

# A 70-MHz 1.2- $\mu$ m CMOS 16-Point DFT Processor

RICHARD W. LINDERMAN, MEMBER, IEEE, CARL G. SHEPHARD, KENT TAYLOR,  
PAUL W. COUTEE, PAUL C. ROSSBACH, MEMBER, IEEE, JAMES M. COLLINS,  
AND ROBERT S. HAUSER, MEMBER, IEEE

**Abstract**—A chip architecture designed to compute a 16-point discrete Fourier transform (DFT) using Winograd's algorithm every 457 ns is presented. The 99 560 transistor 1.2- $\mu$ m chip incorporates arithmetic, control, and input/output circuitry with testability and fault detection into a 144-pin package. A throughput of  $2.3 \times 10^{12}$  gate-Hz/cm<sup>2</sup> and 79-million multiplications per second is attained with 70-MHz pipelined bit-serial logic. Combined with similar chips computing 15- and 17-point DFT's, 4080-point DFT's can be computed every 118  $\mu$ s. Using the 16- and 17-point chips, 272  $\times$  272-point complex data imagery can be transformed in 4.25 ms. A 24-bit block floating-point data representation combined with an adaptive scaling algorithm delivers a numerical precision of 106 dB (17.6 bits) after computing 4080-point DFT's.

## I. INTRODUCTION

THE NEED TO efficiently compute the discrete Fourier transform (DFT) arises in many areas of digital signal processing including radar, sonar, speech, seismic, and image signal processing. To satisfy the most demanding of these applications we propose a 1.2- $\mu$ m custom CMOS VLSI chip set using the Good-Thomas prime factor algorithm (PFA) [1] to combine smaller transforms computed using Winograd's transform [2]. By customizing the VLSI chips to the application at hand, we can achieve a throughput of over 8300 4080-point DFT's per second with processors operated at 70 MHz. The PFA processor places four dual-port ECC memories, a PFA controller, and three Winograd Fourier transform (WFT) chips in a hybrid circuit. The three generic WFT chips compute 15-, 16-, and 17-point DFT's, respectively, each accepting a 35-MHz complex input data sampling rate.

Computing large Fourier transforms requires high

numerical precision in addition to high computational throughput. Large transforms also require an intelligent approach to scaling which can prevent overflows while not sacrificing precision to excess guard bits. Our objective is to keep the noise related to numerical rounding and scaling insignificant compared to other noise sources such as sensor and A/D converter inaccuracies.

The WFT16 architecture is distinguished from other implementations [3], [4] in that it implements the entire Winograd transform on-chip. All the arithmetic, control, address generation, testability, and fault-detection circuitry resides on-chip. This provides simple and efficient partitioning of the data and control flow graphs. The arithmetic circuitry consists of 72 serial adders and 36 serial multipliers which compute the preaddition, multiply, and post-addition operations required by the WFT algorithm. To simplify memory partitioning, off-chip I/O is word parallel. The information is converted to serial bit streams once inside the WFT processor.

The WFT16 arithmetic circuitry extends the 24-bit input word data format to 32 bits. Choosing 26-bit fixed coefficients for the multipliers allows an output precision of more than 20 bits (120 dB). On-chip hardware supports an adaptive scaling algorithm to maximize precision by tracking the magnitude of the data while also preventing overflow.

Fault detection is achieved by using two on-chip methods, parity error and watchdog checking, in conjunction with the off-chip memory ECC.

## II. DFT ALGORITHMS

The DFT is a basic tool of digital signal processing which allows us to move back and forth between time-domain and frequency-domain representations of a signal. As such, it is probably one of the most frequently computed algorithms. Several excellent texts have been written on a variety of "fast" algorithms to compute the DFT [5], [6]. By far the most common approach has been to adopt the Cooley-Tukey fast Fourier transform (FFT). However, the PFA and the WFT have been proven to require fewer multiplications. Placing an algorithm in silicon requires more than just an examination of the number of oper-

Manuscript received August 11, 1987; revised November 2, 1987. This work was supported by the Defense Advanced Research Projects Agency, the Air Force Office of Scientific Research, and the Air Force Wright Aeronautical Laboratories.

