

Complementary GaAs Technology for High-Speed VLSI Circuits

Richard B. Brown, Bruce Bernhardt*, Mike LaMacchia**, Jon Abrokwah***, Phiroze N. Parakh, Todd D. Basso, Spencer M. Gold, Sean Stetson, Claude R. Gauthier, David Foster**, Brian Crawforth**, Timothy McQuire**, Karem Sakallah, Ronald J. Lomax, and Trevor N. Mudge

Electrical Engineering and Computer Science Dept.
University of Michigan
1301 Beal Avenue
Ann Arbor, MI 48109-2122, USA

* Motorola
Communication Products Lab
2100 East Elliot Road
Tempe, AZ 85284

** Motorola
Space and Systems Technology Group
8201 E. McDowell Road
Scottsdale, AZ 85252-1417

*** Motorola
Phoenix Corporate Research Labs
2100 East Elliot Road
Tempe, AZ 85284-5934161

Abstract—A self-aligned complementary GaAs (CGaAsTM) technology developed at Motorola for low-power, portable, digital and mixed-mode circuits is being extended to address high-speed VLSI circuit applications. The process supports full complementary, unipolar (pseudo-DCFL), source-coupled, and dynamic (domino) logic families. Though this technology is not yet mature, it is years ahead of CMOS in terms of fast gate delays at low power supply voltages. Complementary circuits operating at 0.9V have demonstrated power-delay products of 0.01 μ W/MHz/gate. Propagation delays of unipolar circuits are as low as 25 ps. Logic families can be mixed on a chip to trade power for delay. CGaAs is being evaluated for VLSI applications through the design of a PowerPCTM-architecture microprocessor. This paper touches on the major aspects of the project, process technology, circuit design, packaging, architecture, CAD tools and software, with an emphasis on application of the CGaAs technology.

I. INTRODUCTION

Microprocessors have had a profound impact on both the scientific world and on our personal lives. Through astonishing advances in performance, they have replaced traditional mainframes and supercomputers with microprocessor-based workstations and servers [1]. The remarkable decrease in cost vs. performance for microprocessors has made computing ubiquitous in our society.

Processor trends can be seen by surveying the microprocessors presented at ISSCC over the past ten years: gate delays have improved at 12% per year; clock frequencies have increased at 40% per year; and transistor-counts have grown at 40% per year [2]. The performance of systems made from these processors, as measured by the integer SPEC benchmarks, has been improved at a compounded rate of 59% per year, resulting in an increase in computing power of more than 100 fold over the decade.

The disparity in improvement between gate delay and clock frequency is accounted for by noting that some of the additional transistors made available during these years have been used to pipeline processors, reducing the number of gates between latches. However, with present-day processors having on the order of ten pipeline stages, additional pipeline

depth provides diminishing performance returns, and cannot be expected to maintain the steep increase in clock frequency seen in the past decade.

The growing transistor budget also supported the addition of on-chip cache memory, which reduced load/store latency to memory. But again, with large first- and second-level instruction and data caches on chip, the performance return for enlarging cache memories beyond their present sizes (122KB total in the DEC Alpha 21164) is modest.

To keep new processors on the performance curve, architects also invested their additional transistors in multiple functional units for concurrent instruction execution (superscalar architectures). Unfortunately, the benefits of parallelism also diminish with scale in general purpose machines. For example, the designers of the Intel Pentium Pro reduced the initial instruction-issue width of four to three in order to reduce chip size and complexity, and enable an increase in the initial clock frequency from 100 to 150 MHz [3]. A general industry trend toward simpler CPUs running with higher clock frequencies was noted at the 1995 Hot Chips conference [4].

In the past decade, much of the performance gain in high-end microprocessors has derived from deeper pipelining, shorter cache latencies, and multiple instruction issue; but staying on the performance curve will require other solutions, as existing processors have already taken these techniques to their points of diminishing return. Future computers will be more dependent on high clock speeds, resulting from technology advances, for their performance improvements. Scaling of CMOS to channel lengths of 0.35 μ m, together with planarization techniques which can produce five levels of fine-pitch metal interconnect have pushed the clock speed of commercial CMOS processors to as high as 500 MHz [5]. A number of semiconductor manufacturers are investing heavily in fabrication technology in order to develop 0.25- μ m CMOS processes in 1997. The importance of semiconductor technology to future high-end computer performance warrants evaluation of processes such as complementary GaAs, which has the switching speed of HEMTs with many of the circuit advantages of CMOS.




Fig. 1: CGaAs process cross-section.

