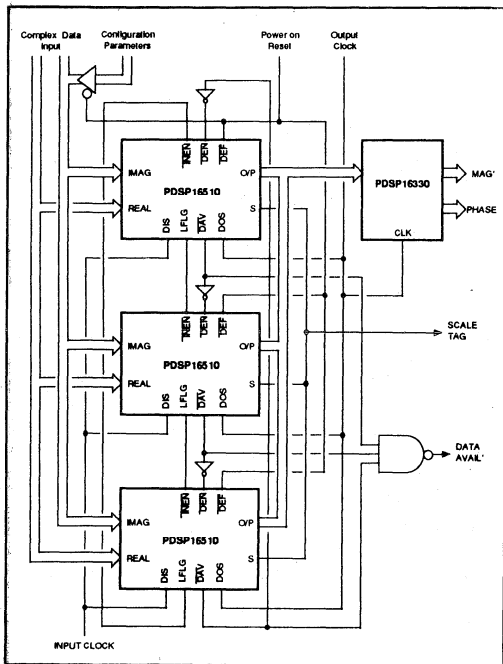


Figure 9. Multiple Device Configuration

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 inactive 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 third rising DIS edge after DEF has gone inactive. 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 ie 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 $\overline{\text{DAV}}$ output from one device to drive the $\overline{\text{DEN}}$ 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 $\overline{\text{DEN}}$ input from the inverted $\overline{\text{DAV}}$ output from a previous device requires a change to the single device $\overline{\text{DAV}}$ and $\overline{\text{DEN}}$ operation. If $\overline{\text{DEN}}$ is active at the end of a transform in a multiple device system, the $\overline{\text{DAV}}$ output will go active when the output circuit has been primed by the

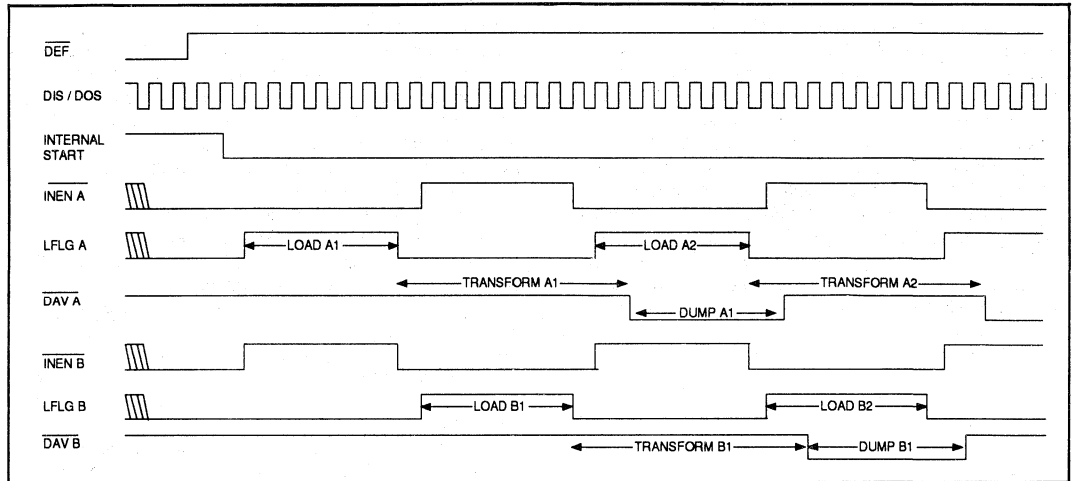


Figure 10. Two Device System with Concurrent Load, Transform, and Dump Operations

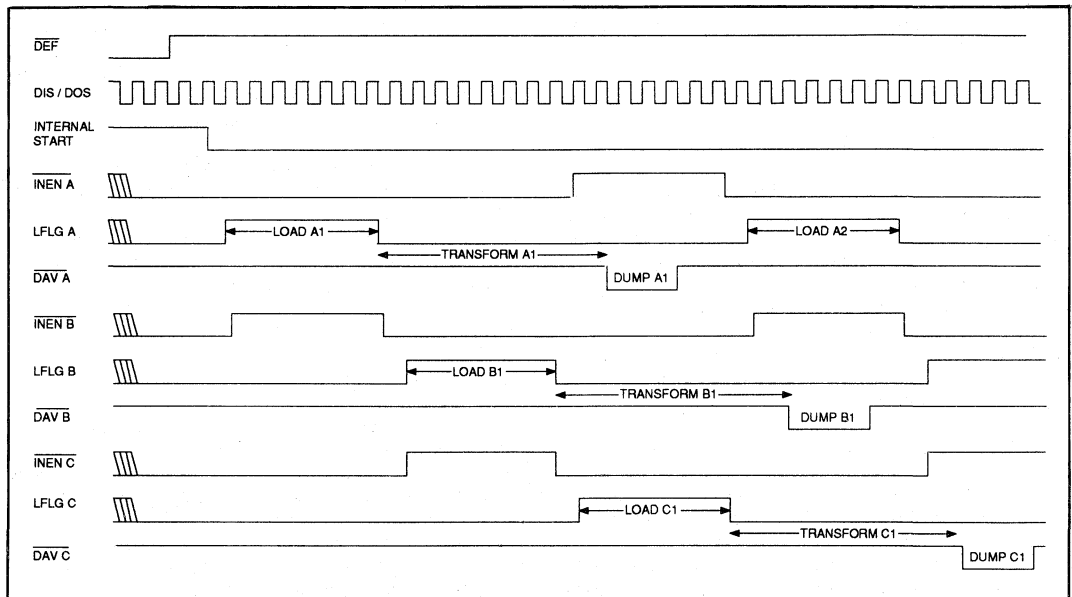


Figure 11. Three Device System with Separate Load, Transform, and Dump Operations

DOS strobes. This operation is identical to that provided for a single device system, and is transparent to the user as long as \overline{DEN} and DOS are active. If \overline{DEN} is not active, however, the \overline{DAV} output will not asynchronously go active as happens in a single device system. Instead \overline{DAV} will only go active when \overline{DEN} eventually goes active. Since \overline{DEN} is the inverted \overline{DAV} output from a previous device, it is thus never possible for two devices to be actively outputting data. The \overline{DAV} active going edge remains synchronised to the DOS strobe since the \overline{DEN} input will only go active when a previous \overline{DAV} goes inactive. A further change to the output circuitry ensures that the output buffer is primed even though \overline{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 occur, until \overline{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 \overline{DAV} output by OR'ing together all the individual, active low, \overline{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 \overline{DAV} goes inactive 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 inactive going \overline{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 affect 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, and Figure 10 shows a typical timing sequence. 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 $F\emptyset$ where F and \emptyset are as defined previously. Within this restriction the theoretical performance is given by;

$$NnS > PK+4W, \text{ or } 0.5NnS > PK+4W, \text{ or } 0.25NnS > PK+4W$$

for 0%, 50%, or 75% overlapping. N is the number of devices and n , S , P , K , and W were defined for a single device.

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 the system clock.

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 inactive after 25%, 50%, or 100% of the block has been loaded. When multiple transforms are performed concurrently (for example 4×64) 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 \overline{DAV} from the first device must drive the \overline{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 $\overline{DAV}/\overline{DEN}$ 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 $\overline{DAV}/\overline{DEN}$ 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. As with Mode 1, additional output processors are needed in order 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:

$$\begin{aligned} NnS &> nS + PK + D \text{ for no overlapping} \\ NnS &> 2X [nS + PK + D] \text{ for 50\% overlapping} \\ NnS &> 4X [nS + PK + D] \text{ for 75\% overlapping} \end{aligned}$$

where N, n, S, P, K, and D are as previously defined. The load time can be derived from the number of samples multiplied by the sample clock period. The dump time is the number of required outputs (plus four as discussed previously) multiplied by the output clock period. The latter is any value defined by the user, down to the minimum system clock period.

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. 11 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 overlapping groups of transforms are done simultaneously. It is used, for example, when 4 x 64 transforms are overlapped by 50% and simultaneously transformed. LFLG will go inactive after a complete group of data blocks have been loaded, regardless of the overlap selected. The device, however, continues to load a further block, or pair of blocks with real transforms. Thus, for example, in the 4 x 64 complex 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.

The full benefits are only obtained if more than one output processor is provided, but an extra processor is not needed for every device. Sampling rates up to the system clock rate are possible. The equations defining the sampling rates become:

$$\begin{aligned} (N - 1)L &> 2PK + 2D \text{ for 50\% overlaps} \\ (N - 1)L &> 4PK + 4D \text{ for 75\% overlaps} \end{aligned}$$

where L is the time needed to load ALL the blocks which are simultaneously transformed, but not the extra block.

When real transforms are to be performed an external FIFO is needed to provide pairs of data blocks, which are loaded simultaneously into the real and imaginary inputs. The FIFO produces a delay of 256 sampling clocks in either the 2 x 256 or 8 x 64 modes.

The amount of internal RAM prohibits the use of this mode when performing overlapped 2 x 1024 point real transforms. The 1024 point complex mode should then be used with the

imaginary inputs forced to logical zero after DEF has gone inactive.

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.

When a real transform is specified by Bits 2:0 together with 50% or 75% overlapping, two blocks of real data must be loaded simultaneously using both the real and the imaginary inputs. With no overlapping only the real input must be used unless Bits 8:6 = 101. This option allows both inputs to be

PDSP16510

used, and supports data from two independent sources. Two additional decodes allow the DIS input to be divided by two or four, with 50% and 75% overlapping respectively. These options allow the DOS and DIS input pins to be supplied from a common source. The frequency of this source would be dictated by the output rate requirement, which must be greater than the input rate when data block overlapping is needed and there is only one output processor..

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

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
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

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

When this bit is set the PDSP16510 will not generate \overline{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 \overline{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 \overline{INEN} input is a simple load enable signal. When it is reset an \overline{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 \overline{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 \overline{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, \overline{DAV} will remain valid until the next \overline{INEN} input, and will indicate that the transformed data still remains in the internal buffer. As soon as the next \overline{INEN} is received the transformed data will be overwritten. Whilst \overline{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 ie 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

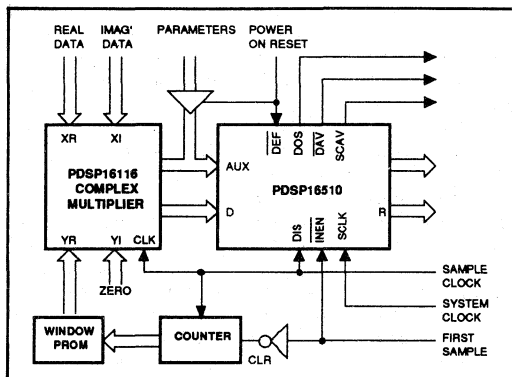


Fig. 12. External Window Generator

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 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;

$$A - B\cos x + C\cos 2x \text{ where } x = \frac{2\pi n}{N} \quad n = 0 \text{ to } N-1$$

For the Hamming window A = 0.54, B = 0.46, C = 0

For the Blackman-Harris window, 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. 12 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.

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

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 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 [ie 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 affect whether the smaller signal 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.

ABSOLUTE MAXIMUM RATINGS [See Notes]


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 pin I_C (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 θ_{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 reliability.
4. Current is defined as positive into the device.

STATIC ELECTRICAL CHARACTERISTICS
Operating Conditions (unless otherwise state)

$T_{AMB} = 0^\circ C$ to $+70^\circ C$.
 $V_{CC} = 5.0V \pm 10\%$



ORDERING INFORMATION

PDSP16510 C0 AC (Commercial - PGA Package).
Call for availability on HIGH Reliability parts and MIL-STD-883c screening.

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	$I_{OH} = 4mA$ $I_{OL} = -4mA$
Output low voltage	V_{OL}	-		0.4	V	
Input high voltage	V_{IH}	2.0		-	V	$GND < V_{IN} < V_{CC}$
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	μA	
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	μA	$GND < V_{OUT} < V_{CC}$ $V_{CC} = Max$
Output S/C current	I_{SC}	10		300	mA	

SWITCHING CHARACTERISTICS

Characteristic	Symbol	Min	Max	Units	Conditions
Clock Frequency	\emptyset	DC	40	MHz	Less than 1024 points or Multiple Devices Mode 1. 1024 points or Multiple Devices Modes 2 and 3.
Clock High Period	T_{CH}	13		ns	
Clock Low Period	T_{CL}	10		ns	
Max DOS, DIS Frequency	\emptyset_D		F \emptyset	MHz	
Note $F = \frac{4}{6 + 0.001\emptyset T_{CL}}$					
Max DOS, DIS Frequency	\emptyset_D		\emptyset	MHz	



PDSP16520

QUAD - PORT SYNCHRONOUS RAM

The PDSP16520 contains 1K by 16 bits of Dual Port Static RAM with separate read and write address ports. All memory and I/O operations are synchronous to a user supplied system clock, with rates up to 20MHz. Two independent 16 bit input ports, plus two independent 16 bit output ports, allow simultaneous read and write operations on two data words every clock cycle. This structure is optimal for FFT processors and the throughput matches the requirements of the PDSP16112 or PDSP16116 Complex Multipliers when configured as butterfly processors.

The RAM is internally partitioned into four 256 by 16 bit blocks, with configurable routing between the blocks and the device I/O pins. Two devices provide the full memory requirements for a not-in-Place, 512 point, complex FFT transform. One device would be dedicated to real data, and the other to imaginary data, as shown in Fig. 8. The provision of register clock enables, and output tristate controls, allows the memory requirements of any transformation size to be supported with additional devices.

To simplify the design of pipelined systems, the PDSP16520 incorporates a user programmable delay on the write addresses and control lines. This provides compensation for the pipeline delay incurred by the processor performing the arithmetic calculations and allows one counter to provide both read and write addresses.

FEATURES

- 16K Registered RAM
- Configured as four 256 by 16 Dual-Port Blocks
- 20MHz Operation
- Simultaneous Read and Write operations
- Fully registered I/O and control for use in data flow architectures such as FFT processors
- Independent inputs and outputs plus separate Read and Write Address Ports
- User programmable delay between read and write addresses to compensate for processing delays
- 144 PGA Package

ASSOCIATED PRODUCTS

- PDSP1601** ALU & Barrel Shifter
- PDSP16112** 16 by 12 Complex Multiplier
- PDSP16116** 16 by 16 Complex Multiplier
- PDSP16318** Complex Accumulator
- PDSP16330** Pythagoras Processor




Fig. 1 Simplified Block Diagram

Signal	Description
A15:0	Input data bus A. Data presented to this input is latched by the rising edge of CLK
B15:0	Input data bus B. Data presented to this input is latched by the rising edge of CLK.
C15:0	Data output bus C. New data appears on this output after the rising edge of CLK.
D15:0	Data output bus D. New data appears on this output after the rising edge of CLK.
RA7:0	RAM read address bus. The address is latched by the rising edge of CLK. A 2 cycle pipeline delay occurs between latching the address and valid data appearing on the device output pins.
WA7:0	RAM write address bus. The address is latched by the rising edge of CLK. The address may be delayed by an amount defined by DL3:0.
$\overline{\text{WE}}4:1$	Active low write enables for RAM quadrants 4-1. These enables are latched by the rising edge of CLK, and the write occurs internally after the next rising edge. Their action may be delayed by an amount defined by DL3:0.
CLK	Device clock input. All I/O operations are initiated by the rising edge.
$\overline{\text{CE}}\text{A}$	A bus input register enable. When low new data is loaded to the input register on each rising clock edge. This signal is not internally latched.
$\overline{\text{CE}}\text{B}$	B bus input register enable. When low new data is loaded to the input register on each rising clock edge. This signal is not internally latched.
MUX7:4	Mux lines 4-7 control the input data multiplexers on RAM quadrants 1-4 respectively. A high level on the mux line routes data from the A input data register and a low level routes data from the B input data register. The mux control signals are latched by the rising edge of CLK.
DL3:0	Delay select pins. To facilitate the design of pipelined systems the PDSP16520 incorporates a user programmable delay between the read and write addresses. This internal delay may be programmed from 0 to 15 clock cycles and is loaded into the device on the rising edge of CLK. DL3 is the most significant bit.
MUX3:2	C output data select. Mux signals 3-2 allow the routing of RAM quadrant outputs 1-4 to the C15:0 output data bus. The mux signals are latched by the rising edge of CLK. See Table 1.
MUX1:0	D output data select. Mux signals 1-0 allow the routing of RAM quadrant outputs 1-4 to the D15:0 output data bus. The mux signals are loaded on the rising edge of CLK. See Table 2.
$\overline{\text{OEC}}$	Active low tri-state enable. This signal is latched by the rising edge of CLK and controls the C15:0 output 2 cycles later. This delay matches the pipeline delay of the device read operation.
$\overline{\text{OED}}$	Active low tri-state enable. This signal is latched by the rising edge of CLK and controls the D15:0 output 2 cycles later. This delay matches the pipeline delay of the device read operation.
$\overline{\text{RESET}}$	Active low power on reset. The clock must be stable before the end of reset.
RDY	Device ready signal. When the output goes high the device is ready for normal operation.




Fig. 2 Device Pinout Bottom View




Fig. 3 Block Diagram

FUNCTIONAL DESCRIPTION

When the input register enables are held low, data presented to the A and B input ports will be loaded into the respective input registers on the rising edge of the clock. Four input multiplexers allow the data from any port to be applied to any RAM quadrant. Each multiplexer has its own control pin; a logical one will route data from the A port into the associated RAM and a logical zero input will route data from the B port.

The read and write address ports specify the particular RAM location used within each quadrant. Each quadrant will receive the same read and write address. A read operation will always be performed when the clock is present, but the write operation to each quadrant is controlled by its own write enable signal. The internal write operation occurs one clock cycle after the write information is latched at the device

inputs. This assumes that no additional delays are present in the write path.

Four control inputs allow the user to select from 0 to 15 additional clock delays between specifying a write operation, and the operation actually occurring. The write address is also delayed by the same amount. When the PDSP16520 is used to support an FFT butterfly processor, the delay allows the user to compensate for the pipeline delay through the data path. This greatly simplifies the design of the address generator, since the address delay can be matched to the data delay.

The routing between the RAMs and the output ports is controlled by two output multiplexers as shown in Fig. 3. Table 1 and 2 show the relationship between the control pins and the quadrant selected. The data specified by a read address will

appear on the output pins after three clock edges (assuming the address met the required set up time with respect to the first edge). To the user the RAM thus appears to functionally contain three clock delays; caused by an input latch, an intermediate latch, and an output latch.

Old data is read before new data is written, if read and write addresses are the same, thus the earliest that data can be read is on the cycle following the write address. The data read will be present on the output pins after two further clock edges, as shown in Fig. 4.

The output multiplexer controls, together with output enables, are normally externally decoded from higher order read address bits. These signals are thus internally delayed by an amount which compensates for this three clock read delay. The possibility of differential delays between the use of various address bits is thus avoided. See Fig. 5.

DEVICE INITIALISATION

To ensure correct operation the device must always be reset after power up. The $\overline{\text{RESET}}$ signal is registered onto the chip and must therefore meet the normal setup and hold requirements. To reset the device signal $\overline{\text{RESET}}$ must be taken low for at least one clock cycle. Once $\overline{\text{RESET}}$ has returned high the device will indicate that it is ready for normal operation by taking signal RDY high. This will occur after 1541 clock cycles.

When $\overline{\text{RESET}}$ is taken low the device will respond by taking RDY low following the next clock rising edge.

M3	M2	RAM Quadrant
1	0	Q1
0	0	Q2
1	1	Q3
0	1	Q4

Table 1 Multiplexer Control for C Output Port

M1	M0	RAM Quadrant
0	0	Q1
1	0	Q2
0	1	Q3
1	1	Q4

Table 2 Multiplexer Control for D Output Port




Fig. 4 Input Timing Diagram




Fig. 5 Output Timing Diagram

APPLICATIONS

The PDSP16520 is ideally suited to be the data memory in an FFT application. In particular the PDSP16520 supports a constant geometry FFT butterfly system, implemented with the aid of the PDSP16116 and the PDSP16318. The system is illustrated in Fig. 8, and described in application note AN50.

Each quadrant provides a separate block of RAM which can accommodate, either the 16 bit real components, or the 16 bit imaginary components of the complex data. The two left hand quadrants are configured to store even data points, and the right hand quadrants store odd data points. Quadrants Q1 and Q3 accommodate data with locations greater than $(N/2-1)$, where N is the transform size. Quadrants Q2 and Q4 accommodate data with locations less than or equal to $(N/2-1)$. Fig. 6 shows how the RAM is internally configured through the input and output multiplexer controls.

The mode of addressing is simple with each quadrant receiving the same address, the address sequence remaining unchanged throughout each pass of the transform. Fig. 7 shows the storage of data points for a 16 point transform. The data points are stored in each of the four RAM quadrants at the addresses and in the quadrants indicated.

The 16 data points are addressed two at a time by an 8 term address sequence. When reading data from the RAM,

quadrants enabled alternate between left and right pairs, and the address sequence used is a simple increment. By using an $n+1$ bit counter where n is size of the required address, it is possible to use the least significant bit to control the output multiplexer and the most significant bits for the address. This count sequence on the address port will result in data points 0 and 8 read on the first cycle, 1 and 9 on the second cycle, 2 and 10 on the third cycle, 3 and 11 on the fourth cycle and so on.

When writing to the RAM the lower pair quadrants are enabled for the first half of the pass and the upper pair are enabled for the second half.

The write address sequence used can be obtained from the same counter used to supply the read address. The least significant bit controls the write enables. This results in data points 0 and 1 being written to the RAM on the first cycle, data points 2 and 3 written on the second cycle and so on. In the second half of the pass the same procedure occurs with the upper quadrants enabled. The write address must be delayed to match the delay through the arithmetic processor.

This addressing sequence described above repeats for each pass throughout the transform.




Fig. 6 Memory Configuration for Constant Geometry Algorithm

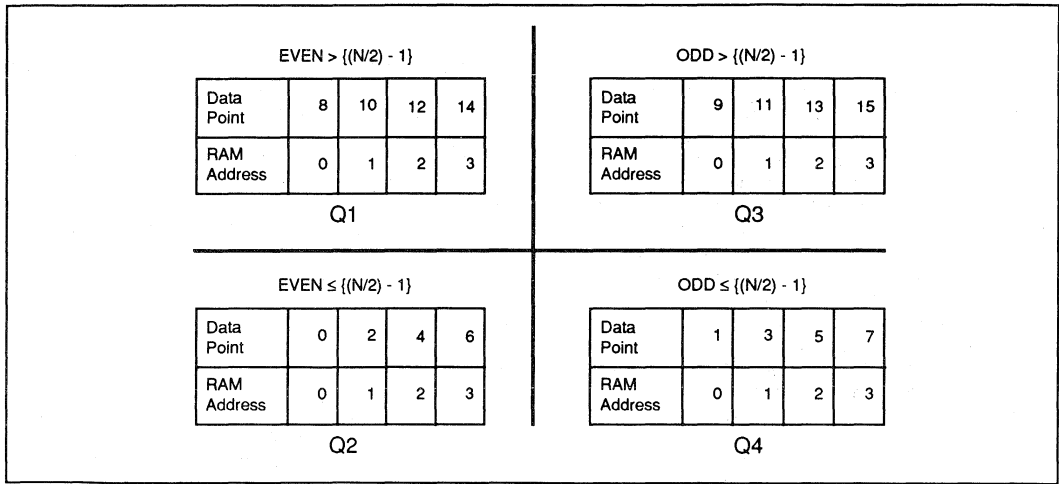


Fig. 7 Storage of Data Points for 16 Point Transform

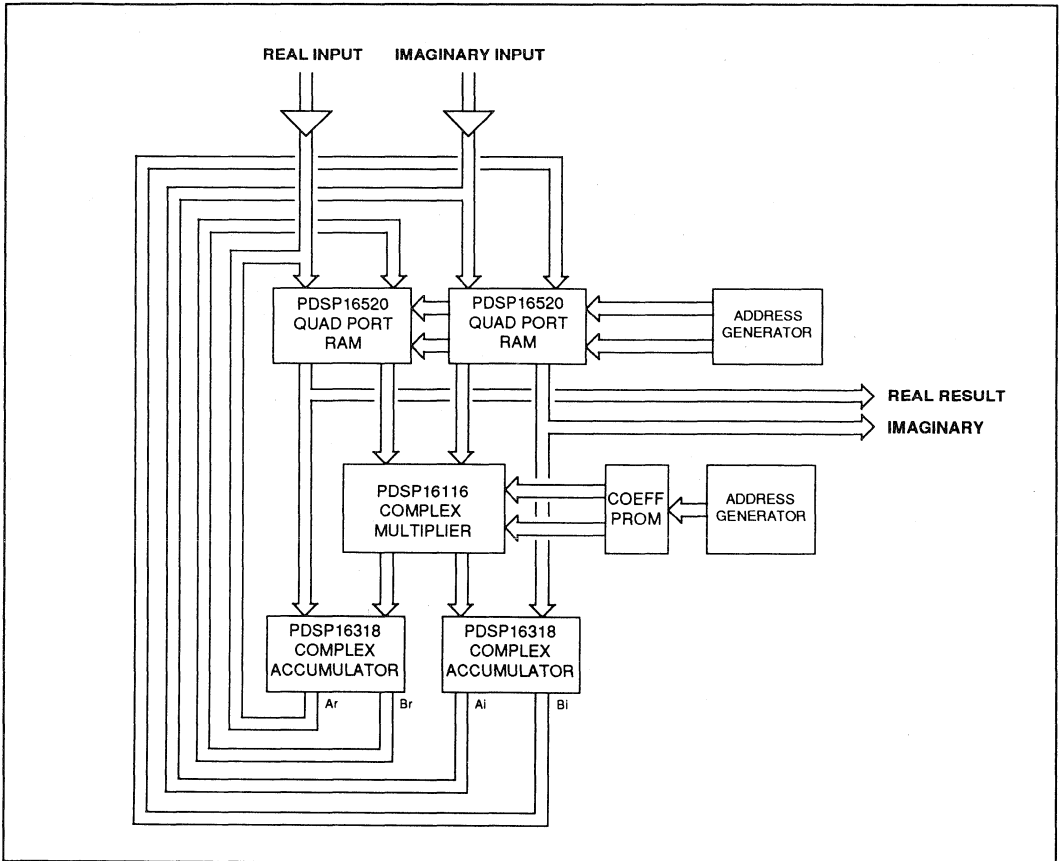


Fig. 8 Constant Geometry Butterfly Processor

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 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 θ_{JC}	5°C/W

NOTES

- Exceeding these ratings may cause permanent damage. Functional operation under these conditions is not implied.
- Maximum dissipation or 1 second should not be exceeded, only one output to be tested at any one time.
- Exposure to absolute maximum ratings for extended periods may affect device reliability.
- Current is defined as positive into the device
- $V_{CC} = \text{Max}$, Outputs Unloaded, Clock Freq = Max
- The θ_{JC} data assumes that heat is extracted from the top face of the package.

ELECTRICAL CHARACTERISTICS**Operating Conditions (unless otherwise stated)**

Industrial:	$T_{AMB} = -40^\circ\text{C}$ to $+85^\circ\text{C}$	$T_{J(MAX)} = 110^\circ\text{C}$	$V_{CC} = 5.0V \pm 10\%$	Ground = 0V
Military:	$T_{AMB} = -55^\circ\text{C}$ to $+125^\circ\text{C}$	$T_{J(MAX)} = 150^\circ\text{C}$	$V_{CC} = 5.0V \pm 10\%$	Ground = 0V

Characteristic	Symbol	Value			Units	Conditions
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	$I_{OH} = -4\text{mA}$ $I_{OL} = 4\text{mA}$
Output low voltage	V_{OL}	-		0.4	V	
Input high voltage	V_{IH}	2.0		-	V	GND < V_{IN} < V_{CC}
Input low voltage	V_{IL}	-		0.8	V	
Input leakage current	I_{IN}	-10		+10	μA	
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	μA	
Output S/C current	I_{OS}	10		300	mA	

SWITCHING CHARACTERISTICS

Characteristic	Symbol	Value				Units	Conditions
		Industrial		Military			
		Min	Max	Min	Max		
Clock period		50	-	50	-	ns	
Clock high time		20	-	20	-	ns	
Clock low time		20	-	20	-	ns	
A/B15:0 setup to clock rising edge	SETUP	10	-	15	-	ns	
A/B15:0 hold after clock rising edge	HOLD	0	-	0	-	ns	
RA7:0,WA7:0,WE4:1,MUX7:0,DL3:0, RESET setup to clock rising edge	SETUP	10	-	15	-	ns	
RA7:0,WA7:0,WE4:1,MUX7:0,DL3:0, RESET hold after clock rising edge	HOLD	0	-	0	-	ns	
CEA,CEB setup to clock rising edge		10	-	15	-	ns	
CEA,CEB hold after clock rising edge		0	-	0	-	ns	
OEC,OED setup to clock rising edge		10	-	15	-	ns	
OEC,OED hold after clock rising edge		0	-	0	-	ns	
Clock rising edge to C15:0,D15:0	Tout	5	20	5	25	ns	30pF load
Clock to O/P low impedance	Tiz	-	20	-	25	ns	See Fig. 9
Clock to O/P high impedance		-	20	-	25	ns	See Fig. 9
Vcc current	Icc	-	100	-	100	mA	See Note 5

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




Fig. 9 Three state delay measurement load.

