High-speed gallium arsenide transversal filter circuit for clutter rejection systems

Spencer W. White, Rafael J. Betancourt-Zamora

Hughes Aircraft Company - Electro-Optical and Data Systems Group
El Segundo, California 90245

ABSTRACT

A high-speed GaAs Transversal Filter IC that has been designed at Hughes Aircraft Company and fabricated at AT&T Bell Laboratories is described. This circuit, which contains in excess of 41,000 transistors, implements a finite impulse response (FIR) filter for infrared data stream signal-to-noise ratio (SNR) enhancement in clutter rejection systems. The die is 9.1 mm on a side and dissipates 5.1 Watts. In addition to a multiply/accumulate processor, the circuit contains an address controller for interfacing with external memory and a register file for storing programmable filter coefficients (tap weights). Very-high-speed operation, as evidenced by a predicted clock rate of 160 MHz, is achieved through the use of a pipelined architecture and an advanced GaAs processing technology. The high-performance characteristics of this device allow it to be applied as a single serial processor to systems normally requiring parallel processing techniques. As such, systems not directly requiring bandwidths achievable with gallium arsenide circuits can utilize this device to reduce power consumption and parts count while at the same time increasing reliability and radiation hardness. The processing technology utilized employs molecular beam epitaxy (MBE) to grow starting material for the fabrication of enhancement and depletion mode SARGIC/HFET (self aligned refractory gate integrated circuit / heterojunction field effect transistor) devices. Minimum gate lengths of 1.0 μm, and two layers of Ti/Pt/Au interconnect metallization are featured. An application specific integrated circuit (ASIC) design methodology was followed utilizing a macrocell library and a completely-integrated computer aided design system known as HCAD.

1. INTRODUCTION

Transversal (non-recursive) filters have long been used for SNR enhancement in digital signal processing (DSP) applications. Space-borne infrared sensor systems in particular, produce tremendous amounts of data which must be processed to eliminate the clutter inherent in the sensed data. Filtering techniques are necessary to improve the SNR to the point where the extraction of information for the identification and tracking of targets is achievable.

Increasing levels of device integration have made possible cascadable single-chip digital filter implementations capable of running in the 20 - 50 MHz range (depending upon resolution) when fabricated in custom CMOS technology[1,2]. Dramatic advances in circuit and processing technologies have not been limited to silicon processes, however. Digital GaAs integration levels are now less than an order of magnitude behind the most advanced CMOS technologies as evidenced by the latest commercial semi-custom offerings[3,4]. Although they have lagged the integration levels offered by MESFET technologies, IC's based upon heterostructures have progressed from MSI to LSI levels of complexity[5-7]. As a result, very-high-performance single-chip DSP circuit implementations can be realized using semi-custom design approaches; this enables such a device to easily and quickly be tailored to a specific system application.

Space-borne and most military systems incorporating signal processing have very stringent power, size, and radiation hardness requirements. Existing DSP systems often exploit parallel processing architectures to achieve the high data rates required. Unfortunately, parallel processing techniques often involve a substantial penalty in terms of added hardware. This paper will describe an alternative to parallel processing in which heterostructure GaAs integrated circuit technology is exploited to implement a serial architecture. Our FIR filter implementation calculates a single tap at a time, yet still provides the data rates required without the size and power penalties often encountered with the use of parallel processors. In addition, the inherent radiation hardness exhibited by gallium arsenide devices makes this an ideal solution for military and space-borne systems.

2. ARCHITECTURE

A transversal filter performs clutter rejection on an incoming data stream through a discrete implementation of the convolution integral. As such, the basic algorithm utilized is the following:

$$Y(k,p) = \sum_{i=0}^{n-1} h_i * X(k-i,p)$$

where Y represents the filter's output, X represents the filter's data input, k represents the sample time or frame index in units of sample time, h represents the filter's impulse response (tap weight function), n represents the number of filter taps, and p represents the pixel index when multi-pixel frame sizes are used. It can be seen that the filter's output is a simple weighting of present and previous values of the input. Tap weight 0 is multiplied by the current frame's data, tap weight 1 by the previous frame's data, etc. The key operations required of such a filter are multiplication and accumulation.

