

PDSP16515A

STAND ALONE FFT PROCESSOR WITH ENHANCED INTERNAL ACCURACY

The PDSP16515A performs Forward or Inverse Fast Fourier Transforms on complex or real data sets containing up to 1024 points. Data and coefficient input are both represented by 16 bits. Data is expanded internally to 18 bits and subject to Block Floating Point arithmetic to preserve a greater 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 PDSP16515A 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 PDSP16515A can be directly connected to an A/D converter to perform continuous transforms. The user can choose to overlap data blocks by either 0%, 50%, or 75%. Inputs and outputs are synchronous to the 40MHz system clock used for internal operations.

A 1024 point complex transform can be completed in some 98µs, which is equivalent to throughput rates of 450 million operations per second. Multiple devices can be connected in parallel in order to increase the sampling rate up to the 40MHz system clock. Six devices are needed to give the maximum performance with 1024 point transforms.

Either a Hamming or a Blackman-Harris window operator can be internally applied to the incoming real or complex data. The latter gives 67dB side lobe attenuation. The operator values are calculated internally and do not require an external ROM nor do they incur any time penalty.

The increased internal bus size together with block floating arithmetic produce up to 85dB of noise rejection.

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.




Fig. 1. Block Diagram

FEATURES

- Completely self contained FFT Processor
- Pin and functionally compatible with the PDSP16510A
- Expanded width internal RAM supports up to 1024 complex points
- 18 bit internal data bus with block floating point arithmetic for increased dynamic range
- 450 MIP operation gives 98 microsecond transformation times for 1024 points
- Up to 40MHz sampling rates with multiple devices.
- Up to 85dB noise rejection
- A choice of internal window operators with no external ROM provide up to 67dB side lobe attenuation.
- 84 pin PGA or 132 pin surface mount package

ASSOCIATED PRODUCTS

PDSP16540	Bucket Buffer
PDSP16330	Pythagoras Processor.
PDSP16256	Programmable FIR Filter.
PDSP16350	I/Q Splitter / NCO
PDSP16510A	Stand Alone FFT Processor





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



Pin Out for 84 PGA Package (AC84) - bottom view

PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC	PIN	FUNC
1	VDD	23	AUX13	45	GND	67	D8	89	GND	111	GND
2	GND	24	VDD	46	VDD	68	D7	90	R3	112	S1
3	I7	25	AUX12	47	SCLK	69	D6	91	VDD	113	GND
4	I8	26	GND	48	GND	70	D5	92	R4	114	DOS
5	I9	27	AUX11	49	GND	71	GND	93	GND	115	DOS
6	I10	28	VDD	50	DAV	72	VDD	94	R5	116	VDD
7	VDD	29	GND	51	GND	73	D4	95	R6	117	S2
8	I11	30	AUX10	52	INEN	74	GND	96	R7	118	GND
9	GND	31	AUX9	53	VDD	75	D3	97	R8	119	S3
10	I12	32	AUX8	54	DEF	76	VDD	98	GND	120	GND
11	VDD	33	AUX7	55	GND	77	D2	99	VDD	121	VDD
12	I13	34	VDD	56	DIS	78	GND	100	R9	122	I0
13	GND	35	AUX6	57	VDD	79	D1	101	VDD	123	I1
14	I14	36	VDD	58	D15	80	VDD	102	R10	124	GND
15	VDD	37	AUX5	59	D14	81	D0	103	R11	125	I2
16	I15	38	GND	60	GND	82	LFLG	104	R12	126	I3
17	GND	39	AUX4	61	D13	83	GND	105	R13	127	I4
18	DEN	40	AUX3	62	D12	84	R0	106	GND	128	GND
19	AUX15	41	AUX2	63	D11	85	GND	107	R14	129	VDD
20	GND	42	VDD	64	D10	86	R1	108	R15	130	I5
21	AUX14	43	AUX1	65	VDD	87	VDD	109	DISAB	131	I6
22	GND	44	AUX0	66	D9	88	R2	110	S0	132	VDD

Pin Out for 132 Leaded Chip Carrier (GC132)

PDSP16515A

SIGNAL	TYPE	DESCRIPTION
D15:0	I	Data input during real only mode. The real component in complex data mode.
AUX15:0	I	When DEF is active AUX15:0 are used to define the operating mode as defined in Table 3. When DEF is in-active AUX15:0 either provide the 16 bit imaginary component of complex input data, or a second set of real only inputs.
R15:0	O	These pins output the real component of the transformed data when DAV and DEN are active. Otherwise they are high impedance.
I15:0	O	These pins output the imaginary component of the transformed data when DAV and DEN are active. Otherwise they are high impedance.
DEF	I	The high going edge of DEF is used to internally latch the contents of AUX15:0, which then define the operating mode. In the simplest system DEF is a power on reset. When DEF is low the internal control logic is reset.
SCLK	I	System clock used for internal computations.
S3:0	O	These pins indicate the number of shifts towards the binary point which have occurred as the result of the conditional scaling logic. When the data path right shift is restricted to 2 places per pass, state 15 is used to indicate an overflow and only a total of 14 shifts is possible.
LFLG	O	This flag indicates that data is being loaded into the device. It goes active in response to an INEN input, and may be programmed to go in-active after the complete, one quarter, or one half a data block has been loaded.
INEN	I	The use of this input is mode dependent. It is either used as an active low, load enabling, signal for the DIS strobe, or it is used to initiate a new block load operation.
DIS	I	The rising edge of this input is used to load data into the device.
DOS	I	The rising edge of this input is used to dump data from the device. In most applications it may be tied to the DIS input, even if the output rate must be higher than the input rate because of overlapped data blocks. The DIS input is then internally divided down.
DAV	O	An active low signal that indicates that a transform is complete. Transformed data will then be output in normal sequential order using DOS. It may be optionally programmed to be delayed by 24 DOS strobes to match the delay through a PDSP16330.
DEN	I	This input is used to enable the data dump operation when DAV has gone active. If it is tied low the device will automatically dump data when DAV goes active. Otherwise the device will wait for the enabling signal to go low before the dump operation commences.
DISAB	I	Only available in the 132 pin GC package. When high the block floating logic is disabled.
VDD	P	+5V pins
GND	P	Ground pins

NOTE. All references to DEF, INEN, DAV, and DEN within the text do not contain the bar designator, signifying an active low signal. This is considered to be implied by the signal name and is not meant to imply a change in the signal function.

FUNCTIONAL OPERATION

The PDSP16515A performs decimation in time, radix 4, forward or inverse Fast Fourier Transforms. Data is loaded

into an internal workspace RAM in normal sequential order, processed, and then dumped in the correct order. With real only input data the processing time can approximately be halved for a given transform size. Two real inputs then replace a single complex input, and are processed in parallel.

Either a Blackman-Harris or a Hamming window can be

generated internally, and applied to the incoming real or complex data with no time penalty. No external ROM is needed to support these windows. The Blackman-Harris window gives improved dynamic range over the Hamming window when two closely spaced frequencies are to be detected, and one is of smaller magnitude than the other. It does, however, reduce the actual frequency resolution, and the Hamming window may then be preferable.

Data in and out of the device is represented by 16 bit real and imaginary components, with 16 bit sine and cosine values contained in an internal ROM. Conditional scaling, coupled with word growth through the butterfly data path, gives increased dynamic range. Transforms can be computed with sample sizes of either 256 or 1024 data points. The 256 point option can alternatively be used to simultaneously execute either four 64 point transforms, or sixteen 16 point transforms. The 16 point mode can only be used with a rectangular window, and no overlapping of data blocks is possible.

The device can be configured, either, to perform continuous transforms in a real time application, or as slave processor to a more general purpose signal processing system. In the continuous mode, with transform sizes of 256 points or less, it contains three internal control units which simultaneously allow new data to be loaded, present data to be transformed, and previous results to be dumped. Additional, external, input/output buffering is not needed. The internal input buffer also allows data blocks to be overlapped by either 50% or 75%, apart from the mode with no overlaps.