ORDERING INFORMATION

PDSP16520 BOAC (Industrial - PGA Package)

PDSP16520 AOAC (Military- PGA Package)

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



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 FIFOs. This wide configuration supports the 16 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 unprocessed.

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 programmable 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 Package

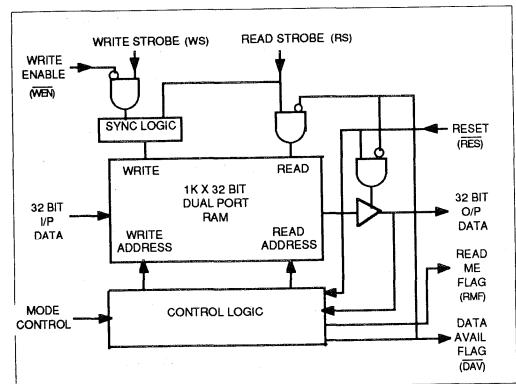



Figure 1. Simplified Block Diagram

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 inactive for one read strobe period every time new data is written to the RAM, and stays inactive 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 $\overline{\text{DAV}}$ 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 new words, 512, 256, or the number defined in groups of 32 words by D9:5 during reset. 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.
$\overline{\text{RES}}$	I/P	Reset, active low. 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 used for power on reset.
GND	I/P	Four ground pins. All must be connected
VCC	I/P	Four +5 volt pins. All must be connected



Pin Out Diagram - Bottom View

FUNCTIONAL DESCRIPTION

The PDSP16540 is designed for use in synchronous data flow systems in which the transfer between system elements is controlled 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 input 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 enabling 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 available 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 synchronisation 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 interrupted for one read strobe period. The flag indicating that data is available goes inactive for this strobe period and the next system element should not accept data during this period.

The correct operation of the write synchronisation 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 inactive 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 Available 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-emphasising 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 \overline{RES} input is active. A tri-state buffer is needed on the outputs which is only enabled during \overline{RES} , 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 (\overline{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 \overline{DAV} will go inactive when the complete block has been read. The \overline{DAV} output will also go inactive 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 \overline{DAV} output must be used as write enable for the next device. All \overline{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 \overline{RES} 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. RMF goes inactive at the same time \overline{DAV} goes inactive.

The gap between the RMF and \overline{DAV} outputs should be sufficient to ensure that the next system component can immediately accept data once \overline{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 irrelevant 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 transforms are to be done (See the next section).

The \overline{DAV} output is used to drive the \overline{INEN} input on the PDSP16510 and RMF is not used. The PDSDP16510 must be used in the mode in which \overline{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

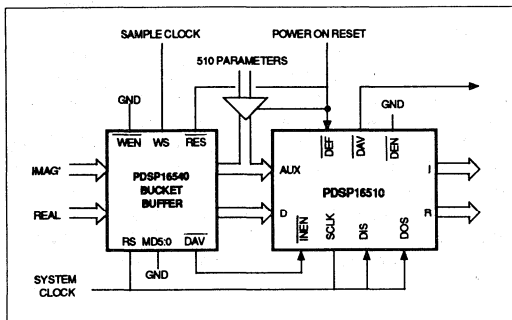


Figure 1. Typical 1024 Point FFT System

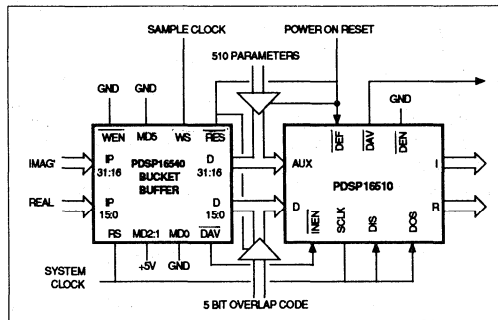


Figure 2. System with Non Standard Overlaps

PDSP16510 is not used by the PDSP16540, since internal logic computes the starting address for the read operation.

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 inputted on D9:5. This high impedance state also allows the PDSP16510 control parameters to be inputted 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;

$$NS > 1024B + \frac{1024B}{S} + T + D$$

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 be 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 is 149 ns. This equates to a 6.7 MHz writing rate when blocks are not

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		
WEN set up wrt WS going high	10ns		WEN going active or in-active
WEN Hold wrt WS going high	2ns		
Data In Set Up wrt to RS going high	15ns		
Data In Hold Time wrt rs going high	0ns		
Delay from RS going high to O/P Data		12ns	
DAV, RMF transition wrt to RS going high	2ns	15ns	Going active or in-active
Time to go Low Z wrt to RS going high		12ns	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

PDSP16540

overlapped, 3.35 MHz with 50% overlaps (512 new words), or 1.675MHz with 75% overlaps (256 new words).

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 transforms

ABSOLUTE MAXIMUM RATINGS [See Notes]

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 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 θ_{JC}	5°C/W

NOTES ON MAXIMUM RATINGS

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

STATIC ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise state)

$$T_{AMB} = 0^\circ\text{C to } +70^\circ\text{C.}$$

$$V_{CC} = 5.0V \pm 10\%$$

are to be performed. It does not support block overlapping in this mode.

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.

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

ORDERING INFORMATION

PDSP16540 C0 AC (Commercial - PGA Package).

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

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	$GND < V_{IN} < V_{CC}$
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	$V_{CC} = \text{Max}$

Application notes



A 50ns BUTTERFLY PROCESSOR

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.

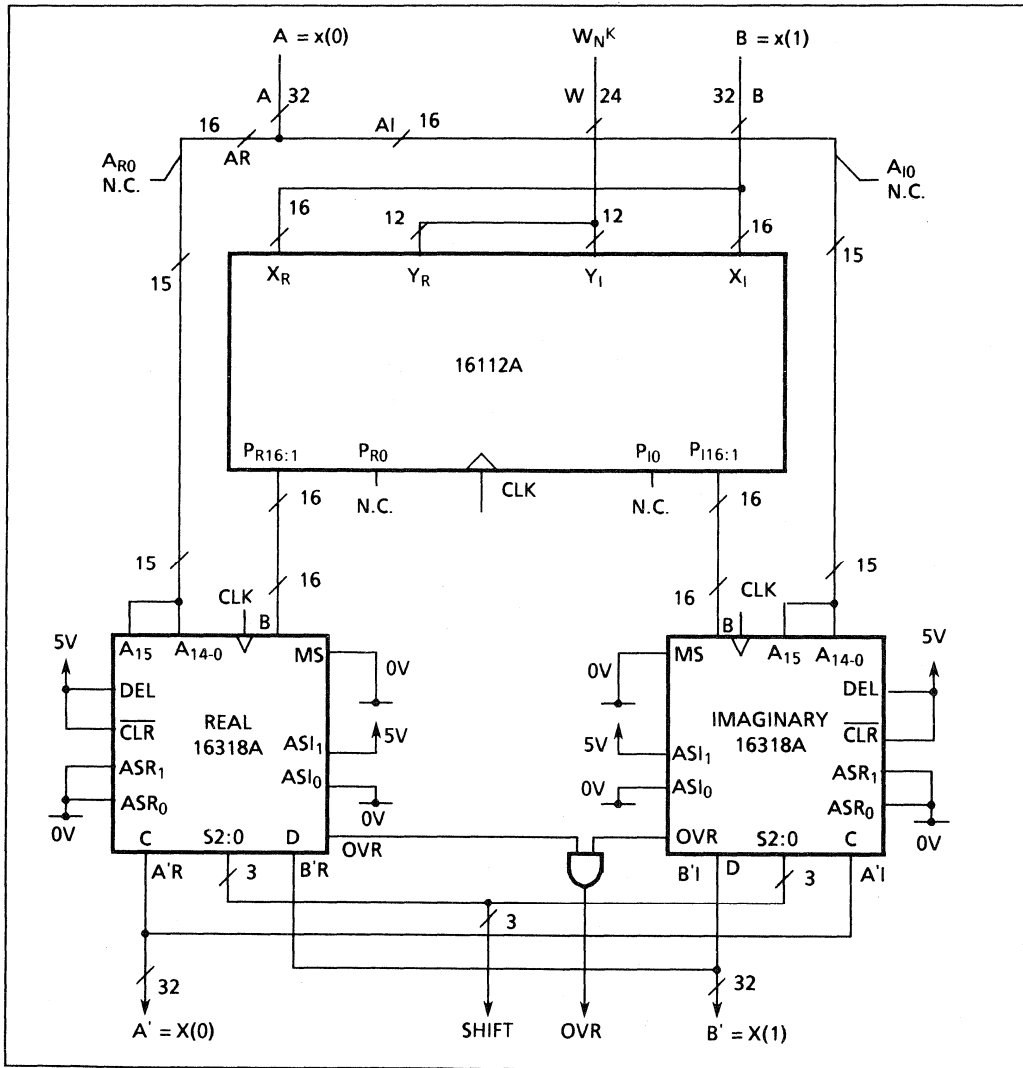


Fig. 2

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.

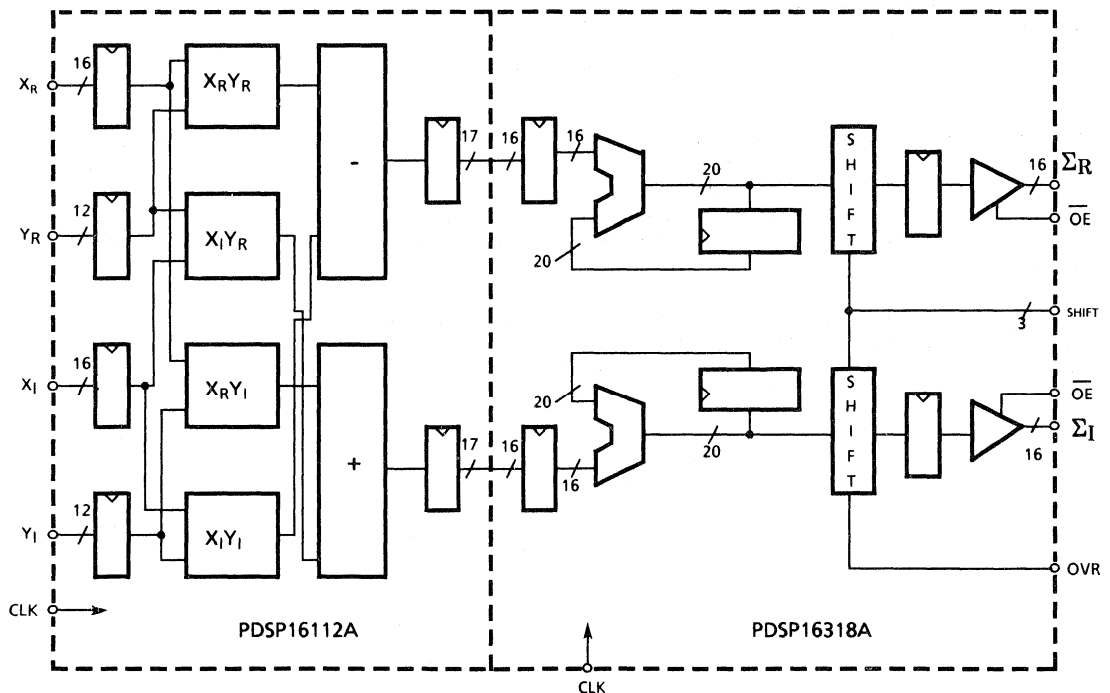


Fig. 1

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.

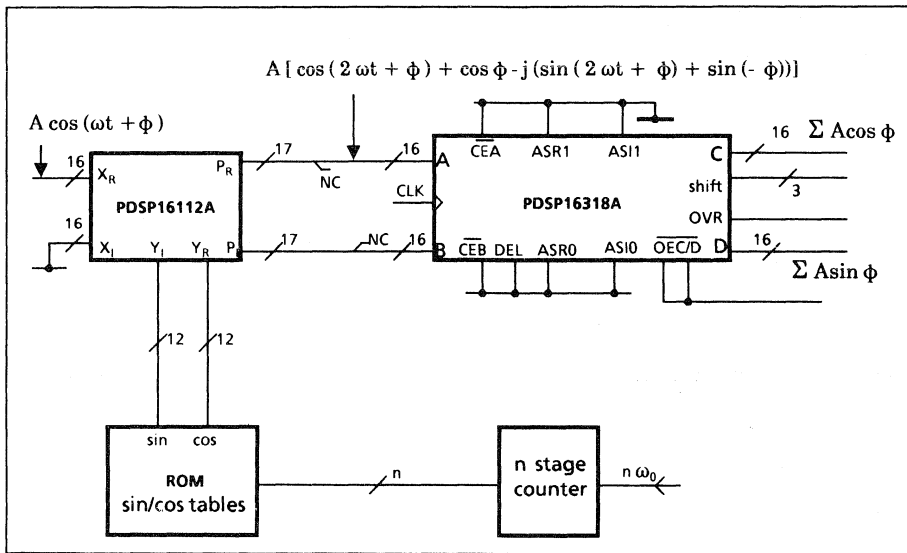


Fig. 2 Digital Demodulator

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 $2e^{-j\omega t} (= \cos \omega t - j \sin \omega t)$ to give an output which, when integrated over the signal period in the accumulator, gives an output $\Sigma A \cos \phi + j \Sigma A \sin \phi$ 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'.

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 PDSP16330 Block Diagram




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 2's 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.

FIR FILTERING WITH THE PDSP16112 AND PDSP16318

The Plessey 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 an 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 an 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 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:

$$y_n = b_0x_n + b_1x_{n-1} + \dots + b_{N-1}x_{n-N+1}$$

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.

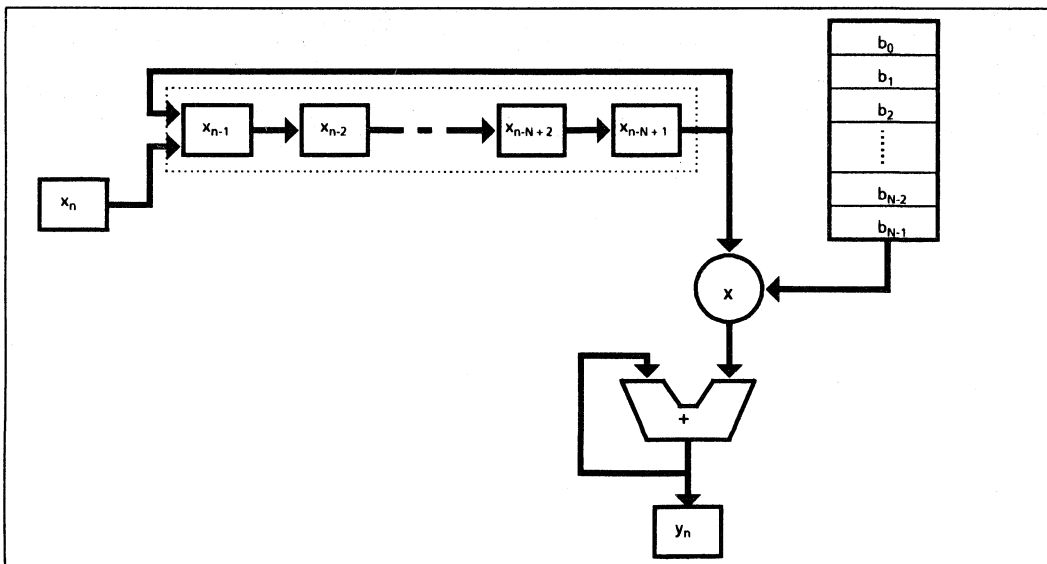


Fig 1 Direct Form Implementation of FIR Filter




Fig 2 FIR filter block diagram using PDSP16112/16318

An 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 an FIR filter using the PDSP16112 and PDSP16318 is shown in fig 2. This circuit also shows the use of a FIFO, such as the Plessey MV66030, as the data store, which may be cascaded to produce the depth and width of registers required, and the PDSP1640 Address Generator to provide the correct address sequence for the coefficient PROM.

INTERFACING THE PDSP FAMILY

INTRODUCTION

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

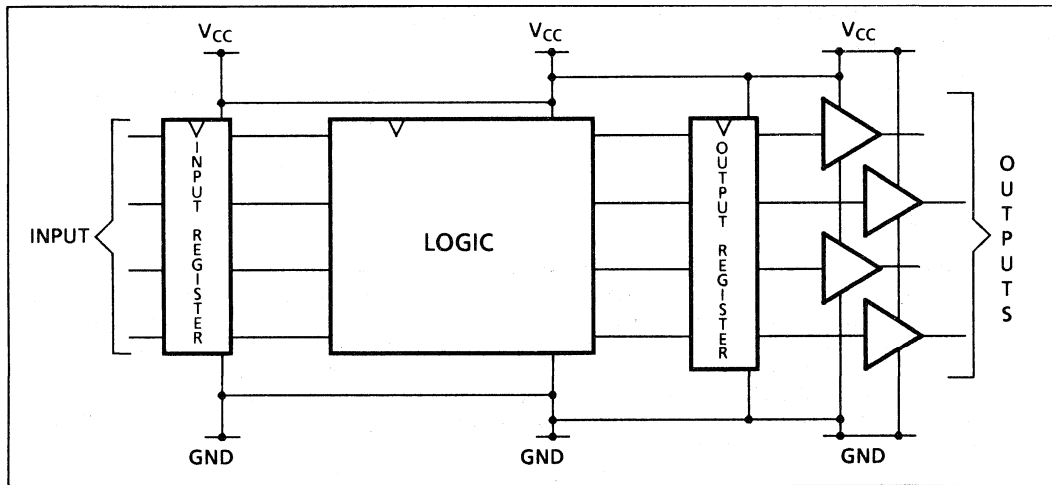


Fig. 1 PDSP Structure

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.

AB16

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 0.47 μ F 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 1601s or 1640s - the long track lengths and heavy capacitive 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 18ohms). 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.

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

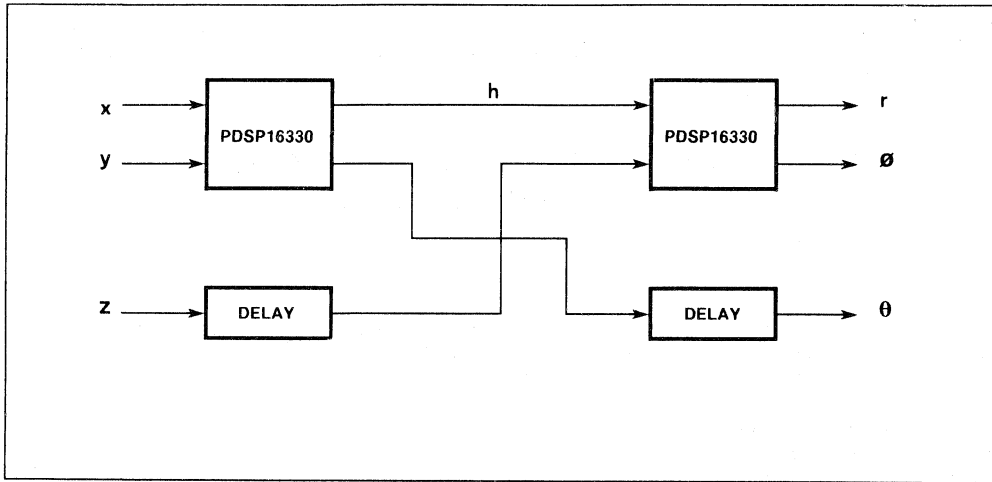


Fig. 2




Fig. 3



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 Plessey Semiconductors PDSP16112 and PDSP16318 in the evaluation of this algorithm.

1.1 The DFT algorithm

The DFT of a limited sequence of values $\{x(n)\}$, $0 \leq n \leq (N - 1)$ is defined as:-

$$X(K) = \sum_{n=0}^{(N-1)} x(n)e^{-j(2\pi/N)nK}, \quad K = \{0,1,2,3,\dots,(N-1)\} \quad 1$$

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

$$X(K) = \sum_{n=0}^{(N-1)} x(n) W_N^{nK}, \quad \text{where } W_N = e^{-j2\pi/N} \quad 2$$

If we split the sequence $\{x(n)\}$ into its even and odd numbered points then:-

$$X(K) = \sum_{n=0}^{(N-1)} x(n) W_N^{nK} + \sum_{n=0}^{(N-1)} x(n) W_N^{nK} \quad 3$$

even only odd only

or

$$X(K) = \sum_{n=0}^{(N/2-1)} x(2n) W_N^{2nK} + \sum_{n=0}^{(N/2-1)} x(2n+1) W_N^{(2n+1)K} \quad 4$$

Now,

$$W_N^2 = [e^{-j(2n/N)}]^2 = e^{-j2n/(N/2)} = W_{N/2} \quad 5$$

and $x_{\text{even}}(n) = x(2n)$, $x_{\text{odd}}(n) = x(2n+1)$

then

$$X(K) = \sum_{n=0}^{(N/2)-1} x_{\text{even}}(n) W_{N/2}^{nK} + W_N^K \sum_{n=0}^{(N/2)-1} x_{\text{odd}}(n) W_{N/2}^{nK}$$

the original transform has now divided into two separate smaller transforms combined in the following way:

$$X(K) = X_e(K) + W_N^K X_o(K) \quad 6$$

where $X_e(K)$ and $X_o(K)$ are the $N/2$ point DFTs of $x_{\text{even}}(n)$ and $x_{\text{odd}}(n)$ 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 $K=0,1$ are ± 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'.

Fig. 1 shows the decomposition of an N point DFT into two N/2 point DFTs and twiddles.

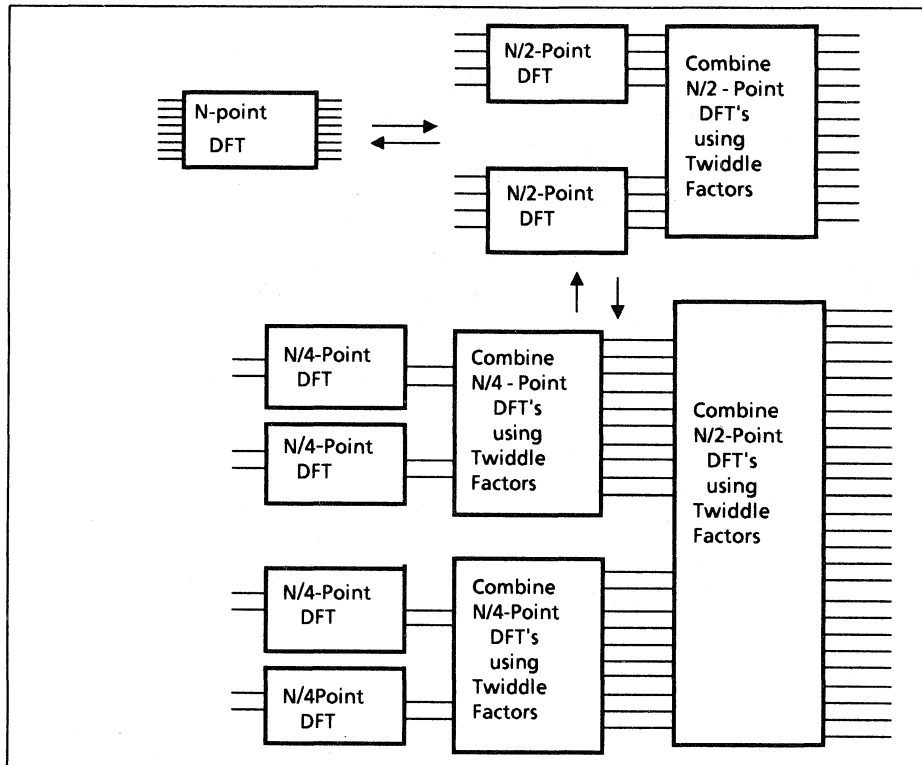


Fig. 1 Typical decomposition for radix 2 FFTs

As an example, Fig. 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 $N/2 \log_2 N$, each Butterfly requiring only one multiplication (W_N^k can be calculated and stored for use twice). Contrast this with the $(N-1)^2$ multiplications required by the direct DFT:

N	$(N-1)^2$	$N/2 \log_2 N$ (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

Fig. 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).

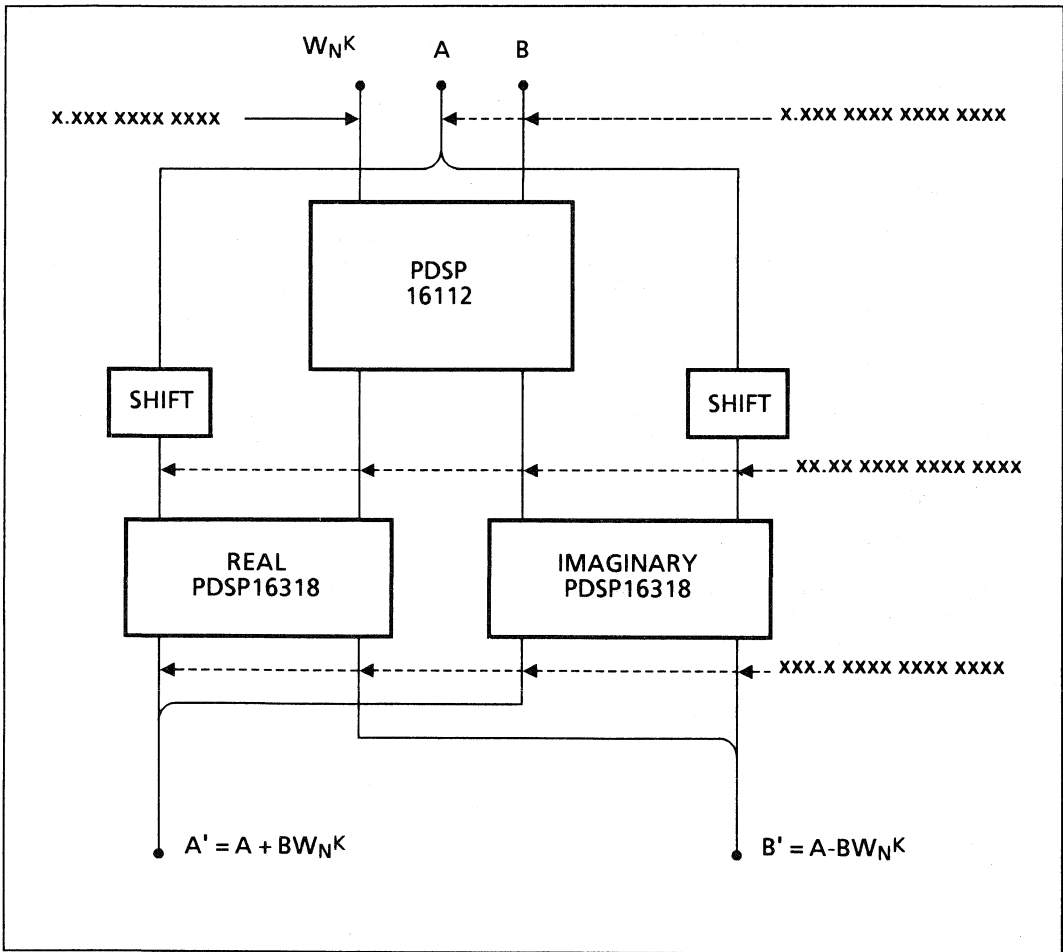


Fig. 4 Basic arrangement of butterfly processor

3. HARDWARE

Figs. 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 eight 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.