R. W. Linderman is with the Air Force Institute of Technology, Wright-Patterson Air Force Base, OH 45433.

C. G. Shephard is with the Air Force Logistics Command, McClellan Air Force Base, CA.

K. Taylor is with the Engineering Installation Division, Tinker Air Force Base, OK.

P. W. Coutee and J. M. Collins are with the Aeronautical Systems Division, Air Force Systems Command, Wright-Patterson Air Force Base, OH 45433.

P. C. Rossbach is with the U.S. Army Signal School, Ft. Gordon, GA. R. S. Hauser is with the Arnold Engineering Development Center, Tennessee.

IEEE Log Number 8719255.

ations. Due to the pin limitations, we must find an algorithm which partitions efficiently between chips while capitalizing to the fullest extent on the inherent computing capacity of the technology at hand. A careful balance must be struck between I/O capacity and computing hardware since the algorithm requires  $N$  inputs and outputs and only  $O(N \log N)$  multiplications. Since either could become a bottleneck, we seek to achieve 100-percent utilization of both buses and ALU components.

The theory of "area-time optimal" DFT processors must also be scrutinized because of the important effect that functional partitioning has on the constant factor which the theory ignores. In fact, a direct comparison between an area-time optimal mesh-connected systolic array FFT processor [7] and the WFT/PFA processor implemented in the same technology and computing 4096- and 4080-point transforms, respectively, shows the WFT/PFA approach to have twice the throughput, in addition to dramatic advantages in reliability, size, and modularity [8]. Part of the reason for this result is the ability of the WFT/PFA approach to make the fullest use of the 1.2- $\mu$ m CMOS technology.

The DFT is written in matrix form as

$$V = Wv$$

where  $v$  is the vector of input data points,  $V$  is the vector of output points, and  $W$  is the matrix of DFT coefficients. Winograd's algorithm breaks the  $W$  matrix up into three matrices:

$$V = CDAv$$

where  $A$  is a matrix of preadditions performed on  $v$ ,  $D$  is a diagonal matrix of constants containing both trivial and nontrivial multiplications, and  $C$  is a matrix of post-additions. The WFT chips use this approach to compute 15-, 16-, and 17-point transforms. The prime factor algorithm specifies a method to compute transform lengths which are the products of these three relatively prime factors such as 240, 255, 272, and 4080, by successively computing the smaller transforms. It maps the longer one-dimensional transforms into two- or three-dimensional transforms using the Chinese remainder theorem [5].

Winograd has shown that his approach requires the fewest number of multiplications [2], and we could use this approach for the larger sizes, however, the ALU components would grow much too large to partition effectively among chips. We chose the 15-, 16-, and 17-point Winograd transforms because they require 36, 36, and 72 real serial multipliers, respectively. The multiplier cells have been designed to accommodate the stack of 72 multipliers on a single chip.

Compared with the Cooley-Tukey FFT, the PFA/WFT approach lacks flexibility in selecting DFT size. We feel this is more than compensated for by the increased throughput derived from a superior algorithm and a greater opportunity for on-chip parallelism in ALU components. The DFT sizes are limited to the base sizes (15, 16, and 17), sizes applicable to image processing (240, 255, and

272), and the 4080-point DFT for high-resolution applications such as radar.

### III. 4080-POINT PFA PROCESSOR

A 4080-point PFA was chosen as a representative system for implementation. In our approach, the algorithm is divided at the Winograd transform level. We chose to put the entire WFT on a single chip as opposed to partitioning the three stages into multiple chips [3], [4]. It uses three Winograd modules pipelined together with block lengths of 16, 15, and 17, respectively. The order is based on the numerical accuracy of the modules, the 16-point having the highest precision, the 15-point less, and the 17-point having the least of the three. Other PFA sizes of 240, 255, and 272 use two Winograd modules. The 4080-point PFA architecture is shown in Fig. 1. There are separate chips for each of the Winograd modules, the four dual-port memories, the PFA controller, and each of the clock generators. In each of the three stages, three identical WFT processors are used. One processor operates in active mode and the other two operate in watchdog mode comparing their results to those of the active. In the most general case, the input and output host need not be the same.