When 1024 point transforms are to be calculated, without loss of incoming data during the transform time, it is necessary to use an input buffer. This requirement is satisfied by a single PDSP16540 support device.

In any of the real or complex modes it is possible to obtain higher performance by connecting devices in parallel. It is then possible to increase the sampling rate to that of the system clock used for internal operations.

The mode of operation of the device is controlled by 16 bits in a control register. These are loaded through the AUX15:0 port when a control signal DEF is active low. This port is also used to provide the imaginary component of complex input data, and, if complex transforms are to be performed, an external tristate buffer will be needed to isolate the control information. This should only be enabled when DEF is active. DEF is also used to initialise the internal circuitry, and can be a simple power on reset if control parameters need not be subsequently changed.

DATA PRECISION

During each pass of a radix-4 fast Fourier transform it is possible for either component of a particular result to grow by a factor of up to four in the first pass, and 5.242 in subsequent passes. This is between two and three bits in each pass and the data path must allow for this word growth to avoid any possibility of overflow. At the end of the data path the word is preserved at 18 bits and stored in the internal RAM. Any unnecessary 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 PDSP16515A




Fig. 3 One of Four Data Paths

compromises by allowing a 2 bit word growth during the butterfly calculation in the first pass. This is equivalent to ignoring the most significant bit of the 19 bit final result, which is assumed to be an extra sign bit, and then selecting the next 18 bits for storage. In subsequent passes a Control Register Bit allows the user to continue to select these 18 bits, or instead to use the 18 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.

PDSP16515A

If the 2 bit option is selected there is a possibility of overflow occurring in one of the passes. The prediction of overflow is mathematically difficult, and only occurs with specific complex square waves. Scaling down the inputs cannot be guaranteed to prevent overflow because of the block floating point shifting scheme, which is discussed later. Overflow can NEVER occur if the 3 bit option is chosen, but at the expense of worse dynamic range.

When overflow does occur a flag is raised which can be read by the user (see later discussion on scale tag bits), and the results ignored. In addition all frequency bins are forced to zero to prevent any erroneous system response.

Even with only 2 bit word growth poor dynamic range can result and becomes worse when the incoming data does not fully occupy all the bits in the word. These problems are overcome in the PDSP16515A, 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 18 bit word, and maximum data accuracy is preserved.

Up to four shifts are possible before every pass after the first, with a total of fifteen for the complete transform. At the end of the transform the number of left shifts that have occurred is indicated on S3:0. Lack of pins prevents a separate output being available to indicate that overflow has occurred in the 2 bit word growth option. For this reason the maximum number of compensating left shifts in this mode is restricted to 14. State 15 is then used to indicate that overflow has occurred.

The first step in the butterfly calculation multiplies 18 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. Eighteen 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




Fig. 4. RAM Organization with 256 Data Points




Fig. 5. RAM Organization with 1024 Point Transforms

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 independent input strobe (DIS), allows data to be loaded without the need for additional external buffering. An independent output strobe (DOS) is also provided. DIS and DOS can thus be tied together, this being particularly useful when the device is performing the inverse transform back to the time domain. Transfer of data occurs internally from DIS to SCLK, so although they can be of different frequencies, they must be synchronous to each other. In the same way transfer of data also occurs from SCLK to DOS, so while DOS can also be independent of SCLK it must also be synchronous to it. Inputs and outputs are both supported by flag and enabling signals which allow transfers to be properly co-ordinated with the internal transform operation.

In many applications the DIS and DOS inputs can be tied together and fed by the sampling clock. If the output rate must be higher than the input rate, as with multiple devices supporting overlapped data samples, both strobes can still be connected together. The clock supplied should then be twice or four times the sampling clock, and an internal divider can be used to provide the correctly reduced input rate. The provision of a separate DOS pin does, however, allow the output rate to be different to the input rate, and therefore faster than strictly needed. Further output processing at higher rates is then possible if this is advantageous to system requirements.

The internal workspace is double buffered when 256 point transforms are to be performed. A separate output buffer is also provided. These resources, together with separate input and output buses, allow new data to be loaded and old results to be dumped, whilst the present transform is being computed. Additional, external, input buffering is not needed to prevent loss of incoming data whilst a transform is being performed.




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

When block overlapping is required, internally stored data will be re-used, and a proportionally smaller number of new samples need be loaded. Note that the internal window operator still functions correctly since it is actually applied during the first pass, and not whilst data is being loaded. The internal RAM organisation is shown in Fig. 4. It should be 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 PDSP16515A always expects to receive 1024 words, regardless of the amount of block overlapping. Data stored internally cannot be re-used when block overlapping is required, and data from the external buffer must be re-read as necessary.

Fig. 6 illustrates a typical 1024 point system with an input buffer which supports complex input data. The input buffer

can be provided by a PDSP16540 Bucket Buffer without the need for any external control logic. It supplies RAM for 1024 x 32 complex words, and allows transfers to the FFT Processor at the full system clock rate. The PDSP16540 also supports the standard 50% and 75% data block overlapping, but in addition allows the user to define the amount of overlap to within 32 words.

If no incoming data is to remain un-processed, the user must ensure that the time taken to acquire sufficient data to instigate a new transform is greater than or equal to the transformation time itself. The latter can be calculated from Table 4, once the system clock rate has been defined. When 1024 point transforms are performed, both the time to read data from the input buffer, and also the time to dump data, must be included in the calculation to determine the minimum time in which data can be loaded into the external buffer.

The peak transfer rate is limited by the characteristics of the I/O circuits, but can be greater than the sampling rate which is determined by the transform time. When load and dump operations are not concurrent with transform operations (as in the 1024 point modes), then the maximum I/O rate is equal to the system clock rate, Φ . When other transform sizes are specified, the sampling rate, S , is reduced by a factor F . This is defined below where Φ is in MHz and L is the system clock low time in nanoseconds:




Table 1. Advanced Timing Information with Continuous Inputs.

PDSP16515A

$$S = F\phi, \text{ where } F = 4 / (6 + 0.001\phi L)$$

F is typically 0.66 and applies to all transforms except for those of 1024 points, even if INEN is driven such that concurrent operations do not actually occur (Note also that S must be synchronous to SCLK). If this causes a system limitation in a single device application, then the device can be configured for pseudo, Mode 2, multiple device operation. Separate load, transform, and then dump operations will then always occur, but DEN must be low when a transform is complete or DAV will never go active. See the section on multiple device operation.

LOADING DATA

Data loading is controlled by three signals; DIS an input strobe, INEN a load enable, and LFLG an output flag. Detailed timing information is given in Table 1. Once sufficient data has been acquired, a transform will automatically commence. This is normally after a complete block has been loaded, except when a single device is performing overlapped transforms of 256 points or less. With 75% overlapping, transforms will commence after 25% of a new block has been loaded, and with 50% overlapping transforms commence after 50% of the data has been loaded. The remainder of the block is provided by data already stored in the internal RAM.

The data strobe is used to load data into the internal workspace RAM, and data must meet the specified set up and hold times with respect to its rising edge. DIS can be a continuous input since the device only loads data when an input enabling signal is active.

An internal synchronisation interval is necessary between the last sample being loaded with the DIS strobe and transforms being started with the system clock. This can be up to twelve system clock periods when data transfers and transforms are overlapped. The transform times given later in Table 4 are maximum values, and include these twelve periods.

The way in which the INEN signal controls data loading is dependent on whether a single or multiple device is to be implemented, and the status of Control Register Bit 12.