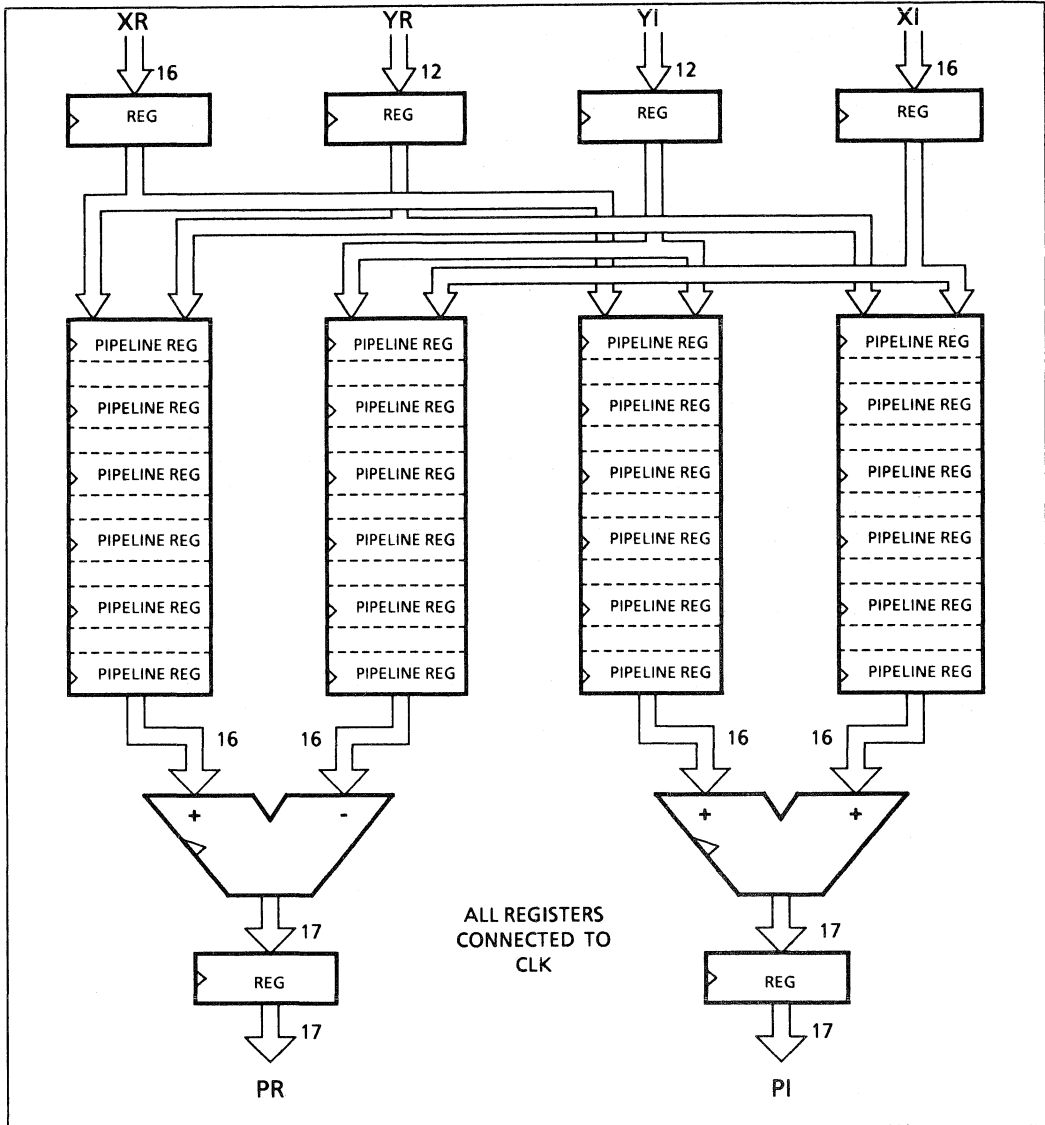


Fig.5 Pipelined multiplier structure of complex multiplier




Fig.6 Block Diagram of Complex Accumulator

ASR or ASI		ALU Function
ASX1	ASX0	
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 that 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 occurring within the accumulator. If overflow occurs an incorrect result is output from the Complex Accumulator, however the overflow flag of the Complex Accumulator 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	0
1	1	0							15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
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 output data is padded with zeros

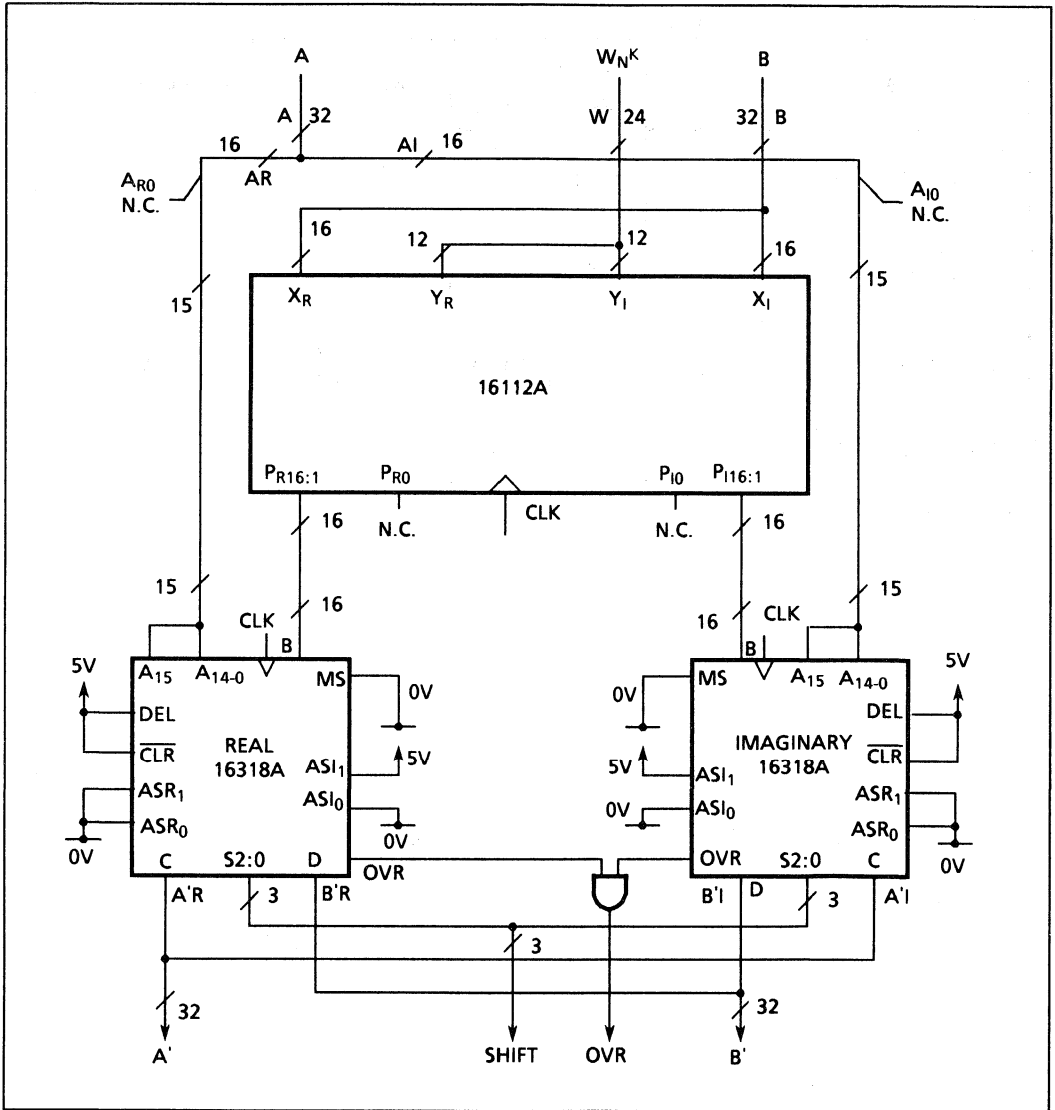


Fig.7

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):

$$V_{cc} = 5V \pm 10\%, T_{amb} = -40^{\circ}C \text{ to } +85^{\circ}C$$

Parameter	Value		Units	Conditions
	Min.	Max.		
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

COMPLEX SIGNAL PROCESSING WITH THE PDSP16000 FAMILY

The PDSP16112 Complex Multiplier, PDSP16318 Complex Accumulator and PDSP16330 Pythagoras Processor are DSP building block parts with the unique capability of operating directly on complex data. Using conventional building blocks it is possible to operate on complex rather than real data, but at the expense of substantially reduced throughput.

Complex signal processing (where the signal comprises both in-phase and quadrature components) is necessary in signal processing applications where the phase of the signal is important. The following example illustrates such a case, and shows how the PDSP16000 devices fit into a typical DSP processor.

Consider the demodulator configuration of Fig. 1.

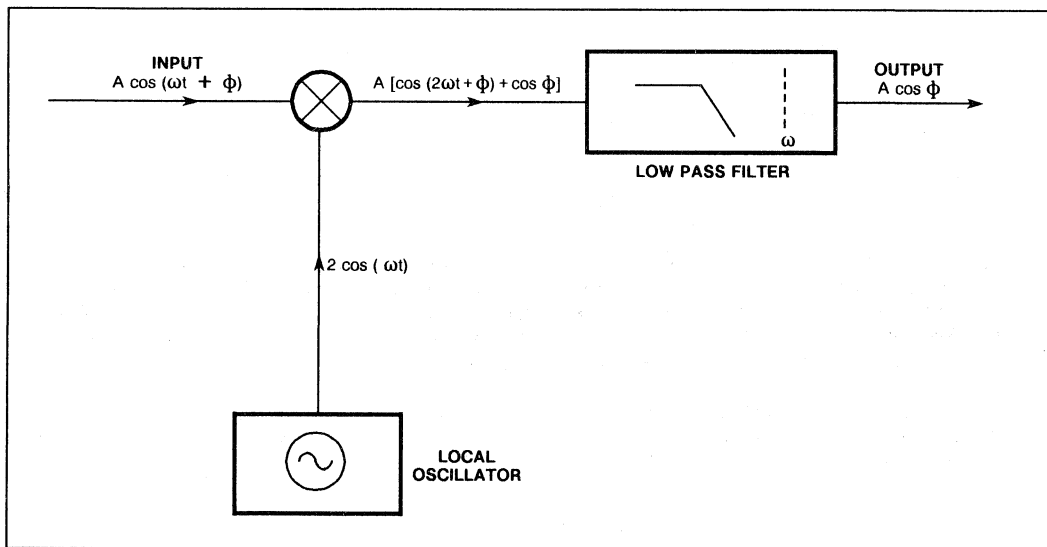


Fig. 1 Simple demodulator

The incoming signal $A \cos(\omega t + \phi)$ is mixed with the output of a local oscillator running at the same frequency ω . The output from the multiplier $A[\cos(2\omega t + \phi) + \cos\phi]$ is low pass filtered to give the baseband output $A \cos\phi$. This is all well and good if the relative phase of the incoming signal to the local oscillator is zero, but unfortunately this cannot be guaranteed. If the phase ϕ is $\pi/2$ then the demodulator output will be zero, which is not very useful! In a real system the relative phase between received signal carrier and local oscillator will vary, giving rise to a variety of undesirable effects from signal fading to noise on the output.

By using complex arithmetic it is possible to demodulate the received signal without phase coherence problems. Fig. 2 shows an I/Q Demodulator.

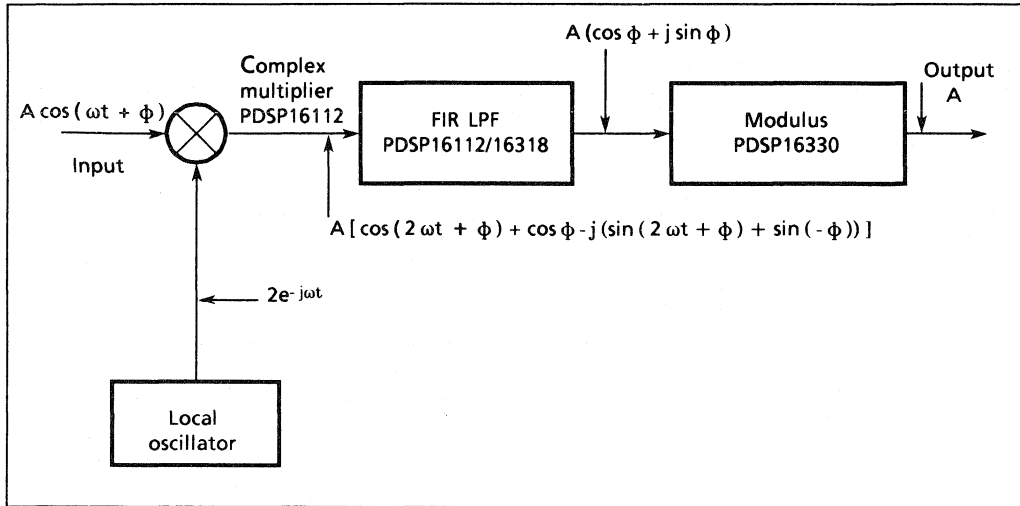


Fig. 2 I/Q demodulator

The output of the local oscillator is now a complex number $2(\cos \omega t - j \sin \omega t)$ which is multiplied by the incoming signal in a Complex Multiplier such as the PDSP 16112. The complex result is then filtered to remove non-baseband signals, leaving a result $A(\cos \phi + j \sin \phi)$ at the output of the filter. This filter must have a cut-off frequency equal to or less than half the channel spacing in the incoming signal to maintain selectivity. The filter will be an FIR filter based on a Complex MAC comprising a PDSP16112 and PDSP16318.

Fig 3 shows a PDSP16112 and PDSP16318 connected as a complex MAC. For maximum accuracy the top 12 bits of the X inputs are used, and the LSB of the outputs discarded. Word growth in the accumulator is handled by the use of the output shifter on the PDSP16318

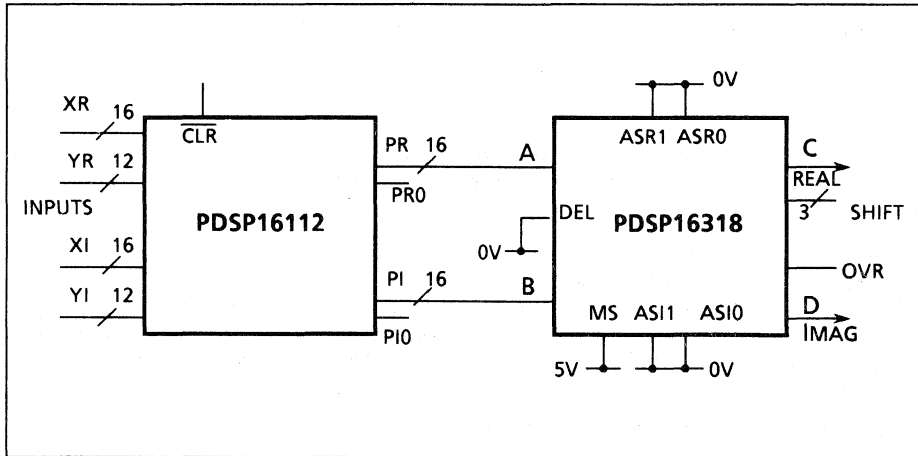


Fig. 3 Complex Multiplier Accumulator

Extraction of the modulation from the complex signal $A(\cos\phi + j\sin\phi)$ is achieved by the use of PDSP16330 Pythagoras Processor. This device calculates $\sqrt{X^2 + Y^2}$, if $X = A\cos\phi$ and $Y = A\sin\phi$ then the MODULUS output from the PDSP16330 will be $\sqrt{A^2\cos^2\phi + A^2\sin^2\phi} = A$, the original modulation signal. In this case the relative phase of the carrier and local oscillator has no effect on the demodulated output.

Similar phase related difficulties are found when carrying out correlation and convolution operations where complex processing is required to preserve the integrity of the signals being operated upon.

COMPLEX CORRELATION

The complex cross-correlation of two signals A and B is

$$R_{AB}(m) = \frac{1}{N} \sum_{k=0}^{N-1} \bar{A}(k) \cdot B(k+m)$$

Where A bar indicates complex conjugation, either the A or B term may be conjugated. The parameter N is the number of points over which the correlation is carried out, k and m are the sample indices.

This function can be carried out by a PDSP16112 and PDSP16318 connected as a complex MAC, k and m may be generated by PDSP1640 Address generators.

CONVOLUTION

The complex convolution of two signals A and B is defined as:

$$C_{AB}(m) = \frac{1}{N} \sum_{k=0}^{N-1} A(k) \cdot B(m-k)$$

The convolution operation is, as can be seen, very similar to the correlation operation except that there is no complex conjugation involved and one of the signals (B in this case) is time reversed.

A convolver and an FIR filter are effectively the same thing. The output of a filter can be considered to be the convolution of the input and the impulse response of the filter. Because of this commonality, FIR filters, correlators and convolvers can all be performed by the same hardware systems with very little redundancy (only the complex conjugation is redundant when filtering or convolving). The time reversal required in convolution is easily achieved by programming one of the data address generators to count backwards.

APPLICATIONS OF CORRELATION / CONVOLUTION

Delay Measurement (Fig. 4)

A common application of correlation in sonar, radar and geological work is the estimation of delay times. By correlating the outgoing signal with the return signal it is possible to estimate the delay between the two.




Fig.4

$T(t)$ is the transmitted signal, $R(t)$ is the returned reflection after a delay t_d . The value of t_d can be determined by the correlation of $T(t)$ and $R(t)$. This function will have a well defined peak at $t = t_d$ even if the signal has suffered considerable distortion on the outward and return path.

Pitch Determination (Fig. 5)

One of the most difficult problems in the analysis of speech or musical signals is the determination of pitch. For simple consistent waveforms, it is quite easy to determine the period - all that is necessary is to time between successive zero-crossings. With complicated waveforms this approach is no longer reliable, but autocorrelation offers a workable technique.




Fig.5

Estimating the period of the waveform $X(t)$ by zero-crossing timing clearly will not work as there are a number of zero-crossings per cycle. The autocorrelation function of $X(t)$, $A_{xx}(t)$ exhibits very pronounced peaks every cycle.

Noise Cancellation

A common difficulty in many signal processing systems is recognizing the desired signal in the presence of large amounts of masking noise. Correlation can be used to extract required signals from background noise if the appropriate characteristics of the signal are known.

If a signal in the presence of background noise is correlated with an estimate of the signal, the required signal will correlate strongly with the estimate and weakly with the noise. The improvement on signal - to - noise ratio known as Correlation Gain, obtained by an N-point correlation is:

$$G = 10\log_{10} N \text{ dB}$$

Pulse Compression (Fig. 6)

In radar systems there is a design conflict between range resolution and operational range. In order to maximise operational range the average radiated power must be as high as possible. However most transmitters have limited peak power delivery which means that for maximum range, the transmitted pulse length must be as large as possible. The range resolution of the radar is $\frac{1}{2}tc$, where t is the pulse duration and c the velocity of light. For a pulse duration of $10\mu\text{s}$, the range resolution will be 1500m, that is, two targets within 1500m of each other would not be distinguished!

Pulse compression radar operates by transmitting not a simple 'burst' but a special waveform known as a 'chirp'.

The important characteristic of a chirp is that its autocorrelation function is substantially shorter in duration than the chirp itself, the greater the swept bandwidth of the chirp, the shorter the autocorrelation pulse.

In a Radar system, the transmitter sends out a linearly frequency modulated pulse of such a duration needed to obtain the required operating range. In the receiver, the signal is correlated with a replica of the transmitted chirp to produce a short duration pulse, which can give good range resolution.




Fig.6

Discrete Fourier Transform (DFT)

The equation for a DFT is:

$$X(k) = \frac{1}{N} \sum_{n=0}^{N-1} x(n) e^{-j2\pi nk/N} \quad k = 0, 1, \dots, (N-1)$$

Contrast this with the function for convolution above:

$$C(m) = \frac{1}{N} \sum_{k=0}^{N-1} A(k).B(m-k)$$

It can be seen that the two operations are identical if B is made equal to $e^{-j2\pi nk/N}$

Usually a DFT will be evaluated by the familiar Cooley-Tukey radix 2 FFT algorithm, but there are applications where the straightforward DFT is preferable. If the transform length is relatively short, a direct DFT using the PDSP16112 and PDSP16318 can be both cost effective (the hardware is much simpler than for FFT) and (because of the high speed of the devices) still capable of acceptable transform speeds. Unlike the Cooley-Tukey Algorithm a DFT may be of any number of points and being non-recursive, can achieve greater accuracy.



A FAST FFT PROCESSOR USING THE PDSP 16000 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 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 multiply, an addition and a subtraction to complete. If we set a target of $256\mu\text{s}$ 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:

$$\begin{aligned} & 256/(N/2 \log_2 N) \mu\text{s} \quad (\text{where } N = 1024) \\ & = 50\text{ns}. \end{aligned}$$

The complex multiplication operation itself requires four real multiplications, an addition and a subtraction, so that the complex Butterfly requires four multiplications, 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, require 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

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:

$$X_R \cdot Y_R; X_R \cdot Y_I; X_I \cdot Y_R; X_I \cdot Y_I.$$

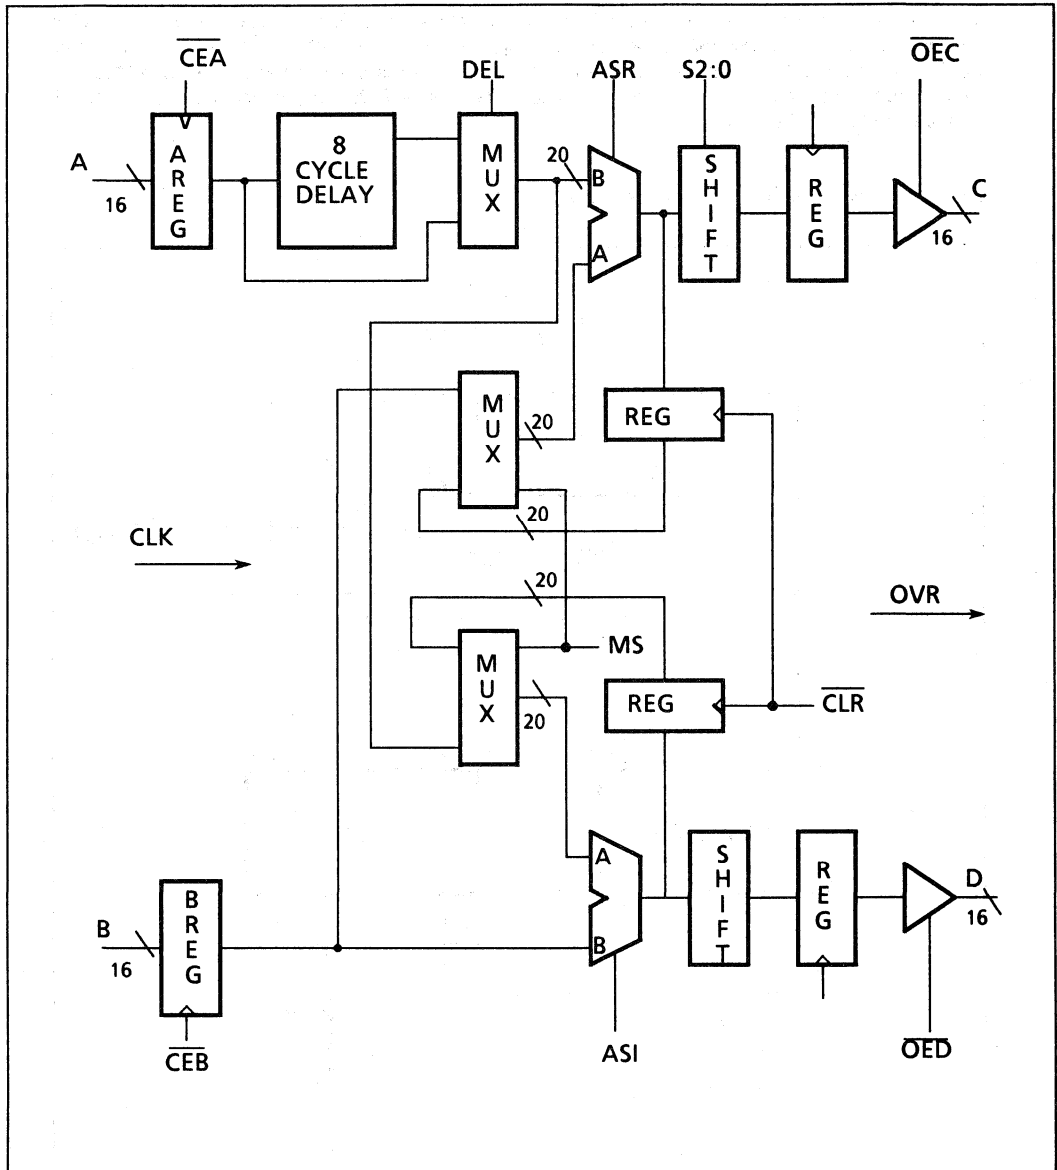


Fig. 3 PDSP16318A 20MHz Complex Accumulator

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 26 μ s!

Application Note AN 47 '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 $n = 0$ to $(N/2-1)$

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 $n = 0$ to $(N/2 - 1)$,

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 quadrant enabled
4	and	12 on the 5th cycle with the left quadrant enabled
5	and	13 on the 6th cycle with the right quadrant enabled
6	and	14 on the 7th cycle with the left quadrant enabled
7	and	15 on the 8th cycle with the right quadrant 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 block 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.

ADDRESSING

The addressing sequence of the data RAMs can be totally satisfied by a PDSP1640 20 MHz Address Generator. The PDSP1640 (see Fig. 8) integrates an eight bit add-latch loop with an on chip comparator and five user programmable registers. The PDSP1640 occupies a 28 pin package which in LCC form offers the smallest footprint of any Address Generator.

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:

$$K = \text{Bit Reversed } [0, \dots, (2^{(m-1)}-1)]$$

this sequence is repeated $\frac{1}{2} N \cdot 2^{m-1}$ 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 $(16/2) / (2 \text{ to the power } 3) = \text{once}$.

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. The Address sequence is easily generated by a pair of PDSP1640s which will address a 16 bit field at 20MHz. The output of the PDSP1640s is wired in bit reversed order before being applied to the Coefficient RAMs.

FFT PROCESSOR CONFIGURATION

The architecture of the complete FFT processor is as shown in Fig 9. Each of the two data RAMs is addressed by its own PDSP1640 address generator, as is the coefficient RAM. A configuration using 50ns access 256 word RAM devices, each addressed by a single 8 bit wide PDSP1640, will in conjunction with PDSP16112A and PDSP16318A Arithmetic processors execute a 1024 point Complex FFT in just 256 μ s, 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

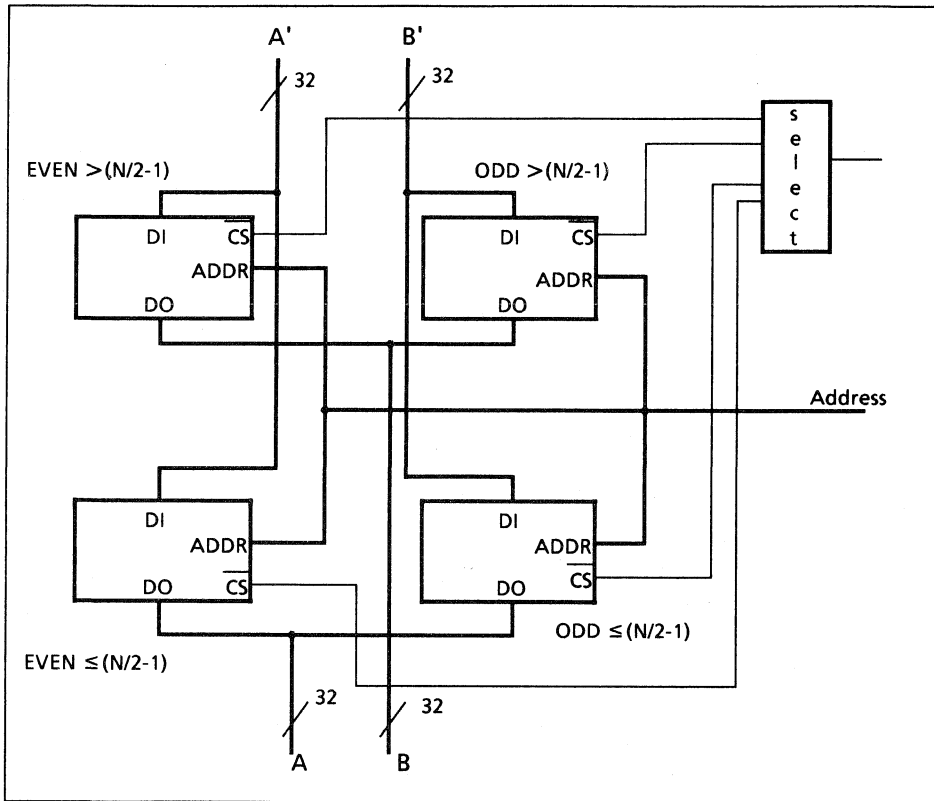


Fig. 6 Memory Configuration For Constant Geometry Algorithm

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.