Handshaking between the various parts of the PFA processor is simple and asynchronous. Three signals are used between the PFA controller and the host and between the PFA controller and the Winograd modules. The host and the WFT processors communicate with the PFA controller only at the beginning and end of each 4080-point DFT. The handshaking signals consist of an *OPERATE* signal, a *LOAD* signal, and a *DONE* signal.

Simple handshaking is possible because of the straightforward flow of data through the processor. Data are initially written to the first memory by the input host. The WFT16 reads the data and computes the transform, writing the result into the next memory. The data continues through the remainder of the pipe in the same manner until it is read from the last memory module by the output host. Unlike other approaches, this approach allows word-parallel data transfers instead of bit-serial transfers.

Each memory module contains data for the real and imaginary parts of the transform. The memory is divided in half so that while one side is being written into the other side is being read from. For the 4080-point DFT, the memory must hold 4080 32-bit words for both the real and imaginary parts on each side. In each 32-bit word, 24 bits are used for the data and 8 bits are used for error correction and detection. Each memory bank must therefore store 512K bits. Three control signals are used to operate the memory modules. A left/right signal from the PFA controller determines which side of the memory is being written to or read from and a READ/WRITE enable is used to initiate a memory transfer for each side.

For the prototype PFA processor, a dynamic RAM cell was chosen over a static cell because of size constraints. The trend in memories indicates that a static RAM with enough storage and speed will be commercially available in

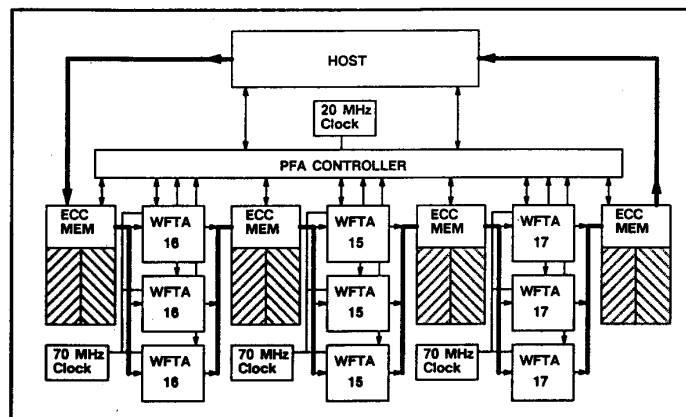


Fig. 1. 4080-point PFA architecture.

the near future. For our purposes, a DRAM was designed to hold enough data for a 272-point DFT. The chip was designed without refresh because of the speed at which the memory is used. For the WFT's operating at 70 MHz, the memory cycles at 35 MHz causing cells to be rewritten every 7.8  $\mu$ s.

The PFA controller chip is an application-specific processor supervising the flow of data through the chip, communicating with the host, and monitoring the status of faults in the pipeline. To begin a DFT, the PFA controller receives information from the host while *OPERATE* is low. At this time, the host may examine or change any register in the PFA controller. Information transferred before a DFT includes time-out information, configuration information, and error summary information. After the host raises *OPERATE*, the PFA controller flips the memory banks and may reconfigure. Reconfiguration is necessary if an active WFT has just made an error or if the host requests a different configuration. Finally, the controller drives the scale factors to each of the three stages and raises the WFT *OPERATE* signal. It then waits for the three WFT stages to finish and the input host to write the next set of data points into the first memory. After the four *DONE* signals are received, the controller latches the scale factors from the three stages and drops *OPERATE*. It then accumulates the scale factors for the DFT and passes the scale factors to the next stage in the pipeline. At this time, the PFA looks at the watchdog signals and uses a voting circuit to determine in which processor, if any, the fault occurred. If the active processor in any stage was at fault, the pipeline is reconfigured so that one of the watchdogs becomes the active processor. If one of the watchdogs was at fault, an error register for that processor is incremented. The controller then sends its *DONE* signal to the host.