The standard parallel processing approach to a transversal filter design involves cascading a number of filter processing elements (PEs) which are often implemented in CMOS. Each PE is used as a single tap of a multi-tap filter and contains its own multiplier and adder. Such a filter configuration is shown in Figure 1(a). A data stream flows from PE to PE with an appropriate delay time added at each stage. Each PE can be thought of as a single pipeline stage in the overall filter. In the case in which a multi-pixel frame of size $a \times b$ is scanned to provide serial input data to the filter, the delay time is simply equal to $(a \times b) \times T$, where T represents the serial data rate. It can be seen that for large frame sizes, this technique requires a large FIFO memory device associated with each PE.
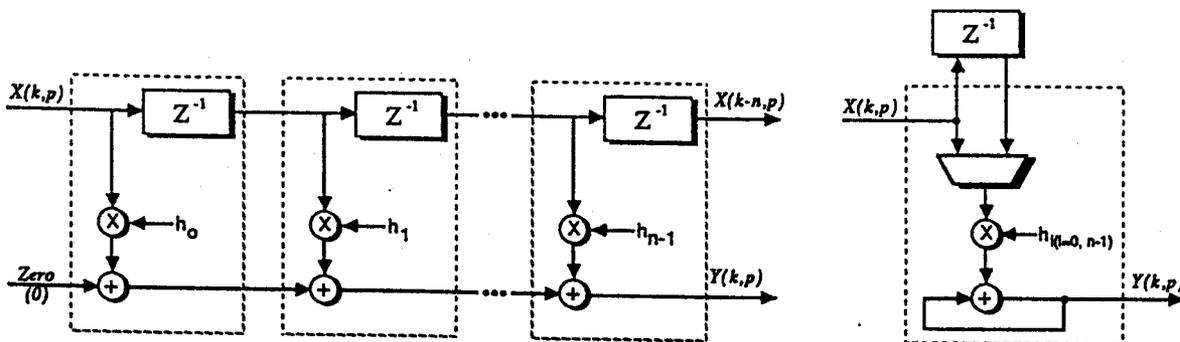

Figure 1. (a) Cascaded Multi-PE Filter Implementation

(b) Single GaAs PE Filter Implementation

The parallel approach described above has often been used out of necessity to ensure that a filter can meet its required data rate. With the advent of significantly faster integrated circuit technologies, simpler architectures can be utilized to achieve the same throughput. By reverting to a serial filter implementation utilizing high-speed GaAs technology, a single PE chip can perform the processing of several CMOS PE chips. Figure 1(b) shows a diagram of a filter using a single GaAs PE. In this GaAs PE, the multiplier generates products for all tap weights rather than a single one, the adder is a true accumulator, and the data delay element has grown to encompass the combined storage capacity of all PE memories. The size of this combined memory requires it to be implemented off chip.

The circuit's architecture was designed to maximize performance while keeping complexity down to a manufacturable level. The final design contains 5.2K gates (41,546 transistors) with simulations predicting a clock rate of 160 MHz. A block diagram of the chosen architecture is shown in Figure 2. The 16-tap filter is implemented using a pipelined multiply/accumulate processor that performs all arithmetic operations. Additional support circuitry includes a register file for filter coefficients (tap weights) and an external memory interface.

