BACKGROUND

The PDSP16510 FFT Processor contains 1K x 32 bits of internal RAM, enough to provide working memory for up 1024 point complex transforms. Once this memory is loaded with data no further external intervention is needed. On every pass of the transform, data is read from the RAM using the correct address sequence, and then written back after the butterfly operation. When the data has been completely transformed, it must be read out of the RAM and transferred to the next system element before new data is loaded.

All this, of course, takes many clock cycles, and in the meantime new data is being collected by the acquisition system. If this data is not to be lost it must be stored somewhere for future processing.

For transform sizes up to 256 points this presents no problem to the PDSP1510; it contains sufficient RAM to provide both working storage and input buffering for new data. In fact it contains sufficient RAM to also provide output buffering. This allows a single FFT Processor to handle higher sampling rates than might otherwise be expected.

The input and output buffers allow the time taken to transfer data in and out of the device to be effectively lost at the system level. Thus, whilst the working RAM is being used to transform a set of data, the output buffer can be dumping data previously transformed and the input buffer can be acquiring data to be next transformed. At the system level data is being continuously transformed, assuming, of course, that the time taken to do a transform is no greater than the time taken to load a new set of data.

When 1024 point transforms are to be undertaken no additional internal buffering is possible. Concurrent load, transform, and dump operations are thus not possible, and incoming data must be externally buffered if no information is to lost whilst a transform is in progress. For continuous transforms, the time taken to load this buffer must be greater than or equal to the sum of the time taken to read data from the buffer into the PDSP16510, then to transform it, and finally to transfer to the next device.

At first sight this buffer could be a simple FIFO, albeit a very wide 32 bit FIFO. It would also needs read rates of 40 MHz if maximum throughputs are to be possible. Once this FIFO were full additional logic would have to ensure that at least one location was transferred to the FFT Processor before the next word was written.

Many DSP applications, however, need to overlap data sets in the time domain before they are transformed to the frequency domain. This is the result of the need to apply a window operator to the data before it is transformed. Since only a finite segment of a signal can be observed at any time, discontinuities at the edge of the segment will introduce spectral errors. These are minimized by applying a window operator which weights the data more in the middle and less at the edges. There is thus a danger of missing some information at the edge of the segment, and this is avoided by overlapping the segments. Typically segments need to be overlapped by 50% or 75% to avoid loss of information, but the greater the overlap the less is the data sampling rate that can be achieved.

Overlapping data sets implies that old data must be re-read before new data is appended; an impossible task with a FIFO. For this reason the PDSP16540 Bucket Buffer has been introduced to support the FFT Processor. It allows data sets to be overlapped in 32 word increments, and requires no supporting logic. Although primarily designed to support 1024 point transforms, it can in fact help in smaller cases. The PDSP16510, itself, then supports 50% or 75% overlapping, but the PDSP16540 can be used when different amounts of overlapping are needed. This is discussed later.
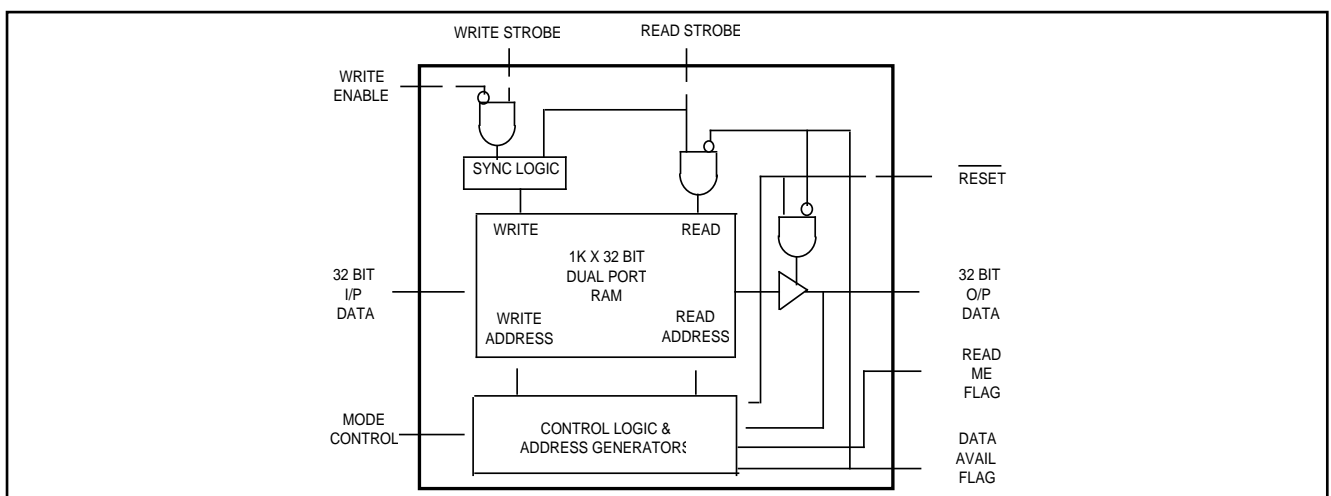