Another justification for exploring digital CGaAs is the growing portable (battery-powered) computer market, which could benefit from complementary GaAs speed at low power supply voltages. The state-of-the art in CMOS low-power processors is exemplified by the IBM 401GF, which consumes just 40 mW (typical) from a 2.5-V supply at 25 MHz (100-MHz parts are also available) [6]. CGaAs, operating with a 1-V supply, would have a dynamic-power advantage of 1:6.25 over this 2.5-V CMOS process, assuming parasitic capacitances were equivalent. This paper describes the characteristics of CGaAs based on test circuits and simulations, and projects its application to a multichip microprocessor.

II. COMPLEMENTARY GAAS TECHNOLOGY DESCRIPTION

CGaAs, a complementary heterostructure-insulated-gate FET technology, has been described elsewhere [7-9]. A sketch of the device structure is shown in Fig. 1. CGaAs integrates an enhancement-mode P-channel HFET with a high performance N-channel HFET. While holes in III/V materials do not enjoy an intrinsic mobility advantage over those in silicon, as do electrons, the pseudomorphic P-channel HFETs in this process have three to five times higher transconductance at given gate dimensions than their silicon counterparts. The wafers are more expensive than silicon, but this process




Fig. 2: SEM cross-section of metal interconnect.

Table 1: CGaAs Device Parameters.

Parameter	NFET ($0.7 \times 10 \mu\text{m}$)	PFET ($0.7 \times 10 \mu\text{m}$)
V_{th} (V)	+0.55	-0.55
I_{dss} (mA)	1.8	0.5
g_m (mS/mm)	280	60
Beta ($\text{mA}/\text{V}^2\text{-mm}$)	270	50
Subth slope (mV/dec)	75	90
Subth Current (nA) at $V_{gs}=0\text{V}$	< 1	< 10

requires only 13 masks through three levels of interconnect, thus compensating in process efficiency for the more costly starting material.

The process currently has three levels of Al-Cu interconnect, a cross-section of which is seen in Fig. 2, and a fourth layer is being added. Standard gate lengths are now $0.7 \mu\text{m}$, but experimental circuits are using 0.5 and $0.35 \mu\text{m}$ gate lengths. A design having 171K transistors has been demonstrated in a high speed signal processor, and a 32Kb SRAM has been fabricated. The process is compatible with RF and other digital processes running through the same fabrication line. Recent improvements to the epi structure make the process resistant to single-event upset as well as tolerant of a large total dose radiation. Typical parameters for $0.7\text{-}\mu\text{m}$ channel-length devices having $\pm 0.55\text{V}$ thresholds, and measured with $V_{dd}=1.5\text{V}$, are given in Table 1. As seen in the table, both N and P channel devices have good output conductances and pinch-off characteristics. The original device thresholds of $\pm 0.55\text{V}$ were selected because they yielded the optimum power-delay product in complementary circuits; unfortunately, they produced a drain-current ratio between N and P devices of about 4:1. Experimental circuits are now being fabricated with $\pm 0.35\text{V}$ thresholds for higher performance.

The low threshold voltages and high transconductance of CGaAs result in good performance at low voltages. Fig. 3 shows unloaded ring oscillator delays versus supply voltage for several logic families (thresholds of $\pm 0.55\text{V}$). The delay of $1.0\text{-}\mu\text{m}$ CGaAs is less than that of $0.5\text{-}\mu\text{m}$ CMOS or Thin-Film Silicon-on-Insulator (TFSOI), and the $0.5\text{-}\mu\text{m}$ CGaAs shows delays below 100ps with a 1.2V power supply. Power dissipation is not indicated in the figure. Lower threshold voltages will make the CGaAs circuits faster yet.

Two key parameters of concern in complementary heterostructure FET devices are gate leakage and sub-threshold drain-source leakage [7,10], which determine the stand-by power dissipation of complementary circuits. Unlike Si CMOS, which has an SiO_2 gate insulator, the CGaAs gate is a Schottky diode to AlGaAs. Substantial gate current flows for gate voltages in excess of about a volt. Gate leakage current




Fig. 3: Propagation delay of CGaAs, CMOS, and TFSOI versus Supply Voltage. Gates are driving one load.

depends on the Schottky barrier height and band offsets. The large valence band offset (about 0.55 V) of high mole-fraction AlGaAs, as used in these devices, improves the gate leakage of PFETs. Typically, the PFET gate-diode turn-on voltage, defined as the gate voltage resulting in 1 A/m^2 gate current at $V_{ds}=0$, is -2V . NFETs have a turn-on voltage of 1.75 V . The gate turn-on voltages are also influenced by implant straggle effects, which should be minimized. Drain-induced barrier lowering increases gate current when the drain-to-source voltage is high (as a logic input changes state).