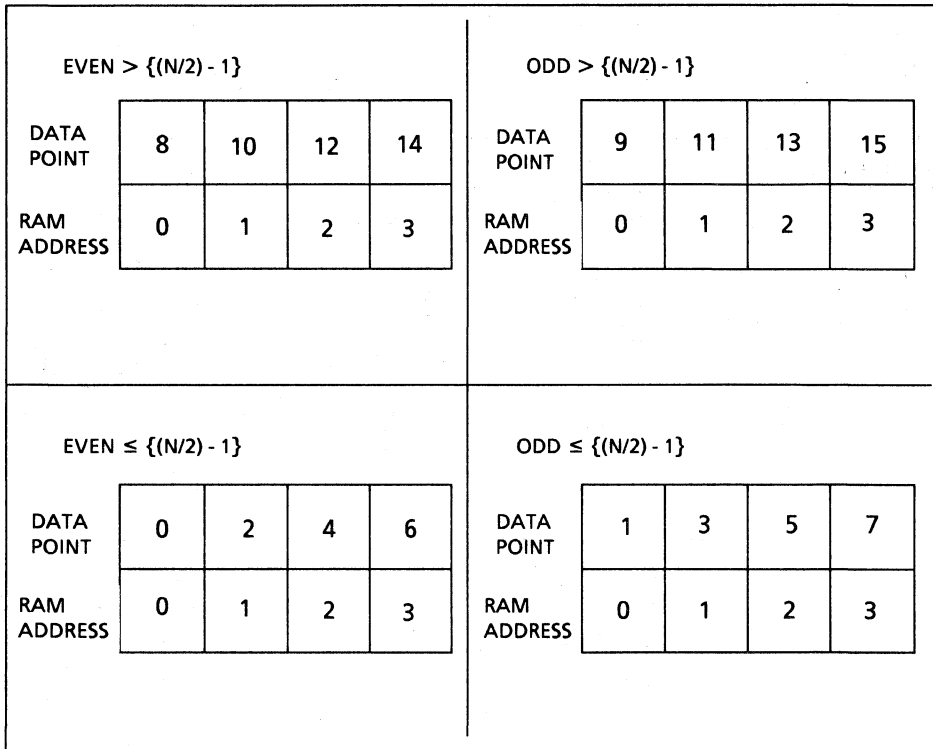


Fig. 7 16 point Transform Example

The data points of a 16 point transform are stored in the four RAM quadrants at the addresses and in the quadrants indicated.

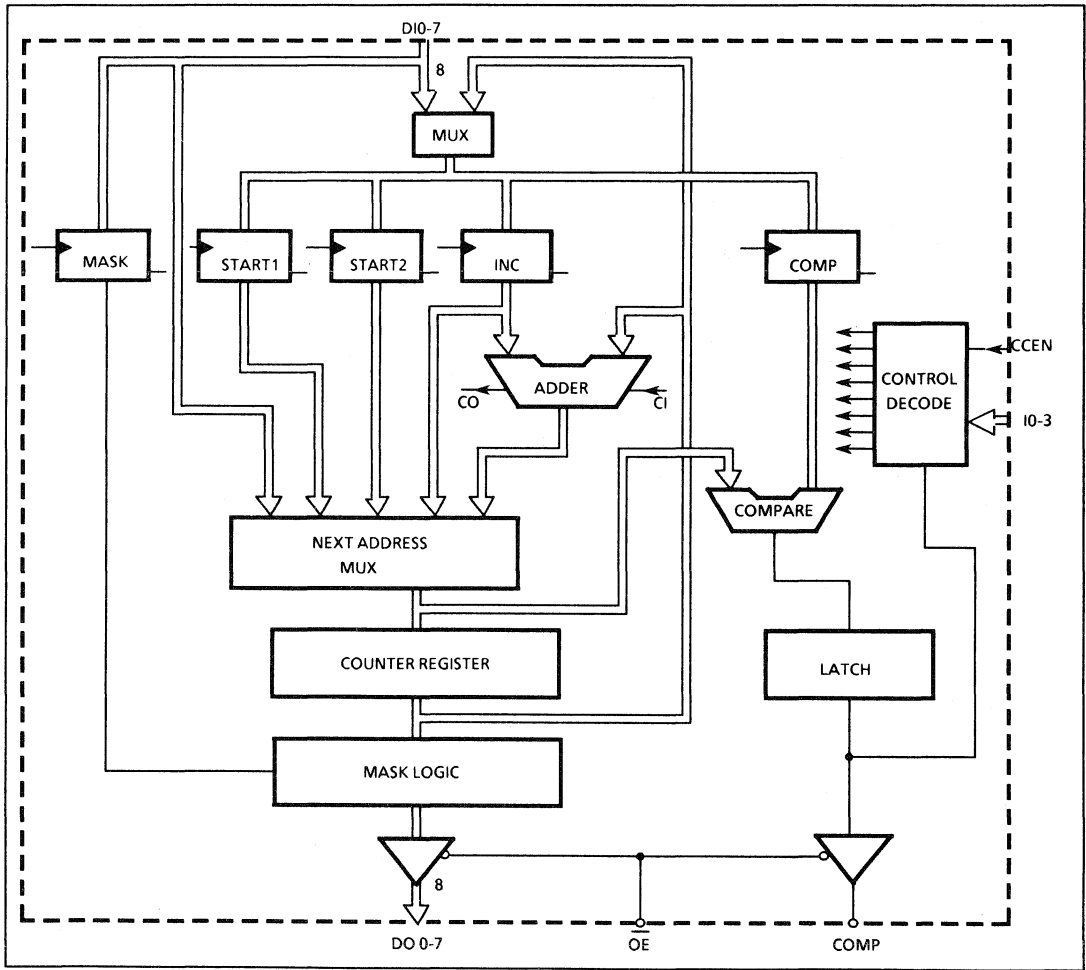


Fig. 8 PDSP1640 20MHz Address Generator

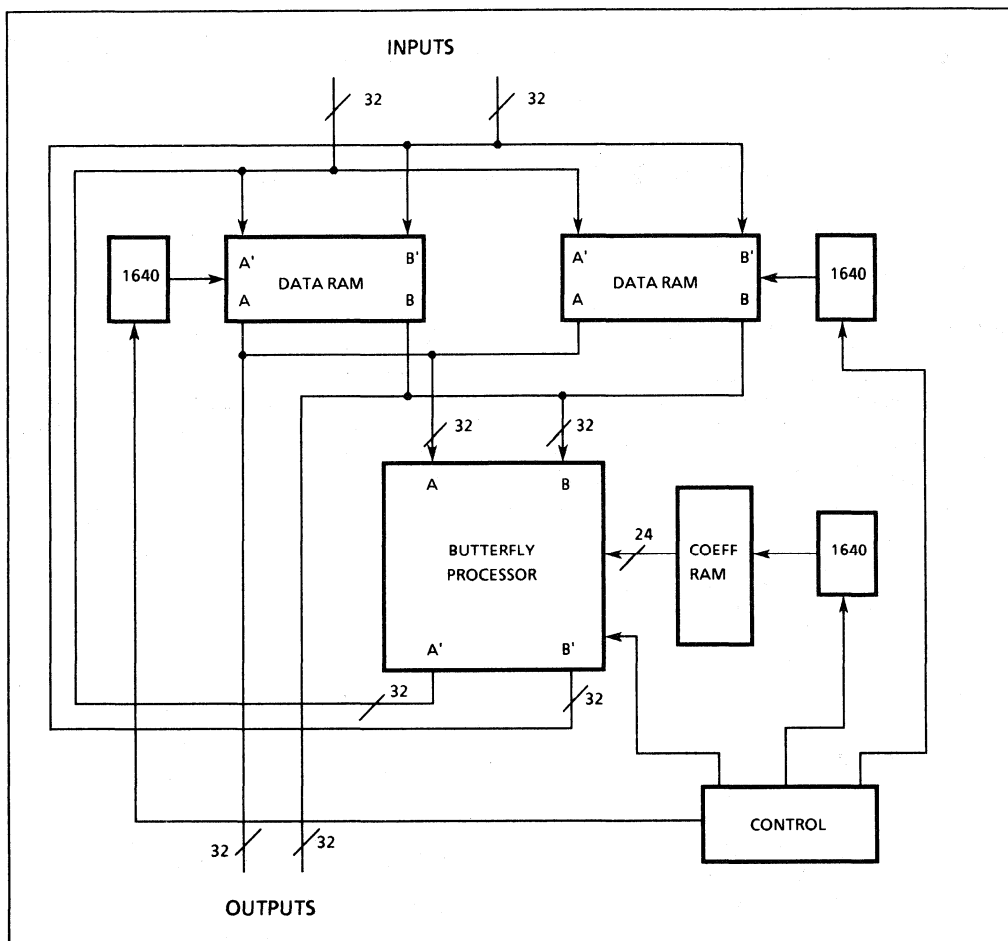


Fig. 9 FFT Processor architecture



FFT ADDRESS GENERATION USING THE PDSP1640

INTRODUCTION

Fast Fourier Transforms (FFTs) are used in a wide variety of applications as a means of calculating the Discrete Fourier Transform of a signal, in order to estimate the signal's spectral energy. There are many different algorithms for computing FFTs, each designed to exhibit particular characteristics, but each of which also have certain disadvantages. The in-place radix 2 DIT algorithm shown in Fig. 1 puts the outputs from each stage into the same memory location from which the inputs are read. This minimises the system memory requirements, but has the disadvantage of requiring a different register addressing sequence at each stage and also requires four memory accesses per butterfly cycle, necessitating the use of a very fast RAM.

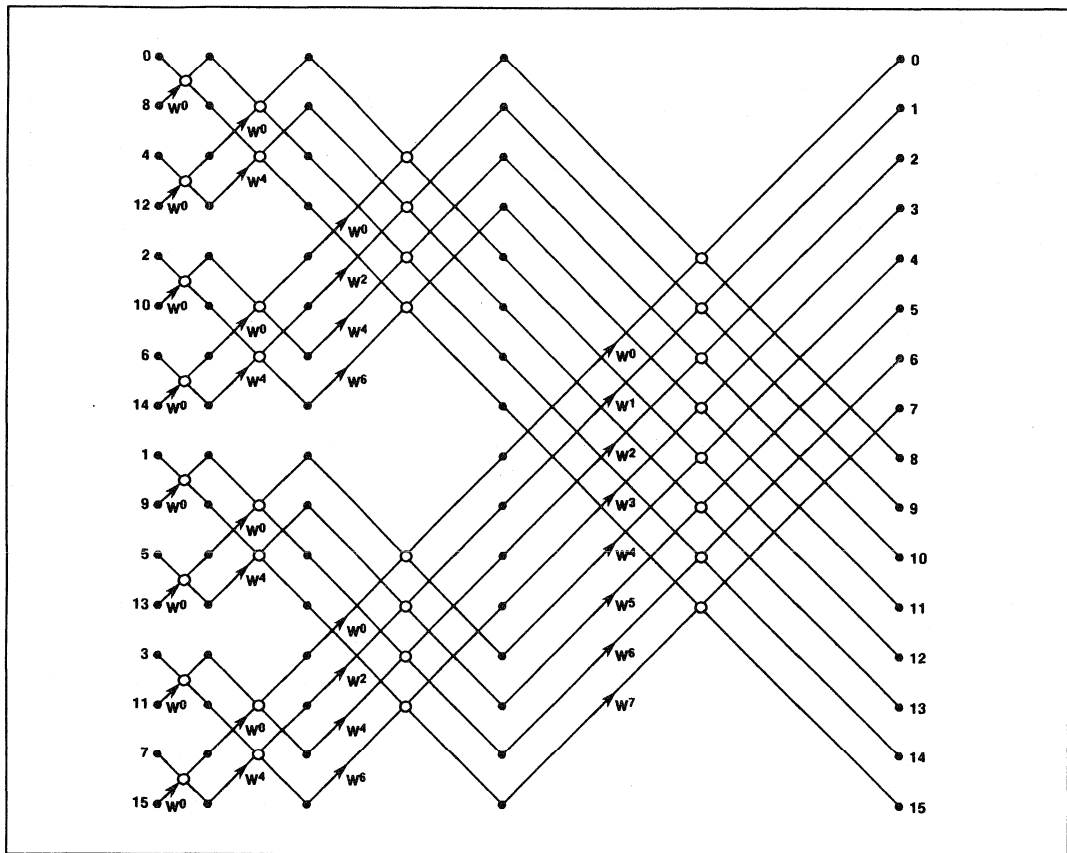


Fig.1 16-point in-place radix 2 DIT FFT algorithm




Fig.2 16-point constant geometry radix 2 DIT FFT algorithm

The main advantage of the 'constant geometry' algorithm, shown in Fig. 2, is that the data memory may be configured such that any particular RAM device need only be accessed once during each butterfly cycle, thus allowing the use of slower devices. The algorithm also has the feature that the memory read address sequence and the memory write address sequence remain the same from one stage to the next. However, the algorithm is not in-place and therefore requires twice as much RAM storage as an in-place algorithm and the coefficients are not in a simple order.

These two algorithms are, perhaps, the most common of FFT algorithms and as such are used in this Application Note as examples of the FFT address generation capabilities of the PDSP1640. The variety of FFT algorithms available shows that the data and coefficient address sequencing is a major consideration in the implementation of any one, requiring in nearly all cases some form of programmable address generation. The principles and techniques explained here for the use of the PDSP1640 may be used as the basis for the implementation of address generation for many other algorithms.




Fig.3 PDSP1640 Address Generator block diagram

THE PDSP1640 ADDRESS GENERATOR

The Plessey Semiconductors PDSP1640 is an 8-bit programmable address generator capable of operation at speeds up to 20MHz. It may be cascaded with other devices to produce wider address fields, for example operating at up to 10MHz for a 24 bit address and is ideally suited to many situations requiring high speed address generation, both for FFT computation and for other digital signal processing applications.

A block diagram of the PDSP1640 is shown in Fig. 3. It consists of five user-programmable registers, an 8-bit address counter, a comparator and instruction decode logic, together with an output multiplexer and mask logic. The various instructions allow each of the registers to be loaded from a number of sources or for the address counter register to be incremented, via the 8-bit adder, by the value specified in the increment register. Carry In and Carry Out signals are available to implement the cascading of the adder with those from other devices. The increment instructions may be conditional such that if a particular address count is reached, a jump to another specified address will occur. By programming the device with a suitable sequence of instructions and data, it is possible to generate the desired address sequence.

For cascaded devices, the instructions input to each device are the same; however, the data associated with these instructions will vary. For two PDSP1640s, for example, the devices may be regarded as a single unit with the data being 16 bits wide.

IN-PLACE FFT ALGORITHM

Data Memory Addressing

Since the data output at each stage is put back into the same memory location from which the input data is read, the write address sequence is exactly the same as the read address sequence, but delayed by the number of clock cycles required to execute the FFT butterfly. Using the Plessey Semiconductors PDSP16112 Complex Multiplier and the PDSP16316 Complex accumulator, this delay is 9 clock cycles. It is therefore possible to use a single address generator for both the read and write sequences.

Two words of data must be read from memory and two words written back again during each butterfly clock cycle, requiring the use of fast dual port RAM and some additional buffering to present the two input words to the butterfly processor at the same time.

From Fig. 1, it can be seen that the register address sequence for the first stage is: 0&1; 2&3; 4&5;....etc. This can be achieved by simply incrementing the address counter register of the PDSP1640 by one after each memory access from zero until the last data register is reached, as indicated by the COMP output from the PDSP1640.

For the second stage, the data address sequence is slightly more complicated being 0&2; 1&3; 4&6; 5&7; 8&10;....etc., which is rather difficult to generate in this jumbled order. Since the order in which the butterflies are generated within each stage is not important, this order may be re-arranged into even and odd sequences of 0&2; 4&6; 8&10;....etc. and 1&3; 5&7; 9&11;....etc. These two half-length series each have an increment of two between consecutively generated addresses, which is easily achieved. The only remaining problem is to get from the end of the first sequence to the start of the second. One way of achieving the jump between the two series is to set the COMPARE register to the end of the first half series (i.e. to N-2) and the START1 register to the beginning of the second (i.e. to 1). The CCJS1 instruction is then used to increment the address counter until the comparator flag output (COMP) becomes active and a jump takes place to the location specified in the START1 register. At this point the COMPARE register must be reloaded with the end of the second half-series (N-1) and counting continued using CCJS2, the START2 register having been previously set to zero for the start of the next stage. This method is obviously rather cumbersome, particularly for subsequent stages where more than two sub-series are necessary.

An alternative method achieves the step between the series automatically, but requires some redundancy in the address bits actually used. In the 16-point FFT example, only four address bits are required. If the generated address from the PDSP1640 is shifted to the left by three positions such that the three least significant bits and the most significant bit are not used, the increment required at each step for each sub-series then becomes 10_H (0 0010 000) instead of 02_H (0000 0010). After seven increments of 10_H, an increment of 18_H is necessary, which causes the address counter to jump from address 70_H to 88_H (or 08_H as MSB is not used, i.e. from 0 1110 000 to X 0001 000). This may be achieved by specifying an actual increment for each step of 11_H, which will produce the 16-step count sequence shown in Table 1(a). The unused bits of the output address may be masked off by use of the MASK register.

	Address sequence	Count sequence		Address sequence	Count sequence
Increment = 11 _H	0	0 (0000) 000	Increment = 22 _H	0	0 (0000) 000
	2	0 (0010) 001		4	0 (0100) 010
	4	0 (0100) 010		8	0 (1000) 100
	6	0 (0110) 011		12	0 (1100) 110
	8	0 (1000) 100		1	1 (0001) 000
	10	0 (1010) 101		5	1 (0101) 010
	12	0 (1100) 110		9	1 (1001) 100
	14	0 (1110) 111		13	1 (1101) 110
	1	1 (0001) 000		2	0 (0010) 000
	3	1 (0011) 001		6	0 (0110) 010
	5	1 (0101) 010		10	0 (1010) 100
	7	1 (0111) 011		14	0 (1110) 110
	9	1 (1001) 100		3	1 (0011) 000
	11	1 (1011) 101		7	1 (0111) 010
	13	1 (1101) 110		11	1 (1011) 100
	15	1 (1111) 111		15	1 (1111) 110

Table 1(a)

Table 1(b)

Table 1 PDSP1640 address count sequence using bit redundancy

AN54

A similar technique may be used for subsequent stages. The third stage, for example, requires registers to be addressed in the sequence 0&4; 8&12; then 1&5; 9&13 and so on. By specifying an increment at each step of 22_H, the count sequence of Table 1(b) is produced.

By cascading PDSP1640 devices together, this technique may be extended to any size of FFT. For the general case, the number of redundant least significant bits required for an N-point FFT is $(\log_2 N - 1)$. The increment at each stage of the FFT is given by:

$$\text{Increment, } i = 2^{(m-2)} (N-1),$$

where $m (= 1, 2, 3, \dots, \log_2 N)$ is the number of the stage of the FFT. The fractional part of the result for the first stage may be ignored.

For the 16-point FFT of Fig. 1, the instruction sequence for the whole FFT is shown in Table 2. It consists of initialising the counter register and the START1 register to address 00_H, loading an initial increment for the first stage of 08_H into the increment register and setting the compare register to the final address required for the first stage (78_H). The MASK register is optionally loaded to mask off the unused address bits and counting started. At the end of each stage, the increment register is simply loaded with the new increment value and the compare register with the final address of the next stage. This technique enables each stage to be executed without a break in the sequence of addresses generated.

Cycle No.	Mnemonic	Op. code	Data 1,2	Operation
1	CLRCR	7 _H	XX	CLEAR COUNT/MASK REGISTERS
2	LIRDI	C _H	XX	LOAD INCREMENT REGISTER
3	LCPDI	E _H	08 _H	LOAD COMPARE REG WITH END ADDRESS
4	LS1DI	8 _H	78 _H	LOAD START1 REG WITH BRANCH ADDRESS
5	LMRDI	3 _H	00 _H	LOAD MASK REGISTER
6	CCJS1	1 _H	87 _H	COUNT BY INC OR GO TO SR1
7	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
21	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
22	LIRDI	C _H	XX	LOAD COMPARE REG WITH END ADDRESS
23	LCPDI	E _H	11 _H	LOAD START1 REG WITH BRANCH ADDRESS
24	CCJS1	1 _H	FF _H	COUNT BY INC OR GO TO SR1
25	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
39	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
40	LIRDI	C _H	XX	LOAD COMPARE REG WITH END ADDRESS
41	LCPDI	E _H	22 _H	LOAD START1 REG WITH BRANCH ADDRESS
42	CCJS1	1 _H	FE _H	COUNT BY INC OR GO TO SR1
42	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
57	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
58	LIRDI	C _H	XX	LOAD COMPARE REG WITH END ADDRESS
59	LCPDI	E _H	44 _H	LOAD START1 REG WITH BRANCH ADDRESS
60	CCJS1	1 _H	FC _H	COUNT BY INC OR GO TO SR1
61	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
75	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1

NOTES

1. XX = don't care
2. Data is input on cycle following relevant instruction

Table 2 Data addressing instruction sequence for in-place algorithm

Coefficient Addressing

The simplest way of achieving the sequence of coefficients is to program a PROM with the correct sequence as required by the complete algorithm. The address generator is then simply programmed to count up to the maximum number of coefficients used. This, however, is a very inefficient method in terms of memory usage since each coefficient occurs several times. An alternative method requires each coefficient to be stored only once and makes use of the output masking capability of the PDSP1640.

The coefficients are held in PROM in numerical order (i.e. $W^0, W^1, W^2, \dots, W^{(N/2-1)}$) and the PDSP1640 programmed to cycle through all coefficient addresses during each stage of the FFT algorithm. This means that the address sequence generated by the counter register is the same for all stages. By considering the required coefficients for each stage and the order in which the FFT butterflies are executed, it is possible to use the MASK register to select a sub-set of the generated addresses. For the first stage only W^0 is required. This can easily be achieved by loading the MASK register with all 1s, which freezes all outputs and produces address 00_H for each cycle. For the second stage, W^0 and W^4 are required in the order $W^0, W^0, W^0, W^0, W^4, W^4, W^4, W^4$, to match the order of execution of the butterflies. If the MASK register is loaded with the data FB_H , only bit 2 is enabled whilst all the others are frozen at logic 0. This has the effect of producing the address 00_H for half of the sequence and 04_H for the other half at the device outputs. Similarly, for the third stage, $F9_H$ is loaded into the MASK register and $F8_H$ for the last stage. Table 3 shows the complete instruction sequence for generating the coefficient addresses. It should be noted that only one instruction is required between stages for the coefficient addressing whereas two were required for the data addressing. Since the coefficient address generator runs at half the clock speed of the data address generator, the two count sequences remain in step.

Cycle No.	Mnemonic	Op. code	Data ^{1,2}	Operation
1	CLRCR	7 _H	XX	CLEAR COUNT/MASK REGISTERS
2	LIRDI	C _H	XX	LOAD INCREMENT REGISTER
3	LCPDI	E _H	01 _H	LOAD COMPARE REG WITH END ADDRESS
4	LS1DI	8 _H	07 _H	LOAD START1 REG WITH BRANCH ADDRESS
5	LMRDI	3 _H	00 _H	LOAD MASK REGISTER
6	CCJS1	1 _H	FF _H	COUNT BY INC OR GO TO SR1
7	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
13	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
14	LMRDI	3 _H	XX	LOAD MASK REGISTER
15	CCJS1	1 _H	FB _H	COUNT BY INC OR GO TO SR1
16	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
22	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
23	LMRDI	3 _H	XX	LOAD MASK REGISTER
24	CCJS1	1 _H	F9 _H	COUNT BY INC OR GO TO SR1
25	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
31	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
32	LMRDI	3 _H	XX	LOAD MASK REGISTER
33	CCJS1	1 _H	F8 _H	COUNT BY INC OR GO TO SR1
34	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
40	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1

NOTES

1. XX = don't care
2. Data is input on cycle following relevant instruction

Table 3 Coefficient addressing instruction sequence for in-place algorithm

CONSTANT MEMORY ALGORITHM

Data Memory Addressing

The memory for this algorithm may be configured into two blocks, one for reading data and one for writing data, which exchange roles at the end of each stage of the FFT. Each memory block may be further sub-divided into four quadrants, each quadrant being a separate 32-bit wide 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 hand 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 address values less than or equal to $(N/2 - 1)$. A full discussion of this memory configuration is given in Application Note AN50 – A Fast FFT Processor using the PDSP16000 Family.

The same location address can then be supplied to each quadrant of a block, the appropriate quadrants being enabled or disabled in pairs according to the above rules. For the 16-point FFT example of Fig. 2, the read address sequence is 0, 0, 1, 1, 2, 2, 3, 3, which accesses data points in the order 0&8; 1&9; 2&10; 3&11;etc., and the write address sequence is 0, 1, 2, 3, 0, 1, 2, 3, which stores data in the order 0&1; 2&3; 4&5;etc. The storage of data points within each quadrant is shown diagrammatically in Fig. 4. Although the read and write address sequences in themselves are easy to generate, the situation is complicated by the fact that the read and write cycles alternate for each block. If dual port RAM devices are used, this presents no problem; however, when using devices with single address buses, each cycle must be carefully controlled to maintain the correct timing with respect to the flow of data through the butterfly processor. For example, after a read cycle, the write cycle for either block must be delayed by a time equivalent to two processor pipeline delays to allow for the other block to complete writing data from the previous stage before reading data for the current stage. Address generators for this configuration may be allocated either one for write addresses and one for read addresses, or one address generator per block of memory, each generating read and write sequences.

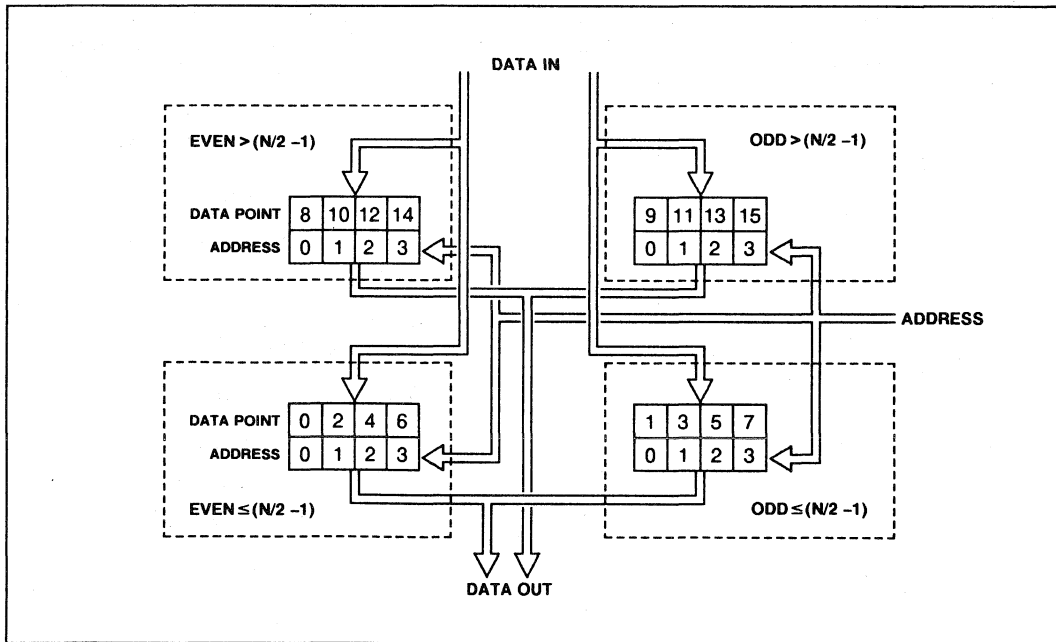


Fig.4 Data storage for constant geometry algorithm

An alternative method if dual port RAM devices are used allows the use of a single address generator for both read and write blocks of memory. The address generator is programmed to count from 0 to $(N/2 - 1)$ for each stage. The write addresses are taken from the least significant output bits (i.e., not using the MSB of the address) whereas the read addresses are shifted one place to the left (i.e., not using the LSB) and thus increment at half the rate of the write addresses. The write address would, of course, have to be delayed by the pipeline delay of the butterfly processor before being presented to the memory devices.