#### IV. 16-POINT WINOGRAD PROCESSOR

The WFT processors use word parallel I/O to simplify and slow down the interface to off-chip RAM's. However, the arithmetic engine makes use of bit-serial adder/sub-




Fig. 2. SNR versus DFT size.

tractor modules and pipelined bit-serial multipliers modified from those suggested by Lyon [9] to incorporate hardwired coefficients. The elements of the  $D$  matrix are known in advance which helps meet the vertical pitch necessary to stack 70 multipliers on a chip. We chose bit-serial hardware because it provides a higher throughput per area per second, it is pipelined, and it requires simpler control. The throughput advantage depends upon high clock rates. We have tested 3- $\mu$ m CMOS macrocells to speeds in excess of 60 MHz and Hatamian and Cash [10] have reported achieving 70 MHz in 2.5- $\mu$ m CMOS. The multipliers and adders themselves could clock much faster but problems are encountered with clock distribution.

The 24-bit inputs are extended to 32 bits allowing sign extensions to prevent overflow and increase numerical accuracy. By using an adaptive rescaling algorithm, high accuracy is preserved in the results. The precision of various length transforms that can be performed is shown in Fig. 2. The simulations were run using random inputs averaged over 10 to 100 sets, properly scaled. The 4080-point transform achieves a signal-to-noise ratio of 106 dB which is the equivalent of 17.6 bits of precision.

The architectures required to compute a 15-, 16-, or 17-point DFT are essentially the same. The few differences in the 15- and 17-point chips relate to the internal data representation (30 bits for the 15-point and 34 bits for the

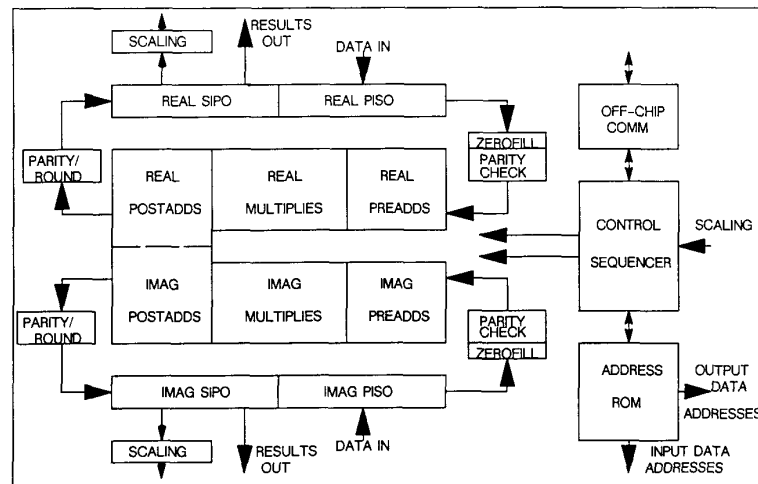


Fig. 3. Generic Winograd processor.

17-point versus 32 bits for the 16-point) and the number of sign extensions required to prevent arithmetic overflow. The internal data representation changes the time to compute the base DFT's, but the longer DFT's are computed in the same time by each processing chip since each accepts data inputs at a 35-MHz rate.

The data flow for the 16-point chip is shown in Fig. 3. As seen in Fig. 3, the circuits of the 16-point processor may be grouped into three categories: 1) input/output (I/O), 2) control, and 3) arithmetic.

The parallel-in/serial-out (PISO) registers take inputs from the data memory. The serial-in/parallel-out (SIPO) registers send results to the data memory. There are two sets of each type of register, one set for real data/results and one set for imaginary data/results. There are two types of control circuits, one for on-chip control of timing and address generation, another for off-chip communication (scaling control, handshaking, etc.). The arithmetic circuits reflect the three components of the Winograd modules: a preaddition matrix, a post-addition matrix, and a multiplication matrix. Adder/subtractors and multipliers compose most of the arithmetic circuitry. There are also several reset circuits for clearing intermediate results before a new 16-point DFT is started. The signal flow through the 16-point processor is described below.

1) The PFA controller sends an *OPERATE* signal to the 16-point processor to start the DFT computation.

2) Sixteen 24-bit data words are shifted into the real and imaginary PISO registers from the input memory.

3) After all 16 data words are loaded, the PISO registers latch the words into the serial output portion of each register.

4a) The input data words are serially shifted into the arithmetic circuitry, after passing through a parity check cell which flags any parity errors on the input data.

4b) While the first set of data words are being shifted out, the next set of data words are being loaded, in parallel, into the PISO register from the data memory.

5) The data pass through the preadders, multipliers, and post-adders to compute the 16-point DFT.

