AB37

Increase the performance of the PDSP16510

Application Brief

AB37 - 1.0 February 1994

BACKGROUND

A single PDSP16510 FFT Processor ,using a 40 MHz system clock, will support sampling rates up to 12.3 MHz when doing 256 point complex transforms, and up to 6.8 MHz when doing 1024 point complex transforms. These rates can be increased to 40 MHz when several devices are connected in a ring arrangement.

In such a system one device is loaded with a complete block of data, and then starts a transform operation using its internal RAM. In the meantime incoming data is loaded into another FFT Processor which will then start its transform operation when all the data is available. Sufficient devices are needed to ensure that the first device has finished before its turn comes round again for new data.

The PDSP16510 actually supports two multiple device modes of operation. One mode always does separate load, transform, and then dump operations regardless of the actual transform size. The other mode does concurrent load, transform, and dump operations, but cannot be used to perform 1024 point transforms. It can in some circumstances allow less devices to be used in order to achieve a given sampling rate.

Note that an input buffer is not needed in either mode of operation (even when doing 1024 point transforms), and interdevice flags support block overlapping. With the standard 50% and 75% block overlapping no external logic is needed.

Note also that this arrangement is only intended to increase the sampling rates possible with the transform sizes supported by a single device. If larger transform sizes are needed see Application Note AB35.

GENERAL CONSIDERATIONS

Figure 1 illustrates the basic ring arrangement, using three devices for convenience. It can, of course, be expanded to any required number of devices. It shows that both inputs and outputs are commoned together. A block of data is loaded into the first device, then the next block is loaded into the second device, and so on. Sufficient devices are provide to ensure that continuous data can be supported without any loss.

The LFLG output and INEN input are used to co-ordinate the splitting of incoming data between the devices. This requires