For the 16-point constant geometry example of Fig. 2 and using dual port RAM with a single PDSP1640, the instruction sequence to program the device is shown in Table 4. Although no new instructions are required between stages, one dummy instruction may be required to keep the data address generation in step with the coefficient addressing.

Cycle No.	Mnemonic	Op. code	Data ^{1,2}	Operation
1	CLRCR	7 _H	XX	CLEAR COUNT/MASK REGISTERS
2	LIRDI	C _H	XX	LOAD INCREMENT REGISTER
3	LCPDI	E _H	01 _H	LOAD COMPARE REG WITH END ADDRESS
4	LS1DI	8 _H	07 _H	LOAD START1 REG WITH BRANCH ADDRESS
5	CCJS1	1 _H	00 _H	COUNT BY INC OR GO TO SR1
6	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
(36)	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1

NOTES

1. XX = don't care
2. Data is in out on cycle following relevant instruction

Table 4 Data addressing instruction sequence, constant geometry algorithm

Coefficient Addressing

From Fig. 2, it can be seen that the coefficient addressing sequence for each stage is a count from 0 to $(N/2 - 1)$ in bit-reversed order or a repeated sub-set of it. The sequence for each stage may be generated by programming the PDSP1640 to count through the full address sequence and at each stage masking different output bits. The address generator outputs are then connected in bit-reversed order to a PROM holding all the coefficient terms.

For the second stage, the least significant bit is unmasked; for the third stage the two least significant bits are unmasked and so on, producing the output address sequence shown in Table 5. The instruction sequence to generate these addresses is as shown in Table 6.

ADDR CNTR REG	1ST STAGE MASK = FF _H		2ND STAGE MASK = FE _H		3RD STAGE MASK = FC _H		4TH STAGE MASK = F8 _H	
	OUTPUT ADDR	BIT RVRSD ADDR	OUTPUT ADDR	BIT RVRSD ADDR	OUTPUT ADDR	BIT RVRSD ADDR	OUTPUT ADDR	BIT RVRSD ADDR
00	00	00	00	00	00	00	00	00
01	00	00	01	80	01	80	01	80
02	00	00	00	00	02	40	02	40
03	00	00	01	80	03	C0	03	C0
04	00	00	00	00	00	00	04	20
05	00	00	01	80	01	80	05	A0
06	00	00	00	00	02	40	06	60
07	00	00	01	80	03	C0	07	E0

Table 5 Coefficient output address sequence for constant geometry algorithm

Cycle No.	Mnemonic	Op. code	Data ^{1,2}	Operation
1	CLRCR	7 _H	XX	CLEAR COUNT/MASK REGISTERS
2	LIRDI	C _H	XX	LOAD INCREMENT REGISTER
3	LCPDI	E _H	01 _H	LOAD COMPARE REG WITH END ADDRESS
4	LS1DI	8 _H	07 _H	LOAD START1 REG WITH BRANCH ADDRESS
5	LMRDI	3 _H	00 _H	LOAD MASK REGISTER
6	CCJS1	1 _H	FF _H	COUNT BY INC OR GO TO SR1
7	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
13	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
14	LMRDI	3 _H	XX	LOAD MASK REGISTER
15	CCJS1	1 _H	FE _H	COUNT BY INC OR GO TO SR1
16	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
22	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
23	LMRDI	3 _H	XX	LOAD MASK REGISTER
24	CCJS1	1 _H	FC _H	COUNT BY INC OR GO TO SR1
25	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
31	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
32	LMRDI	3 _H	XX	LOAD MASK REGISTER
33	CCJS1	1 _H	F8 _H	COUNT BY INC OR GO TO SR1
34	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1
40	CCJS1	1 _H	XX	COUNT BY INC OR GO TO SR1

NOTES

1. XX = don't care
2. Data is input on cycle following relevant instruction

Table 6 Coefficient addressing instruction sequence for constant geometry algorithm

2-D EDGE DETECTOR BOARD AP16401

The AP16401 is a real time two dimensional edge detector module for generating edge presence, magnitude and direction information from 8 bit digitized video.

Constructed on a single Eurocard, the module incorporates a PDSP16401 2-D Edge Detector with two dual port RAMS and all necessary address generation and support logic.

Sync regeneration facilities are provided to maximise video dynamic range, and a blanking pulse output is provided to allow for easy reconstruction of a composite video signal.

A 10MHz crystal controlled clock is provided, though external clocks up to 15MHz may be used.

FEATURES

- PDSP16401 Edge Detector Chip
- Two Line Stores
- Programmable Line Store Delay
- Sync Recovery Circuit
- Black Level Period Circuit
- On-board or External Threshold Setting
- Single Eurocard Format
- 5V Power Supply
- External Clock up to 15MHz

APPLICATIONS

- Robotic Vision Systems
- Pattern Recognition
- Video Effects Generation
- Video Bandwidth Compression




Fig.1 Block diagram of 2-D Edge Detector Board AP16401

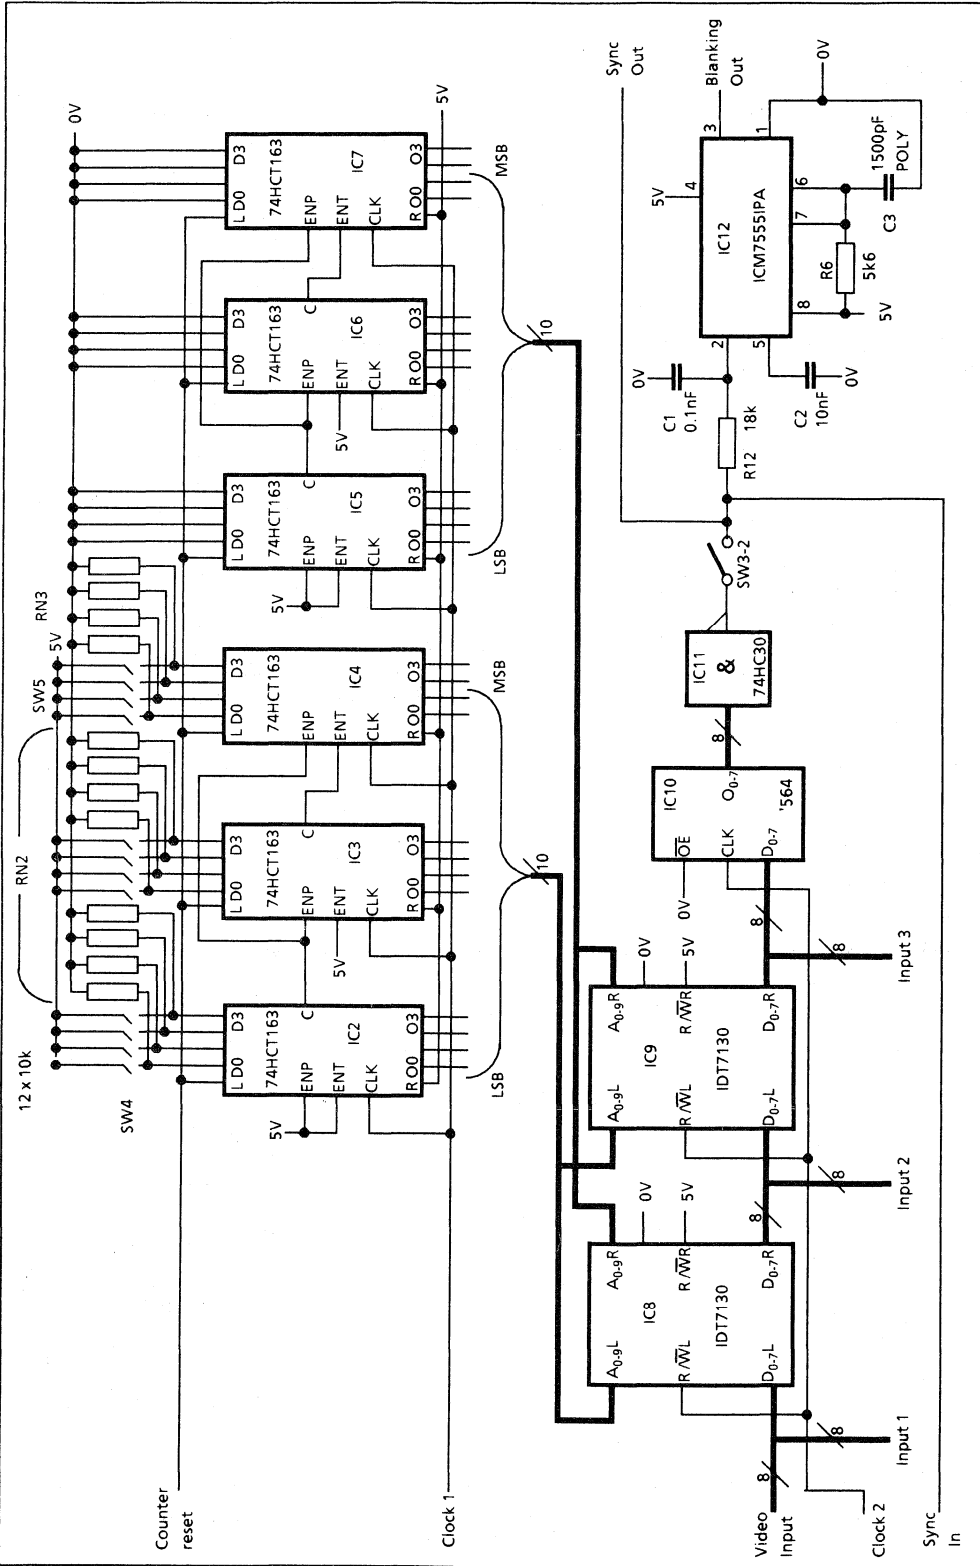


Fig. 2 Delay Lines & Sync Recovery

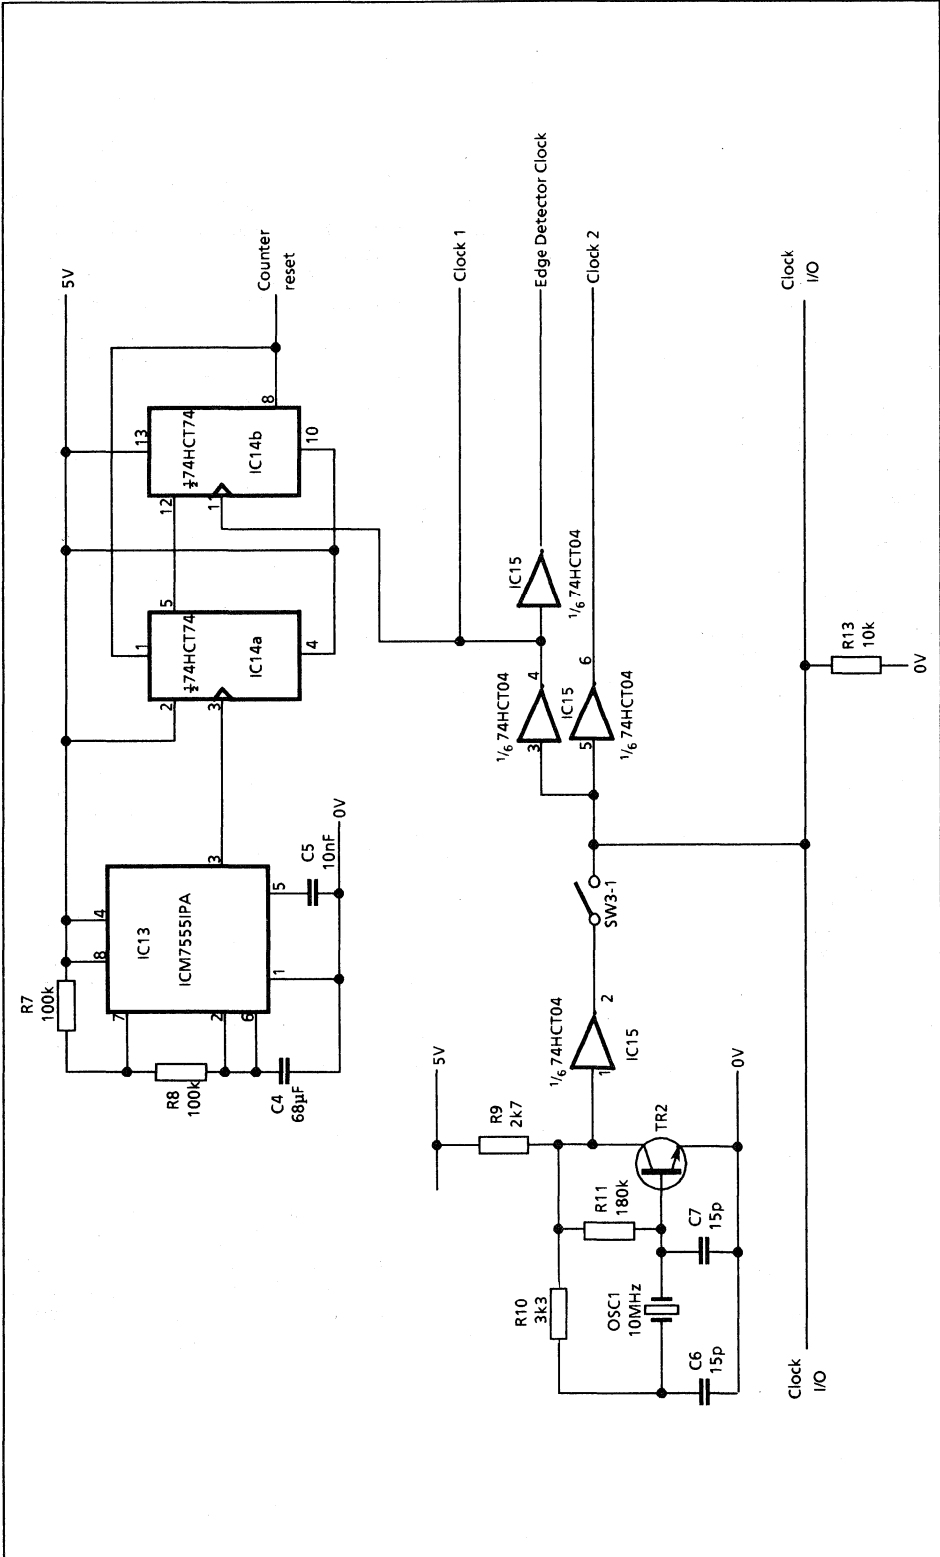


Fig. 3 Clock & counter reset circuits

REAR (DIN41612) CONNECTOR**EDGE DIRECTION OUT**

	EDG1	EDG2	EDG3
Pin	29b/c	30a	30b/c

VIDEO INPUT

Bit	7	6	5	4
Pin	23a	24a	25a	26a
Bit	3	2	1	0
Pin	23b/c	24b/c	25b/c	26b/c

The data on this input must be valid on the rising edge of the clock.

EDGE MAGNITUDE OUT

Bit	12	11	10	9	8	7	6
Pin	19a	20a	20b/c	21b/c	22b/c	15b/c	16b/c
Bit	5	4	3	2	1	0	-
Pin	17b/c	13a	21a	16a	13b/c	11a	-

THRESHOLD INPUT

Bit	9	8	7	6	5
Pin	19b/c	18b/c	12a	9b/c	12b/c
Bit	4	3	2	1	0
Pin	8b/c	11b/c	6b/c	10b/c	7b/c

Note that if the threshold inputs are used, the rotary switches SW1 and SW2 must both be set to zero.

SYNC OUT

6a

BLACK LEVEL PERIOD (blanking)

4b/c

THRESHOLD LATCH

14a

This is pulled high on the board by a 10kΩ resistor.

THRESHOLD EXCEEDED

14b/c

CLOCK I/O

5b/c

FRONT EDGE CONNECTOR

This is a gold plated 30 way connector, with a 0.1inch contact spacing. The contacts on the underside of the board are used for signals. Pin 2 on the upper side of the board is ground. Digitised video can alternatively be input via this connector. The dual DIL switch, SW3, determines whether the clock and sync lines are inputs or outputs. SW3-1 controls the clock line and SW3-2 controls the sync line. The lines are outputs when the switches are closed.

The connections to the front of the card are:

PIN	FUNCTION	PIN	FUNCTION
1	N/C	16	D6
2	N/C	17	D7 MSB
3	N/C	18	N/C
4	N/C	19	N/C
5	N/C	20	N/C
6	N/C	21	N/C
7	N/C	22	SYNC
8	N/C	23	N/C
9	N/C	24	N/C
10	D0 LSB	25	CLOCK
11	D1	26	N/C
12	D2	27	N/C
13	D3	28	N/C
14	D4	29	N/C
15	D5	30	N/C

LINE PERIOD

The video line period is set on SW4 and SW5. The value is calculated by dividing the period of a line by the period of the 10MHz clock and subtracting 3. For example, in the UK, this value is $(64\mu\text{s} / 0.1\mu\text{s}) - 3 = 637$. This number is converted to binary and set on the switches. The value of each switch position together with the common settings are:

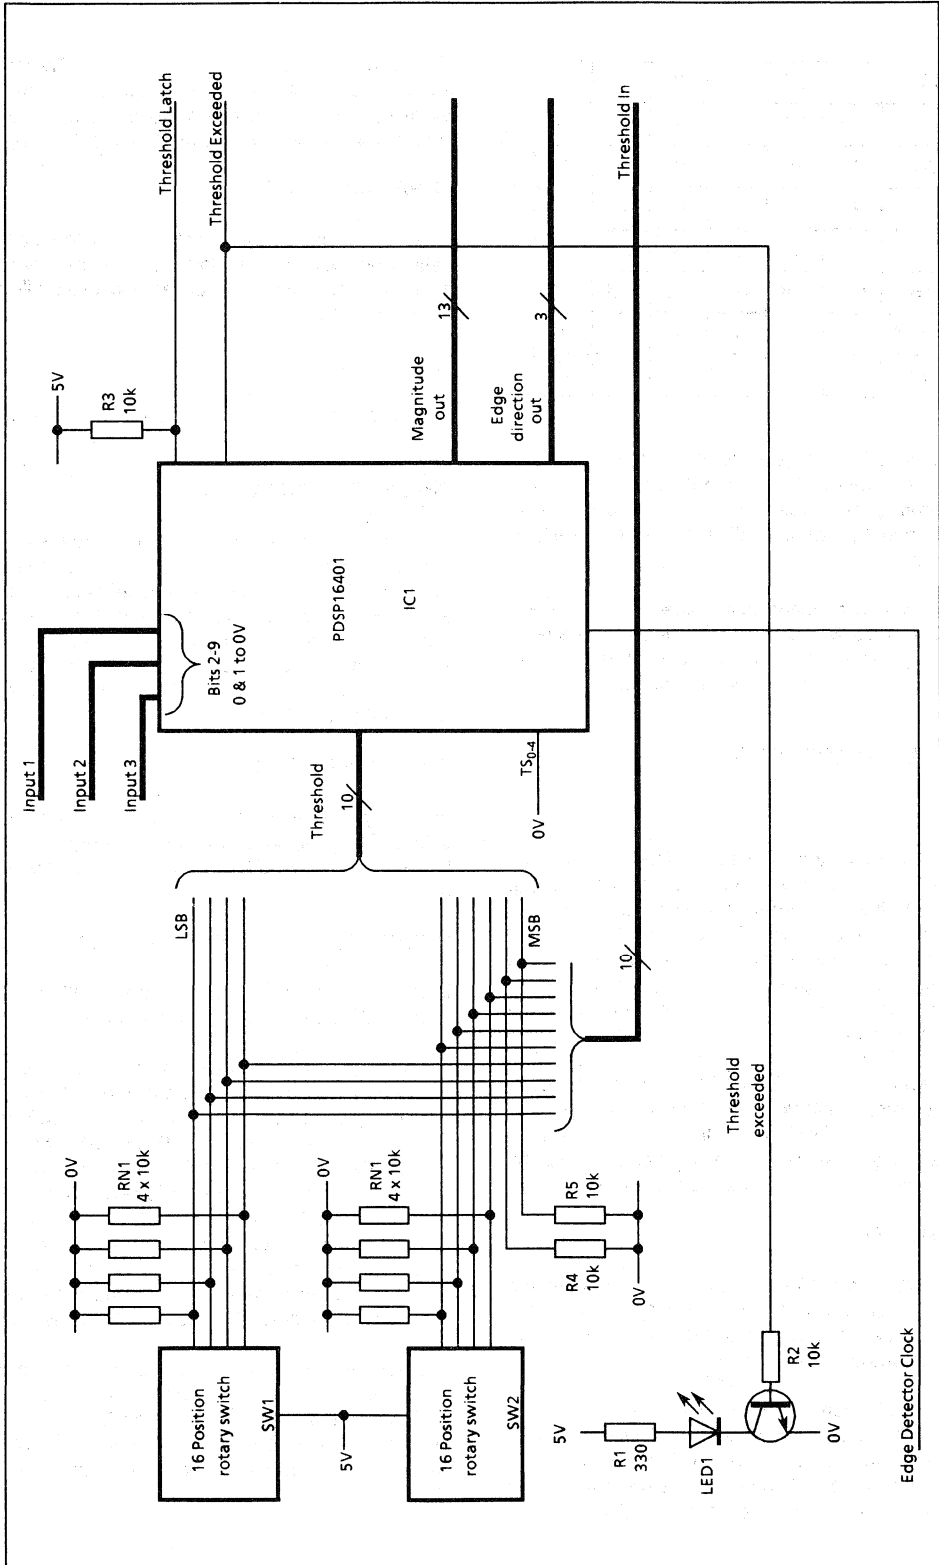


Fig. 4 PDS16401 connections

Value	2 ¹¹	2 ¹⁰	2 ⁹	2 ⁸	2 ⁷	2 ⁶
Switch	5-1	5-2	5-3	5-4	4-5	4-6
625/50Hz*	0	0	1	0	0	1
525/60Hz*	0	0	1	0	0	1
Value	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
Switch	4-7	4-8	4-1	4-2	4-3	4-4
625/50Hz*	1	1	1	1	0	1
525/60Hz*	1	1	1	0	0	0

* These switch settings assume normal interlacing. In this case, note that the edge detector will actually be operating on three horizontal by six vertical screen pixels, which will lose some vertical resolution.

The settings are periodically reloaded by a pulse derived from IC13 and IC14. This is done to ensure that the system will set itself up again in the event of a power supply glitch upsetting the counters.

SIGNAL INTERFACING

VIDEO

Video is input to the card in digitised form on the rising edge of the clock. The word width can be up to 8 bits. Either the front or rear edge connectors can be used. If it is desired to convert the output of the card back into composite video form, it is advantageous to digitise the video so that it is all zeros during sync pulses. This is detected by a circuit at the output of the second line store and converted back into a negative going sync pulse. This is available at the SYNC OUT pin on the rear of the card. Also available when SYNC OUT operates is the BLACK LEVEL PERIOD output. This is a positive going pulse which can be used to clamp the reconstructed composite video to black level.

THRESHOLD

The threshold is a 10 bit number which is compared with the magnitude of the current edge. If the edge magnitude exceeds the threshold then the threshold exceeded output will be active.

The threshold can be set by using the two rotary switches, or input from another card using the rear connector. On the card, the ten threshold lines are pulled low by 10k Ω resistors, with the rotary switches, which only operate on the 8 least

significant bits, taking the lines to 5V. Therefore if the rear connector is used, both of the switches **must** be set to zero, otherwise the external device driving the card will be shorted to the 5V supply.

THRESHOLD LATCH

Threshold latch enables clocking in of the threshold data when high. It is pulled high on the card by a 10k Ω resistor. This pin is pulled low externally to fix the threshold value.

MAGNITUDE

This is a measure of the magnitude of the current edge.

THRESHOLD EXCEEDED

Either all of the data from the edge magnitude output can be used, or only data which exceed the threshold. The threshold exceeded signal indicates when this occurs and can be used, for example, to enable a gate passing the magnitude data. Alternatively, if just an outline is required, it is only necessary to use the threshold exceeded output which indicates where the edge is.

EDGE DIRECTION OUT

Each edge has a direction which takes one of eight values. These specify whether the edge is vertical, horizontal or diagonal, and which side of the edge is brighter.

CONSTRUCTING COMPOSITE VIDEO

If composite video is required, the SYNC OUT and BLACK LEVEL PERIOD outputs are used.

For example, the simplest method of displaying edges on a monitor is to use the threshold exceeded output to produce black lines on a white background, see Fig. 5. To do this, THRESHOLD EXCEEDED and BLACK LEVEL PERIOD are passed through a NOR gate. A resistor network is used to add this to SYNC OUT so that the magnitude of the sync pulse is 0.66V below black and the video white level is 1.33V above black. This is buffered with a 75 Ω driver for a monitor with a 75 Ω input. If the monitor input is unterminated, the above voltages are halved.

Another possible video effect is to have a white background, and use the edge magnitude with a DAC to draw the edges in varying levels of grey. If the greatest magnitude edges produce the blackest lines the effect is very much like a moving pencil sketch.

CIRCUIT DESCRIPTION

The PDSP16401 requires three digitised video inputs. Inputs 2 and 3 must be delayed by one and two line periods respectively. This is implemented by the circuitry in Fig. 2. IC2, IC3 and IC4 are periodically loaded with the value set on the switches SW4 and SW5, which set the time delay per line. IC5, IC6 and IC7 are loaded with all zeros. These counters operate the address inputs of the dual port RAM chips, IC8 and IC9. The left ports are written under control of the address generated by IC2, IC3 and IC4 and the right ports are read under control of the address generated by IC5, IC6 and IC7. Thus there is a delay between a byte being written into the left port of one of the RAMs and read from the right port. These two delayed video signals, along with the current video signal, are fed into the PDSP16401.

Also in Fig. 2 is a digital sync detector, which operates by decoding bytes with all zeros in the digitised video. A monostable is also provided for the purpose of clamping the video for the period

of the sync pulse and black level. Under some circumstances an analogue sync detector may be more successful than the digital one. A suitable circuit is shown in Fig. 6. This circuit has the advantage of also providing a DC level clamped video signal which can be input to the ADC.

Fig. 3 shows the system clock, which is a simple crystal controlled transistor oscillator with inverters to provide buffering and different phases.

The reset circuit periodically provides a pulse which reloads the original starting values into the counters. This was provided to make the system restart itself in the event of a glitch on the power supplies.

Fig. 4 shows the connections to the PDSP16401 and the method of setting the threshold input, which can be done using either the rotary switches or the edge connector. The LED is provided to give an indication of the activity of the threshold exceeded output.