When Bit 12 is set in a SINGLE device system the INEN signal is simply used as an enable for the DIS strobes. When INEN is low, and provided the relevant set up and hold times have been satisfied, data will be loaded with the rising edge of the DIS strobe. If no gaps occur within the incoming data, INEN can be tied permanently low, provided that the sampling rate has been chosen such that transforms are completed before a new block of data is loaded. For transforms of less than 1024 points, data will then be continually processed without any loss of information. In the 1024 point modes the device will cease loading data when 1024 samples have been loaded, and even if INEN remains low no more data will be accepted until the previous results have been dumped.

In a multiple device system an edge is ALWAYS needed to commence a load operation, and Bit 12 has a different purpose. The edge is provided by INEN going low. Loading will cease when a complete block (or group of blocks with multiple concurrent transforms) of data has been loaded, even if INEN remains low. INEN must go high at some point after the minimum hold time has been satisfied, and then return low AFTER ALL DATA HAS BEEN LOADED, before a new load operation can commence. Low going edges which occur before all data has been loaded will be ignored.

The INEN edge mode is actually provided for the correct operation of multiple device systems, but if Bit 12 in the Control Register is reset in the SINGLE device mode, the edge activated operation will still be possible. With all but 256 point complex transforms, the single device edge mode of operation is identical to that of a multiple device system. With 256 point transforms, and their concurrent derivatives, the location of the low going edge in the data stream is dependent on the amount of block overlapping. The low going edge transition must be provided after 64 samples have been loaded with 75% overlapping, and after 128 samples have been loaded with 50% overlapping. With no overlapping the edge must be provided after 256 samples have been loaded.

In a single device system with Bit 12 set, INEN can be taken high to inhibit the load operation when gaps occur in the data stream. In the INEN edge activated mode gaps in the data stream can only be accommodated if the DIS clock is externally inhibited. Taking INEN high will not inhibit the loading of data in this mode.

With gaps in the data stream the peak sampling rates can be higher than continuous sampling rates. When data loading is not coincident with transform operations the peak rate can equal that of the system clock, otherwise it is reduced by the factor, F, given on the opposite page.

When Control Register Bit 12 is set in any multiple device mode, the DEF high going edge will also initiate a load operation after it has been internally synchronised to the rising DIS edge. If the first device in a multiple device system is programmed in this manner, the transform sequence will automatically start when DEF goes in-active. The other devices need the INEN edge as usual, and must have Bit 12 reset. A fuller explanation of the use of Bit 12 in a multiple device mode is given in the section on I/O In Multiple Device Systems. Note that the use of Bit 12 in a single device system (Control Register Bits 10:9 = 00) is completely different to its use in a multiple device mode.

The LFLG output goes active in response to the DIS rising edge used to load the first data sample, and indicates that a load operation is occurring. In an edge activated system the LFLG output will go high as the result of the first high going DIS edge after INEN has gone low. In the simple INEN enabling mode, internal logic counts the number of valid inputs and detects when the programmed block length has been reached. LFLG then goes low and will go high again in response to the next valid DIS strobe. LFLG will go low when DEF is active and will go high in response to the first INEN enabled DIS edge after DEF has gone in-active.

The active going LFLG edge does not normally have any system significance, but in the block overlapping modes the in-active going edge will occur when 50% or 75% of the data has been loaded. By driving the INEN input on one device with the LFLG output from a previous device, this edge can be used to partition data between several devices in a multiple device system. It can also be used to provide an address marker for a user defined input buffer, when executing 1024 point transforms with a single device. It is not needed, however, when the input buffer is provided by the PDSP16540.

DUMPING DATA

Data output is controlled by an output strobe [DOS], a dump enable signal [DEN], and a Data Available signal [DAV]. The DAV signal is used to indicate that the internal output

buffer contains transformed data, and the DEN input is used to control the outputting of that data. The output buffer within the device is clocked by the DOS input, and must be primed with a number of DOS strobes (see "user notes - stopping DOS") once a transform is complete in order to transfer data to the output pins. DAV will not go active until this priming has occurred.

The state of the DEN input at the end of a transform is used to control the transition of the active going edge of the DAV output with respect to the DOS strobes. The latter are then used to transfer data from the device to the next system component. If the DEN input is tied low in a single device system, the active going DAV transition will be internally synchronised to the rising edge of a DOS clock. If DEN is not tied low it must be guaranteed to be low at the end of the internal transform operation for this synchronization to occur. Since there is no external indication of this event, the user must take care to only allow DEN to go high whilst DAV is active, if this DAV synchronous mode is needed.

SYNCHRONIZED DAV OPERATION

In the DAV synchronised mode the first rising edge of the DOS clock, after DAV has gone active, must be used to transfer the first transformed sample from the output pins to the next system component. It should be noted that the output buffer will have been primed before the active DAV transition, since DOS must be a continuous clock, and there is then no delay before the first output becomes valid. The DAV output can be used as a clock enable for this next device, and transfers will continue in normal sequential order until the required data has been dumped. DAV will then go inactive in

response to the last DOS edge which was used to transfer data to the next device.

This mode of automatically dumping data when it is ready finds applications in real time data flow systems, and detailed timing is given in Table 2. It should be noted that the DOS input MUST be continually present before DAV goes active. If this is not the case the DAV output will not go active at the correct time, and the internal output circuitry will not be primed. Once DAV is active, however, it is possible for DOS to be irregular, and DEN can be used to inhibit the action of the output strobe as discussed previously. For the correct operation of the device the user must ensure that DOS becomes continuous and DEN remains low once DAV goes in-active.

When continuously transforming data such that new outputs are internally available before the previous block has been completely dumped, then DAV would normally stay active and give no indication that one block dump had been finished and another block started. Additional internal circuitry is, however, provided to ensure that DAV goes inactive for one DOS high time, thus supplying an inter block marker.

ASYNCHRONOUS DAV MODE

If DEN is not active in a single device when the transform is complete, then the device will wait for DEN to go active before any data is dumped. This mode is suitable for applications in which output processing is under the control of a remote host, such as a general purpose digital signal processor. The DAV output will then go active as soon as the output buffer is full, and will not be synchronised to the DOS edge. In such systems the DOS strobe may not necessarily be present at this time. Table 3 gives the relevant timing




Table 2. Output Timing with DEN tied low. (Advanced Data)

PDSP16515A

information.

In this host controlled dump mode the PDSP16515A waits for the host to activate the DEN input after DAV has gone active. DEN then functions as an enable for the host produced data strobes on the DOS pin. DEN may either stay active for the complete transfer, or may be used to enable each DOS input. When DEN and DOS are both active an internal read operation occurs, and an address generator is incremented. DAV goes in-active in response to the DOS edge needed to read the last output, unless Bit 15 in the Control Register is set. In this case DAV goes in-active when the next INEN edge is received for reasons given later.

In host controlled systems the time to dump data could be longer than the transform time. The dump time in such a system will dictate the maximum sampling rate that can be used without the loss of incoming data. In the 1024 point mode, when the loss of data is not important, the PDSP16515A is designed to not accept new data until the previous results have been dumped. Such a system needs no input buffer, and INEN can be permanently tied low if the edge activated mode is not in use. If the loss of data is to be avoided an input buffer is needed and the host must have received all the results before a new block of data has been loaded into the buffer.

For 256 point transforms, with host controlled dumping, it is still possible to overlap load and dump operations. The maximum dump times, however, must be less than the load times to avoid data corruption. Previously converted outputs will be actually corrupted, rather than inputs simply not being used.

If the loss of incoming data is not important, the device can be forced to do separate load, transform, and then dump operations. The corruption of results will then never occur, no matter what dump time is taken. This can be achieved by ensuring that INEN is not active between loading a block of data and completing the dump of the results from that data. The same ends can be achieved if the INEN edge activated mode (Bit 12 reset) is used, and the inverted DAV edge is used to drive the INEN input. This then initializes a new load operation only when the previous dump has been completed.