Methods for significantly reducing the leakage current of both N- and P-channel devices are being investigated at Motorola and other laboratories. A stacked epitaxial scheme could accomplish this, but would add too much process complexity. Further bandgap engineering of the AlGaAs/InGaAs interface is expected to reduce the leakage only slightly. A large improvement in NFET gate leakage could be achieved using a GaInP/InGaAs heterostructure, but this would be at the expense of PFET leakage.

A metal-insulator-semiconductor (MIS) gate structure is the most promising approach for reducing gate leakage [11]. In such a structure, a wide-bandgap insulator (dielectric or semiconductor such as AlN) is interposed between the gate and channel. The major material requirements for MISFET gate insulators are large bandgap, low interfacial charge density with the active HIGFET channel, and good chemical and thermal stability.

III. CGAAS LOGIC FAMILIES

CGaAs technology offers designers a range of circuit options with excellent power-delay products in each regime. CGaAs has the flexibility to implement full complementary, unipolar (pseudo-Direct Coupled FET Logic), Domino Logic, and Source Coupled FET Logic. Fig. 4 shows a NAND gate schematic in each of these logic families. In addition, pass-

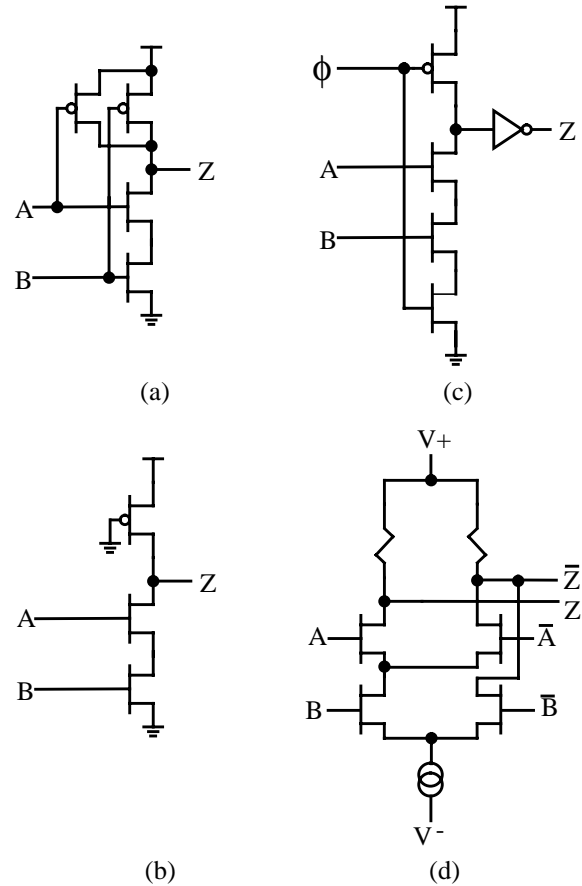


Fig. 4: CGaAs circuit configurations: (a) complementary, (b) unipolar, (c) Domino, and (d) Source-Coupled FET Logic.

gate logic can be used in conduction with complementary and unipolar logic families. These styles can be mixed on an integrated circuit, giving the designer a range of speed and power options from complementary logic to SCFL. The power and speed can also be adjusted by varying the power supply voltage.

A. Complementary CGaAs

The complementary circuit style offers CMOS-like logic and memory elements in CGaAs that can operate with power supply voltages of 0.9 to 1.5 Volts . In most applications, the static power dissipation is negligible, though it is higher than in CMOS due to higher leakage currents. The measured power delay product for two-input complementary gates driving a fan-out of four ranges from 0.01 to $0.1 \mu\text{W}/\text{MHz}/\text{Gate}$, depending on the power supply voltage.

As in CMOS, complex logic gates with stacked transistors can be built. There are several interesting differences between CGaAs and CMOS. In the CGaAs process, latchup is not possible, there are no wells, and no substrate or well contacts are needed. P- and N-transistor drains can be abutted, and their ohmic contacts can be merged. Compared to CMOS, in which device types must be separated, this feature allows sim-

pler, more efficient layout topologies. Existing CMOS designs can be mapped directly to complementary CGaAs using high level CMOS EDA tools, including synthesis and place and route tools. The only modification typically required involves the smaller supply voltage (limits on device stacking) which would also be an issue in low-voltage CMOS processes. The transconductance ratio of P to N devices calls for superbuffers to drive large loads.