Fig. 5 Example circuit to produce composite video from threshold exceeded output

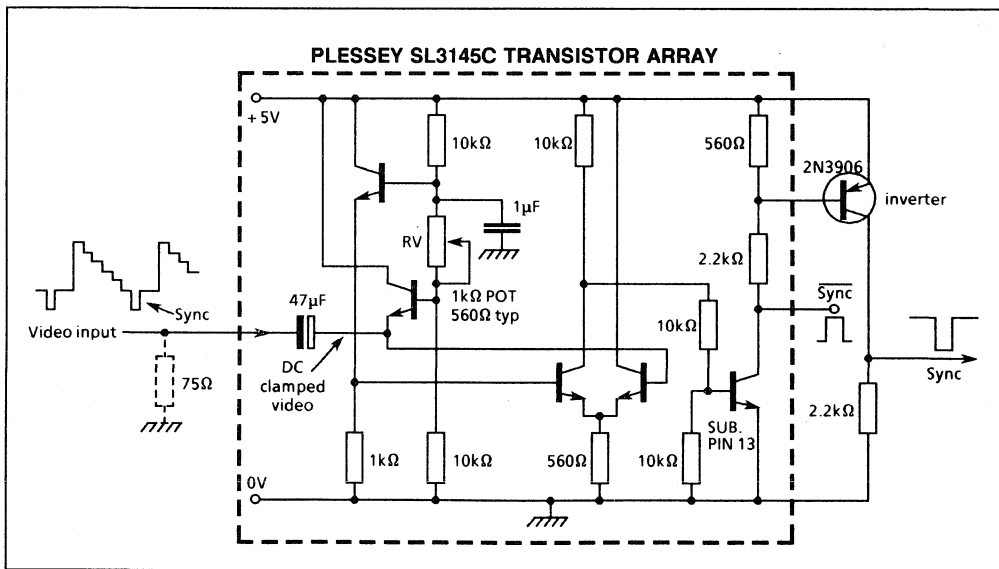


Fig. 6 Sync separation and DC level clamping circuit

SOBEL v. PDSP16401 OPERATORS

The Sobel operator in common use is given by:

$$S(x) = \left| (x) * \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix} \right| + \left| (x) * \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix} \right|$$

The four-operator system used by the PDSP16401 gives a magnitude output which is the greatest of the four convolutions of the operators below with the input array:

$$1.5 * \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix} \quad 1.5 * \begin{bmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{bmatrix}$$

$$\begin{bmatrix} 2 & 1 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & -2 \end{bmatrix} \quad \begin{bmatrix} 1 & -1 & -2 \\ 1 & 0 & -1 \\ 2 & 1 & 0 \end{bmatrix}$$

The characteristic of an ideal edge enhancement operator is that it gives a constant degree of enhancement irrespective of the orientation of the edge. To test this characteristic, a test edge of constant gradient and variable angle is used:

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \quad \text{(i)} \quad \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{(ii)}$$

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{(iii)} \quad \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{(iv)}$$

$$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{(v)} \quad \text{etc.}$$

The Table below summarises the effect of the Sobel and Plessey operators on the above test edge.

Orientation	Sobel	Plessey
i	3	4.5
ii	4	4
iii	3	4.5
iv	4	4
v	3	4.5

As can be seen, the variation in output magnitude with edge orientation with the operator set from the PDSP16401 is much smaller than that experienced with the Sobel operator.



A HIGH RESOLUTION FFT PROCESSOR USING THE PDSP16116/A

The PDSP16116/A has been designed with an integral Block Floating Point system which can be used, in conjunction with other 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 259 μ s
- 512 point complex radix-2 transform in 118 μ s
- 256 point complex radix-2 transform in 53 μ s

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 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' ($\log_2 N$ passes for an N -point transform), each pass consisting of $N/2$ 'butterfly' operations:



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 butterfly 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 occurring: 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:

```

0101110101110100
+ 0110100110101001
-----
01100011100011101

```

Wordgrowth does not occur:

```

0101010111011010
+ 0010100010101110
-----
00111111010001000
↑
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:

$$\begin{array}{r}
 \downarrow \text{1 bit sign extension} \\
 0011010001010110 \\
 * 0010001101000101 \\
 \hline
 00000111001101011101110110100101110 \\
 \uparrow\uparrow\uparrow\uparrow \text{4 bit sign extension}
 \end{array}$$

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 occurred 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 wordgrowth 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.XXXXXXXXXXXXXXXXXX	word tag = 00
XXX.XXXXXXXXXXXXXXXXXX	word tag = 01
XXXX.XXXXXXXXXXXXXXXXXX	word tag = 10
XXXXX.XXXXXXXXXXXXXXXXXX	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.XXXXXXXXXXXXXXXXXX	word tag = 00
XXXXX.XXXXXXXXXXXXXXXXXX	word tag = 01
XXXXXX.XXXXXXXXXXXXXXXXXX	word tag = 10
XXXXXXX.XXXXXXXXXXXXXXXXXX	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.XXXXXXXXXXXXXXXXXX

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:

XXXXXXXXXXXX.XXXXXX

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

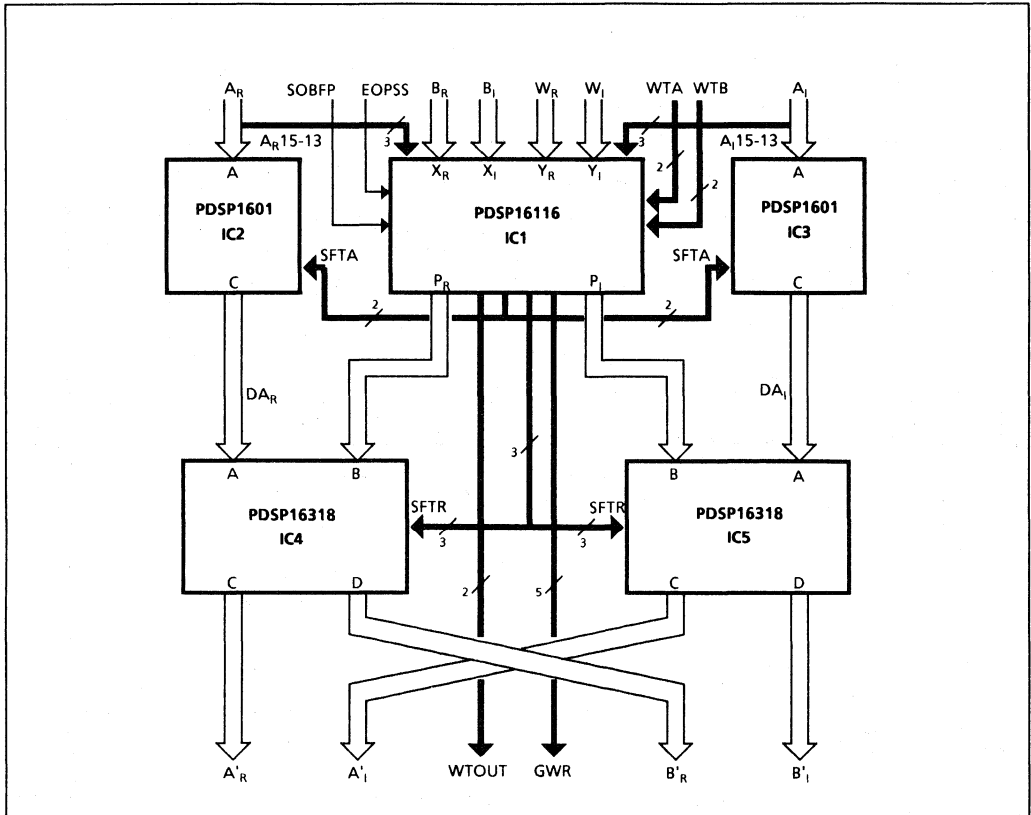


Fig. 4 Block Floating Point FFT Butterfly

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 Plessey PDSP16318s and two Plessey 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).

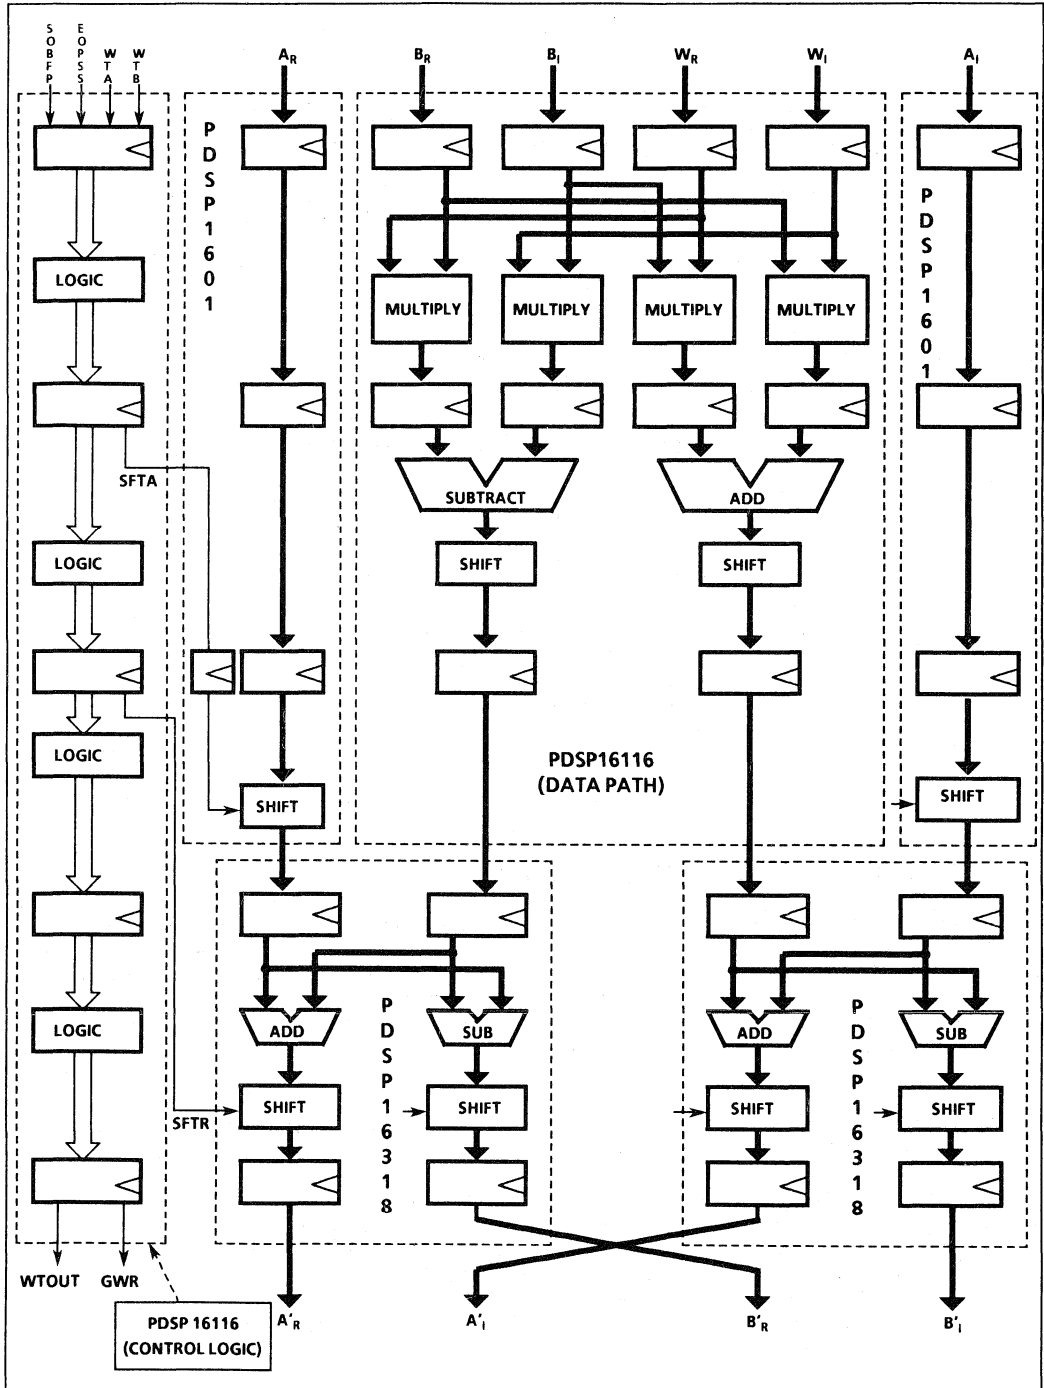


Fig. 5 BFP Butterfly Detail




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.

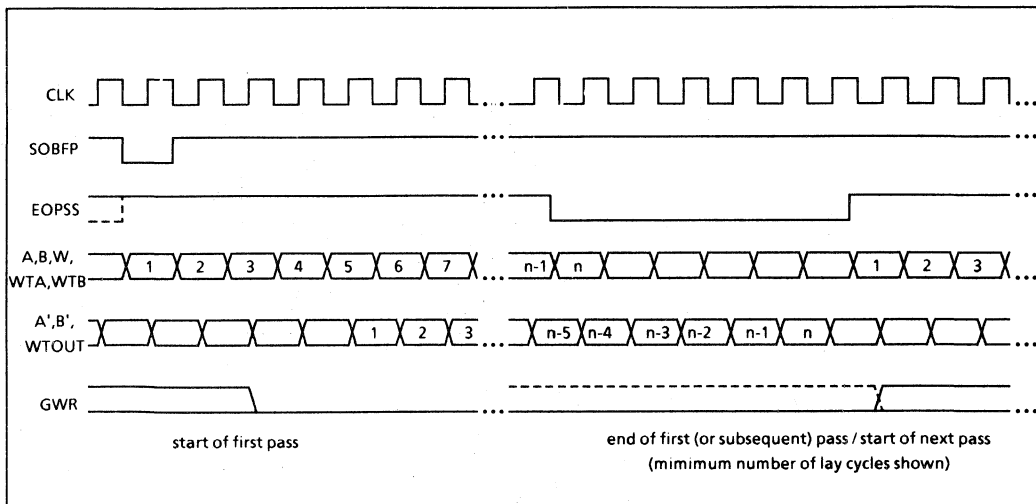


Fig. 7 Use of BFP Control Signals

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. Plessey Semiconductors makes 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. The PDSP1640 is ideal for generating the data and coefficient addressing sequence. Output normalisation is a simple matter using the PDSP1601's adaptable barrel shifter and the PDSP16330 Pythagoras Processor to convert Cartesian to polar co-ordinates.

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.

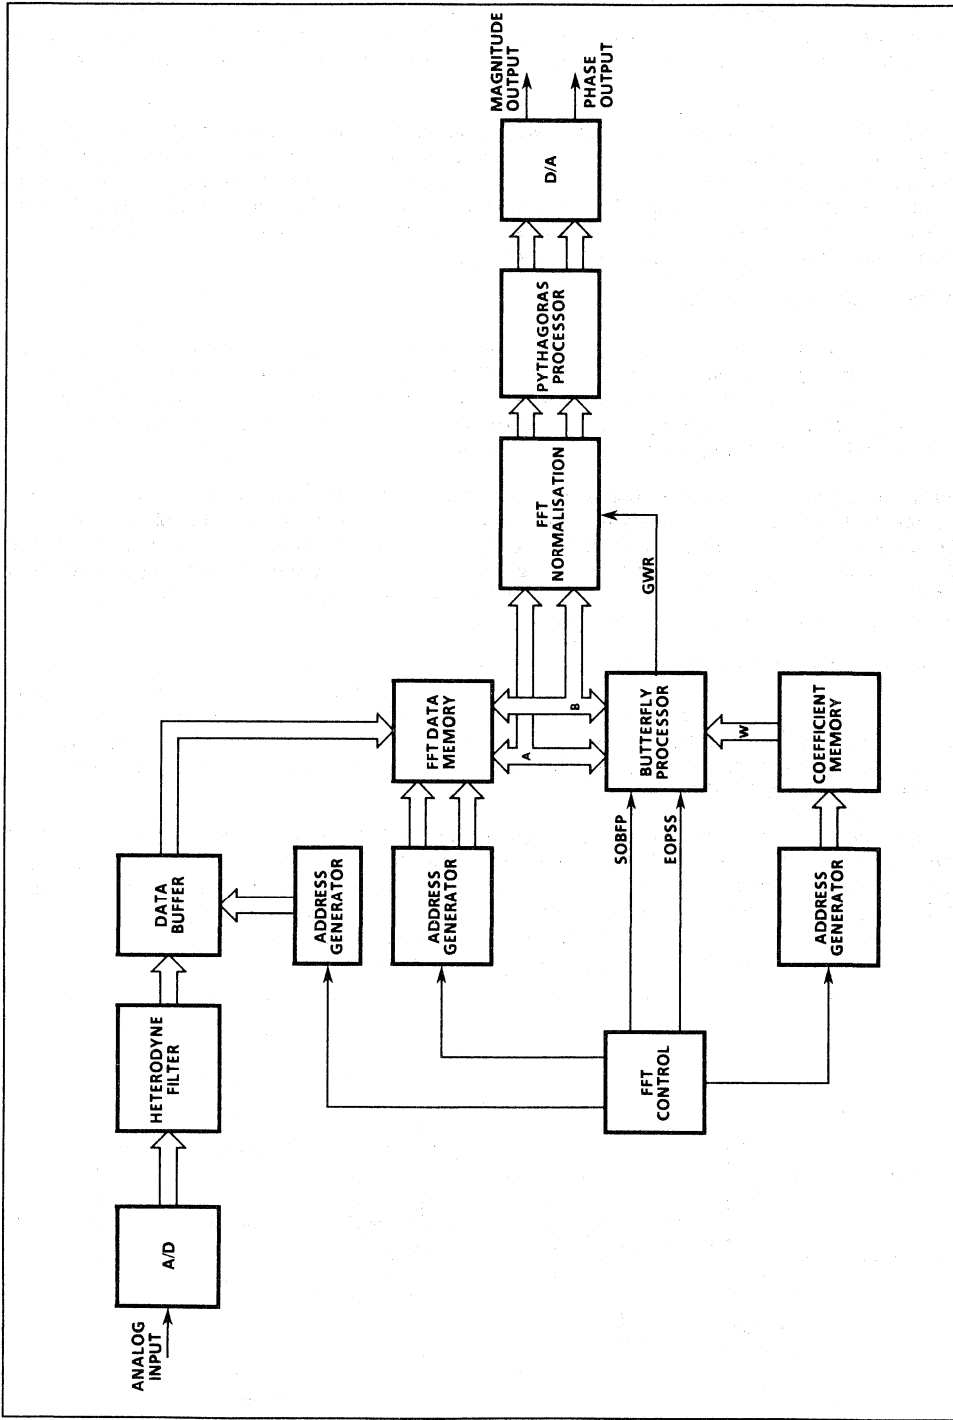


Fig. 8 Typical FFT System

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)

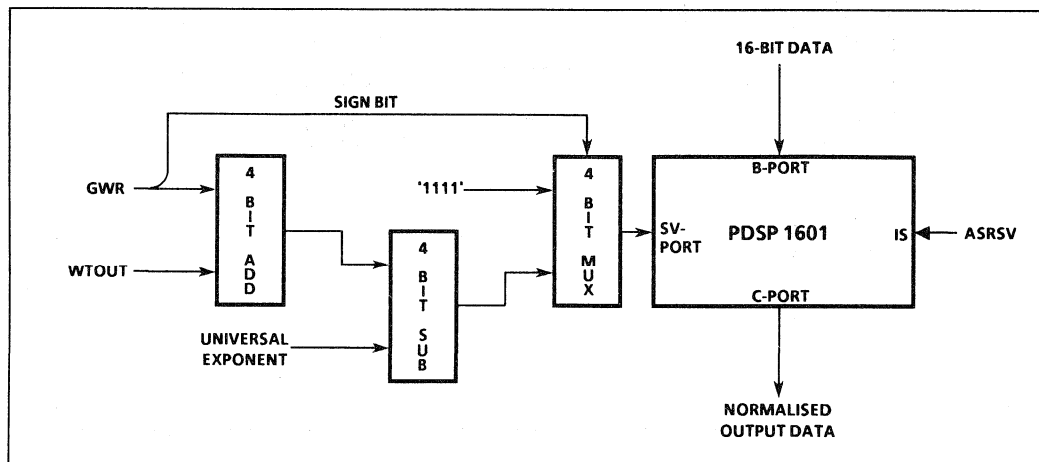


Fig. 9 Output Normalisation Circuitry

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	OE1		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	Y115	W115	external i/p
J3	Y114	W114	external i/p
J1	Y113	W113	external i/p
K1	Y112	W112	external i/p
J2	Y111	W111	external i/p
K2	Y110	W110	external i/p
K3	Y19	W19	external i/p
L1	Y18	W18	external i/p
L2	Y17	W17	external i/p
M1	Y16	W16	external i/p
N1	Y15	W15	external i/p
M2	Y14	W14	external i/p
L3	Y13	W13	external i/p
N2	Y12	W12	external i/p
P1	Y11	W11	external i/p
M3	Y10	W10	external i/p
N3	X10	B10	external i/p
P2	GND	GND	0V supply rail
R1	VDD	VDD	+ 5V supply rail
N4	X11	B11	external i/p
P3	X12	B12	external i/p
R2	X13	B13	external i/p
P4	X14	B14	external i/p
N5	X15	B15	external i/p
R3	X16	B16	external i/p
P5	X17	B17	external i/p
R4	X18	B18	external i/p
N6	X19	B19	external i/p
P6	X110	B110	external i/p
R5	X111	B111	external i/p
P7	X112	B112	external i/p
N7	X113	B113	external i/p
R6	X114	B114	external i/p
R7	X115	B115	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	EOP5S	EOP5S	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 : PDSP16318 - 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'17	external o/p
C2	D8	B'18	external o/p
B1	D9	B'19	external o/p
C1	D10	B'110	external o/p
D2	GND	GND	0V supply rail
D1	VDD	VDD	+ 5V supply rail
E3	D11	B'111	external o/p
E2	D12	B'112	external o/p
E1	D13	B'113	external o/p
F2	D14	B'114	external o/p
F3	D15	B'115	external o/p
G3	C15	A'115	external o/p
G1	C14	A'114	external o/p
G2	C13	A'113	external o/p
F1	C12	A'112	external o/p
H1	VDD	VDD	+ 5v supply rail
H2	GND	GND	0V supply rail
J1	C11	A'111	external o/p
K1	C10	A'110	external o/p
J2	C9	A'19	external o/p
L1	C8	A'18	external o/p
K2	C7	A'17	external o/p
K3	C6	A'16	external o/p
L2	C5	A'15	external o/p
L3	C4	A'14	external o/p
K4	C3	A'13	external o/p
L4	C2	A'12	external o/p
J5	C1	A'11	external o/p
K5	C0	A'10	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	M5		tie low
L6	AS11		tie high
L8	AS10		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'10	external o/p
A4	D1	B'11	external o/p
B4	D2	B'12	external o/p
A3	D3	B'13	external o/p
A2	D4	B'14	external o/p
B3	D5	B'15	external o/p
A1	D6	B'16	external o/p

REFERENCES

For a general introduction to FFTs the following texts are recommended :

1. Rabiner and Gold, '*Theory and Application of Digital Signal Processing*' (Prentice - Hall, 1975)
2. Oppenheim and Shafer, '*Digital Signal Processing*' (Prentice - Hall, 1975)

Other 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'
- AN54 FFT Address Generation using PDSP1640
- 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.

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 development of Plessey's PDSP family of Complex Digital Signal Processing Building Blocks, and the bit widths selected for the data and coefficient paths. Plessey 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 Plessey PDSP family of complex signal processing building blocks.

The relevant sections of an FFT system that determine the dynamic range 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 time 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-emphasis 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	1024	FFT Size			
			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 additional 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 only 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 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 converter 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 Butterfly processor that supports block floating point arithmetic automatically allowing FFT System 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 Plessey 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 (sampled 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 signal 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.

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'.

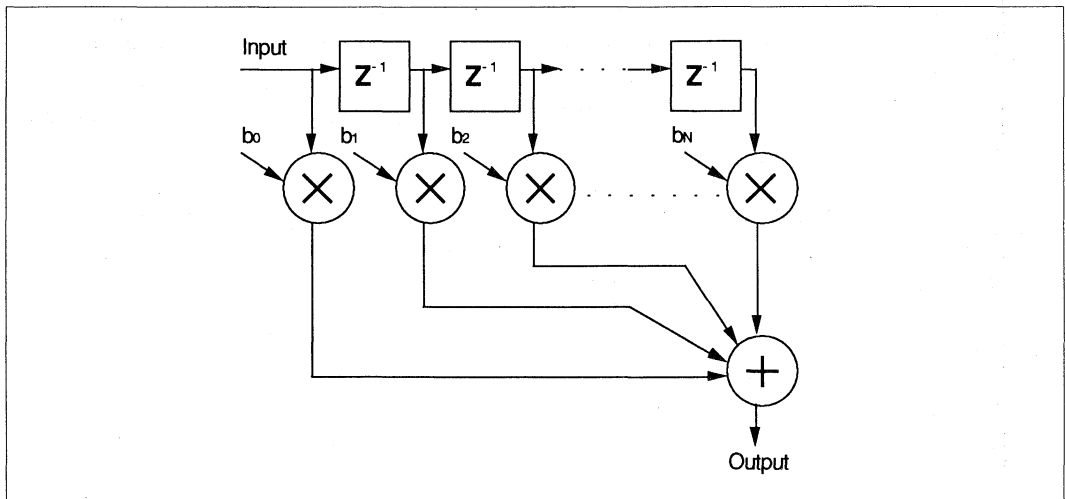


Fig 1 FIR Filter Structure.

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:

$$y_n = b_0x_n + b_1x_{n-1} + \dots + b_{N-1}x_{n-N+1}$$

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.

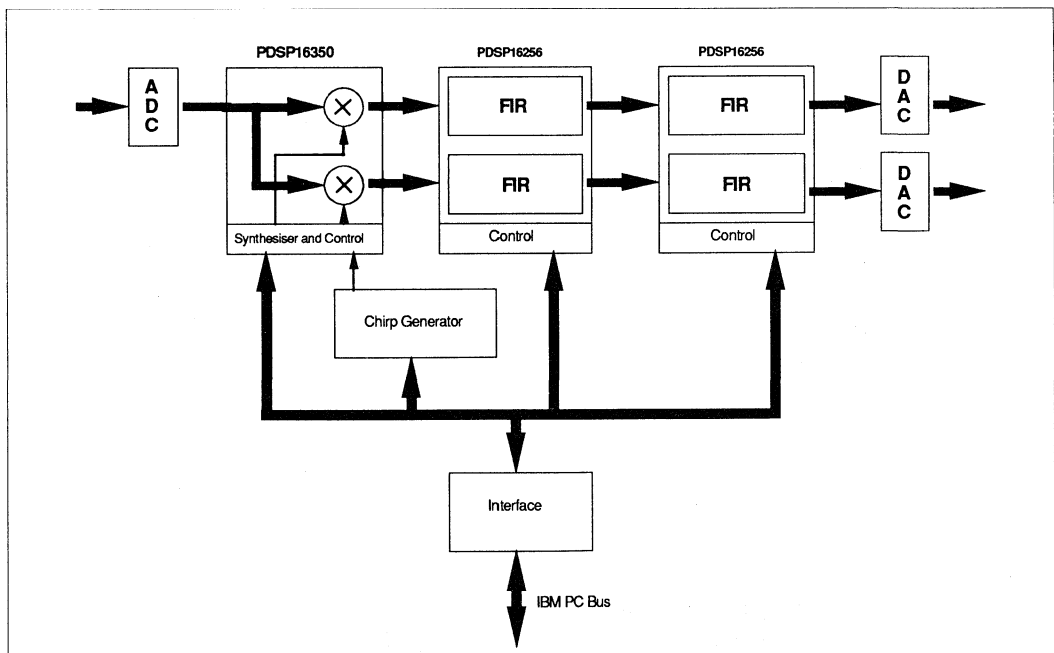


Fig 2 Schematic of Development Board.

The development board is arranged as shown in Fig. 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 PDSP16256s. 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 DACs. 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.

Fig. 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.

Fig. 4 shows the frequency response of a filter designed to the same specification as the one shown in Fig. 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 its simplest it can be configured to auto load the filter

coefficients directly from EPROM on power up and be directly connected to ADCs and DACs (Fig. 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).

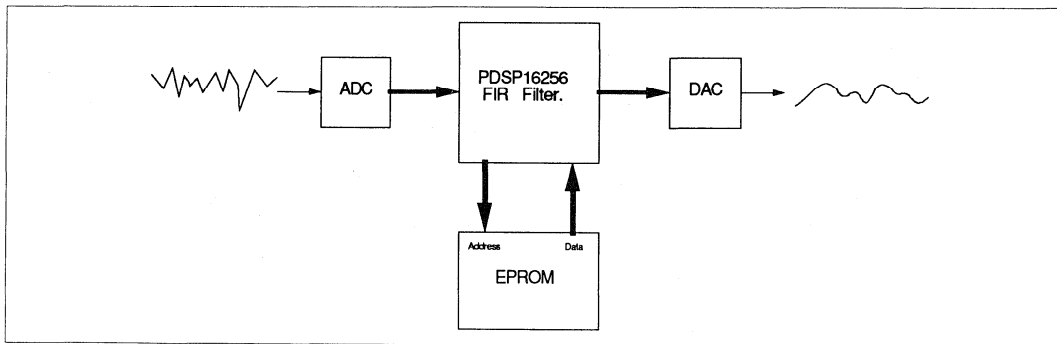


Figure 6 Simple Auto Load Configuration.

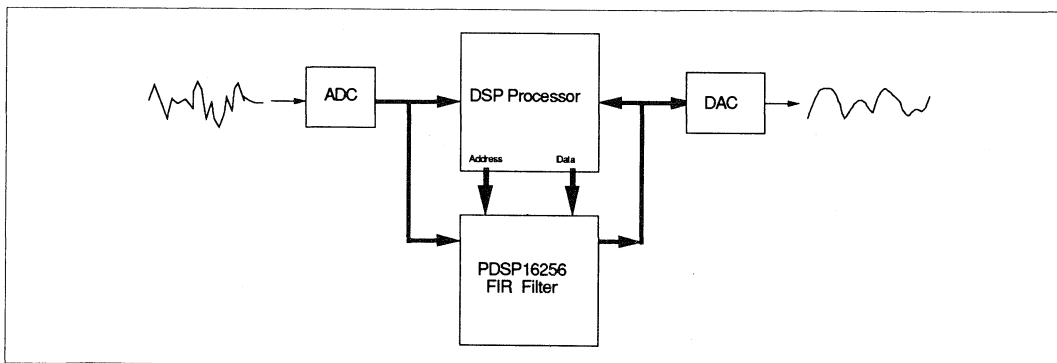


Figure 7 Slave Processor Configuration.

Support tools

PDSP Demonstrator

The PDSP Demonstrator is a microprogram development tool for the PDSP device family offering device and system simulation facilities based upon functional models of the devices. It includes a powerful line editor for microprogram file preparation, a simulator with user control of program execution and an interactive trace facility which incorporates a print and plot function. It runs on an IBM-PC or any compatible machine under the MS-DOS operating system, the minimum configuration comprising one 360Kb floppy disc unit and 512Kb of RAM. It is therefore a convenient and powerful method of introducing the new user to the devices by 'animating' the traditional data sheet, giving a real ability to experiment with particular algorithms and speed breadboard system development.




Fig.1 Demonstrator functional block diagram

FEATURES

- DSP Microprogram Development Environment
- Device/System Simulator with Programmable Breakpoint Control
- Interactive Trace Facility Including Print Function
- Graphical Representation of I/O Data
- Powerful Microprogram Editor
- High Level Program Construction using Macros
- User Friendly Interface
- Standard Application Microprogram Library
- IBM-PC Compatible Software on a Single Floppy Disc

APPLICATIONS

- Application Microcode Development
- Debugging Aid
- Test Vector Generation
- Training Aid

THE PDSP DEMONSTRATOR RANGE

The Demonstrator is available in several versions offering facilities ranging from the simulation of individual devices to the simulation of complete system solutions to common DSP problems.

The following programs are currently available in the PDSP Demonstrator range:-

Device Function Demonstrator 1

Supports programming and simulation of the following devices:

- **PDSP1601/A** ALU and Barrel Shifter
- **PDSP1640** 20MHz Address Generator
- **PDSP16112/A** Complex Number Multiplier
- **PDSP16318/A** Complex Accumulator

Each model is a faithful reproduction of the real device and is programmed in accordance with its data sheet specification. All four models are programmed as individual devices only.

COMPLEX ARITHMETIC DEMONSTRATOR

Includes the same four devices as the Device Function Demonstrator 1 but supports the programming and simulation of these devices in the following system configurations:

- Dual 1640 demonstrating 16 bit addressing
- Fast Fourier Transform Processor comprised of one PDSP16112 and two PDSP16318 devices
- Complex Arithmetic Processor comprised of one PDSP16112 and one PDSP16318 device

The PDSP1601 device model is retained for programming and simulation as an individual device.

GETTING STARTED

Loading the ANSI Device Driver

The ANSI device driver is an independent software module used by the operating system to drive the display screen. It incorporates facilities for cursor movement and screen erasure which the PDSP Demonstrator requires.

To enable MS-DOS to load the ANSI device driver the CONFIG.SYS file must contain the statement:-