Filter coefficients are programmed into the device prior to data processing operations by writing them into the register file. This register file can accomodate up to 16 coefficients each of which is an 8-bit 2's complement value. To minimize the number of I/O signals required, the pixel data bus is used during the programming sequence to hold both address (tap number) and filter coefficient data. A separate controller is used to properly access the filter coefficients for arithmetic processing.
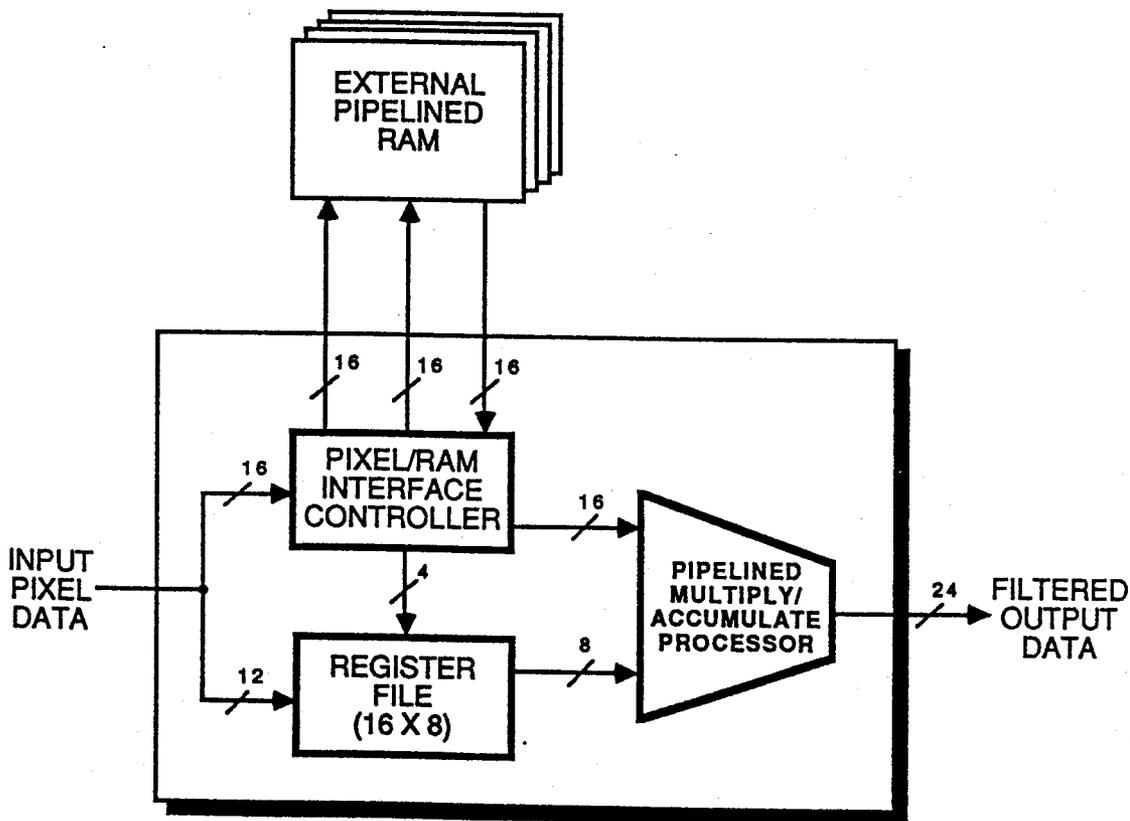

Figure 2. GaAs Transversal Filter Architecture

The GaAs filter PE circuit is designed to interface with systems which provide data that is serially scanned a single pixel at a time for each frame. Frame synch and data valid input signals are provided to enable proper synchronization. Each data word consists of a 16-bit 2's complement value. Because up to 16 multiply/accumulate operations are required for each pixel, the circuit's clock rate must be 16 times the input data rate. A single clock signal and reset signal are provided to control the sequential circuitry.

Fifteen frames of pixel data must be stored in external memory to enable the calculation of the 16-tap filter's output. This external memory is assumed to be fully-pipelined operating at the same clock rate as the transversal filter device. Specifically, pipeline registers for input data, address, write enable, and output data are assumed to exist on the memory device. It has been shown that incorporating pipeline registers directly on the memory components maximizes the performance of GaAs signal processing systems operating at very high data rates[8]. Such a design minimizes the inter-chip speed penalties that can otherwise prove to be a system's performance bottleneck.

The circuit is designed to accomodate up to 4096 pixels per frame. As such, both address and data buses between the circuit and external memory are 16 bits wide. The 4 least significant bits of the address bus represent the frame address. An address controller ensures that the current frame's pixel data overwrites the oldest frame's data and that the starting frame count is decremented whenever a new frame is encountered. The 12 most significant bits of the address bus represent the pixel address within a frame. This address value is incremented for each new data point and is reset for each new frame. Finally, it can be seen that the required external memory depth is equal to 16 times the number of pixels for a maximum of 64K words.