Memories implemented in CGaAs are similar to CMOS memories. Static power dissipation is limited to that which derives from the small drain-source and gate leakages. The small power supply voltage, which is responsible for low power dissipation in CGaAs, does make sense-amplifier design challenging.

Complementary CGaAs is finding applications where the required circuit speed is greater than CMOS can provide at a given power dissipation, and where additional speed is required compared to CMOS. Significant power savings over CMOS can be achieved in systems having high clock rates and a high percentage of nodes switching each cycle. Its low power and radiation hardness suit full complementary CGaAs circuits particularly well to space applications.

B. Unipolar CGaAs

CGaAs offers a superset of E/D MESFET logic styles, as CMOS does for E/D NMOS. Just as CMOS technology can implement pseudo-NMOS, CGaAs can implement Direct-Coupled FET Logic (DCFL) using a grounded-gate P-transistor as a load. The P device can be switched on and off using complementary logic levels for power management. CGaAs DCFL offers higher frequency performance than complementary CGaAs logic, but at higher power.

CGaAs DCFL differs from E/D MESFET DCFL in that it has a wider logic swing, less gate current, and improved noise margins due to the high bandgap (1.7 eV) of AlGaAs. E/D MESFET technology does derive much of its speed from its small thresholds and small logic swings, but some of its speed advantage over silicon is lost because it is a NOR-only logic family. In contrast, CGaAs DCFL allows stacked transistors, making NAND functions and complex gates practical. NAND structures tend to have asymmetrical propagation delays (t_{PHL} vs. t_{PLH}) because the pull-up has to be weak compared to the pull-down transistors in order to guarantee a valid output-low voltage. The low gate current of CGaAs means that much wider NOR functions can be implemented without compromising noise margins. The wider logic swing of CGaAs DCFL does generate more internal noise, but it has wider noise margins than E/D MESFET DCFL. The rich library of logic gates in CGaAs DCFL reduces the number of gates required to implement a given function, thereby reducing critical path lengths.

The simulated power-delay product of CGaAs DCFL based on a two-input NOR driving a fan-out of four is $0.075 \mu\text{W}/\text{MHz}/\text{Gate}$. The simulated gate delay is 113 ps.

C. Domino CGaAs

The datapaths of high-performance CMOS processors make extensive use of Domino logic for its high performance at reasonably low power levels. CGaAs can also implement Domino logic. As in CMOS, the efficiency of Domino logic is very dependent upon circuit topology: when the logic is multiple levels deep, the precharge time (during which no computation is being done) is amortized over many gates. In CGaAs Domino logic, AND structures and complex gates are supported as well as OR gates. The evaluation time of a Domino gate is limited to the time required to discharge internal nodes, and to pass the signal through an output inverter. Domino gates evaluate quickly because the logic is implemented in the faster N-channel transistors (unlike in complementary circuits), the gates do not spend part of their drive current overcoming the load current (as in unipolar logic), and driven gates present smaller loads, only one unratiod transistor per input, compared to two transistors per input in complementary circuits, or one larger (ratioed) transistor in unipolar logic.

Domino logic gates consume power only during the precharge phase, as they charge nodes that were discharged during the previous evaluation phase. This characteristic does dissipate more dynamic power for signals that remain low for multiple clock signals than would be dissipated in complementary logic. Drain leakage on the dynamic node requires the use of a small keeper transistor. Domino logic requires that a precharge clock be routed to each dynamic logic gate in the circuit. The increased clock load also raises the power cost of dynamic circuits.

Domino logic is non-inverting: the inputs to every logic gate must be low when the evaluation phase begins. The output of each gate is permitted a single rising transition during the evaluation phase. Strict setup requirements for inputs to Domino gates make mixing inverting logic with Domino logic difficult. For each level of inverting logic, a separate precharge pulse is required. Inverting logic may be inserted before the dynamic logic (without the separate precharge) provided that the inputs to the dynamic logic block stabilize before the evaluation phase begins.

To summarize, in many circuits, Domino logic offers higher speed than complementary logic, at lower power than unipolar logic. CGaAs Domino test circuits are in fabrication. The simulated results (ignoring clock power and precharge time) indicate that simple Domino gates driving a fan-out of four should have power-delay products ranging from 0.03 to $0.08 \mu\text{W}/\text{MHz}/\text{Gate}$. The evaluation delay of such a two-input Domino AND gate is simulated at 115 ps. OR structures and complex gates having larger fan-in show Domino logic to its best advantage. An eight-input OR gate has a simulated evaluation delay of 70 ps. A three-level logic four-bit carry look-ahead circuit has a simulated delay of 250 ps.