Results are transferred from the device with the rising edge of the DOS strobe when DEN is active. This is consistent with using the device in a data flow architecture, as is commonly employed in data processing systems. In a typical microprocessor based system, however, data is normally expected to become valid before the end of the data strobe produced by the processor. It is thus necessary for the user to provide a 'dummy' data strobe in order to transfer data to the outputs which can then be read by the host during the next data strobe. In addition further 'dummy' strobes are needed each time DAV goes active in order to prime the output circuitry. The actual output sequence is given in Table 3 for a single device system and is described more fully in "user notes - stopping DOS".

GENERAL DUMP CONSIDERATIONS

The tri-state drivers on the output buses are only enabled when both DAV and DEN are active. When DEN is tied permanently low the output bus will start to become valid from the DOS edge which also generates the DAV output. The next




Table 3. Host Controlled Output Timing. (Advanced Data)




Fig. 7. Host Controlled System

DOS edge can then be used to transfer the first output to the next device. When DEN is driven low in response to the DAV output, the outputs start to become valid when DEN goes low. The Scale Tag outputs become valid at the same time as data, and when enabled will continue to indicate the correct value until all frequency bins have been dumped. If at any time during the dump operation DEN goes 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 DAV going active until INEN again goes active, and a new set of data is loaded. The output tristate drivers, however, normally go high impedance when DAV goes inactive once a dump operation has been completed. In order to support systems in which it may be necessary to read the transformed data more than once, a Control Register Bit is provided which keeps the DAV output active until a further INEN edge is received. The user must then keep track of how many outputs have been dumped before INEN is generated to start a new load operation.

The DAV output can be delayed by an amount equivalent to the pipeline delay through the PDSP16330. This option is invoked by setting a control bit, and allows DAV to indicate that polar data is available at the output of the PDSP16330. When the option is used the tri-state outputs will be enabled when data is actually available and DEN is active, and not when DAV eventually goes active.

Two Control Register Bits allow a range of dump size options to be supported. In some applications the results of interest may only lie in the lower 25 or 50% of the frequency bins, the sampling rate having been chosen to prevent aliasing, and the transform size having been selected to give the required frequency resolution. In other systems it is only necessary to output the second half of a given sized transform. This is useful when filtering is to be performed in the frequency domain using Overlap /Discard Fast Convolutions. With this method FIR filters with N taps can be implemented in the frequency domain using 50% overlapped transforms on 2N samples. After multiplication in the frequency domain with the required frequency response, the inverse transform is performed and the first half of each output is discarded. Since only half the results are dumped, the dump clock need not be twice the rate of the clock used to load data.

FULL CO - PROCESSOR OPERATION




Figure 8. 1024 Point Real Transforms

A single device can be configured as a co-processor to a host system in which both the loading and dumping of data is under the control of the host. Such a system is shown in Figure 7, in which DEN is a host provided enable for host read operations, and INEN is an enable for host write operations. DIS and DOS are host data strobes.

The host loads a block of data into the PDSP16515A, using DIS enabled by INEN, which is then automatically transformed. The DAV output provides a flag indicating that the transform is complete, and results are then read by the host using DOS enabled by DEN. A new set of inputs is not normally loaded until the previous results are complete. If, however, 1024 point transforms are not to be performed, loading new data could coincide with dumping previous results. This, however, would require a host system with separate input and output buses, and which also allowed coincident transfers. As discussed previously, transferring results must take no longer than loading new data to prevent corruption of the outputs.

In the system illustrated by Figure 7, the host also controls the mode of operation of the FFT processor. The DEF signal is produced from an address decode, and the control parameters are loaded from the host bus by connecting the AUX inputs to the data outputs.

REAL ONLY TRANSFORMS WITH A SINGLE DEVICE

In the simplest case real transforms can, of course, be computed by forcing zero levels on the imaginary input pins. The device can, however, be configured to internally perform two simultaneous real transforms instead of a single complex transform. The block floating point logic will then use data from both blocks when it determines the number of shifts to be applied. This dual transform technique is used to increase the maximum permissible sampling rates, but since an additional data pass is required in order to un-scramble the transformed data, the actual performance is not quite double that possible with a complex transform of the same size. The 4 x 64 point complex mode becomes an 8 x 64 real mode, but the change from 16 x 16 complex transforms to 32 x 16 real transforms is not supported.

When a real transform is performed the algorithm produces complex results for each of the incoming data

PDSP16515A

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. 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. If block overlapping is needed, it is always necessary to load pairs of data blocks simultaneously, using both the real and imaginary inputs. With dual sources of data this presents no problem, and Control Bits 8:6 should be set to 110 or 111 for the relevant amount of overlapping. If data is from a single source an external FIFO is needed to provide a simple delay for a block of data. Decodes 001 through 100 from Control Bits 8:6 must be used to select the required overlap.

The output of the FIFO must provide data for the real inputs. Continuous inputs can still be accepted, and each block 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 & 2, and 2.5 is the overlap between 2 & 3.

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 outputs will be in the correct sequence from 1 to 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

Configuration		Clock Periods
16 X 16PT	COMP	456
4 X 64PT	COMP	660
256PT	COMP	852
1024PT	COMP	3943
8 X 64PT	REAL	852
2 X 256PT	REAL	1068
2 X 1024PT	REAL	4735

Table 4. Computation Times in Clock Periods

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 in-active 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 a 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 sizes, which include an allowance for synchronisation

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 : Guide to MAX Sampling rates (in MHz) possible from a single device system. SCLK is 40 MHz. Where sampling rate is asynchronous to SCLK, a pdsp16540 (or similar) is assumed on the input.

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 peak rate is limited by the system clock rate and the factor, F, given previously.

The time taken to dump the transformed data must be no more than the load time, if continuous inputs are to be supported and I/O operations are concurrent with transforms. With block overlapping the dump time must be reduced to the time taken to load the partial block. This dump time must include four extra DOS strobes needed to prime the output circuitry when a transform is complete. These, in effect, can be added to the transform time such that with concurrent I/O and 0%, 50%, or 75% overlapping;

nS or $(nS)/2$ or $(nS)/4$ must be greater than or equal to $PK + 4W$

where n is the transform size, S is the input DIS period, P is the number of clock periods given in Table 4, K is the system clock period, and W is the DOS period which can be less than S if necessary. Note also that S must be synchronous to SCLK, and if an asynchronous ratio is required then a pdsp16540 input buffer should be used.

When DIS and DOS are produced from a common source the minimum allowable sampling period must be increased to allow for the extra dumping time. Thus when DIS and DOS have equal periods and, for example, there is no overlapping;

$(n - 4)S$ must be greater than or equal to PK




Figure 9. Multiple Device Configuration

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 PDSP16515A at a user defined rate. The time taken to load this external buffer must be at least equal to the sum of the time to transfer data in and out of the FFT processor and the transform time itself. When data blocks are overlapped by 50% or 75%, no more than one half or one quarter of the block, respectively, must have been loaded in the same time. In the 1024 point modes the dump time can be any user defined value, and need not be increased to allow for block overlapping. The dump time, however, does directly effect the maximum sampling rates which can be accommodated without loss of incoming data.

The maximum sampling rates for 1024 point transforms at any load and dump rate can be calculated from the following relationship:

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

for 0%, 50%, or 75% overlapping respectively. S , P , and K were defined opposite. B is the clock period in which data is read from the input buffer and loaded into the device, D is the total dump time allowing for the four extra DOS periods. The periods of the load and dump clocks cannot be less than the system clock period. The maximum sampling rates given in Table 5 assume that a 40 MHz I/O rate is used, and that all results are dumped.

MULTIPLE DEVICE SYSTEMS

In real time applications several devices may be used in parallel in order to increase the sampling rate, but not to increase the transform size. When all outputs are commoned together, and feed a single output processor, then the data dump time must always be less than or equal to the time taken to load the data block (or 50% or 25% of the time with block overlapping). In most configurations with block overlapping the dump rate requirements will limit the maximum input rate, if only one output processor is provided. This can be avoided if the system provides separate output processors for every device. The system clock used for internal calculations then ultimately imposes a limit on the maximum sampling rate possible.