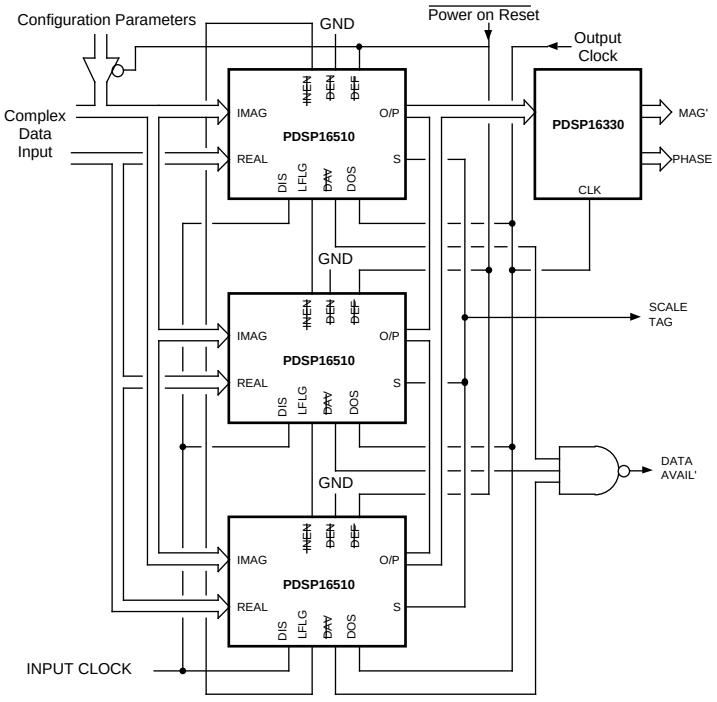

Figure 1. The Basic Multiple Device Arrangement

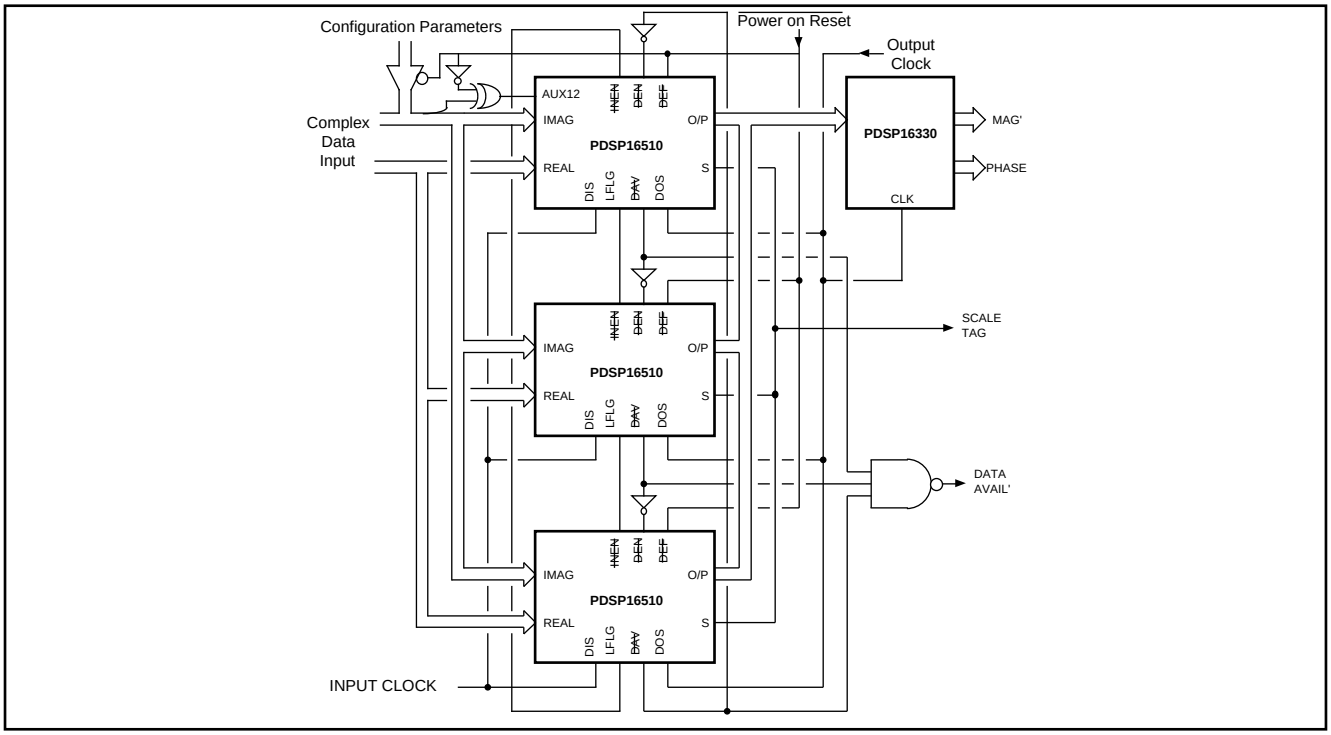

Figure 2. Adaptation of the Basic System

that the input data enabling signal (INEN) is no longer a simple DIS strobe enabling input. Whenever the Control Register specifies a multiple device mode of operation, it becomes a global ,edge activated, block enabling input. When it goes low a new block load operation commences, which will then continue until the programmed number of samples have been loaded. INEN can go back high at any time without causing inputs from the current block to be inhibited (early silicon requires that INEN stays low at least until all the block has been loaded).

The Load Flag (LFLG) goes active high after the first sample has been loaded, and goes low when the programmed number of samples have been loaded. This low going LFLG edge from a previous device provides the low going INEN edge to the next device. Thus once one device has received its full complement of data the next one will start loading data. As a further enhancement the LFLG will go in-active after half the block has been loaded if 50% overlapping is in use, or after a quarter the block has been loaded if 75% overlapping is in use. Thus, if for example, 50% overlapping is in use, the second device will also start loading data half way through the load operation into the first device. Two devices thus receive the same data in order to mechanize the overlap requirement.

By tying all the outputs together the system need only provide one output processor. The actual processing needed on the outputs is application dependent, and is beyond the scope of this article. It is assumed, however, that unless speed requirements make it impossible, the use of only one output processor will provide some economic advantage. It is also assumed that the reason for using multiple devices was to continuously process incoming data, without any loss of information. Thus once results are available they must be passed on to the rest of the system.