```
DEVICE = ANSI.SYS
```

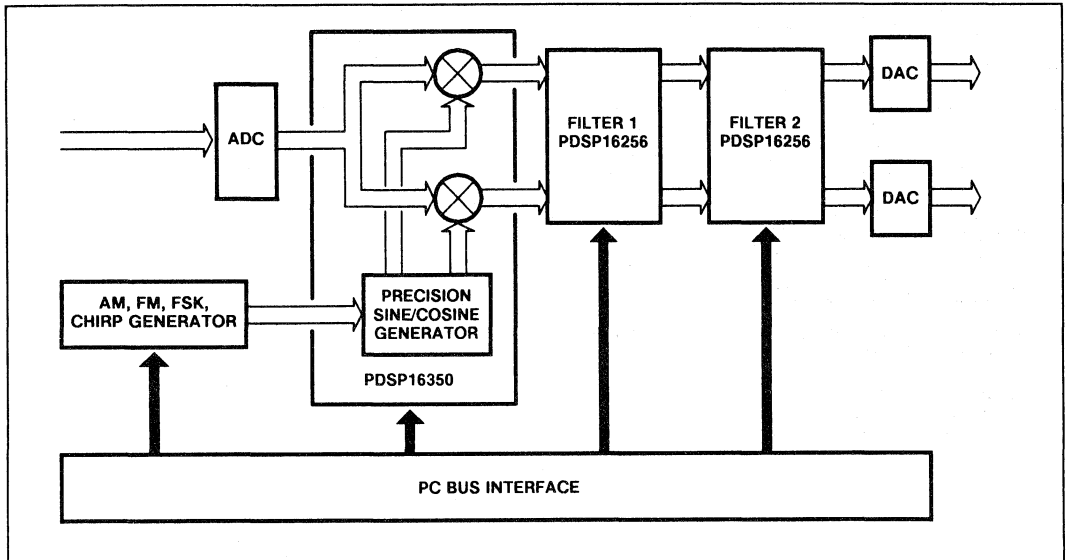
You can use the line editor EDLIN to modify the CONFIG.SYS file in the root directory as follows where <CR> implies carriage return:-