A multiple device system performing complex transforms with a single output processor is shown in Figure 9. The INEN/LFLG signals are used to co-ordinate the segmentation of data between devices. The in-active going edge of LFLG instigates the load procedure in the next device, and, since this edge can be programmed to occur either 25%, 50%, or 100% through the load operation, it can cause the next device to commence loading before the previous one has finished. In this manner data block overlapping is achieved. When multiple concurrent transforms are performed (for example 4×64 or 8×64) two LFLG transitions are sometimes needed to support block overlapping. This is fully explained in the section on Mode 1 sampling rates.

In any of the multiple device modes an INEN edge

transition is needed to start a new load procedure when the previous one has finished. When the LFLG output from the last device is fed back to the INEN input of the first device, continuous transforms will be executed. This continuous sequence can be started by the rising edge of DEF if Control Register Bit 12 is set in the first device (see section on Loading Data). This bit must not be set in the other devices. Since all devices are supplied from a common input bus and have a common source of control parameters, this Bit 12 inversion is best mechanized with an Exclusive OR gate in the AUX12 input line of the first device. The input can then be inverted when DEF is active but otherwise not be effected. Once the first device has been started with the DEF edge, the sequence will continue automatically using the LFLG /INEN connection between devices.

In many applications data is transformed continuously after power on, and the concept of a first data sample does not exist. If, however, the opposite is true, the first data sample must be present on the input pins such that it can be loaded with the second rising DIS edge after DEF has gone in-active. The data must meet the set up and hold times given in Table 1, and DEF itself must meet the parameters normally met by the INEN rising edge. The latter requirement is necessary to avoid a possible one DIS cycle variance, due the internal DEF synchronization logic. If the position of the first data sample is not important, it is not necessary for DEF to have any set up specification.

Without the feedback from the last device, the first device would wait for another externally supplied initialising pulse. In such a system with N devices in parallel, then N continuous transforms must be executed before the first device can wait for a new INEN input.

When only one output processor is provided the data outputs from all devices are connected together, and internal

logic will enable the tri-state outputs when a device is ready to output data i.e. DAV goes active. When data blocks are overlapped it is possible that the output rate requirements will limit the input sampling rate (see section on Multiple Device Sampling Rates). Additional output processors will remove this restriction, and the correct choice of multiple device operating mode will optimise the sampling rates that can be achieved with a given number of devices.

The synchronisation intervals, necessary to co-ordinate input and output operations with the transform operation, lead, in effect, to some uncertainty in the time needed to complete a transform. Thus a particular device in a multiple device system can effectively complete a transform in less system clock periods than another device in the same system. To prevent one device turning on its output bus before the previous one has finished, it is either necessary to use a faster output rate than would otherwise be required, or to use the inverted DAV output from one device to drive the DEN input of the next. The latter option allows DIS and DOS to be connected together, and ensures that the second device will not output data until the first device has finished.

This method of driving the DEN input from the inverted DAV output from a previous device requires a change to the single device DAV and DEN operation. If DEN is active at the end of a transform in a multiple device system, the DAV output will go active when the output circuit has been primed by the DOS strobes. This operation is identical to that provided for a single device system, and is transparent to the user as long as DEN and DOS are active. If DEN is not active, however, the DAV output will not asynchronously go active as happens in a single device system. Instead DAV will only go active when DEN eventually goes active. Since DEN is the inverted DAV output from a previous device, it is thus never possible for two devices to be actively outputting data. The DAV active going




Figure 10. Three Device System with Separate Load, Transform, and Dump Operations

edge remains synchronised to the DOS strobe since the DEN input will only go active when a previous DAV goes in-active. A further change to the output circuitry ensures that the output buffer is primed even though DEN is not active. The first word, however, only progresses as far as the final output latch. The output bus is not enabled, and address increments do not occur, until DEN is finally received. This modification to the internal control logic ensures that the output buffer does not impose unnecessary gaps between consecutive transforms. These gaps would, in turn, force the required DOS frequency to be greater than the DIS frequency (or greater than twice or four times the frequency with 50% and 75% overlaps).

The system illustrated by Figure 9 produces a common DAV output by OR'ing together all the individual, active low, DAV outputs. This is not guaranteed to give an indication when one transform has finished, and the next one has started, since it may simply glitch as one DAV goes in-active and the next one goes active after some delay. This glitch will not cause system problems since it occurs at a point clear of the high going edge of the DOS strobe. To provide a marker for the end of a transform each in-active going DAV edge should set its own latch, which is then reset by a subsequent DOS edge. The output of the latches can then be OR'd together if necessary.

Three multiple device operating modes are actually provided, and are selected with Control Register Bits 10:9. The choice of a particular mode is application dependent, and will effect the maximum sampling rate achievable with a given number of devices.

MULTIPLE DEVICE SAMPLING RATES

MODE 1. (BITS 10:9 = 01)

In this mode transfers in and out of the device are concurrent with transform operations. This mode must not be used for 1024 point transforms due to internal memory size restrictions. When real transforms are performed in this mode, only the real data input is used, regardless of the amount of block overlapping.

The increase in performance is directly related to the number of devices provided, but the input and output rates are limited to F/Φ where F and Φ 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, n is the transform size, S is the DIS strobe period, P is the number of system clock periods given in Table 4, K is the system clock period, and W is the DOS strobe period. Note that DIS should be synchronous to SCLK, and also that DOS should be synchronous to SCLK.

If an output processor is provided for every device, two devices with 50% block overlapping or four devices with 75% block overlapping will give the same sampling rates as a single device with no overlapping. If only one output processor is provided, the two or four times increase needed in the output rate over the input rate, usually imposes a limit on the input rate, since the output rate is limited to a factor, F , of SCLK.

In this operating mode the DIS and DOS strobes can often be tied together, since a faster DOS strobe gives no improvement in the sampling rates possible. This remains true even when the output rate must be twice or four times the input

rate due to block overlapping. Options can then be used which internally divide the DIS strobe by two or four, and thus allow the input to be driven by the faster DOS strobe.

In this mode the LFLG goes in-active after 25%, 50%, or 100% of the block has been loaded. When multiple transforms are performed concurrently (for example 4×64) a LFLG transition occurs at the relevant point whilst the first block in the group is being loaded. LFLG then goes high again and returns low at the overlap point in the last block. This double LFLG transition allows two devices to support 50% block overlapping, since the first transition from the first device can be used to initiate the load procedure in the second device. The second transition from the second device then initiates a new load procedure in the first device. The additional edges from each device have no effect since they occur when the device they are driving is already doing a load operation.

In such a two device system supporting 50% overlaps the inverted DAV from the first device must drive the DEN input of the second device. The data dumping time is then shared equally between both devices. The second device only outputs data when the first has finished, but both dumps must be finished in the time taken to load the group of blocks if only one output processor is provided. Without the DAV/DEN connection one device would only have had the time needed to load half of one sub block in which to dump its data.

In a similar manner four devices will handle 75% overlaps when concurrent multiple transforms are to be computed. The second, third, and fourth devices make use of the first transition, and ignore the second. The first device uses the second transition from the last device, and ignores the first. With the DAV/DEN connection each device will have one quarter of the load time to dump its data when a single output processor is provided.

More than two devices will provide increased performance for multiple transforms with 50% overlapping, and more than four devices will increase the performance with 75% overlapping. External logic is then needed to ensure that each device only uses the correct LFLG transition. Any device should only use the negative LFLG transition from a previous device if its own LFLG is low, and the LFLG output from the previous device plus one is low.