This can be achieved by tying all the Data Enable (DEN) inputs low, as illustrated in Figure 1. The Data Available Flag (DAV) will then be synchronized to the output strobe, and will go active when the first result is available on the output pins. The output bus goes low impedance with respect to the same output clock edge which causes DAV to go active. Figure 1 in fact shows all the DAV flags OR'ed together to give a common indicator for the rest of the system. In practice this only makes sense if the dump time is less than the load time. The combined output would then go in-active between individual devices for a period equivalent to the difference between the load and dump times.

For convenience we have so far indicated that all the DEN inputs are tied low, but there is some uncertainty in the time taken to complete a transform once all the data has been loaded. In fact the device uses an internal 12 cycle sequence, which will lead to a 12 system clock variation in the time needed to complete a transform. The number of system clock cycles needed to complete each transform, as given in Table 4 of the data sheet, and are worst case numbers. This uncertainty makes tying DEN low potentially dangerous.

If device two effectively completes its transform in less time than device one, then it could start dumping its results before device one had finished. If DEN is tied low this can only be prevented if the dump time (plus 4 DOS periods) is at least 12 system clock periods less than the load time. DIS and DOS can then not be tied together, and it is usually convenient to derive DOS by dividing down from the system clock.

An alternative approach has, however, been provided which still allows DIS and DOS still to be commoned together.

Whenever the control register specifies a multiple device node of operation, the operation of the DAV / DEN circuitry is modified. IF DEN is not active (i.e. high) when the device is internally about to generate DAV, then the output pin will not "asynchronously" go active as happens in the single device mode. Instead the DAV output pin will stay in-active until DEN goes active low. The output bus will also stay high impedance.

By connecting the inverted DAV output from one device to the DEN input of the next, this second device cannot commence its output sequence until the first one has finished, and its DAV has gone in-active. The DAV output from the second device will go active as soon as its DEN input goes low, but it will still be effectively synchronized to DOS since it is derived from the previous DAV signal going in-active. This arrangement is shown in Figure 2.

When DOS is physically connected to DIS, the combined Data Available Flag shown in Figure 1 will only glitch in-active as one device finishes and the next one starts. This glitch would safely occur after a DOS active edge, but even so the flag would only usefully indicate the initial delay from start up before valid results are obtained. When a reliable flag is needed to indicate the end of a set of results, then each PDSP16510 should be provided with a D type latch. This is set by the inactive going edge of DAV and reset when DOS again goes low.

The circuit shown in Figure 1 needs an edge after power on in order to initiate the load procedure into the first device. Rather than providing external logic to generate a start pulse which is OR'ed into the first INEN line, an alternative scheme is supported. By setting bit 12 in the Control Register contained within the first device, it is possible to cause the power on reset signal (DEF) to initiate a load procedure.

All Control Registers are loaded from the common AUX port whilst DEF is active, and we now require one bit in one of the registers to be different from the others! The easiest way to mechanize this is to include an EX-OR gate in the AUX12 input to the first device. The other input is driven from the DEF signal such that it causes a logical inversion as the Control Register is loaded, but not when imaginary data is loaded. This is also shown in Figure 2.

SAMPLING RATES POSSIBLE WITH SEPARATE LOAD TRANSFORM, AND THEN DUMP (MODE 2)

This mode can be used with all transform sizes and the maximum DIS and DOS rates can theoretically be equal to the system clock rates used. The DIS rate is the data sampling rate and can, of course, only be equal to the system clock rate if sufficient devices are provided in the ring to make this sensible at the system level. The DOS rate can be any rate greater than or equal to the system clock rate, and would normally be limited by the capabilities of the output processor.

The number of devices, N, needed to achieve a sampling period of S with a block size of n, is governed by;

$$NnS > nS + PK + D \qquad \text{where D is less than } nS$$

and K is the system clock period, P is the number of system clock periods needed to complete a transform as given in Table 4 of the datasheet, and D is the total dump time allowing for the 4 extra DOS periods needed for the internal output circuitry.

With 50% block overlapping the above equation becomes;

$$NnS > 2(nS + PK + D) \qquad \text{where D is less than } nS/2$$

With 75% overlapping it becomes;

$$Nns > 4(nS + PK + D) \qquad \text{where D is less than } nS/4$$

Table 1 gives the maximum sampling rates possible with 3, 4, 5, or 6 devices and output rates of 20 MHz and 40 MHz. It covers transform sizes of 256 and 1024 complex points.