D. Source-Coupled FET Logic

SCFL, which uses differential current steering logic analogous to ECL or CML, is the highest speed logic family in CGaAs. Like DCFL, SCFL dissipates static power. CGaAs SCFL circuits are well-suited for use in fast communications circuits such as ATM and SONET applications. Data rates as high as 5 Gigabits per second are currently achievable in CGaAs SCFL technology. Differential CGaAs circuits share the low noise susceptibility and good load-driving attributes of other differential logic families. The N-device-rich circuits allow 10/20 mA outputs to drive 50/25 Ω loads and transmission lines. No separate high current buffering is required for system I/O.

CGaAs SCFL can implement 3 full levels of differential logic with a 2.5 V supply, compared to the 4 to 5.2 V supply typically required in silicon ECL, or even more in GaAs HBT circuits. In CGaAs SCFL, all 3 levels of logic operate within a 1V region, allowing the remaining 1.5 volts to be used for the current source and the output logic swing. The power-delay product has been demonstrated at 0.16 $\mu\text{W}/\text{MHz}/\text{Gate}$ operating at 1 GHz. The higher power dissipation of SCFL compared to the other CGaAs logic families may keep it from being used to implement circuits such as microprocessors.

A great advantage of CGaAs is realized when these logic families are mixed on an integrated circuit. For example, the combination of SCFL for high serial data rates and rate buffering, with complementary-style CGaAs for processing the information at lower power, allows a designer to construct in a single integrated circuit what would otherwise require multiple technologies.

IV. MICROPROCESSOR PROJECT

CGaAs is being evaluated for VLSI applications at the University of Michigan through the design of a microprocessor called PUMA. Feedback from this project is helping to guide the development of the CGaAs process so that it will better suit VLSI digital applications. In addition to process and circuit design studies, the project includes architecture, operating system, compiler, packaging, and CAD tool aspects, which will be briefly described.

The general architecture of the PUMA processor follows that of the PowerPC [12], but the instruction set and microarchitecture are being tailored to fit the CGaAs technology. The microarchitecture is evolving as circuit and architecture studies proceed. As seen in Fig. 5, current CGaAs integration levels would force a partitioning of a full PowerPC processor onto multiple chips: Fixed-Point Execution Unit (FXU), Floating-Point Unit (FPU), Memory Management Units (D-MMU and I-MMU), a clock generator circuit, and memory chips. The FXU in this figure includes a small L1 I-Cache, an eight-cache line stream buffer, and instruction and data memory queues, in addition to the fixed-point datapath, register file, and control circuits. The FPU datapath would execute floating point instructions in lock-step with the fixed-point

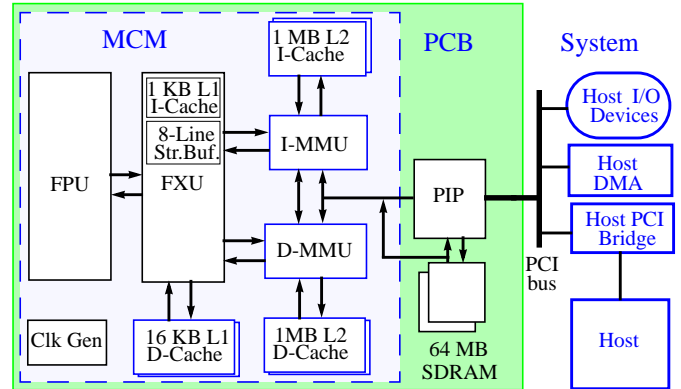


Fig. 5. Diagram of proposed prototype computer, showing microprocessor on MCM, mounted on a PC Board, which is hosted by workstation.

execution unit. Separate memory management chips for the instruction and data streams would interface the processor to CMOS L2 caches (1MB each) and, through a PCI interface chip (PIP), to synchronous silicon DRAM and to the host computer's memory and I/O.

Architectural work in this project emphasizes high bandwidth interconnect to memory, which is supported by the advanced packaging scheme described below. The characteristics of CGaAs (short gate delay and comparatively low integration levels) dictate that the datapaths and control be simple. Since the integration level will not support much parallelism, the processor's performance must come from high clock frequency, which argues for a deep pipeline (currently at ten stages in the fixed-point unit and nine stages in the floating-point unit). Deep pipelines require good branch prediction because of the high performance cost of flushing the pipe on mispredicted branches.