MODE 2 (BITS 10:9 = 10)

This mode is suitable for all transform sizes, since separate load, transform, and then dump operations occur. More devices than required by Mode 1 are necessary to achieve a given sampling rate, but the input and output rates can be any value up to the full system clock rate. As with Mode 1, additional output processors are needed to avoid the sampling rate restriction imposed by block overlapping.

The number of devices, N , needed to achieve a given sample rate can be derived from the following formula:

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

N is the number of devices, n is the transform size, S is the DIS strobe period, P is the number of system clock periods given in Table 4, K is the system clock period, and D is the total dump time including 4 extra DOS periods as discussed previously. The DIS and DOS periods are any value defined by the user,

PDSP16515A

down to the system clock period. Note that DIS should be synchronous to SCLK, and also DOS should be synchronous to SCLK.

In this mode increasing the output clock frequency will allow a greater continuous input rate. The provision of separate DIS and DOS pins allows this to be mechanized, and the DOS frequency can be increased to that of the system clock used internally. When the sum of the dump time (including four extra DOS periods for output priming) plus 12 system clock periods (the transform time variation caused by input synchronization) is less than the load time, one device will be guaranteed to have finished dumping before the next one starts. The inverted DAV to DEN connection between devices is then not needed, and all DEN inputs can be grounded.

The LFLG transitions occur at the same times as Mode 1, except that the double transition does not occur with multiple concurrent transforms. Fig. 10 illustrates a timing sequence with three devices. Real transforms still only use the real inputs regardless of the amount of block overlapping.

MODE 3 (BITS 10:9 = 11)

Multiple device Mode 3 is provided in order to improve the performance when block overlapping is needed, and separate output processors are provided. In this mode transfers in and out of the device are never concurrent with transform operations. The device will actually load extra data such that the required data to perform two overlapped transforms is stored internally. The amount of internal RAM prohibits the use of this mode when performing overlapped 1024 point transforms.

LFLG will go in-active after a normal data block have been loaded, regardless of the overlap selected. The device, however, continues to load more data. Thus, for example, in the 4 x 64 mode, five 64 point blocks will be loaded. This technique allows each device in the system to complete two or four overlapped transforms (depending on the amount of overlap) before any new data is needed. When doing a straightforward 256 point transform the device will load 256 + 128 data points.

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

$$(N - 1)L > 2PK + 2D \text{ for 50\% overlaps}$$

$$(N - 1)L > 4PK + 4D \text{ for 75\% overlaps}$$

where L is the time needed to load a normal block of data but not including the extra data, P is the number of system clock periods given in Table 4, K is the system clock period, and D is the total dump time including 4 extra DOS periods. As before, both DIS and DOS must be synchronous to SCLK.

When real transforms are to be performed on single sourced data, an external FIFO is needed to provide pairs of data blocks. These are loaded simultaneously into the real and imaginary inputs. See the section on real transforms.

OPERATING MODES

The operating mode of the PDSP16515A is determined

by the condition of 16 bits in an internal Control Register. The status of these bits is defined by the inputs present on the AUX15:0 pins when the DEF input is active. The DEF input can be a simple power on reset if the operating mode is fixed once power is supplied. The AUX pins are also used to provide the imaginary component of the complex input data. Thus, if complex inputs are needed, the mode definition must be implemented through a tri-state buffer which is only enabled when DEF is active. The imaginary input data must be disabled during this time.

Table 6 lists the functionality of each of the bits in the mode control register, and further explanations are as follows:-

BITS 2:0

These bits define one of 7 options for the sample size and type of data. In the 1024 point options the device will assume the non concurrent operating mode, regardless of whether a single or multiple device system is specified. The internal control logic will then ensure that data is loaded, transformed, and dumped in sequential operations.

For other data set sizes, loading, transforming, and dumping, can all occur simultaneously with a single device; the actual overlap will be dependent on the relative occurrences of the INEN input. Only in Mode 1 can concurrent operations be done with multiple devices.

BIT 3

This bit determines the number of right shifts built into the data path. In either condition only two right shifts occur during the first pass. If the bit is reset, three shifts occur in subsequent passes and the block floating point scheme allows up to fifteen compensating left shifts. If it is set, two shifts occur in every pass and overflow is possible. This is indicated by reducing the number of compensating left shifts to fourteen, and using scale tag value fifteen to indicate that overflow has occurred.

BITS 5:4

These bits define the choice of window operator. If other windows are needed they must be applied externally. The fourth option is used to specify the inverse transform, which does not require the use of a window operator. When 16 x 16 complex transforms are specified by Bits 2:0, only the rectangular window can be used. The use of any of the other options will cause the device to enter an internal test mode.

BITS 8:6

These bits define 0%, 50%, or 75% data block overlapping, and the division factor on the DIS input. Overlapping must not be specified with 16 x 16 complex transforms.

Two decodes allow the DIS input to be divided by two or four, when 50% and 75% overlapping is respectively needed. These options allow the DOS and DIS input pins to be still supplied from a common source, even though the output rate must be faster than the input rate. The frequency of this source would be dictated by the output rate requirement, with the input rate internally reduced by the correct amount.

Special decodes are provided to support real only transforms from dual sources, using both the real and auxiliary inputs. When data is from a single source, and no overlaps are needed, only the real input should be used. If 50% or 75%

overlaps are needed from a single source of real data, the device always expects blocks to be simultaneously loaded. An external FIFO is then needed to supply data to the real inputs after a delay of one block. Each block is thus loaded twice, firstly through the Auxiliary inputs and then through the Real inputs.

BIT 10:9

These bits define a single device system, or one of three multiple device possibilities. The choice between the first and second multiple device mode is dependent on the transform size and the sampling rate needed. The third mode should only be used when overlapped multiple transforms with less than 1024 points are to be performed simultaneously. It

BITS	Dec'	OPTION
2:0	000	16 x 16 COMPLEX
	001	4 x 64 COMPLEX
	010	256 COMPLEX
	011	1024 COMPLEX
	100	8 X 64 REAL
	101	2 X 256 REAL
	110	2 X 1024 REAL
	111	NOT USED
3	0	SHIFT 3 PLACES AFTER PASS1
	1	ALWAYS SHIFT 2 PLACES
5:4	00	RECTANGULAR
	01	HAMMING WINDOW
	10	BLACKMAN-HARRIS
	11	INVERSE TRANSFORM
8:6	000	NO OVERLAP
	001	50% OVERLAP
	010	50% OVERLAP AND DIS + 2
	011	75% OVERLAP
	100	75% OVERLAP AND DIS + 4
	101	DUAL SOURCE, NO OVERLAP
	110	DUAL SOURCE, 50% OVERLAP
	111	DUAL SOURCE, 75% OVERLAP
10:9	00	SINGLE DEVICE
	01	N DEVICES, CONCURRENT I/O
	10	N DEVICES, LOAD-TRANS-DUMP
	11	SPECIAL MULTIPLE TRANSFORM
11	0	DAV NOT DELAYED
	1	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

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 PDSP16515A will not generate DAV until 24 DOS clocks after data was actually valid. In this case the output tri-state drivers will be enabled at the correct time, even though the DAV signal was not externally valid. Host controlled dumping should not be used.

BIT 12

When this bit is set in the single device mode, the INEN input is a simple load enable signal. When it is reset an INEN edge is needed at the end of a load sequence before a new one can commence.

When it is reset in a multiple device mode it has no action, but when it is set it will cause the DEF high going edge to also initiate a load operation.

BIT 14:13

These bits allow four dump size options to be provided. Individual frequency bins are not accessible.

BIT 15

Under normal circumstances DAV would be expected to go invalid when a transform has been dumped. In some applications, however, it may be necessary to read the outputs more than once. When this bit is set, DAV will remain valid until the next INEN input, and will indicate that the transformed data still remains in the internal buffer. As soon as the next INEN is received the transformed data will be overwritten. Whilst DAV remains active the output tri-states will be enabled.