6) Prior to arriving at the SIPO registers, the real and imaginary results are sent through a parity/rounding cell which rounds the 32-bit arithmetic results to 23-bit output results and computes a parity bit to be appended to the 23-bit result.

7) After all 23 data bits and the single parity bit have been shifted into the SIPO registers, the result words are latched into the parallel portion of the SIPO.

8) The 24-bit results are shifted out of the SIPO register, through the scaling cell (which checks the most significant seven bits of each output data word to find the smallest number of sign extensions for all 4080 words), and to the output memory.

Once the pipeline fills, it is kept 100 percent busy until the end of the assigned transform length (i.e., 4080, 240, 272, etc.). In addition, the input and output data buses are kept 100 percent busy performing 35-MHz I/O. There is only a short delay between transforms while the PFA processor performs its functions which include flipping the memory banks.

The 16-point processor continues to compute DFT's until it has exhausted all the addresses for that particular DFT block length (i.e., 255 16-point DFT's will be computed in the 4080-point PFA system). On-chip control detects the completion of the DFT, shuts down the computation, and raises a *DONE* line to interrupt the PFA controller.

The 16-point WFT processor generates the addresses required to communicate with the memories in the PFA pipeline using an on-chip ROM which stores the required sequence of addresses. The ROM was automatically generated with a silicon compiler [11]. Since the output addresses are just the input addresses delayed by the latency of the pipeline, we use one ROM to generate addresses for both external RAM banks. Addresses are fetched four at a time and the circuitry alternately fetches input addresses

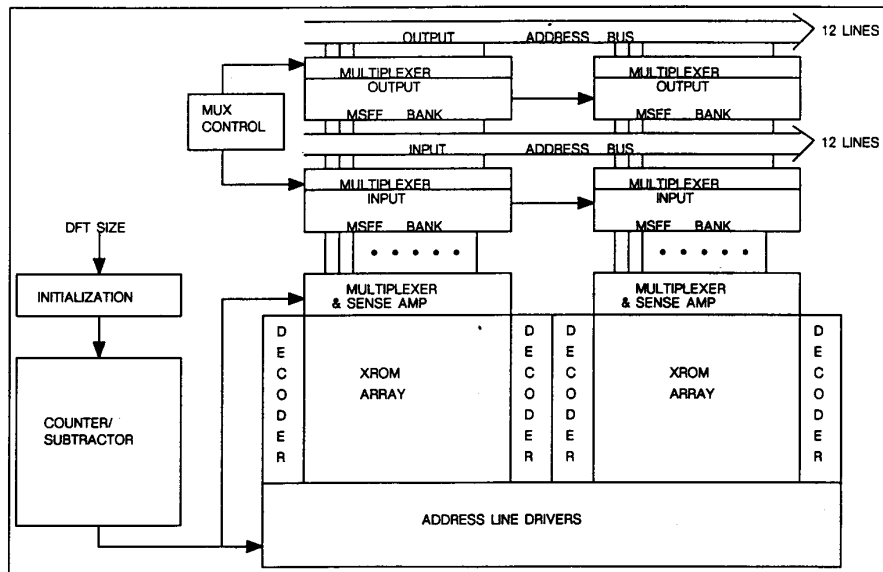


Fig. 4. Address-generation hardware.

and output addresses. The addresses are buffered in flip-flops and then multiplexed out onto the address buses. A block diagram of the address generation circuitry is given in Fig. 4.

To keep the control interface simple and asynchronous, an on-chip control sequencer guides the arithmetic and I/O circuitry through the execution of the WFT. This allows the simple and asynchronous *OPERATE/DONE* handshaking with the PFA controller. The control sequencer is a high-speed PLA driven by a ring counter to increase speed. To generate and then distribute the 70-MHz control signals, all outputs of the PLA are buffered through two master-slave flip-flop (MSFF) stages. These also allow set/scan paths to be woven through the control-sequencer outputs to improve testability.

Scaling control is the final major component. This circuitry detects the number of useless sign extensions on the largest output of the DFT. This information is passed onto the next WFT chip in the pipeline and finally to the host. The WFT chips use the input scale factors to provide fewer guard bits while still assuring no overflow may occur, thus improving numerical precision. This adaptive scaling mechanism can also detect and correct for poorly scaled inputs from the host.