All microarchitectural decisions are based on results from a trace-driven simulator which takes as inputs, streams of instruction- and data-references generated from benchmark programs, and produces plots of performance loss in terms of cycles per instruction (CPI) attributable to various components of the architecture (e.g., cache misses, branch misprediction, data-dependency pipeline stalls, etc.). The simulator has been used to model four major configurations of the FXU pipeline, two versions of the FPU/FXU interface, and many memory organizations (on- and off-chip L1 cache, various L1 and L2 cache sizes, cache line sizes, multi-way stream buffers, and a software-managed TLB). As microarchitectural features such as the size of the branch target buffer or number of forwarding paths on the pipeline are varied, one can see the effect on overall CPI, and determine whether spending the hardware resources to implement a given feature is justified.

Changes in the instruction set or microarchitectural features require modifications to the software. An example of the interdependence of software and architecture is the software prefetch capability in the PUMA processor and compiler designed to reduce the performance loss due to data-cache misses.




Fig. 6. Cross-section of packaging scheme, showing flip-chip mounted CGaAs ICs and edge-mounted silicon SRAMS.

A major focus in the PUMA project is advanced packaging. As Fig. 6 indicates, the CGaAs chips are to be flip-chip mounted using fine-pitch gold bumps in an array pattern, and the CMOS cache memory chips are to be mounted on edge. Stacking the L2 caches will keep the interconnect short, minimizing time-of-flight latency. The high I/O counts provided by area pads will increase bandwidth among processor and memory chips, and allow power to be brought directly to the circuit modules, eliminating the need for wide power rails on chip.

Expected benefits of this packaging approach are smaller integrated circuits for a given functionality, faster logic paths due to less total interconnect, lower power dissipation because of reduced parasitic capacitance, and a higher percentage of the interconnect on the MCM where signal propagation characteristics are better.

The integrated circuits are being designed with the Cascade Design Automation EPOCH [13] circuit compiler. This tool suite is being extended to enable it to distribute pads, pad-drivers and receivers in an array [14] and to be able to run static timing analysis across multiple chips on the MCM. A major packaging advance will come through a new capability to concurrently place integrated circuits on the MCM and modules within the integrated circuits, in order to minimize the total interconnect length, especially on critical nets.

Our IC design flow begins with a circuit description written in the Verilog hardware description language. As pioneers in VLSI digital CGaAs, we have designed and modeled our own cell library, including memories and complementary, Domino and P-loaded DCFL logic circuits. Initial cell layouts are done in Mentor's IC Station, after which the cells are transferred to the EPOCH database using MasterPort, a Cascade tool which translates physical layout to new design rules. The CGaAs design rules have changed often as the process has developed, and the ability to efficiently regenerate the cell library has been extremely valuable. EPOCH compiles, synthesizes, places and routes the circuits automatically from the Verilog description, but its floorplanner allows user interaction, which we employ extensively. Electrical and design-rule checks are then run on the full physical design, and a gate-level Verilog model generated by EPOCH is verified to have functionality equivalent to the original behavioral Verilog specification.




Fig. 7. Maximum frequencies for a 32-bit shift register test circuit operating at 1.0, 1.2, and 1.5 Volts.

Using the design methodology described, together with preliminary device models and architectural specifications, our circuit design group has been able to evaluate both the technology and the proposed architecture by designing and simulating major microprocessor circuit blocks. This work has included a comparison of the various logic styles, and optimization of both leaf cells and complex modules such as adders and shifters in CGaAs.

Our SRAM design experience was a major influence in our recommendation that for digital CGaAs, the threshold voltages be reduced to ± 0.35 V. A PowerPC ALU test chip was also designed in CGaAs, with a major objective being to explore Domino CGaAs logic. The ALU design provided information that has been used to modify the processor architecture.

V. RECENT CIRCUIT RESULTS

Results from recently-processed circuits validate the good performance at very low power of full complementary style CGaAs technology. A 32-bit shift register is used as part of the process control monitor for complementary circuits. Fig. 7 shows the maximum operating frequency for one version of this circuit at three power supply voltages. The supply current at 1.2V is less than 1mA when the circuit is operating at more than 100 MHz. These circuits have been used to verify the power-delay performance reported above of $0.01 \mu\text{W}/\text{MHz}/\text{Gate}$ at 0.9V and $0.1 \mu\text{W}/\text{MHz}/\text{Gate}$ at 1.2V, while operating at more than 500 MHz.

A recent mixed SCFL/Complementary signal processor demonstrated an average power-delay product of $0.4 \mu\text{W}/\text{MHz}/\text{Gate}$ while operating at more than 1 GHz over -35°C to $+110^\circ\text{C}$ with a -4V SCFL supply and a -1.2V complementary supply. The circuit still operated at more than 1GHz when the supply voltage for the SCFL portion was reduced to -2.5V ; with this supply voltage, its power-delay product was $0.16 \mu\text{W}/\text{MHz}/\text{Gate}$. This performance can be compared to the Multithreshold-Voltage CMOS at NTT, which has a $0.3 \mu\text{W}/\text{MHz}/\text{Gate}$ power-delay product when operating at up to 63 MHz with a 2V power supply[15].