WINDOW OPERATORS

Since only a finite segment of a signal can be observed and processed at any one time, it is impossible to obtain pure spectral lines. Discontinuities are introduced at the boundaries of the observation interval which lead to spectral leakage. Windows are weighting functions applied to the data in order to reduce these discontinuities at the boundaries.

In the time domain the signal has to be observed through a finite window as a matter of accord. This is in fact equivalent to multiplying the signal with a set of uniform weights i.e. a rectangular window operator. In the frequency domain the spectrum of the data will be the spectrum of this weighting function shifted to the sinusoidal frequencies of the components in the data.

The rectangular window has a Fourier Transform which is a SINC(X) function. This has sidelobes which are only 13dB down from the main lobe. This severely limits the dynamic range of the system since a second sinusoid in close proximity would have its main lobe swamped by this side lobe. This would occur if its amplitude was a mere 13dB down from the first sinusoid.

Window operators are thus mathematically constructed to cancel these sidelobes as far as possible. Unfortunately this is normally done at the expense of making the main lobe spread over more frequency bins. This reduces the ability of the system to resolve two frequencies, and can only be overcome by using more data samples. This may not always be possible because of other system constraints.

A common rule of thumb defines the resolution of an FFT

PDSP16515A




Fig. 11. External Window Generator

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 PDSP16515A. These are constructed on the fly as needed, and take the general form:

$$A - B \cos x + C \cos 2x \text{ where } x = (2\pi n)/N, n = 0 \text{ to } N-1$$

For Hamming, $A = 0.54$, $B = 0.46$, $C = 0$

For Blackman-Harris, $A = 0.42323$, $B = 0.49755$, $C = 0.07922$

These windows can be applied to any of the transform size options, except the 16×16 complex variant. When the latter is specified the rectangular window option MUST be selected, or the device will be configured in an internal test mode.

If other operators are required these must be applied externally. This can be conveniently achieved with either a PDSP16112 or a PDSP16116, both of which are complex multipliers but with different accuracies. Fig. 11 shows how either one can be configured to perform two separate multiplications with one input common to both. This arrangement is necessary to perform the window function on

complex inputs.

Important features of the windows generated by PDSP16515A, and other commonly used windows, are illustrated in Table 7. The results are obtained from the reference quoted, which should be consulted for a full mathematical treatment. The significance of each parameter is outlined below :

Highest Side Lobe Level

The inherent rectangular window has sidelobes which are only 13dB down from the mainlobe. These severely limit the dynamic range. The object of the window is to improve this situation with better side load attenuation.

Mid-Point Loss

In line with the filter concept it is possible to conceive of an additional processing loss for a tone of frequency mid-way between two bins. This is defined as the ratio of the coherent gains of two tones, one at the mid-point and one at the sample point. It is expressed in dB in Table 8.

Overall loss

An overall figure for the reduction in signal to noise ratio can be obtained by adding the mid-point loss to the reciprocal of the equivalent noise power bandwidth in dB. It is a measure of the ability of the window to detect single tones in broadband noise. The variance between windows is less than 1dB.

6.0dB Bandwidth

This figure, expressed in bin widths, represents the ability of the window to resolve two tones and should be as close to unity as possible. As the highest sidelobe level is reduced, this parameter tends to get worse, and a compromise must be used when choosing a window.

Overlap Correlation

In many practical systems the squared magnitudes of successive transforms are averaged to reduce the variance of the measurements. If, however, a windowed FFT is applied to non overlapping partitions of the sequence, data near the boundaries will be ignored since the window exhibits small

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 PDSP16515A block FP	85	67	65
Full 32 bit Floating point with 16 bit inputs	93	82	67

Table 8. Comparative Dynamic Range Measurements

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 [i.e. the width of the frequency bin] and the window operator that is used. Dynamic range is also determined by the window operator, but in a hardware implementation it is also influenced by the number of bits used to represent the data throughout the calculation.

The hardware effects include the accuracy of the A/D converter, the number of bits representing the window operator and the twiddle factors, and the way the growth in word length is handled as the FFT calculation proceeds. The

obvious way to overcome these limitations is to use floating point arithmetic; but in real life the accuracy of the A/D converter is fixed and the sample size is limited. Floating point arithmetic is thus an overkill solution for the majority of applications. This is especially true for transform sizes up to 1024 points, which is the intended application area.

Figures given for the dynamic range of a system must be carefully interpreted, since there is no exact definition of the measurement. Three different ways of measuring dynamic range have been investigated using 1024 point transforms.

The 'best' dynamic range figures will be obtained with single tone measurements, and these results are often quoted to indicate the need for greater bit accuracies. The measure is the ratio of a full scale sinusoid to the average noise level and the results will be essentially independent of the window operator. The results given by the PDSP16515A 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 PDSP16515A. This can be attributed to the scaling scheme, word growth, and rounding method used within the device.

When two nearby tones are to be capable of detection, the window operator will dictate the performance of the system. The final column in Table 8 illustrates the results obtained using two sinusoids of different amplitudes, with the larger one residing mid-way between two frequency bins, and the smaller 5.5 bins away. The two frequencies are five bins apart to avoid the effects of the mainlobe widths. The dB figures given are the difference in amplitude between the two signals when the smaller one is still just detectable as a separate peak from the larger one.

This technique illustrates the performance of the window, since the amount by which sidelobe structure of the larger signal swamps the mainlobe of the smaller signal will determine if the smaller is detected. The theoretical attenuation of the highest sidelobe levels, with respect to the mainlobe, for the window options provided by the PDSP16515A 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 PDSP16515A is equivalent to using 24 bit data. The Blackman Harris window was used in all cases.

USER NOTES - STOPPING DOS

GENERAL DESCRIPTION

The transform is calculated internally fully synchronous to SCLK. However, as all outputs are referenced to DOS, a transfer has to be made between the two clocks. In addition, some dummy DOS strobes are needed to operate the internal control logic, and to advance data from the internal RAMs to the output pins.

The most simple configuration for the device is to have DOS running continuously and for DEN to be permanently active. When this happens the user will just be aware of data appearing on the output pins on the same DOS cycle when DAV goes active. However, there are many situations where either DOS is not continuously running, or DEN is not permanently active. To help explain how to operate the device in these situations, the internal operation of the output circuits must be described. For those who are not going to be interrupting DOS, the remainder of this section can be ignored.

INTERNAL RAM - GENERAL DESCRIPTION.

For single device operation of transforms less than 1024 points, the internal RAM is shared between three separate operations which enable the device to output old transformed results, calculate the current transform, and input new data ready for the next transform. All these operations, along with the internal control logic, are controlled by a 12-cycle state machine. The RAM operations are :

- (a) 2 cycles in every 12 are dedicated to reading new information in the input buffer and writing it to the RAM.
- (b) 2 cycles in every 12 are dedicated to reading the contents of the RAM and advancing that data to the output buffer.
- (c) 8 cycles in every 12 are dedicated to the read and write operations of the transform currently being calculated.

SEQUENCE OF EVENTS

The sequence of events relating to the output control and data flow is as follows :

(3.1) An SCLK rising edge :

- (a) An internal flag is raised to indicate that the transform has finished and data is available to be dumped. Data will be present in the internal RAM, and the output address generator will be at the correct address. Access to the RAM at this moment, however, has not been made.
- (b) If at this moment the device is programmed to be a single device, and DEN is inactive, then DAV will be made active - ie without the presence of DOS. If DEN is active at this point, or the device is programmed in any multiple device mode, then DAV will remain inactive.

(3.2) Accessing the RAM at this point

At this moment, when DAV has been made active before data appears on the output pins, data is not yet in the output buffer. Internally the precise SCLK cycle at which the RAMs are read and written to the output buffers now has to be waited for. This cycle, as described above occurs 2 in every 12

SCLK cycles, so at worst case 6 SCLK cycles have to elapse until data is guaranteed to be in the output buffer.