	1024 POINT COMPLEX TRANSFORMS						256 POINT COMPLEX TRANSFORMS					
Number of	Dump at 20 MHz			Dump at 40 MHz			Dump at 20 MHz			Dump at 40 MHz		
Devices	0%	50%	75	0%	50%	75%	0%	50%	75%	0%	50%	75%
3	13.7	6.8	-	16.6	8.2	-	15.3	7.6	-	19	9.5	-
4	20.6	10*	-	24.8	16.6	-	22.9	10*	-	28.5	19	-
5	27.4	10*	5*	33.2	20*	8.2	30.6	10*	5*	38	20*	9.5
6	34.3	10*	5*	40*	20*	10*	38.4	10*	5*	40*	20*	10*

* indicates that the sampling rate is limited by the maximum dumping rate

Table 1. Maximum Sampling Rates with separate load transforms and dumps. Where sampling rate is asynchronous to SCLK, a PDSP16540 (or similar) is assumed on the input.

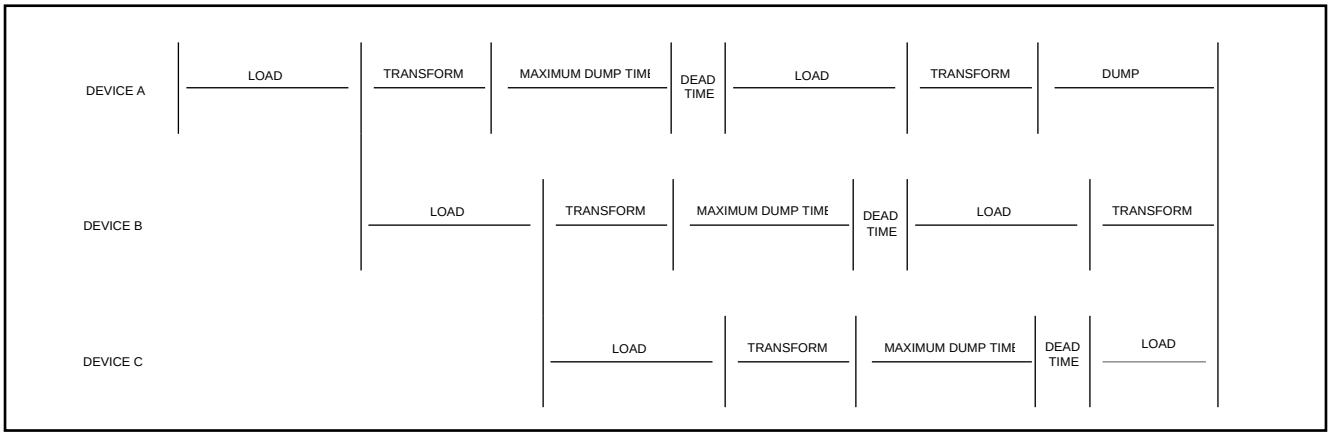

Figure 3. Sequence of events in a 3 device system doing separte load, transforms and then dumps

Figure 3 shows the sequence of events taking place in such a system without overlaps, and Figure 4 shows the sequence in a five device system doing 50% overlaps. It should be noted that the time taken to dump the results must be no more than the time taken to load the block of data (if you put more in than you take out, then something has to give). If 50% block overlapping is used the dump must be complete in half the load time, and one quarter the load time if 75% overlapping is needed. If these criterion are not met, then one device will not have finished dumping before the next one starts. Thus two sets of outputs would be enabled at the same time. Remember that the dump time can easily be made faster than the load time when required.

SAMPLING RATES POSSIBLE WITH CONCURRENT LOAD, TRANSFORM, AND DUMP (MODE 1)