```
Enter EDLIN CONFIG.SYS <CR>
Enter I <CR>
Enter DEVICE = ANSI.SYS <CR>
Press CTRL-C <CR>
Enter E <CR>
```

The CONFIG.SYS file is now prepared for MS-DOS to set the correct running environment for the PDSP Demonstrator. Refer to your PC manual for further details about EDLIN or the CONFIG.SYS file. Re-boot the PC Operating System before using the Demonstrator.

The PDSP Demonstrator is available NOW from Plessey Semiconductors.

**AN IBM PC COMPATIBLE DEVELOPMENT SYSTEM
FOR HIGH PERFORMANCE DIGITAL FILTERING AND SIGNAL GENERATION
USING THE PDSP16256 AND PDSP16350**



The evaluation 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 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 Plessey Semiconductors devices.

APPLICATIONS

The evaluation system is applicable to a wide range of areas, including:

- | | |
|--------------------|----------------------------|
| • Radar | • Sonar |
| • Ultrasonics | • Data communications |
| • Video processing | • Instrumentation |
| • Digital radio | • 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 20MHz
- Digital output ports

Digital Filtering

- Configurable as separate I and Q channels or as a single real only channel
- Cascaded or single filter options
- Filter lengths between 8 and 128 coefficients (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

PDSP16256/PDSP16350 EVALUATION SYSTEM

DEVELOPMENT SOFTWARE

- Filter 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
- Supporting documentation

GEC Plessey Semiconductors wishes to thank Loughborough Sound Images Ltd.
for giving permission to reproduce the following leaflet




Loughborough Sound Images Limited

PDSP16488 Real-Time Digital Image Processor Board



- *Features the Plessey Semiconductors PDSP16488 Digital Image Convolver*
- *Fully Programmable 8 x 8 convolution window over a 512 x 512 (8 bits/pixel) digitised image in **Real-Time***
- *Direct Camera Input and RGB Monitor Output*
- *Linear Operations - Filtering, Edge Detection, Field Differencing*
- *Non-Linear Operations - Contrast Enhancement, Binary Images, Pseudo Colouring, Histograms*
- *PC, XT or AT Plug-In Board*
- *Menu-Driven Software Support*



Applications

Though primarily designed as an evaluation board for the PDSP16488 the board offers a complete or partial solution to many image processing problems in both research, education and industrial applications.

As an evaluation tool the board demonstrates the impressive processing power of the chip. By providing all the additional circuitry for image digitisation, processing and display, the board allows real-time imaging applications to be quickly realised without prototyping delay.

Different filter characteristics can be easily programmed and compared without lengthy computer simulation runs and combined linear/non-linear image processing algorithms can be realistically evaluated on moving images.

The instant response of the board to coefficient changes or alternative transfer functions provides a powerful aid to teaching image processing. Different filter characteristics can be quickly illustrated. Both traditional and new image processing algorithms can be readily programmed.

A number of algorithms can be readily applied to existing problems in image processing such as contrast enhancement, edge and motion detection

and noisy image filtering. The ability of the board to apply these algorithms in real-time offers a complete solution in a number of fields such as surveillance, infra-red imaging, image enhancement, histogram equalisation and others previously limited by processing power or cost.

User Support

The board is supplied complete with user manual and comprehensive software support. The manual contains full technical details of the board's operation including software examples of how to drive the board from a user's own programs. These are illustrated by a simple menu program written in Microsoft 'C' which allows full control of all board facilities at a very basic level. Well documented source code shows how to read and write to the various data and control registers from within a high level language environment.

Interactive Control Program

A full graphics-based control program has been written for the board by Plessey Semiconductors. This program, which requires a minimum of EGA graphics, provides a friendly visual means of controlling board functions using cursor keys and menu selection. The full status of the board is displayed continuously on a single interactive graphics screen.

PDSP16488 Real-Time Digital Image Processing Board

Processor

This high performance chip integrates an array of multiplier-accumulators with 32K bits of RAM configured as line delays to provide a complete two dimensional digital convolver. The chip performs a weighted sum of all pixels within an $N \times N$ 2D window. Each pixel value is multiplied by a signed coefficient and the products summed to provide a result for every pixel clock. Positive coefficients produce averaging effects or -ve coefficients can be used for edge enhancement.

At pixel rates of up to 10 MHz, an 8×8 convolution window can be continuously scanned over a 512 pixel by 512 line image in real-time, achieving in excess of 640 million, 8 bit multiplication-accumulation operations per second.

Coefficient weights, output gain and saturation, multiplier and line delay configurations are all programmable by means of an 8 bit microprocessor interface.

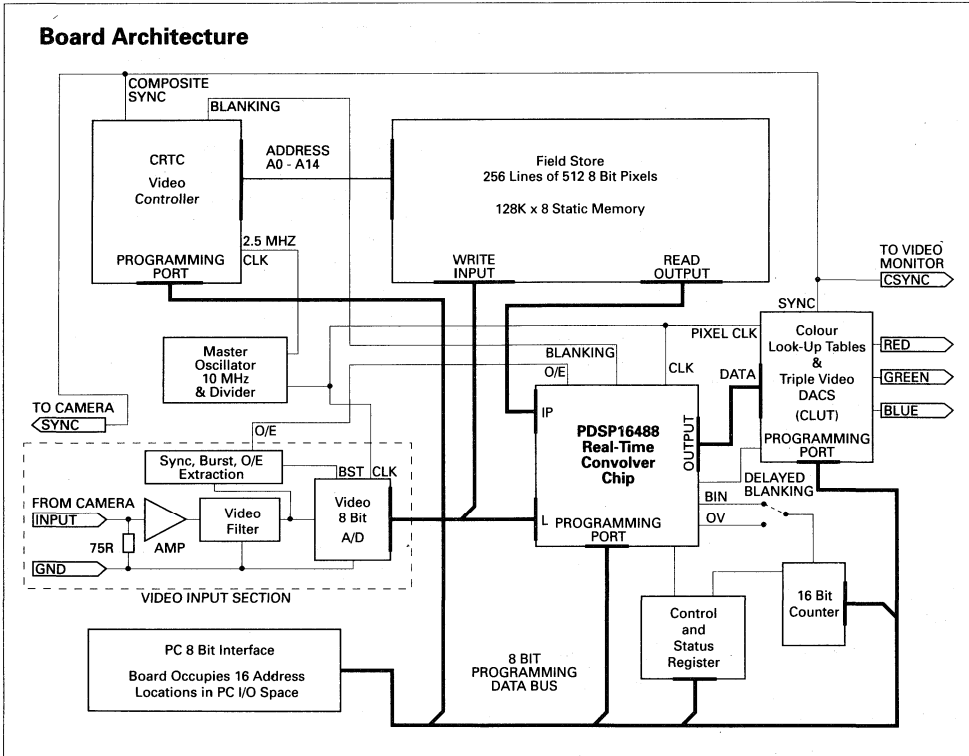
Board Features

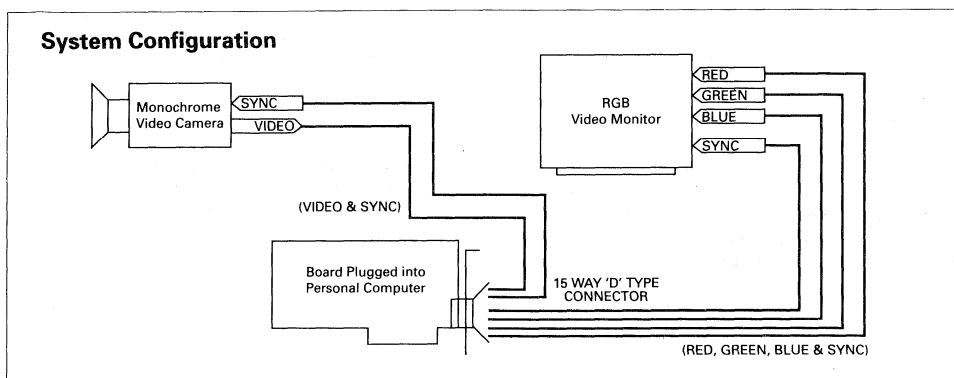
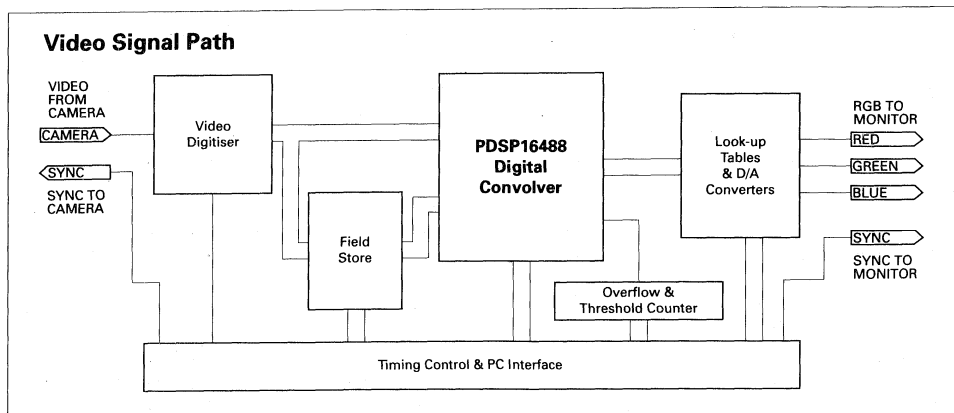
This Processor Board is designed to support and extend the capabilities of the PDSP16488 convolver by adding picture digitisation, a field store to de-interlace the video information, colour look-up

table and video output facilities all under the control of a Personal Computer. Video information from a monochrome camera is digitised with a resolution of 8 bits per pixel. This is fed both directly, and via a field delay, to the convolver chip which continuously operates in real-time. Results are quantised to 8 bit resolution and passed through a triple colour look-up table and video DAC's to give a programmable colour monitor display.

The video screen format is controlled by an on-board CRT Controller which generates sync signals for the camera and video monitor. Thus these are both gen-locked to the board giving a very stable, crystal controlled display. A number of video formats are programmable, including a 512 x 512 display at 50Hz and 512 x 480 display at 60Hz. A non-interlaced sync can be generated to give a 512 x 256 display, which can be captured and held within the field store.

A programmable digital comparator within the convolver, coupled with a 16 bit counter on the board allows histogram information to be obtained on stationary images. This can be used to reprogram the colour palette for contrast enhancement, binary, inverse or pseudo colour images.





Control of the board is split into three main areas:

- Firstly, linear convolution is supported by a display of the 8×8 array of coefficients which may be edited using cursor control and keyboard entry. Single commands will clear the array and conveniently generate circular symmetric filters from a single line of coefficients. Full control of output gain is by means of cursor keys.
- The output function of the colour look-up tables is displayed graphically and can be edited, again using cursor keys, to generate non-linear functions for contrast enhancement and binary imaging. Menu options allow the selection of output saturation modes or the use of pseudo-colouring palettes.

- Finally, the full status of the board can be saved or loaded from disc so that a library of effective routines can be developed. Further menu options allow images to be frozen within the field store, the generation of non-interlaced displays and field differencing, and the display of histogram statistics on stationary images.

Ordering Information

PC/16488 PDSP16488 Real-Time Digital Image Processing Board and Software

We are continually improving our products, and reserve the right to alter the above specifications at anytime, without notice.

Designed and manufactured by:



**Loughborough
Sound Images Limited**

The Technology Centre Epinal Way
Loughborough Leics LE11 0QE England
Telephone: (0509) 231843 Telex: 341409 LUFBRA G
Fax: (0509) 262433

Data Converters for Digital Signal Processing

Somewhere in your system you are almost certain to need to convert between the analog and digital domains.

GEC Plessey Semiconductors has a whole range of data conversion ICs with performance from hundreds of Megahertz down to microsecond cycle times.

Typically, video processing – whether for robotics, radar or any other imaging system – would use fast front end ADCs such as our eight-bit SP94308 video system ADC or the simple but faster SP973T8. Still eight bits, but with the possibility of oversampling for even greater accuracy, is the 100MHz SP97508.

Perhaps you need to drive a graphics display at the back end of your system or maybe you want to synthesise an analog waveform. Either way, check our range of fast DACs, which operate up to 450MHz and includes parts with graphics features.

Finally, for servo control and mechanical measurement at lower speeds, GEC Plessey Semiconductors has a range of microprocessor-compatible ADCs and DACs. Complete technical data for these and many more products is contained in our Data Converters IC Handbook.

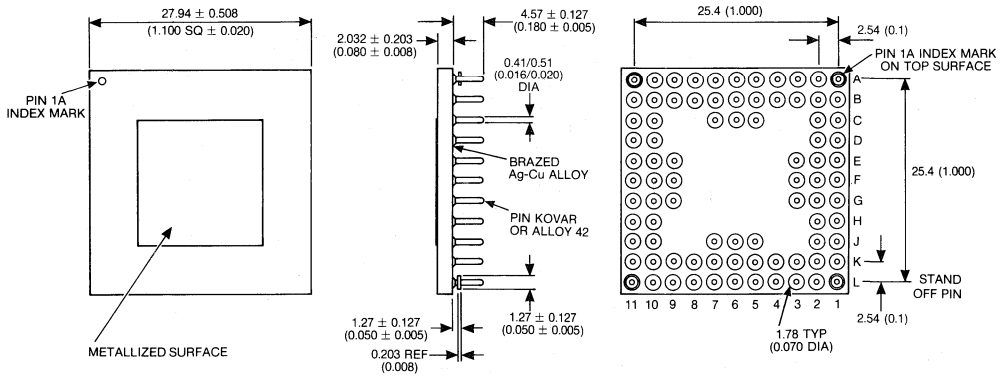
ANALOG TO DIGITAL CONVERTERS

Type	Function	Guaranteed Minimum Clock Rate	Process
SP97508	8-bit flash ADC	100MHz	Bipolar
SP973T8	8-bit flash ADC, TTL/CMOS outputs	30MHz	Bipolar
SP94308	8-bit video system ADC	20MHz	Bipolar

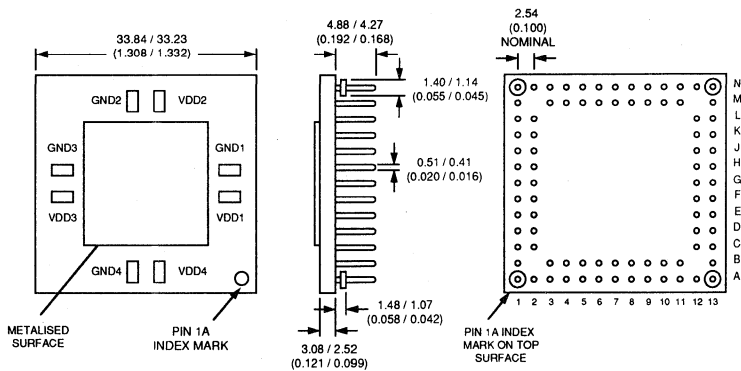
DIGITAL TO ANALOG CONVERTERS

Type	Function	Guaranteed Minimum Clock Rate	DAC Max. Rise Time (10% - 90%)	Process
SP98608	8-bit multiplying DAC	450MHz	800ps	Bipolar
MV95308	8-bit video DAC	30MHz	6ns	CMOS
MV95408	8-bit video DAC	50MHz	5.5ns	CMOS

Package outlines

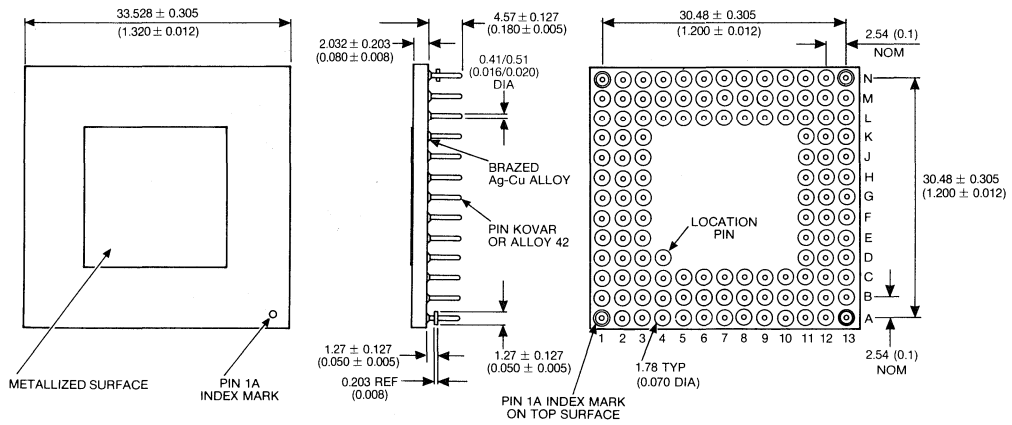


84-PIN GRID ARRAY PACKAGE — AC84




84-PIN GRID ARRAY POWER PACKAGE - AC84

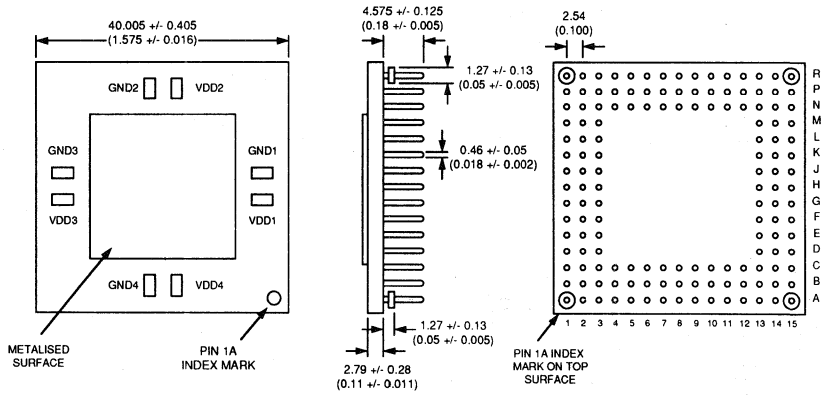
(Used for the PDSP 16340, PDSP16350, PDSP16488 and PDSP16510)




120-PIN GRID ARRAY PACKAGE — AC120




144-PIN GRID ARRAY PACKAGE - AC144




144-PIN GRID ARRAY POWER PACKAGE - AC144
 (Used for the PDSP 16256)




28-LEAD CERAMIC DIL - DG28



**28-PIN LEADLESS CHIP CARRIER - LC28
(HERMETIC)**



68 CONTACT LCC PACKAGE — LC68



**84-PIN LEADLESS CHIP CARRIER - LC84
(HERMETIC)**

Locations

HEADQUARTERS OPERATIONS

UNITED KINGDOM Cheney Manor, Swindon, Wiltshire SN2 2QW, United Kingdom.
Tel: (0793) 518000 Tx: 449637 Fax: 0793 518411

NORTH AMERICA Sequoia Research Park, 1500 Green Hills Road, Scotts Valley,
California 95066, United States of America. Tel:(408) 438 2900 ITT Telex: 4940840 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. Tel: (1) 64 46 23 45 Fax: (1) 64 46 06 07 Tlx: 602 858 F.
Ungererstraße 129, 8000 Munchen 40. Tel: 089/36 0906-0 Fax: 089/360906-55 Tlx: 523980

GERMANY (FDR),
AUSTRIA and
SWITZERLAND Viale Certosa, 49 20149 Milano. Tel: (02) 33 00 10 44/45 Fax: (GR3) 31 69 04 Tlx: 331347

ITALY Sequoia Research Park, 1500 Green Hills Road, Scotts Valley,
California 95066, United States of America. Tel: (408) 438 2900 ITT Telex: 4940840 Fax: (408) 438 7023

NORTH AMERICA 152 Beach Road, #04-05 Gateway East, Singapore 0718.
Tel: 2919291 Fax: 2916455.

SOUTH EAST ASIA Unit 1, Crompton Road, Groundwell Industrial Estate, Swindon,
Wiltshire SN2 5AY. Tel. (0793) 518510 Tx: 444410 Fax: (0793) 518582.

UNITED KINGDOM and
SCANDINAVIA

WORLD-WIDE AGENTS

AUSTRALIA and
NEW ZEALAND
EASTERN EUROPE **GEC Components Group.**, Electronic Division, 2 Giffnock Avenue, North Ryde, Sydney, New South Wales 2113
Tel: (2) 8876222 Tx: AA26080 Fax: (2) 8050272

GREECE **CTL Empexion Ltd.**, Falcon House, 19 Deer Park Road, London SW19 3WX Tel: (081) 543 0911
Tx: 928472 Fax: (081) 540 0034.

JAPAN **Impel Ltd.**, 30 Rodon Str, Korydallos, Piraeus, Greece. Tel: 010 30 1 49 67815 Tlx: 213835
Fax: 01 49 54041.

HONG KONG **Cornes & Company Ltd.**, Maruzen Building, 2-3-10 Nihonbashi, Chuo-ku, Tokyo 103. Tel: 3 272 5771
Tx: 24874 Fax: 3 271 1479.

KOREA **Cornes & Company Ltd.**, 1-Chome Nishihonmachi, Nishi-Ku, Osaka 550. Tel: 6 532 1012
Tx: 525-4496 Fax: 6 541-8850.

MALAYSIA **Microtek Inc.**, Itoh Bldg, 7-9-17 Nishishinjuku, Tokyo 160. Tel: 3 371 1811 Tx: 27466.
Fax: 3 369 5623.

SCANDINAVIA **YES Products Ltd.**, Block E, 15/F Golden Bear Industrial Centre, 66-82 Chaiwan Kolk Street, Tsuen Wan,
N.T. Hong Kong. Tel: 4442416 Tx: 36590 Fax: 4993065.

Denmark **KML Corporation**, 3rd Floor, Bang Bae Station Building, 981-15 Bang Bae, 3-Dong Shucho-Gu, Seoul, Korea,
CPO Box 7981. Tel: 2 588 2011/6 Tx: K25981 Fax: 2 588 2017.

Finland **Plessey Malaysia**, 1602 Pernas International Building, Jalan Sultan Ismail, Kuala Lumpur 50250.
Tel: 3 2611477 Tlx: 30918 Fax: 3 2613385.

Norway **Scansupply A/S**, 18-20 Nannasgade, DK-2200 Copenhagen N. Tel: 31 83 50 90 Tx: 19037 Fax: 31 83 25 40.

Sweden **Scansupply A/S**, Marselisborg Havnevej 36, 8000 Aarhus C. Tel: 45 86 12 77 88 Fax: 45 86 1277 18.

Oy Ferrado AB, P.O.Box 54, SF-00381 Helsinki 38. Tel: 98 0550 002 Tx: 122214 Fax: 98 0551 117.

Skandinavisk Elektronikk A/S, Ostre Aker Vei 99, 0596 Oslo. Tel: 2 64 11 50 Tx: 71963 Fax: 2 643443.

Swedesupply AB, Vastra Vagen 5, P.O.Box 1028, 171 21 Solna. Tel: 08735 81 30 Tx: 13435
Fax: 0883 9033.

SPAIN **Anatronic SA**, Avda de Valladolid 27, 28008 Madrid. Tel: 91 542 4455 Tx: 47397 Fax: 91 2486975.

TAIWAN **Artistex International Inc.**, B2, 11th Floor, 126, Nanking East Road, Section 4, Taipei, Taiwan,
Republic of China. Tel: 2 7526330 Tx: 27113 Fax: 2 721 5446.

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**, Refik Saydam Cad No.89 Kat 5, 80050 Sishane, Istanbul, Turkey. Tel: 0 143 621213 Fax: 0 143 6549.

WORLD-WIDE DISTRIBUTORS

- AUSTRALIA** **GEC Components Group.**, Electronic Division, 2 Giffnock Avenue, North Ryde, Sydney, New South Wales 2113
Tel: (2) 8876222 Tx: AA26080 Fax: (2) 8050272
- AUSTRIA** **Moor-Lackner GesmbH**, Elektrotechnik Elektronik Datentechnik, A-1232 Wien/Austria, Lamezanstrasse 10.
Tel: 222 610620 Tx: 135701 Fax: 222 61062151.
- BELGIUM** **Heynen**, De Koelen 6, B-3530 Koutmalen. Tel: 011/52 57 57 Tlx: 39047 Fax: 011/52 57 77.
- FRANCE** **Mateleco**:
Ile de France, 66 Avenue Augustin Dumont, 92240 Malakoff. Tel: 010 33 1 46 57 70 55 Tx: 203436.
Rhône-Alpes, 2, Rue Emile Zola, 38130 Echirolles. Tel: 010 33 76 40 38 33 Tx: 980837.
- ICC**:
Bordeaux, Rue de la Source, 33170 Gradignan. Tel: 56 31 17 17 Tx: 541539 Fax: 61 48 11 25.
Clermont-Ferrand, 2 bis, Avenue Fonmaure, 63400 Chamalieres. Tel: 73 36 71 41 Tx: 990928.
Z.A. Artisanord 11, 13015 Marseille. Tel: 91 03 12 12 Tx: 441313.
78, Chemin Lanusse, 31200 Toulouse. Tel: 61 26 14 10 Tx: 520897.
- CGE Composants S:A**:
Ile de France-6, avenue Marechal-Juin - Z.I, Grange-Dame-Rose, 92363 Meudon La Forest
Tel: (1) 40 94 84 00 Tx: 632253 Fax: (1) 46 30 01 29.
Aquitaine, Avenue Gustave Eiffel, 33605 Pessac Cedex. Tel: 56 36 40 40 Tx: 571224 F.
Bretagne-9, rue du General Nicolet, 35015 Rennes Cedex. Tel: 99 50 40 40 Tx: 740311 F.
Centre/Pays-de-Loire, Allee de la Detente, 86360 Chasseneuil du Poitou. Tel: 49 52 88 88 Tx: 791525 F.
Est-27 rue Kleber, 68000 Colmar. Tel: 89 41 15 43 Tx: 870569 F.
Midi-Pyrenees 55, Avenue Louis Breguet, 31400 Toulouse. Te: 61 20 82 38 Tx: 530957 F.
Nord, 2 rue de la Creativite, 59650 Villeneuve-d'Ascq. Tel: 20 67 04 04 Tx: 136887 F.
Provence/Cote d'Azur, Avenue Donadei, bat.B-06700 Staint. Laurent-du-Var. Tel: 93 07 77 67
Tx: 461481 F.
Rhône-Alpes, 101, rue Dedieu, 69100 Villeurbanne. Tel: 78 68 32 29 Tx: 305301 F.
- Aquitech**:
Aquitech, 73 Avenue du Chateau d'Eau 33700 Merignac Tel: 56 55 10 30 Tx: 550529 Fax: 56 47 53 20
Aquitech, 25 rue de la Chalotais 35000 Rennes, Tel: 99 78 31 32 Fax: 99 79 21 80
Aquitech, 2 rue Alexis de Tocqueville 92189 Antony, Tel: (1) 40 96 94 94 Tx: 550529 Fax: (1) 40 96 93 00
- Rhonalco**, 3 Rue Berthelot, 69627 Villeurbanne Tel: 33 78 53 00 25 Tlx: 380284 Fax: 33 72 34 67 72.
- GERMANY(FDR)** **Altron GmbH & Co. KG**, Gaussstr. 10, 3160 Lehrte. Tel: 05132 50990 Tx: 922383 Fax: 05132 57776.
API Elektronik Vertriebs GmbH, Lorenz-Brarenstr. 32, 8062 Markt Indersdorf.
Tel: 08136 7092 Tx: 5270505 Fax: 08136 7398.
AS Electronic Vertriebs GmbH, In den Garten 2, 6380 Bad Hamburg 6.
Tel: 06172 458931 Fax: 06172 42000.
Astronic GmbH, Gruenwalderweg 30, 8024 Deisenhofen. Tel: 089 6130303 Tx: 5216187 Fax: 089 6131668.
Micronetics GmbH, Weil de Stadter Str. 45, 7253 Renningen 1.
Tel: 071 59 6019 Tx: 724708 Fax: 071 59 5199.
Weisbauer Elektronik GmbH, Heiliger Weg 1, 4600 Dortmund 1. Tel: 0231 57 95 47 Tx: 822538
Fax: 0231 57 75 14.
- ITALY** **Eurelettronica Spa**, Via E Fermi 8, 20090 Assago Milan. Tel: 2 4880022 Tlx: 350037 Fax: 2 4880275
Adelsi Spa, Via Novara 570, 20153 Milan. Tel: 2 3580641 Tlx: 332423 Fax: 2 3011988.
Fanton S.R.L., Milano-Bologna, Firenze, Roma, Padova, Torino. Tel: 2 3287312 Tlx: 350853 Fax: 2 3287948.
- NETHERLANDS** **Heynen B.V.**, Postbus 10, 6590 AA Gennep. Tel: 8851-96111 Tx: 37282 Fax: 8851 96200.
Tekelek Airtronic BV., PO Box 63, NL 2712 LB Zoetermeer. Tel: 79 310100 Fax: 79 417504
- SWITZERLAND** **Basix fuer Elektronik AG**, Hardturmstr 181, CH-8010 Zurich Tel: 01 2761111 Tlx: 822966 Fax: 01 2761199
Eibatex AG., Hardstr 72, CH-5430 Wettingen. Tel: 41 56 27 51 11 Tx: 826300 Fax: 41 56 27 19 24.
- TAIWAN** **Artistex International Inc.**, B2, 11th Floor, 126, Nanking East Road, Section 4, Taipei, Taiwan,
Republic of China. Tel: 2 752630 Tx: 27113 Fax: 2 721 5446.
- UNITED KINGDOM** **Celdis Ltd.**, 37-39 Loverock Road, Reading, Berks RG3 1ED. Tel: 0734 585171 Tx: 848370
Fax: 0734 509933.
Farnell Electronic Components Ltd., Canal Road, Leeds LS12 2TU. Tel: 0532 790101 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, Trafalgar House, 28 Paradise Circus, Queensway, Birmingham
B1 2BL. Tel: 021 6436365 Tx: 338731. Fax: 021 633 3207
Macro, Burnham Lane, Slough SL1 6LN Tel: 0628 604383 Fax: 0628 66873
RR Electronics Ltd., St. Martins Way, Cambridge Road, Bedford MK42 0LF. Tel: 0234 47188/270777 Tx: 826251.
Fax: 0234 210674.
Semiconductor Specialists (UK) Ltd., Carroll House, 159 High Street Yiewsley, West Drayton, Middlesex UB7 7XB
Tel: (0895) 445522 Tx: 21958 Fax: (0895) 422044
STC Electronic Services Ltd, Edinburgh Way, Harlow, Essex CM20 2DF Tel: 0279 626777 Tlx: 818801 Fax: 0279 441687.
2001 Electronic Components Ltd, Woolners Way, Stevenage, Herts SG1 3AJ Tel: 0438 742001 Tx: 827701
Fax: 0438 742002.

UK EXPORT

(To countries other than those listed)

GEC Plessey Semiconductors Ltd, Unit 1, Crompton Road, Groundwell Industrial Estate, Swindon Wilts, UK SN2 5AY Tel: (0793) 518510 Tlx: 444410 Fax: (0793) 518582.

AMERICAN DESIGN CENTRES

CANADA **Alberta Microelectronics Center**, 3553 31st St., N.W., Calgary, Alberta T2L2K7. Tel: (403) 289-2043.
Microstar Technologies, 7050 Bramelea Rd., #27A Mississauga, Ontario L5S1T1 Canada Tel: (416) 671-8111
COLORADO **Analog Solutions**, 5484 White Place, Boulder, CO 80303. Tel: (303) 442-5083.
ILLINOIS **Frederikssen & Shu Laboratories, Inc.**, 531 West Golf Rd., Arlington Heights, IL 60005. Tel: (312) 956-0710.

NORTH AMERICAN SALES OFFICES

NATIONAL SALES 1500 Green Hills Road, Scotts Valley, CA 95066 Tel: (408) 438-2900, ITT Telex: 4940840, Fax: (408) 438-5576.
METRO NY/NJ 1767-42 Veterans Memorial Hwy., Central Islip, NY 11722 Tel: (516) 582-8070, Tlx: 705922, Fax: (516) 582-8344.
EASTERN 132 Central Street, #216 Foxboro, MA 02035 Tel: (508) 543-3855, Tlx: 316805, Fax: (508) 543-2994.
WESTERN 2727 Walsh Ave. #102, Santa Clara, CA 95051 Tel: (408) 986-8911, Fax: (408) 970-0263.
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.
NORTH CENTRAL 2625 Butterfield Rd. #109N Oak Brook, IL 60521 Tel: (708) 573-7773, Fax: (708) 573-7790.
FLORIDA 668 N. Orlando Ave., Suite 1015 B, Maitland, FL 32751 Tel: (407) 539-1700, Fax: (407) 539-0055.
DIXIE 41 Milton Ave., #104, Alpharetta, GA 30201 Tel: (404) 343-9904, Fax: (404) 343-9972.
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 207 Place Frontenac, Quebec, H9R-4Z7 Tel: (514) 697-0095/96, Fax: (514) 695-9250.

NORTH AMERICAN REPRESENTATIVES

ALABAMA **DHR Marketing, Inc.**, 1580 Sparkman Dr. NW #202, Huntsville, AL 35816 Tel: (205) 722-0440, Fax: (205) 722-0494.
CALIFORNIA **Select Electronics**, 14730 Beach Blvd. Bldg. F #106, La Mirada, CA 90638 Tel: (714) 739-8891, Fax: (714) 739-1604.
CONNECTICUT **Pioneer**, 112 Main St. Norwalk, CT 06851 Tel: (203) 235-1422, Fax: (203) 634-4884.
FLORIDA **DHR Marketing, Inc.**, 1860 Old Okeechobee Rd. #506 West Palm Beach, FL 33409 Tel: (407) 697-9680 Fax: (407) 697-9711.
DHR Marketing, Inc., 417 Whooping Loop St. 1747, Altamonte Springs, FL 32701 Tel: (407) 331-1199 Fax: (407) 331-1263.
DHR Marketing, Inc., 800 W. Platt St. #8 Tampa FL 33606 Tel: (813) 254-2009, Fax: (813) 251-2904.
GEORGIA **DHR Marketing, Inc.**, 3100 Breckenridge Blvd. #145 Duluth, GA 30136 Tel: (404) 564-0529, Fax: (404) 564-1127.
ILLINOIS **Micro Sales, Inc.**, 54 West Seegars Road, Arlington Heights, IL 60006 Tel: (312) 956-1000, Tlx: 510-600-0756, Fax: (312) 956-0189.
INDIANA **Leslie M. DeVoe**, 4371 E. 82nd St., Suite D, 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-Houck**, 10706 Reisters Town Rd., Suite D, Owings Mills, MD 21117 Tel: (301) 356-9500, Fax: 301-356-9503.

MASSACHUSETTS **Stone Components**, 2 Pierce Street, Framingham, MA 01701 Tel: (508) 875-3266, Fax: (508) 875-0537.
Stone Components, 10 Atwood Street, Newburyport, MA 01950 Tel: (508) 875-3266, Fax: (508) 465-3544.
Stone Components, 11 Blueberry Hill Rd., Long Meadow, MA 01106 Tel: (413) 567-9075, Fax: (413) 567-1019.

MICHIGAN **R. P. Urban & Associates**, 2335 Burton Street S.E., P.O. Box 7386, Grand Rapids, MI 49510 Tel: (616) 245-0511, Fax: (616) 245-4083.
R. P. Urban & Associates, 24634 Five Mile Rd., Detroit, MI 48239 Tel: (313) 535-2355, Fax: (313) 535-7109.

MINNESOTA **Electro Mark, Inc.**, Valley Oaks Business Center, 7167 Shady Oak Rd., Eden Prairie, MN 55344 Tel: (612) 944-5850, Fax: (612) 944-5855.

MISSOURI **Lorenz Sales, Inc.**, 10176 Corporate Square Dr. #120, St. Louis MO 63132 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.

PENNSYLVANIA/
NEW JERSEY **Metz Benham Associates, Inc.**, (MBA), 1916 Fairfax Av, Cherry Hill, NJ 08003 Tel: (609) 424-0404 Fax: (609) 751-2160.
Metz Benham Associates, Inc., (MBA), 7 Waynnewood Rd, Wynnewood, PA 19096 Tel: (215) 896-7300 Fax: (215) 642-6293.

METRO NY/NJ
NEW YORK **Metro Logic**, 271 Route 46 West, Suite D-202, Fairfield, NJ 07006 Tel: (201) 575-5585, Fax: (201) 575-8023.
Metro Logic, 554 Polaris St., N. Babylon, NY 11703 Tel: (516) 243-5617, Fax: (516) 242-1670.
Micro-Tech, 1350 Buffalo Road, Rochester, NY 14624 Tel: (716) 328-3000, Fax: (716) 328-3003.
Micro-Tech, 401 South Main St., North Syracuse, NY 13212 Tel: (315) 458-5254, Fax: (315) 458-5919.
Micro-Tech, 10 Guyton Street, Kingston, NY 12401 Tel: (914) 338-7588.

NORTH CAROLINA **DHR Marketing, Inc.**, 211 Six Forks Rd. #105, Raleigh, NC 27609 Tel: (919) 829-1970, Fax: (919) 829-1906.

SOUTH CAROLINA **Pioneer Technologies**, 9401-L South Pine Blvd., Charlotte, NC 28217 Tel: (704) 527-8188, Fax: (919) 522-8564

OHIO **DHR Marketing, Inc.**, 33 Villa Road, B 140, Greenville, SC 29615 Tel: (803) 235-3594, Fax: (803) 271-8712.
Scott Electronics, Inc., 3131 S. Dixie Dr., Suite 200, Dayton, OH 45439. Tel: (513) 294-0539, Fax: (513) 294-4769.
Scott Electronics, Inc., 360 Alpha Park, 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) 882s-0900.
Scott Electronics, Inc., 10901 Reed Hartman Hwy., Suite 301, Cincinnati, OH 45242 Tel: (513) 791-2513, Fax: (513) 791-8059.

TENNESSEE **DHR Marketing, Inc.**, 417 Welchwood, Suite 102, Nashville, TN 37211 Tel: (615) 331-2745, Fax: (615) 331-3453.

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., Suite 218, Austin, TX 78758 Tel: (512) 453-0275, Fax: (512) 453-0088.

WISCONSIN **Micro Sales, Inc.**, 16800 W. Greenfield Av., # 116, Brookfield, WI 53005 Tel: (414) 786-1403, Fax: (414) 786-1813.

CANADA **GM Assoc. Inc.**, 7050 Bramalea Road, Suite 27A, Mississauga, Ontario L5S 1T1 Tel: (416) 671-8111, Fax: (416) 671-2422.
GM Assoc. Inc., 3860 Cote-Vertu, Suite 221, St. Laurent, Quebec H4R 1N4 Tel: (514) 335-9572, Fax: 514- 335-9573.
GM Assoc. Inc., 3525 McBean St., Richmond, Ontario K0A 2Z0 Tel: (613) 838-4480, Fax: 613-838-4479.

NORTH AMERICAN DISTRIBUTORS

ALABAMA **Pioneer Technologies**, 4825 University Square, 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**, 1525 W. University Dr., Suite 103, Tempe, AZ 85282 Tel: (602) 829-1800, Twx: 510-601-1618, FAX: (602) 967-2658.

CALIFORNIA **Insight Electronics (Corp.)**, 6885 Flanders Dr., Unit C, San Diego, CA 92121 Tel: (619) 587-0471, Twx: 183035, Fax: (619) 587-0903.
Insight Electronics, 28038 Dorothy Dr., Suite 2, Agoura, CA 91301 Tel: (818) 707-2100, Fax: (818) 707-0321.
Insight Electronics, 15635 Alton Pkwy., Suite 120, Irvine, CA 92718 Tel: (714) 727-2111, Fax: (714) 727-4804.

NORTH CAROLINA **Pioneer Technologies**, 134 Rio Robles, San Jose, CA 95134 Tel: (408) 954-9100 Fax: (408) 954-9113.
Hammond, 2923 Pacific Ave., Greensboro, NC 27406 Tel: (919) 275-6391, Tlx: 62894645, Fax: (919) 272-6036.
Pioneer Technologies, 9401-L South Pine Blvd., Charlotte, NC 28217 Tel: (704) 527-8188, Fax: (919) 522-8564
Pioneer Technologies, 2810 Meridian Pkwy., Suite 148, Durham, NC 27713 Tel: (919) 544-5400, Fax: (919) 544-5885.

FLORIDA **Hammond**, 1035 Lowndes Hill Rd., Greenville, SC 29607 Tel: (803) 232-0872, Fax: (803) 232-0320.
Jaco Electronics, 384 Pratt Street, Meriden, CT 06450 Tel: (203) 235-1422 Fax: (203) 634-4884.
Pioneer/Norwalk, 112 Main Street, Norwalk, CT 06851 Tel: (203) 853-1515, Twx: 710 468-3373, Fax: (203) 838-9901.
Hammond Ft. Lauderdale, 6600 N.W. 21st Ave., Ft. Lauderdale, FL 33309 Tel: (305) 973-7103 Twx: (510) 956-9401 Fax:(305)973-7601.
Hammond Orlando, 1230 West Central Blvd., Orlando, FL 32805 Tel: (407) 841-1010 Fax:(407)648-8584.
Pioneer Technologies, 674 South Military Trail, Deerfield Beach, FL 33442 Tel: (305) 428-8877, Twx: 510-955-9653, Fax: (305) 481-2950.
Pioneer Technologies, 337 South Northlake Blvd., Suite 1000, Altamonte Springs, FL 32701 Tel: (407) 834-9090, Twx: 810-853-0284, Fax: (407) 834-0865.

GEORGIA **Hammond**, 5680 Oakbrook Pkwy, Suite 160, Norcross, GA 30093 Tel: (404) 449-1996, Fax: (404) 242-9834.
Pioneer Technologies, 3100 F Northwoods Pl., Norcross, GA 30071 Tel: (404) 448-1711, Twx: 810-766-4515, Fax: (404) 446-8270.

INDIANA **Pioneer/Standard**, 9350 N. Priority Way W. Drive., Indianapolis, IN 46240 Tel: (317) 573-0880, Fax: (317) 573-0979.

ILLINOIS **Pioneer/Standard**, 2171 Executive Drive #104, Addison, IL 60101 Tel: (312) 495-9680, Fax: (312) 495-9831.

MASSACHUSETTS **Jaco Electronics**, 1053 East Street, Tewksbury, MA 01876 Tel: (508) 640-0010, Fax: (508) 640-0755.
Pioneer/Standard, 44 Hartwell Avenue, Lexington, MA 02173 Tel: (617) 861-9200, Fax: (617) 863-1547, Fax: (301) 995-6620

MARYLAND **Jaco Electronics**, Rivers Center, 10270 Old Columbia Road, Columbia, MD 21046 Tel: (301) 995-6620, Fax: (301) 995-6032
Pioneer/Tech. Group, Inc., 9100 Gaither Rd., Gaithersburg, MD 20877 Tel: (301) 921-0660, Twx: 710-828-0545, Fax: (301) 921-4255.

MICHIGAN **Pioneer/Standard**, 13485 Stamford, Livonia, MI 48150 Tel: (313) 525-1800, Twx: 810-242-3271.
Pioneer, 4505 Broadmoore Ave., S.E., Grand Rapids, MI 49512 Tel: (616) 698-1800, Twx: 510-600-8456 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 St. Louis**, 2029 Woodland Pkwy. #101, St. Louis, MO 63146 Tel: (314) 432-4350, Fax: 314-432-4854.

NEW JERSEY **Jaco Electronics**, 110 Greenvale Ave., Wayne, NJ 07470 Tel: (201) 942-4000, Fax: (201) 942-0088.
Pioneer/Standard, 14A Madison Rd., Fairfield, NJ 07058 Tel: (201) 575-3510, Fax: (201) 575-3454.

NEW YORK **Jaco Electronics**, Hauppauge, NY 11787 Tel: (516) 273-5500, Twx: 510-227-6232, Fax: (516) 273-5528.
Mast, 710-2 Union Pkwy, Ronkonkoma, NY 11779 Tel: (516) 471-4422, Twx: 4974384, Fax: (516) 471-2040.
Pioneer/Standard, 68 Corporate Drive, Binghamton, NY 13904 Tel: (607) 722-9300, Twx: 510-252-0893, Fax: (607) 722-9562.
Pioneer/Standard, 60 Crossways Park West, Woodbury, NY 11797 Tel: (516) 921-8700, Twx: 510-221-2184, Fax: (516) 921-9189.

OHIO **Pioneer/Standard**, 840 Fairport Park, NY 14450 Tel: (716) 381-7070, Twx: 510-253-7001, Fax: (716) 381-5955.
Pioneer/Standard, 4800 East 131st St., Cleveland, OH 44105 Tel: (216) 587-3600, Twx: 810-421-0011, Fax: (216) 587-3906.
Pioneer, 4433 Interpoint Blvd., Dayton, OH 45424 Tel: (513) 236-9900, Twx: 810 459 1622, Fax: (513) 236-8133.

PENNSYLVANIA **Pioneer Technologies**, 500 Enterprise Rd., 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**, 6034 W. Courtyard, #305-49, Austin, TX 78730 Tel: (512) 467-0800, Fax: (512) 343-2612.
Insight, 1778 Plano Rd., Suite 320, Richardson, TX 75081 Tel: (214) 783-0800, Fax: (214) 680-2402.
Insight, 10500 Richmond, Suite 201, Houston, TX 77042 Tel: (713) 448-0800, Fax: (713) 952-0289.
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, Twx: 910-860-5563, Fax: (214) 490-6419.
Pioneer/Standard, 10530 Rockley Rd., Houston, TX 77099 Tel: (713) 495-5642, Twx: 910-881-1606, Fax: (713) 495-4700.

WASHINGTON **Insight**, 12002 115th Ave. N.E., Kirkland, WA 98034 Tel: (206) 820-8100, Fax: (206) 821-2976.

WISCONSIN **Pioneer Wisconsin**, 120 Bishops Way #134, Brookfield, WI 53005 Tel: (414) 784-3480, Fax: (414) 784-8207.

CANADA EASTERN **Semad**, 1825 Woodward Dr., Suite 303, Ottawa, Ontario K2C 0R3. Tel: (613) 727-8325 Twx: 0533943 Fax: (613) 727-9489.
Semad, 243 Place Frontenac, Pointe Claire, P.Q. H9R 4Z7. Tel: (514) 694-0860 Twx: 05821861 Fax: (514) 694-0965.

CANADA WESTERN **Semad**, 6120 3rd. St. SE, Calgary, Alberta T2H 1K4. Tel:(403) 252-5664 Twx:03824775 Fax:(403)255-0866.
Semad, 8563 Government St,Burnaby, BC V3N 4S9. Tel: (604) 420-9889 Twx: 04356625 Fax:(604)420-0124.
Semad, 85 Spy Court, Markham, Ontario, L3R 4Z4. Tel: (416) 475-8500 Twx: 06966600 Fax: (416) 475-4158.

PRIMARY SEMI-CUSTOM DESIGN CENTRES

AUSTRALIA 2 Giffnock Avenue, North Ryde, Sydney, New South Wales 2113
Tel: (2) 8876222 Tx: AA26080 Fax: (2) 8050272

FRANCE & BENELUX Z.A. Courtaboeuf, Miniparc-6, Avenue des Andes,
Bat.2-B.P. No. 142 91944 Les Ulis Cedex A. Tel: (6) 446 23 45 Tx: 602858F Fax: (6) 446 06 07.

ITALY Viale Certosa, 49, 20149 Milan. Tel: (02) 39001044/5 Tx: 331347 Fax: (GR3) 2316904.

WEST GERMANY Ungererstrasse 129, D 8000 Munich 40. Tel: (089) 3609 06 0 Tx: 523980 Fax: (089) 3609 06 55.

JAPAN Saito Building 6F 6-6-1, Sotokanada Chiyoda-Ku, Tokyo.
Tel: (3) 839 3001 Fax: (3) 839 3005

TAIWAN B2, 11th Floor, 126 Nanking East Road, Section 4, Taipei.
Tel: 2 7526330 Tx: 27113 Fax: 2 721 5446.

UNITED KINGDOM Cheney Manor, Swindon, Wiltshire SN2 2QW.
Tel: (0793) 518000 Tx: 449637 Fax: (0793) 518411.
Hollinwood Avenue, Hollinwood, Oldham, Lancashire OL9 7LB.
Tel: 061 682 6844 Tx: 666001 Fax: 061 688 7898.

UNITED STATES Sequoia Research Park, 1500 Green Hills Road, Scotts Valley, California 95066.
OF AMERICA Tel: (408) 438 2900 ITT Telex: 4940840 Fax: (408) 438-5576.
Two Dedham Place, Suite 1, Allied Drive, Dedham, Massachusetts 02026.
Tel: 617/320-9790 Fax: 617/320-9383.
13900 Alton Parkway #123, Irvine, California 92718.
Tel: (714) 455 2950 Fax: (714) 455 9671.

© GEC Plessey Semiconductors 1990

All rights reserved

Publication No. PS 2252 November 1990

This publication is issued to provide outline information only which (unless agreed by the Company in writing) may not be used, applied or reproduced for any purpose or form part of any order or contract or be regarded as a representation relating to the products or services concerned. The Company reserves the right to alter without notice the specification, design, price or conditions of supply of any product or service.

IBM-PC is a registered trademark of IBM. MS-DOS is a registered trademark of MICROSOFT Inc.



GEC PLESSEY
SEMICONDUCTORS

Marconi Electronic Devices Limited and Plessey Semiconductors Limited have been grouped together to form GEC Plessey Semiconductors.

However, until further notice, contracts, purchase orders, invoicing and payments should continue to be made to Marconi Electronic Devices Ltd and Plessey Semiconductors Ltd.