All arithmetic processing takes place in the multiply/accumulate unit. This unit is implemented as a three-stage pipeline consisting of an 8-bit × 16-bit 2's complement multiplier and a 24-bit 2's complement accumulator. The pipeline is partitioned to achieve maximum performance by closely matching the stage delays. The multiplier is implemented using the modified Booth's algorithm[9] which is used to reduce the number of partial products involved in the operation. The summation of partial products is performed using carry save adders (CSAs) in a Wallace tree configuration[10]. Carry lookahead adders (CLAs) are used to generate the final product. In addition to the 24-bit accumulator output bus, the device has two output status flags. One indicates the presence of an accumulator overflow while the other indicates that the output data is valid and represents a final accumulation value.

A number of test signals are included on the device to ensure that it is testable. All sequential devices on the circuit are linked into 3 serial scan chains controlled by a single scan enable input. This scan path design technique has proven to be helpful in the generation of test vectors for complex sequential circuits. Separate serial input and output signals are provided for these chains. Other test inputs can disable accumulation and initiate a test mode in which the pixel address counter's 4-bit sections count independently.

3. CIRCUIT DESIGN & METHODOLOGY

3.1. Overview

The design process followed a standard cell methodology using the HCAD toolset. HCAD was developed by Hughes to support the design of GaAs IC's for DARPA's Pilot Line III. It is based on Cadence Design Systems' Framework technology and incorporates commercial as well as Hughes designed tools. HCAD is a completely integrated design environment which supports schematic capture, simulation, physical design and verification with a single unified database and a common user interface. It also incorporates a Hughes developed design-for-test toolset known as HTEST. Full custom, standard cell and gate array layout methodologies are all supported.

Included with HCAD is a GaAs macrocell library consisting of 45 custom logic cells developed by AT&T Bell Laboratories. These macrocells are fixed-height standard cells which range in complexity from simple logic gates to high level functions. The use of macrocells results in an increase in speed as well as reductions in area and power dissipation while simplifying the design process by providing MSI-level building blocks. The library includes symbols, circuit schematics, layouts, and simulation models (SILOS, HSPICE, TA, HITS, N.2) for each cell. The logic family used is Source Follower Logic (SFL) which exhibits typical unloaded gate delays of 65 ps with only 0.674 mW power dissipation[11].

3.2. Design flow

Figure 3 depicts the steps followed during the design process. First, the macrocell-based design was captured using HCAD's schematic editor. Using the circuit's specification, a program was written in HCAD's Simulation and Test Language (STL) and executed to generate input vectors and their expected outputs. SILOS logic simulation proceeded using these input vectors, and the results were compared with those generated by the STL code to verify the correctness of the design. A static timing analysis was then run using HCAD's TA to identify critical paths.

After the front-end verification was completed, the physical design was initiated. First, the relative locations of the I/O cells were defined, and these cells were subsequently placed automatically. Then, the nets in the critical path (critical nets) and clock nets were assigned weights to optimize the placement. Two placement regions were defined for the automatic placement of all core logic cells. At this point, a Gap cell was added to each row to define the required center power bus. Also, Cap cells were added at both ends of the rows to facilitate tapping the core's power ring. Figure 4 depicts the circuit's floorplan.
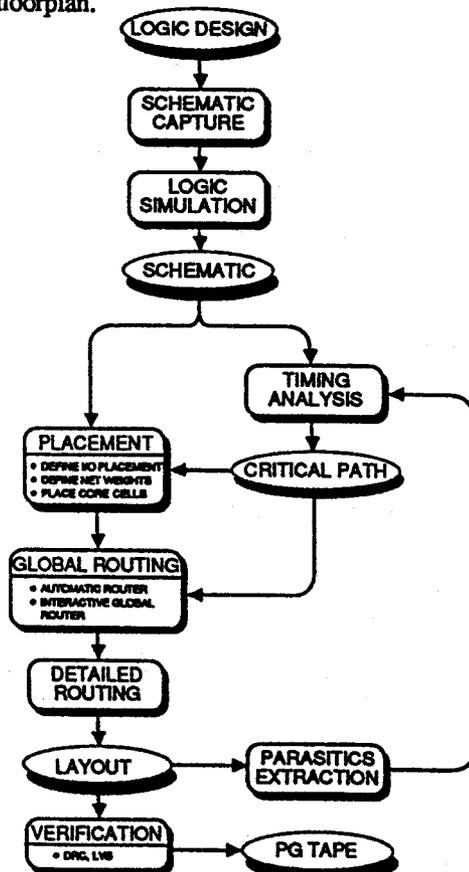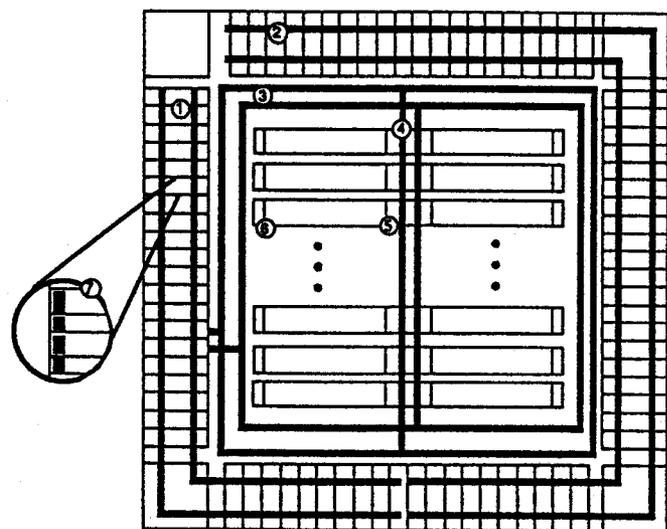