In this mode transfers in and out of the PDSP16510 are concurrent with transform operations. Internal RAM restrictions do not allow this mode to be used with 1024 point transforms. For other sized transforms the sampling rates possible are theoretically much higher for a given number of devices. A limitation is, however, imposed on the maximum I/O rates, which can no longer be increased up to the system clock rate. Instead the rates are reduced to a factor, F of the system clock rate where F is given by;

$$F = \frac{4}{6 + 0.001\text{ØL}}$$ where ØL is the system clock low time

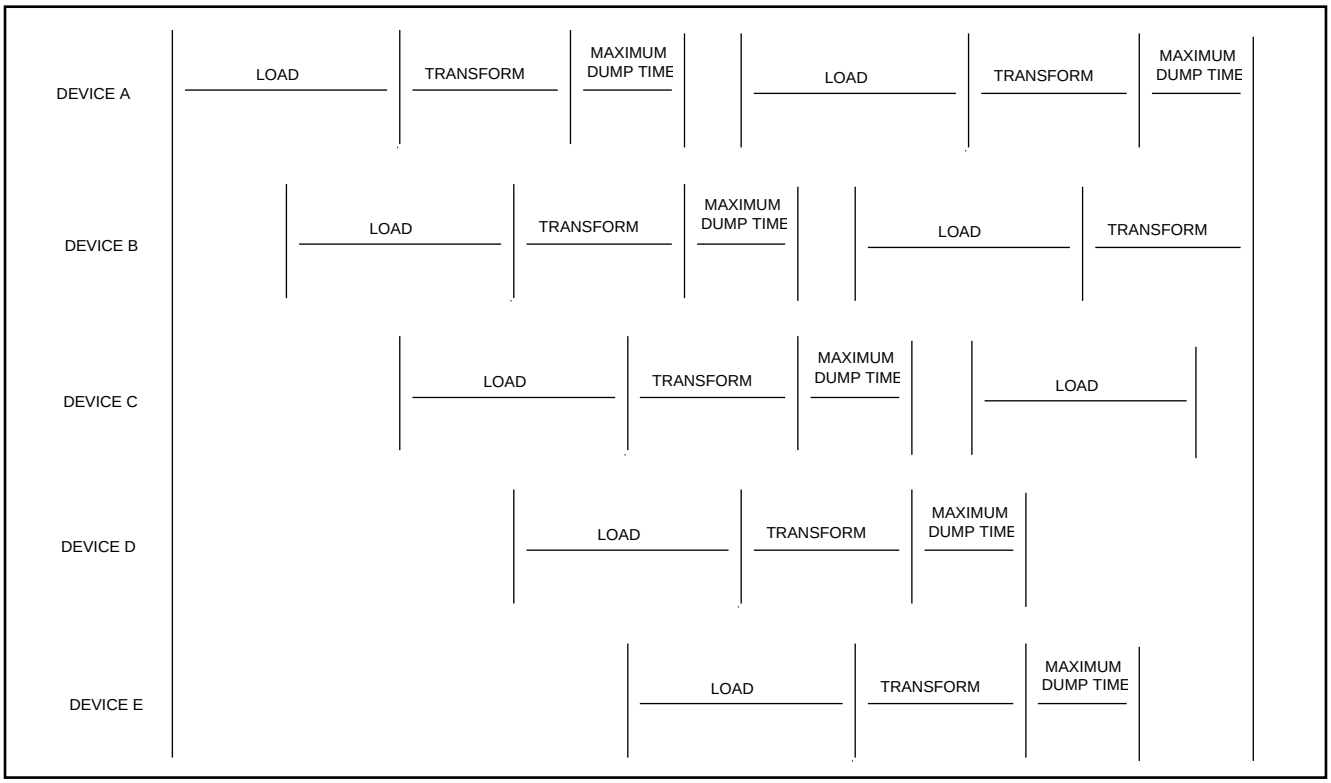

Figure 4. Sequence of events in a 5 device system with 50% overlaps

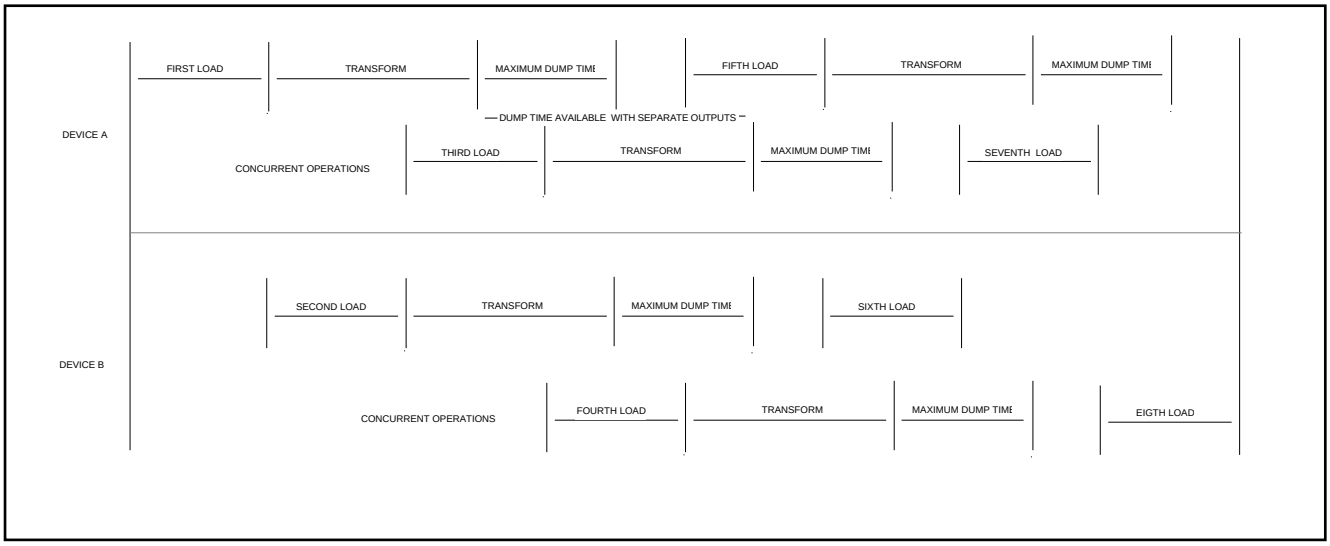

Figure 5. Sequence of events in a 2 device system with concurrent I/O Operations

With a 40 MHz system clock the low time might well be 15 nanoseconds. The factor, F, then becomes 0.66. The maximum theoretical I/O rate is thus 26.6 MHz. In this mode there are no benefits to be gained by increasing the output rate above the input rate (except when necessary because of overlaps). DIS and DOS are thus normally commoned together. In fact the device provides an internal divide DIS by or 2 or 4 feature. Thus the DIS and DOS pins can be externally connected to a source to match the DOS requirements, with the internal DIS strobe internally divided down to the correct frequency for 50% or 75% overlaps.

With N devices the theoretical sampling period, S, ignoring the F factor and with no overlapping is governed by;

$$NnS > PK + 4 \text{ DOS periods}$$

where P is the number of clock periods needed to complete the transform and K is the system clock period.

Figure 5 illustrates the sequence of events which occur with a two device system, with the outputs joined together. With a 40 MHz system clock, and common DIS and DOS, the theorectical maximum input sampling rate given by the above equation is 24.6 MHz when doing a 256 point complex transform. This sampling rate is just less than the limit which would be imposed by the factor F, but all other transform sizes supported by this mode of operation would be limited by the factor F, to sampling rates of 26.66 MHz. If each PDSP16510 has its own output processing circuit, then the outputs would not be joined together and each