If the DOS rate is similar to the SCLK rate, and the user has been immediately applying DOS pulses (on seeing DAV go active) hoping to get data off the chip, then this will not actually happen.

The next internal flag raised is the one which indicates that the output data has been successfully read from the RAMs and is now in the output buffer.

(3.3) The next DOS rising edge (regardless of DEN status) :

The flag indicating that the RAMs have been read is transferred to circuitry operating on DOS. The output enable signal, DEN, does not have to be present at this point.

(3.4) The next DEN-Enabled DOS rising edge (ie the 1st one of this sequence)

The output state machine receives it's first edge.

(3.5) The next DEN-Enabled DOS rising edge (ie the 2nd)

Internal output address generators start to count (ready for fetching the next set of output data).

(3.6) The next DEN-Enabled DOS rising edge (ie the 3rd)

An enable signal is raised for the final data latch in the output buffer.

(3.7) The next DEN-Enabled DOS rising edge (ie the 4th)

- (a) The final data in the output buffer latch clocks-through new data and presents it to the output pads.
- (b) The output pads come out of high impedance.

- (c) If DAV was previously inactive, it is now made active.

OUTPUT SCENARIOS

Considering the above sequence, therefore, some single device situations can now be explained :

(4.1) DOS is continuously present, but DEN is inactive (Transform size less than 1024)

In this case, when the transform is complete, as the device is programmed as a single device and DEN is inactive, DAV will be made active. Even though DOS is running, the status of DAV at this point does not rely on it.

The user can now monitor the status of DAV, and after at least 6 SCLK cycles can initiate some further action, eg by external control force DEN active at some later time when the rest of the system is ready to accept the transformed data. Independently of this external control, the next DOS pulse will start to operate the sequence of events as described above (ie point No. 3.3). When DEN is eventually made active, the remainder of the above sequence (points Nos 3.4 to 3.7) is executed, with 4 DEN-Enabled DOS pulses needed before data is observed on the output pins.

If however the user immediately forces DEN active upon monitoring DAV go active and waiting for the required 6 SCLK cycles, then 5 DOS pulses would have to be issued. The first of these 5 would start the sequence of events as described above (3.3), and the fact that it is enabled by DEN would be irrelevant. The required DEN enabled pulses in this situation would be the 2nd, 3rd, 4th and 5th pulses supplied.

(4.2) DOS is not running, and DEN is inactive. (Transform sizes less than 1024)

In this situation, again as the device is programmed to be a single device and DEN is inactive at the point where the transform is complete, DAV will be made active regardless of the state of DOS. The user can now monitor this event on DAV and after waiting a further 6 SCLK cycles, use it to switch on DOS and to make DEN active.

DOS can now be switched on for at least one pulse (but may be more), and the sequence of events as described earlier (from point No 3.3) will start. DEN can then be made active, whereby a further 4 DEN-Enabled DOS pulses will be required before data is seen on the output pins. This is the situation shown in table 3.

Alternatively, DEN and DOS could be made to operate on the same cycle. In this case data will appear on the output pins on the 5th DOS pulse (the first would not actually require the presence of DEN, but the 2nd, 3rd, 4th and 5th would)

(4.3) 1024 point transforms, single device mode.

In the case of 1024 point transforms, the internal RAM is no longer operated in the manner described in section 2. The RAM is instead totally dedicated to one operation at a time. Thus data for a transform will be loaded, and all 12 out of 12 SCLK cycles will be available for the transfer of input data to the RAMs. During the transform no transfers from the input

to the RAM or from the RAM to the output are possible. This is why DIS and DOS can be equal to SCLK for 1024 point transforms.

If 1024 point transforms are being performed and the device is programmed as a single device, then "asynchronous" operation of DAV is possible as described earlier for transform sizes less than 1024 points. If DEN is inactive at the time the transform has finished calculating, then DAV will be made to go active regardless of the state of DOS. Although 6 SCLK cycles do not have to be waited for as in section 3.2, a transition has to be made from the transform controlling the internal RAM to the output circuits controlling it. This operation plus the time taken to advance data from the RAMs to the output buffer takes exactly 4 SCLK cycles.

Hence the sequence of events is exactly as described in section 3, except that section 3.3 should read 4 SCLK cycles rather than 6. The analysis of sections 4.1 and 4.2 are also true if the 6 SCLK cycle time is substituted with 4 SCLK cycles.

DUMMY DOS STROBES AFTER DEF

In addition to the dummy DOS strobes needed prior to dumping data, it is necessary to provide at least 4 DOS strobes after DEF has gone inactive, but before DAV goes active. These initialise the internal address counters and do not rely on DEN also being active. They are needed every time DEF has been used to change the operating mode.


PDSP16515A

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
Junction Temperature, Commercial	100°C
Junction temperature, Industrial	115°C
Junction Temperature, Military	155°C
Package power dissipation	5000mW

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.



ELECTRICAL CHARACTERISTICS

Operating Conditions (unless otherwise state)

PDSP16515A C0 Tamb = 0°C to +70°C.	$V_{CC} = 5.0V \pm 5\%$
PDSP16515A B0 Tamb = -40°C to +85°C.	$V_{CC} = 5.0V \pm 10\%$
PDSP16515A A0 Tamb = -55°C to +125°C.	$V_{CC} = 5.0V \pm 10\%$

Characteristic	Symbol	Value			Units	Notes
		Min.	Typ.	Max.		
Output high voltage	V_{OH}	2.4		-	V	$I_{OH} = 4mA$
Output low voltage	V_{OL}	-		0.4	V	$I_{OL} = -4mA$
Input high voltage	V_{IH}	2.0		-	V	SCLK, DIS, DOS, DEN need 3V
Input low voltage	V_{IL}	-		0.8	V	DEN needs 0.7V max
Input leakage current	I_{IN}	-10		+10	μA	$GND < V_{IN} < V_{CC}$
Input capacitance	C_{IN}		10		pF	
Output leakage current	I_{OZ}	-50		+50	μA	$GND < V_{OUT} < V_{CC}$
Output S/C current	I_{SC}	10		300	mA	$V_{CC} = Max$

SWITCHING CHARACTERISTICS

Characteristic	Symbol	Min	Max	Conditions
Clock Frequency (MHz)	$\emptyset$	DC	50	Max $\emptyset$ high time is 1msec
Clock High Period (ns)	T_{CH}	8		
Clock Low Period (ns)	T_{CL}	8		
Max DOS, DIS Frequency	$\emptyset_o$		F $\emptyset$	Less than 1024 points or Mult Dev Mode 1 Note $F = \frac{4}{6 + 0.001\emptyset T_{CL}}$
Max DIS Frequency	$\emptyset_o$		40	1024 points or Mult Dev Modes 2 and 3
Max DOS Frequency	$\emptyset_o$		40	D15, DOS must be derived from $\emptyset$

SCLK to DIS/DOS RELATIONSHIP

Both DIS and DOS must be synchronous to SCLK. Ideally they should both be produced from SCLK, in which case the SCLK rising edge would either be first or coincident with the DIS and DOS rising edges.

In any event, the rising edge of SCLK must not fall between 2ns and 10ns after the rising edge of either DIS or DOS

ORDERING INFORMATION

PDSP16515A C0 AC	(Commercial - PGA Package)
PDSP16515A C0 GC	(Commercial - Leaded Chip Carrier)
PDSP16515A B0 AC	(Industrial - PGA Package)
PDSP16515A B0 GC	(Industrial - Leaded Chip Carrier)
PDSP16515A A0 AC	(Military - PGA Package)
PDSP16515A A0 GC	(Military - Leaded Chip Carrier)
PDSP16515A/MA/GCPR	(Military - Screened Leaded Chip Carrier. See separate datasheet for details)