Figure 1. The PDSP16540 Bucket Buffer

THE PDSP16540 BUCKET BUFFER

This device is essentially a 1k x 32 bit synchronous RAM. Being synchronous it requires a continuously available clock, which normally would be the same as the PDSP16510 input and system clocks. It thus has the same 40 MHz maximum rate as the PDSP16510. Note that the data sheet for the PDSP16540 refers to this continuous clock as the read strobe - this does not imply that the strobe should only be present when a read operation is needed.

A write strobe with an enabling signal is needed to write data to the RAM. This write strobe can be asynchronous to the continuous read strobe, and is only needed when data is actually to be loaded. It would normally, however, be the data sampling clock used by the data acquisition circuitry, and thus is expected to be slower than the PDSP16540 read strobe (i.e. PDSP16510 input clock DIS). For example a single PDSP16510 will support continuous 1024 point transforms with sampling rates of some 6.7 MHz, when using 40 MHz clocks. The read strobe for the PDSP16540 is then 40 MHz, and the write strobe is 6.7MHz.

This biased read write ratio makes the use of a true dual port RAM unnecessary. Whenever a write operation is needed the read operation can be interrupted for one cycle, and the write operation actually internally performed with the continuously available read strobe (internal write strobe to read synchronization also takes place).

The device is designed to interface easily to the PDSP16510, and provides comprehensive data overlapping facilities. For correct operation both the block length of the data to be transferred to the PDSP16510, and also the amount of new data in that block must be defined. For commonly used set ups, these two parameters can be defined by tying mode pins high or low. For other alternatives tri-state buffers are needed, connected to up to 16 of the output pins. These are enabled during reset, when these outputs become inputs to an internal latch.

When the programmed amount of new data has been written to the RAM, a Data Available flag (DAV) goes active. This goes in-active for one cycle whenever more data is written to the buffer, and goes permanently in-active when the programmed block length has been transferred. When DAV is active the PDSP16540 will automatically produce new output data on every read strobe edge. The receiving device cannot halt the operation, and it must be dedicated to the transfer task. The DAV signal should be used to provide a clock enable signal for this receiving device.

DAV is the only signal needed to interface to the PDSP16510. For more general applications an additional Read Me Flag is provided. This can be programmed to go active before DAV, and thus warn the receiver that data is about to appear. This signal has no action internal to the PDSP16540.

CONNECTING THE BUCKET BUFFER TO THE FFT PROCESSOR

Figure 2 shows a typical 1024 point system with 50 % block overlapping. Grounding MD0 specifies that 1024 word blocks will be read from the RAM when DAV goes active. Forcing MD2:1 to logical 01 will ensure that DAV goes active when 512 new words have been written to the RAM. Thus the 1024 word block that is transferred to the FFT Processor consists of 512 previously used words and 512 new words. These new words are written to the RAM using the asynchronous Write Strobe, which is also the sampling clock used by the data acquisition circuit. Inputs MD4:3 are really don't care inputs defining when the unused Read Me Flag goes active, but are grounded for electrical reasons rather than leaving them open circuit.