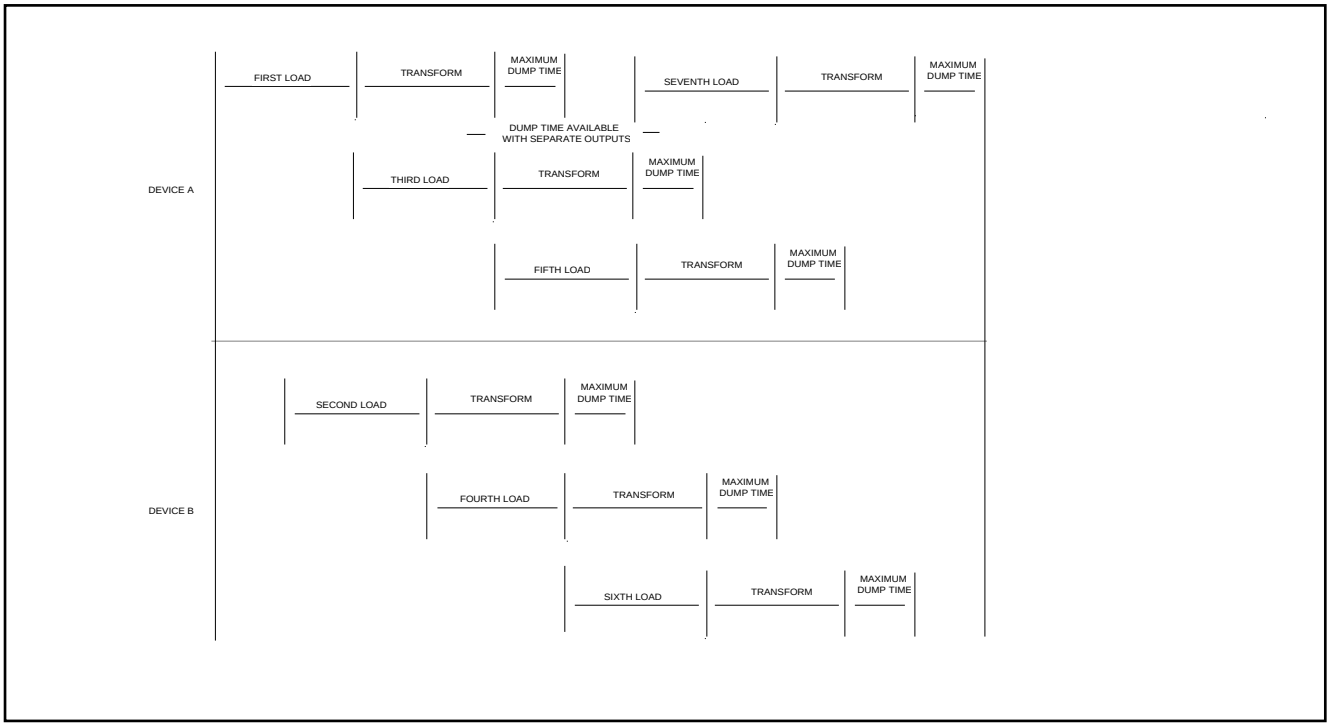

Figure 6. Sequence of events in a 2 device system with concurrent I/O and 50% overlaps

would have a dump time available equivalent to twice the load time.

In this mode two devices can be used to gave the same performance when doing 50% overlaps as a single device with no overlaps. But the dump must be completed in half the load time if the outputs are commoned together. This does not limit the performance when doing 256 point complex transforms, and two devices will handle 50% overlaps at up to 12.3 MHz sampling rates. All other sizes are limited to up to 13.3 MHz sampling rates by the F factor, unless each device has its own output processing circuitry. Figure 6 illustrates the sequence of events in such a system.

As long as each PDSP16510 has its own output processor, then extra devices can be added as needed to increase the overlapped sampling rate up to that defined by the F factor i.e. 26.66 MHz. The minimum sampling period down to this limit is governed by;

$$NnS > 2(PK + 4DOS)$$

With 75% overlapping at least four devices are needed, and the minimum sampling period under the same conditions as above is given by;

$$NnS > 4(PK + 4DOS)$$

When all the outputs are commoned the dump must be completed in one quarter the load time. Since the maximum output rate is 26.66 MHz, then the maximum input rate is 6.66 MHz. For all transform sizes this is much less than suggested by the above equation. Thus separate output circuits must always be provided if the maximum performance is to be achieved.

Note that if you input or output rate chosen is asynchronous to SCLK, then a PDSP16510 (or similar) is assumed at the PDSP16510 interface.

MULTIPLE CONCURRENT TRANSFORMS

The PDSP16510 will support 4 concurrent 64 point complex transforms, or 8 concurrent 64 point real transforms. When the performance is increased by using multiple devices configured in MODE 1, a double LFLG transition is provided to support block overlapping. This is illustrated in Figure 7.

The LFLG output goes high after the first sample is loaded, and then low half way through the first sub block if 50% overlaps are in use. This transition is used to instigate the load procedure into the second device. LFLG then returns high and goes low again half way through the last sub block. This second low going transition is used to instigate a new low procedure in the first device.

Note that once a load procedure has started, the occurrence of a second edge will not effect the device in any way and thus each device only responds to an edge as shown in Figure 7. Note also that in this arrangement the DAV / DEN connection MUST be made, even if DIS and DOS are not commoned. This ensures that each device has an equal amount of time to dump its data.