## V. RELIABILITY AND TESTABILITY

Reliability allows the system to operate in the presence of failures. The PFA system uses several means at different levels to increase reliability. At the PFA controller level, time-out counters, error registers, and voting circuits are used. At the memory level, on-chip ECC is employed. The WFT level uses parity checking and watchdog monitoring. Finally, at the global level, the PFA system may be operated in an active or watchdog mode.

The PFA controller uses a time-out register to provide recovery when the WFT processors do not respond with a *DONE* signal. The PFA also monitors the faults reported in the pipeline. For most reported faults, the PFA updates registers that hold the error counts. However, if an active WFT processor was at fault the PFA controller must also reconfigure the pipeline.

The memory modules include single error correction and double error detection. Eight parity bits are added onto the incoming word and stored. When the word is read, the eight parity bits are used to generate eight syndrome bits which create an error vector which is then combined with the 24 data bits to form the final output word. [12]

The Winograd processors include internal parity checking on the input data. If a memory chip should suffer catastrophic failure, the on-chip ECC would be useless. However, the parity checking inside the Winograd processor would detect this fault.

The Winograd processors can operate in the active or watchdog mode. The PFA controller sends control signals to the WFT processors which configure the processors as either active or watchdog. One possible configuration is to group one active Winograd processor and two watchdog processors in each stage. All three processors in each stage compute the same DFT; however, the watchdog processors only monitor the output data lines of the active processor. An error signal is sent to the PFA controller if the watchdog output data differ from the active. The PFA uses a simple voting procedure to determine which processor was at fault.

For system reliability, the entire PFA system can be operated in the active or watchdog mode in the same manner as the WFT processors. At the output to the PFA system, the system WATCHDOG determines whether the PFA drives the bus or simply monitors the bus and reports

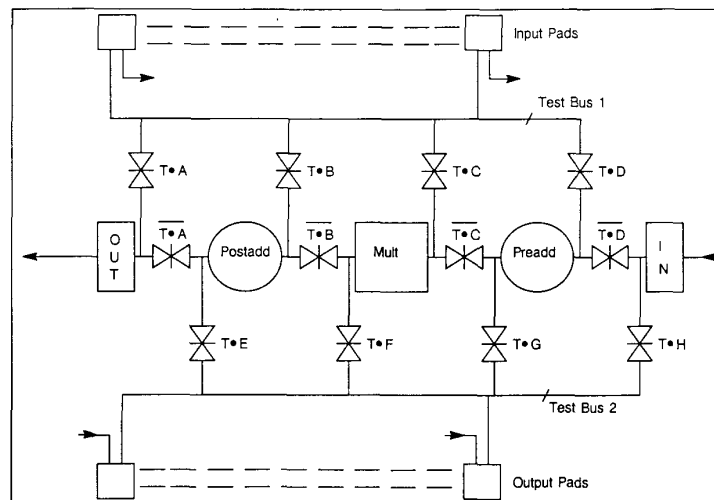


Fig. 5. Arithmetic circuitry with testing hardware.

any discrepancies.

The goals of testability are to gain as much observability and controllability as possible without degrading system performance and without using excessive VLSI chip area for the testing hardware [13]. One of the major design efforts was to incorporate the hardware needed to provide testability without affecting system performance. By using dual-input and dual-output pads, transmission gates, chained MSFF's, separate clocks for arithmetic and control hardware, and a PLA, observability and controllability were achieved for the WFT16.

The WFT16 requires a 144-pin pad frame where only one of the 144 pins is dedicated exclusively to testing. The remaining 143 pins are used for the input and output of data, for the output of addresses, and for the input and the output of control signals. By designing the input pads with two lines into the chip, "normal data and test data," and the output pads with two similar multiplexed inputs, 123 pads are made available to support testing. The test data signals do not affect the normal data signals when using these pads. Fig. 5 shows the dual-purpose input and output pads used in the arithmetic section.