MD5 should be grounded when complex words are to be processed. It should only be tied high if 1024 real transforms are to be performed with no block overlapping (i.e. MD2:0 must be tied low). In this particular case the Bucket Buffer will acquire two blocks of 1024 point real data, through inputs IP15:0, before DAV goes active. These two blocks are then transferred concurrently using all 32 outputs, and the PDSP16510 must be programmed to expect dual real blocks (Control Register Bits 8:6 = 101).
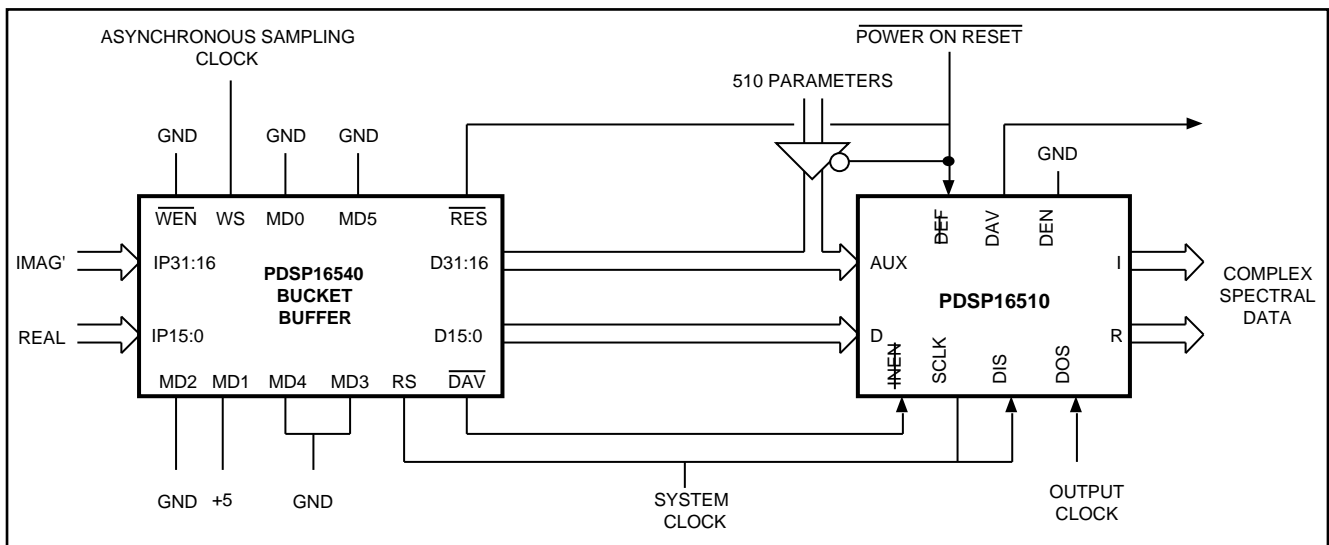

Figure 2. A Typical 1024 Point System with 50% Overlaps

The DAV output is directly connected to the INEN input on the PDSP16510. For correct operation the PDSP16510 must be programmed to use this input as a simple enable i.e. Control Register Bit 12 must be set. The following equation must be obeyed to prevent the loss of any incoming data when doing continuous transforms:

$$NS > 1024B(S/(S - B)) + T + D$$

where N is the amount of new data, S is the input sampling period, B is the read strobe period, T is the transform time which in the case of 1024 point transforms is 3907 system clock cycles, and D is the time to transfer data into the next device.