The 16-channel signal distribution circuit shown in Fig. 8 is implemented in complementary logic, and is composed of




Fig. 8. A 171,000-transistor 16-channel signal distribution circuit.

171,000 transistors. This represents the highest level of integration in HIGFET technology attained to date. On the most recent wafer lot, this circuit achieved 50% yield at wafer probe. The design was done using automated CAD tools. As is seen in Figs. 9 and 10, the CGaAs circuit dissipates less power and has a better power-delay product than 2.7 V CMOS at clock frequencies over 55 MHz. A power-delay product of 0.01 $\mu\text{W}/\text{MHz}/\text{Gate}$ is attainable at 100MHz.

We have also initiated design and characterization of CGaAs analog, MMIC, and RF power circuits. While some device improvements still need to be made for precision analog circuits, applications such as A/D and D/A converters, and ATM can be addressed well with the current transistors. N-channel CGaAs devices are similar to GaAs HFETs, and

Parameter	NFET ($L_g=0.7 \mu\text{m}$)	PFET ($L_g=0.7 \mu\text{m}$)
f_t ($V_{ds}=1.0\text{V}$)	20 GHz	5 GHz

Table 2: RF device characteristics.




Fig. 9. CGaAs vs. CMOS power comparison.

Parameter	NFET ($L_g=0.7 \mu\text{m}$)	PFET ($L_g=0.7 \mu\text{m}$)
f_{max} ($V_{ds}=1.0\text{V}$)	30 GHz	10 GHz
MAG ($V_{ds}=1.0\text{V}$)	20 dB	6 dB

Table 2: RF device characteristics.

can be used in MMIC and power RF circuits in similar ways. Table 2 shows typical RF device characteristics for 0.7 μm gate length devices with $V_{ds} = 1.0 \text{ V}$. A 0.7 μm x 3 μm NFET was used as a power amplifier in class AB operation, with a 3V supply and 26 dBm output power at 1GHz; power-added efficiency was greater than 60%

A variety of complementary, DCFL and Domino test circuits, including simple gates, I/O drivers and receivers, an SRAM, a PowerPC arithmetic logic unit, a register file, and a PLL-based clock generator have been designed to evaluate the technology for high-speed digital applications. These circuits are being fabricated.

VI. SUMMARY

A complementary GaAs technology which has shown promising performance and yield in RF and digital communications applications is being evaluated for high-speed VLSI circuits. For circuits such as microprocessors, a mixture of complementary, DCFL and Domino logic styles allows designers to trade power for speed over a wide range, with power supply voltages of 0.9 to 1.5 V.

The major challenges in CGaAs seem to be developmental, rather than fundamental. Both the process and device models are still being tuned. The interconnect metallization is not competitive with state-of-the art CMOS processes. It is more difficult to make good ohmic contacts to III/V material than to silicon, so contact rules are larger than their counterparts in CMOS. CGaAs transistors have gate current, which is made worse by drain-induced barrier lowering, and they have more drain-source leakage than MOSFETs. CGaAs device scaling is two generations behind CMOS. CGaAs is radiation hard, and much progress has been made recently in CGaAs defectivity (yield) and integration levels.

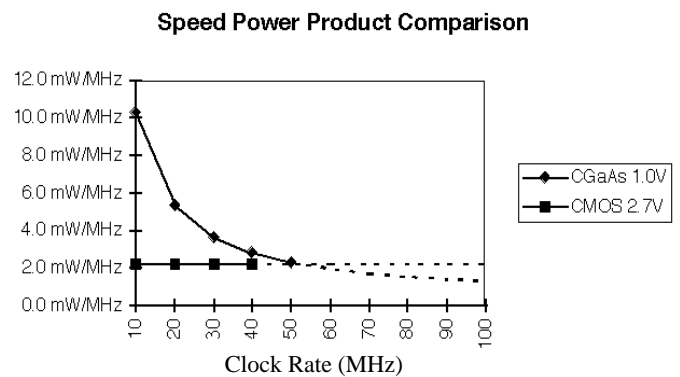


Fig. 10. CGaAs vs. CMOS power-delay comparison.