By using transmission gates in an architecture, we can disconnect different modules from one another. For example, in Fig. 5 transmission gates are used to disconnect the different sections of the arithmetic section. To test the preadd section, data are input through the test port of the input pad onto Test Bus 1, and routed through the  $T \cdot D$  transmission gate, the preadd,  $T \cdot G$  transmission gate, Test Bus 2, and the test port of the output pads. All other transmission gates are turned off, thus preventing contention on the test buses.

Another useful testability tool is the chained MSFF. The control sequencer and address generation sections use MSFF's extensively and to test these sections, the MSFF's are connected via set/scan paths to make them observable and controllable.

The WFT16 clocks are routed so that the arithmetic clock and control clocks can be separated from one another. This method allows the state of the arithmetic hardware to be preserved while a new test vector is shifted into the control section.

To control the testing circuitry, a PLA is used. Fifteen test modes are available. The system is set to TEST by the only data pin dedicated to testing. The testing hardware uses only 2 percent of the total chip area. Testability of the arithmetic, control, and address generation is achieved.

## VI. PERFORMANCE AND FEATURES

The WFT16 is a completely pipelined architecture where the 96 data pads, the 36 serial multipliers, and the 72 serial adders are busy nearly 100 percent of the time. The system delivers an I/O bandwidth of 420 Mbyte/s. Approximately 78-million multiplications and 157-million additions are performed every second. The chip has a throughput of one 16-point DFT every 457 ns, and supports sampling complex input data at 35 MHz.

On-chip address generation to the off-chip RAM is provided. All control is provided with the exception of the handshaking control signals from the PFA controller. Complete testability of each major subsection of the chip is provided.

The WFT16 chip was designed using Caesar, a VLSI layout editor written by John Ousterhout. To check for design errors, the cif file output by Caesar was read into MAGIC, another VLSI layout editor also written by John Ousterhout.

The WFT16 has been extensively simulated with SPICE, a circuit simulator, and ESIM, an event driven switch-level simulator. ESIM was run after circuit extracting on the cif file output from Caesar. ESIM, however, has difficulty in simulating certain nodes where "fighting" takes place be-




Fig. 6. WFT16 chip layout.

cause it has limited understanding of the size ratio between transistors and thus would set node value to unknown where fighting takes place. To solve this problem, mostly found in MSFF's, the feedback loops in the MSFF cells were removed prior to simulation.

The WFT 16-point processor is being fabricated in a 1.2- $\mu\text{m}$  n-well technology through MOSIS. A floorplan is shown in Fig. 6.

## VII. CONCLUSION

The need for high-speed digital signal processing is readily apparent with the current applications of synthetic aperture radar and image recognition among the most noticeable. The DFT is a powerful tool for measuring the spectral content of sampled signals. Limited resources often force system designers to choose a VLSI circuit implementation for DFT computation. The 4080 PFA processor has been designed for this purpose and the WFT16 processor computes the small Winograd DFT's using a 70-MHz clock. On-chip control and address generation along with testability are important features of this custom VLSI processor.

## REFERENCES

- [1] C. S. Burrus and P. W. Eschenbacher, "An in-place, in-order prime factor FFT algorithm," *IEEE Trans. Acoust., Speech, and Signal Processing*, vol. ASSP-29, no. 4, pp. 806-817, Aug. 1981.
- [2] S. Winograd, "On computing the discrete Fourier transform," *Math Comput.*, vol. 32, pp. 175-199, Jan. 1978.
- [3] R. M. Owens and J. Ja'Ja', "A VLSI chip for the winograd/prime factor algorithm to compute the discrete Fourier transform," *IEEE Trans. Acoust., Speech, and Signal Processing*, vol. 34, no. 4, pp. 979-989, Aug. 1986.