The factor S/(S - B) arises because the read sequence is interrupted for one B period every time new data is written to the buffer. It thus requires more than 1024 B periods to transfer 1024 words to the PDSP16510. For example if the read rate (B) is 4 times faster than the write rate (S), every 4th read cycle will be inhibited. Thus only 3 out of every 4 read cycles will actually result in data being transfered from the PDSP16540 to the PDSP16510. To achieve the maximum sampling rate possible (i.e. minimum S) data should be transferred in and out of the PDSP16510 at 40 MHz, and the system clock should also be 40 MHz.

Solving for these values gives;

$$NS \text{ must be } > \frac{25600 \times S}{S - 25} + 123375 \qquad \text{where S is in nanoseconds.}$$

Rearranging the above equation into the standard quadratic form (i.e. $S^2 + pS + q = 0$) and solving for the routes gives the value for S.

For no overlapping N = 1024 and S must be greater than 150 nanoseconds. The maximum sampling rate is thus 6.66MHz.

For 50% overlapping N = 512, and the minimum S period is 296 nanoseconds. The maximum sampling rate is then 3.37 MHz

For 75% overlapping N = 256, and the minimum S period is 589 nanoseconds. The maximum sampling rate is then 1.69 MHz

Suppose system requirements, for example, dictate a sampling rate of 4 MHz and some overlapping is required. One solution, of course would be to use more than one FFT Processor as explained in the data sheet. The PDSP16540 bucket buffer would then not be needed. Since, in this particular example, the sampling rate achievable with 50 % overlapping is close to the 4 MHz requirement, it may be possible to compromise on the actual overlap used.

Solving the above equation for S = 250 results in the need to load at least 606 new samples before DAV goes active. By setting MD2:1 to 11 it is possible to define the required number of new words in multiples of 32. A 5 bit code is then inputted through a tri-state driver connected to the D9:5 outputs, which become inputs during RESET. This binary code specifies up to 31 additional blocks of 32 above the minimum of 32.

Rounding 606 up to 608 (32 x 19) results in the need to load the code 10010 through pins D9:5 to achieve the overlap possible. The actual percentage overlap is then 40.6% and the 4 MHz sampling rate will be possible.

In a similar manner the Bucket Buffer can be used to provide non standard overlapping when transform sizes smaller than 1024 points are required. Such a system is illustrated in Figure 3. By forcing pin MD0 high it is possible to define block sizes of less than 1024 words. A 5 bit code on pins D4:0, during RESET, defines up to thirty one additional 32 word blocks after the basic 32 word block. Thus to define a 256 word block it is necessary to input the code 00111 via a tri-state driver on pins D4:0.

The calculation needed to define the minimum sampling period with a given overlap is different when 1024 point transforms are not being performed. As explained earlier, the time to transfer data in and out of the PDSP16510 can then be effectively lost
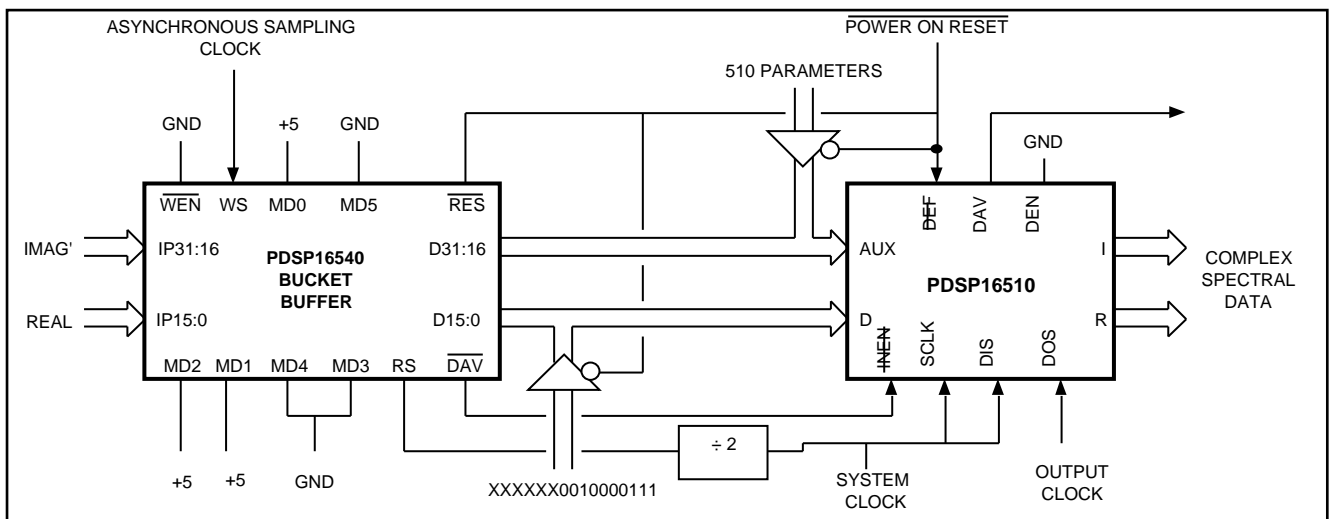