This double LFLG transition also supports 75% overlapping with four devices. The first low going edge occurs 25% through the first sub block, and the second low going edge 25% through the last sub block.

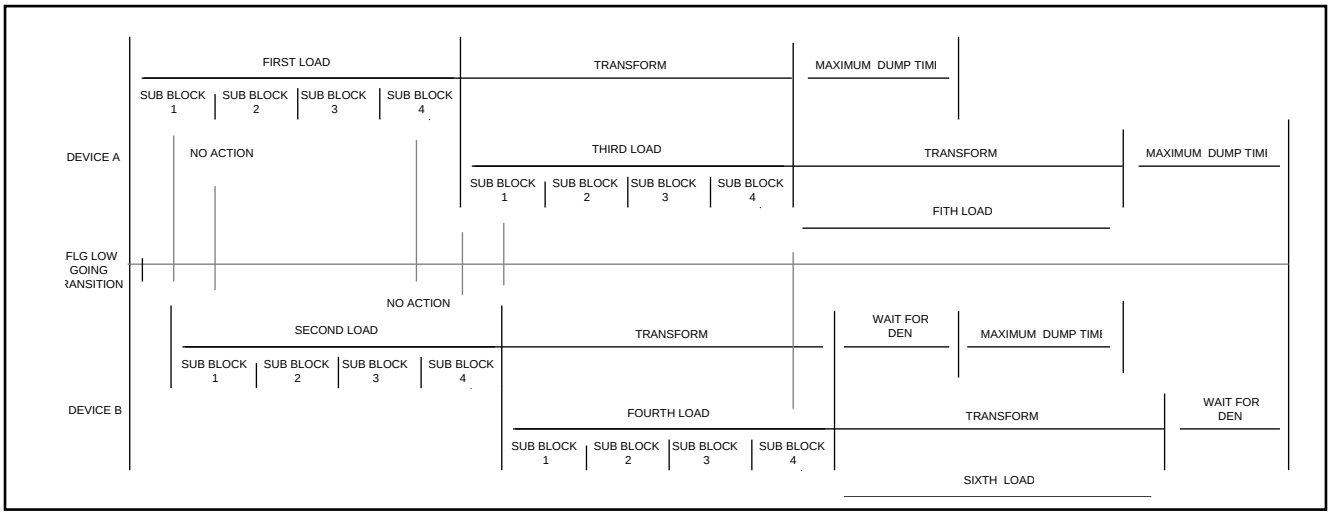

Figure 7. Sequence of events when performing multiple concurrent transforms.

HEADQUARTERS OPERATIONS
MITEL SEMICONDUCTOR
Cheney Manor, Swindon,
Wiltshire SN2 2QW, United Kingdom.
Tel: (01793) 518000
Fax: (01793) 518411

MITEL SEMICONDUCTOR
1500 Green Hills Road,
Scotts Valley, California 95066-4922
United States of America.
Tel (408) 438 2900
Fax: (408) 438 5576/6231

Internet: http://www.gpsemi.com
CUSTOMER SERVICE CENTRES
- **FRANCE & BENELUX** Les Ulis Cedex Tel: (1) 69 18 90 00 Fax : (1) 64 46 06 07
- **GERMANY** Munich Tel: (089) 419508-20 Fax : (089) 419508-55
- **ITALY** Milan Tel: (02) 6607151 Fax: (02) 66040993
- **JAPAN** Tokyo Tel: (03) 5276-5501 Fax: (03) 5276-5510
- **KOREA** Seoul Tel: (2) 5668141 Fax: (2) 5697933
- **NORTH AMERICA** Scotts Valley, USA Tel: (408) 438 2900 Fax: (408) 438 5576/6231
- **SOUTH EAST ASIA** Singapore Tel:(65) 3827708 Fax: (65) 3828872
- **SWEDEN** Stockholm Tel: 46 8 702 97 70 Fax: 46 8 640 47 36
- **TAIWAN, ROC** Taipei Tel: 886 2 25461260 Fax: 886 2 27190260
- **UK, EIRE, DENMARK, FINLAND & NORWAY**
  Swindon Tel: (01793) 726666 Fax : (01793) 518582

These are supported by Agents and Distributors in major countries world-wide.