① Input Buffer's Power Bus   ⑤ Gap Cell

② Output Buffer's Power Bus   ⑥ Cap Cell

③ Core's Power Ring   ⑦ Variable-width
                                     Supply Pads

④ Center Power Bus

Figure 3. Design Methodology Flowchart   Figure 4. Circuit Floorplan

After all cells were placed, the automatic global router was invoked. The interactive global router (*iglr*) was used to route the clock nets and key critical nets. It provides a graphical display which allows manual routing and sizing of nets during the global routing phase. Also, the *iglr* was used to route all power nets. When the global routing was completed, the automatic detailed router was used to create the actual geometries and resize the routing channels where necessary.

With the physical design phase completed, parasitics extraction was run on the layout and back annotated into the schematics to reassess the circuit's performance. Timing analysis was rerun and the effect of parasitics on the critical path was revealed. The weighting of critical nets was subsequently changed and the placement and routing process was repeated. In some cases, the design was modified to add buffering required to drive long nets. The *iglr* also was used to further optimize the routing given the more accurate timing results. This iteration of the physical design cycle resulted in a 20% reduction in the critical path's propagation delay.

The last phase in the design process is physical verification. HCAD provides for three different levels of design rule checking (DRC): macrocell, hierarchical, and flat. All three were used and revealed no design rule violations. Verifications were completed when layout vs. schematics (LVS) checking showed a match between the physical design and the schematics.

3.3. Circuit design issues

In order to minimize the effects of switching noise, the power distribution scheme utilizes three independent power buses isolating inputs, outputs, and the core (logic). Due to the relatively low supply voltage (2.0 V) and noise margin characteristics of GaAs technology, it was important to minimize the resistive voltage drops across these power buses to prevent noise margin degradation. This factor, along with the need to control current densities to avoid electromigration, dictated the size of the buses and the I/O placement. As such, the macrocell power rails are 64 μm while the I/O cell buses are 200 μm. Power buses supplying the core logic consist of a 250 μm ring bisected by a 400 μm center power bus. The *iglr* was used extensively to produce this unique power distribution topology. Finally, no more than two output buffer pads were placed together between supply pads.

Wider supply pads were used to accomodate multiple probes where it was necessary to reduce probe current density. This was done to maximize probe card life while significantly reducing transient noise during testing. On some pads a sense probe can also be accomodated. Finally, the wider pads enable multiple wire bonds from a single pad to the package to minimize current density in the bond wires.