Figure 3. A 256 Point System with Non Standard Overlaps

at the system level. The required equation then simplifies to:

NS > T   if transfer times in and out of the PDSP16510 are less than T.

N, S, and T are as previously defined

When concurrent load, transform, and dump operations occur in the PDSP16510, it is not possible for the input clock rate (DIS) to be same as the system clock rate. The actual input rate must be reduced by the factor F from the system clock rate, where F is given by;

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

Thus if the system clock rate is to be 40 MHz, the clock low time would be say 12 nanoseconds, and the factor F would be 0.665. The maximum input rate is thus 26.6 Mhz.

In practise this is much greater than that needed to ensure that the time, to transfer 256 words from the Bucket Buffer to the FFT Processor, is less than the transform time. From the PDSP16510 data sheet the time taken to perform a 256 point transform, with a 40 MHz system clock, is 20.4 microseconds. Thus the input clock period needed to load 256 points in that time is 79 nanoseconds: or an input rate of 12.65MHz.

In this example the most convenient approach is to simply divide the system clock by two in order to provide the read strobe for the Bucket Buffer. There is, however, a relationship between the read strobe rate and the maximum write strobe rate. The write strobe period must be at least twice the read strobe period plus 10 nanoseconds. Thus with a 20 MHz read strobe the maximum write strobe rate is 9 MHz.

This writing rate is only achievable with read overlaps up to 50 %. Beyond this a second read rate requirement comes into effect. The write strobe period must also be greater than the read period multiplied by L/N, where L is the read block length and N is the amount of new data. This is another way of saying that the time taken to read the complete block must be no more than the time taken to load the required amount of new data.

In practise with a 20 MHz read strobe these considerations will not limit the writing rate in any way, and the maximum rates will be solely governed by the transform time in the FFT processor. Suppose, for example, we need to support a sampling rate of 7MHz (144 ns period) when doing 256 point transforms with some overlap. Then;

N X 144 must be > 20400 (the transform time)

Thus N must be greater than 142 for 7MHz sampling rates. Since N must be rounded up to a multiple of 32 it is thus necessary to load 160 (32 x 5) new samples in the 256 word block. This requires the code 00100 to be present on pins D15:10 during RESET, and gives 37.5% block overlapping.

Note   :   In all the above equations any requirement for the input clock (DIS) to be asynchronous to the system clock (SCLK) of the PDSP16510, will have to be modified in a practical system. The PDSP16510 has a requirement that its input and system clocks must be synchronised to each other. It may be possible to burst into the PDSP16510 data at a higher rate than the equations specify, such that DIS is synchronous to SCLK, but that on average the required DIS rate is achieved. When using the PDSP16540 on the input, however, it is very easy to burst data into the PDSP16510.

MITEL

SEMICONDUCTOR