The PUMA project at the University of Michigan, an effort to build a CGaAs PowerPC-architecture microprocessor, is providing feedback to the process engineers regarding changes that are needed to better suit CGaAs to digital VLSI applications. The area and performance effects of changing various design rules is evaluated on memory and logic circuits from the processor design. Examples of the results of such experiments are the shift of threshold voltages in the process from +/- 0.55 V to +/- 0.35 V, the addition of a fourth level of metal, and scaling of channel lengths in experimental circuits to 0.35 μm .

CGaAs overcomes many of the weaknesses of E/D MES-FETs for digital circuits, and provides the flexibility of CMOS technology with a better power-delay product. The most promising applications for digital CGaAs are very high-performance circuits (500 to 1000 MHz) and low-power circuits operating at speeds over 100 MHz.

ACKNOWLEDGMENT

This work is funded by the U.S. Defense Advanced Research Projects Agency, under ARO DAAH04-94-G-0327.

REFERENCES

- [1] Linley Gwennap, "The Death of the Superprocessor," *Microprocessor Report*, vol. 9, no. 13, p. 3, October 2, 1995.
- [2] Michael D. Upton, *Architectural Trade-offs in a Latency Tolerant Gallium Arsenide Microprocessor*, Ph.D. Dissertation, University of Michigan, Ann Arbor, 1996.
- [3] David B. Papworth, "Tuning the Pentium Pro Microarchitecture," *IEEE Micro*, vol. 16, no. 2, pp. 8–15, April 1996.
- [4] Ron Wilson, "Superscalar sunset seen at Hot Chips?" *Electronic Engineering Times*, issue 862, pp. 1 and 108, August 21, 1995.
- [5] Linley Gwennap, "Digital's 21164 Reaches 500 MHz," *Microprocessor Report*, vol. 10, no. 9, pp. 1 and 12, July 8, 1996.
- [6] Jim Turley, "401GF is Coolest, Cheapest PowerPC," *Microprocessor Report*, vol. 10, no. 8, pp. 9–10, June 17, 1996.
- [7] B. Bernhardt, M. LaMacchia, J. Abrokwah, J. Hallmark, R. Lucero, B. Mathes, B. Crawforth, D. Foster, K. Clauss, S. Emmert, T. Lien, E. Lopez, V. Mazzotta, B. Oh, "Complementary GaAs (CGaAsTM): A High Performance BiCMOS Alternative," GaAs IC Symposium Technical Digest, San Diego, CA, Oct. 29–Nov. 1, 1995, pp. 18–21.
- [8] J. Abrokwah, J. Huang, W. Ooms, C. Shurboff, J. Hallmark, R. Lucero, J. Gilbert, B. Bernhardt, and G. Hansell, "A Manufacturable Complementary GaAs Process," GaAs IC Symposium Technical Digest, San Jose, CA, Oct. 10-13, 1993, pp. 127-130.
- [9] P. O'Neil, B. Bernhardt, F. Nikpourian, C. Della, Y. Abad, G. Hansell, "GaAs Integrated Circuit Fabrication at Motorola," GaAs IC Symposium Technical Digest, San Jose, CA, Oct. 10-13, 1993, pp. 123-126.
- [10] J. Hallmark, C. Shurboff, W. Ooms, R. Lucero, J. Abrokwah, and J. Huang, "0.9V DSP Blocks, A 15ns 4K SRAM and a 45ns 16-bit Multiply/Accumulator," *GaAs IC Symposium Technical Digest*, 1994, pp. 55-58.
- [11] F. Ren, M. W. Hong, W. S. Hobson, J. M. Kuo, J. R. Lothian, J. P. Mannaerts, J. Kwo, Y. K. Chen, and A. Y. Cho, "Enhancement-Mode p-Channel GaAs MOSFETs on Semi-Insulating Substrates," *IEDM Technical Digest*, San Francisco, CA, Dec. 8-11, 1996, pp. 10.7.1-10.7.3.
- [12] Shlomo Weiss and James E. Smith, *Power and PowerPC: Principles, Architecture, Implementation*, Morgan Kaufmann, San Francisco, 1994.
- [13] Cascade Design Automation, *EPOCH User's Manual, ver. 3.2*, Bellevue, WA 98006, 1995.
- [14] Ray Farbarik, Xiaowen Liu, Mark Rossman, Richard B. Brown, Phiroze Parakh, and Todd Basso, "CAD Tools for Area Distributed I/O Pad Packaging," Proceedings of the IEEE Multi-Chip Module Conference, Santa Cruz, CA, Feb. 4–5, 1997, pp. 125-129.
- [15] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki, S. Shigematsu, and J. Yamada, "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," *IEEE J. of Solid-State Circuits*, vol. 30, no. 8, pp. 847-854, August 1995.