4. DEVICE FABRICATION

Circuits were fabricated (see Figure 5) in DARPA's Pilot Line III which features SARGIC/HFET technology developed for digital applications. This process utilizes molecular beam epitaxy for growing AlGaAs/GaAs heterostructures on semi-insulating GaAs wafers in order to achieve high electron mobility and carrier confinement. This MBE process yields depletion and enhancement mode HFET's - the latter being formed through the selective removal of GaAs cap and AlGaAs etch stop layers. Device isolation is achieved with an oxygen implant which has been shown to limit inter-device leakage currents to 5 nA or less at 2 V. The refractory gate material is tungsten silicide which is patterned with a reactive ion etch to form typical gate lengths of 1.0 μm. Source and drain regions are formed using self-aligned ion implantation followed by a capless furnace anneal for N$^+$ activation. Ohmic contacts to these regions employ Au/Ge based metallization. Finally, two levels of Ti/Pt/Au metallization, patterned with a lift-off process, are used to interconnect the devices. Both bottom and top level metal runners have minimum design rules of 2.0 μm-wide lines and spacings, and use SiON as the underlying dielectric. Interconnect sheet resistances have been characterized at less than 0.05 Ω/square.
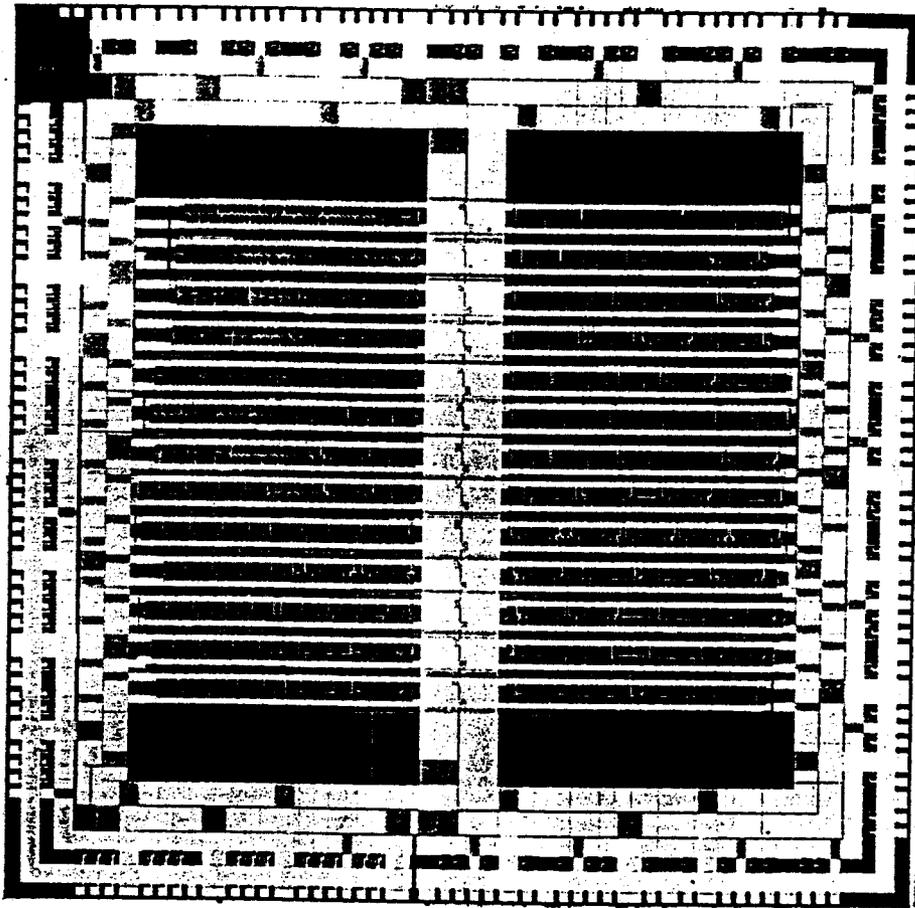
Figure 5. Fabricated GaAs Transversal Filter

5. TEST RESULTS

Wafer-probe testing was performed on an Advantest T-3340 test system using a Microprobe 50 Ω 222-pin probe card at speeds up to 40 MHz. Each fabricated wafer is 3 inches in diameter containing 31 circuit sites along with process control monitors (PCMs). PCM measurements indicate that EFET devices have a typical threshold voltage of 0.29 V with a transconductance of 225 mS/mm while DFET devices have a typical threshold voltage of -0.41 V and a transconductance of 170 mS/mm.