- [4] J. P. Costello, "Digital beamforming radar VLSI processor," *SPIE Real Time Signal Processing VIII*, vol. 564, pp. 66-72, 1985.
- [5] R. E. Blahut, *Fast Algorithms for Digital Signal Processing*. Reading, MA: Addison-Wesley, 1985.
- [6] H. J. Nussbaumer, *Fast Fourier Transform and Convolution Algorithms*. New York: Springer-Verlag, 1981.
- [7] P. P. Reusens, "High performance VLSI digital signal processing architectures and chip design," Ph.D. dissertation, Cornell Univ., Ithaca, NY, Aug. 1983.
- [8] R. W. Linderman, "High performance VLSI technologies, integrated circuits, and architectures for digital signal processing," Ph.D. dissertation, Cornell Univ., Ithaca, NY, Aug. 1984.
- [9] R. F. Lyon, "Two's complement pipeline multipliers," *IEEE Trans. Commun.*, vol. COM-24, pp. 418-425, Apr. 1976.
- [10] M. Hatamian and G. L. Cash, "A 70 MHz 8-bit  $\times$  8-bit parallel pipelined multiplier in 2.5- $\mu\text{m}$  CMOS," *IEEE J. Solid-State Circuits*, vol. SC-21, no. 4, pp. 505-513, Aug. 1986.
- [11] P. C. Rossbach, R. W. Linderman, and D. M. Gallagher, "An optimizing XROM silicon compiler," in *Proc. IEEE Custom Integrated Circuits Conf.*, 1987, pp. 13-16.
- [12] S. Lin and D. Costello, *Error Control Codes: Fundamentals and Applications*. Englewood Cliffs, NJ: Prentice Hall, 1983.
- [13] H. Fujiwara, *Logic Testing and Design for Testability*. Cambridge, MA: M.I.T. Press, 1985.



**Richard W. Linderman** (S'81-M'84), Captain, U.S. Air Force, received the B.S.E.E., M.S.E.E., and Ph.D. degrees from Cornell University, Ithaca, NY, in 1980, 1981 and 1984, respectively.

He is currently assigned as an Assistant Professor of Electrical and Computer Engineering at the Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, where he teaches graduate-level courses in computer architecture and VLSI. His research interests include VLSI architecture and design, computer systems architecture, digital signal processing, and submicrometer VLSI device fabrication.

Dr. Linderman is a member of Tau Beta Pi and Eta Kappa Nu.



**Carl G. Shephard**, Captain, U.S. Air Force, received the B.S.E.E. degree from University of Texas, Austin, in 1980 and the M.S.E.E. degree from the Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, in 1985.

He is currently assigned to Air Force Logistics Command, McClellan Air Force Base, CA, as a Systems Engineer.



**Kent Taylor**, Captain, U.S. Air Force, received the B.S.E.E. degree from the University of Arizona, Tucson, in 1979, and the M.S.E.E. degree from the Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, in 1985.

He is currently stationed at the Engineering Installation Division, Tinker Air Force Base, OK, as a Communications Engineer.

Captain Taylor is registered as a Professional Engineer in the state of Oklahoma.



**Paul W. Coutee**, Captain, U.S. Air Force, received the B.S.E.E. degree from the Fenn College of Engineering, Cleveland State University, Cleveland, OH, in June 1977, and the M.S.E.E. degree from the Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, in 1985.

He is currently assigned to the Air Force Systems Command, Aeronautical Systems Division, Wright-Patterson Air Force Base, OH, Deputy for Engineering as a Computer Systems Engineer.




**Paul C. Rossbach** (M'86), Captain, U.S. Army, received the B.S. degree from the U.S. Military Academy, West Point, NY, in 1980 and the M.S.E.E. degree from the Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, in 1985.

He is currently assigned to the U.S. Army Signal School, Ft. Gordon, GA, where he teaches courses on the theory and application of artificial intelligence. His teaching and research interests include VLSI design, knowledge engineering,

and expert systems.


Captain Rossbach is a member of Phi Kappa Phi, Tau Beta Pi, and Eta Kappa Nu.



**James M. Collins**, Captain, U.S. Air Force, received the B.S.E.E. degree from Norwich University, Northfield, VT, in May 1981 and the M.S.E.E. degree in December 1985 from the Air Force Institute of Technology.

He is currently with the Engineering Directorate at the Aeronautical Systems Division, Wright-Patterson Air Force Base, OH.

Captain Collins is a member of Eta Kappa Nu.



**Robert S. Hauser** (S'82-M'87), Second Lieutenant, U.S. Air Force, received the B.S. degree in computer engineering from the University of Illinois, Urbana, in 1986, and the M.S. degree in computer engineering from the Air Force Institute of Technology, Wright-Patterson Air Force Base, OH, in 1987.

He is now at the Arnold Engineering Development Center, Tennessee.