The filter circuit was found to be functional on the first lot of GaAs wafers. At V_{DD} = 2.0 V, the average dynamic current (measured with and without 50 Ω termination at the output) was 2.54 A and 1.64 A respectively. This corresponds to an average power dissipation of 0.66 mW per logic gate. The measured I/O noise margins averaged NM_L = 0.34 V and NM_H = 0.29 V. Device performance has been characterized using a 17-stage *loaded* ring oscillator. Each stage is a buffered inverter loaded with a fanout of 3 and approximately 1 mm of routing. At V_{DD} = 2.0 V, the average delay is 237 ps/gate. The layout for this ring oscillator was implemented in the same manner as the remainder of the circuit using automatic placement and routing of standard cells. This ensures the legitimacy of its use as a vehicle for predicting the speed of the functional circuitry. Initial tests show measured delays are within 20% of simulated values.

Technology	1.0 micron HFET GaAs
No. of Transistors	41,546
No. of I/O Pins	105
Chip Area	$9.1 \times 9.1 \text{ mm}^2$
Predicted Clock Rate	160 MHz
Power Dissipation	5.1 W @ V_{DD} = 2.0 V
Word Length	coefficients = 8 bit input data = 16 bit output data = 24 bit
No. of Filter Taps	16 (max.)

Table 1. Circuit Characteristics

Functional parts will be packaged in a Triquint 196-pin (128 signals) leaded chip carrier. This is a multilayer ceramic package specifically developed for very high performance ICs and includes 50 Ω controlled impedance signal lines as well as internal power and ground planes. High-speed testing (up to 200 MHz) of packaged devices will be performed using an HP 82000 test system.

6. CONCLUSION

In summary, we have developed a high-performance GaAs transversal filter IC. Important circuit characteristics and capabilities appear in Table 1. The high speed achieved with HFET GaAs technology enables the implementation of filter modules which use a serial architecture. The benefits of this architecture are clear when comparing it to an existing parallel architecture which implements a 16-tap, 256-pixel filter module based upon CMOS components. In this case, the serial GaAs filter implementation dissipates 40% less power and requires 69% fewer parts (1 PE chip and 4 SRAM devices in the GaAs implementation vs. 16 PE chips in the silicon implementation). This clearly results in a significant decrease in required board area along with an increase in reliability. Finally, it is expected that the GaAs filter module will exceed the radiation hardness of the silicon filter by two to three orders of magnitude.

7. ACKNOWLEDGEMENTS

8. REFERENCES

1. L. R. Tate, "Sparse Canonical Signed Digit FIR Filter," CICC Proceedings, pp. 215-218, 1987.

2. T. Baji et al., "A 20-ns CMOS Micro-DSP Core for Video Signal Processing," *IEEE Journal of Solid State Circuits*, vol. 23 No. 5, pp. 1203-1210, October 1988.

3. "Vitesse Doubles Gate-Count Limit of GaAs Array Family," *ECN*, p. 37, January 1990.

4. T. Wong, "Can ASICs Remain Technology Drivers?," *EE Times*, p. 51, October 2, 1989.

5. A. Lee, S. W. White, and S. M. Lau, "The Design, Fabrication, and Test of Four Standard Cell Circuits in HFET GaAs," GaAs MANTECH Proceedings, pp 105-109, 1988.

6. K. Kajii et al., "A 40-ps High Electron Mobility Transistor 4.1K Gate Array," *IEEE Journal of Solid State Circuits*, Vol. 23 No. 2, pp. 485-489, April 1988.

7. A. I. Akinwande et al., "A 500-Mhz 16 × 16 Complex Multiplier Using Self-Aligned Gate GaAs Heterostructure FET Technology," *IEEE Journal of Solid State Circuits*, Vol 24 No. 5, pp. 1295-1300, October 1989.

8. B. K. Gilbert et al., "Signal Processors Based Upon GaAs IC's: The Need for a Wholistic Design Approach," IEEE Computer, pp. 29-43, October 1986.

9. M. Annaratone et al., "The Design of an LSI Booth Multiplier: nMOS vs. CMOS Technology," *Report from the Department of Computer Science, Carnegie-Mellon University (CMU-CS-84-150)*, 1984.

10. C. S. Wallace, "A Suggestion for a Fast Multiplier," *IEEE Transactions on Electronic Computers*, pp. 14-17, February 1964.

11. A. I. Faris, P. J. Robertson et al., "500 MHz Gallium Arsenide Macrocell Library for High-Speed, Low-Power LSI Digital ASICs," SPIE Optical Engineering and Photonics in Aerospace Sensing Proceedings, 1990 (these proceedings).