

HIGH POWER KEEPS COOL

Mission-Specific Processing Techniques Enable Low-Power, High-Performance Surveillance Sensors

*Michael Lucas, Naresh Shanbhag, Kaushik Roy,
Fadi Kurdahi, and John Fagan*

Advanced U.S. Department of Defense (DoD) surveillance sensors are continually pushing the need for ever-increasing electronics performance. Full-time, large-area ground surveillance capability, for instance, has been proven to be of great importance in recent conflicts; it drives the requirement for the signal processing of growing numbers of radar channels with increasing bandwidth. At the same time, the surveillance platforms carrying these sensors, typified by unmanned aerial vehicles (UAVs) or other small aircraft or spacecraft, are being pushed to be smaller, lighter, cheaper, and to have longer mission durations. On the Defense Advanced Research Projects Agency (DARPA)-sponsored Mission-Specific Processing (MSP) project, the team of Northrop Grumman Electronic Systems, Purdue University, University of Illinois at Urbana-Champaign, Morpho Technologies, and Atmel Corporation developed and demonstrated low-power, high-performance application-specific integrated circuit (ASIC) cells and computation techniques to provide 10× throughput performance per watt improvements to meet next-generation DoD sensor-processing applications. Northrop Grumman has defined a set of cells that form the core of the digital signal processing (DSP) functions needed for emerging military radio-frequency (RF) applications. The University of Illinois has developed a fluid intellectual property (IP) core generator that produces fully laid out optimized DSP function cells

using transistor and voltage scaling to maximize ASIC performance per watt. Purdue University developed low-computation complexity algorithms and advanced-circuit techniques to achieve additional improvements in operational performance and input/output (I/O) throughput per watt.

Morpho Technologies analyzed application of these MSP techniques to their reconfigurable core cell, which provides on-the-fly reconfiguration with reduced power and improved performance. MSP cells were integrated into the design process of a commercial ASIC foundry (Atmel Corporation) standard cell library, allowing designers to achieve ASIC performance levels normally associated with full custom designs, yet using a standard-cell design flow to reduce the design cost by a factor of ten. Northrop Grumman defined a wideband adaptive processing ASIC chip set incorporating these cells and projected their performance in a multichannel real time surveillance radar processor, exceeding the 10× improvement goals. An overview of the project is shown in Figure 1.

KEY ASIC METRICS FOR DOD SENSORS

Emerging DoD sensors need more sensor channels, wider bandwidth, and adaptive cancellation of interference, jamming, and clutter. Projections of sensor signal processing throughput needs are on the order of tens of trillion operations per second (TO/s). Many of the sensor platforms have severe payload constraints on the size, weight, and power



© EYEWIRE

available for this processing. The use of commercial off-the-shelf (COTS) programmable processing falls several orders of magnitude short in providing the required throughput in the available physical envelope. The vast majority of the processing throughput required in these advanced sensors is in dedicated functions that are specific to the sensor, so optimized mission-specific ASICs are an effective solution to meet the processing requirements within the platform constraints. A key metric on the project is the ASIC performance in terms of operations per second per watt (O/s/W). This parameter is extremely critical to military systems, as they frequently must operate over wide temperature ranges with severely constrained cooling abilities, as compared to commercial computing hardware placed in air-conditioned rooms. Also, since the bulk of the targeted functions are closely coupled to high data rate sensors, low-power, high-bandwidth, chip-to-chip I/O measured in gigabits per second per watt (Gb/s/W) is also a key metric.

SENSOR REQUIREMENTS ANALYSIS

MSP ASICs are targeted for the mission-specific application domain, which is typically in the front end of the signal processing sequence. These functions are usually tightly coupled to the sensor configuration and require very high arithmetic throughputs and data rates. Northrop Grumman performed system analysis and processing requirements projections for emerging DoD RF sensors to identify which functions provided the highest pay-

**Surveillance platforms
are being pushed to be
smaller, lighter, cheaper,
and to have longer
mission durations.**

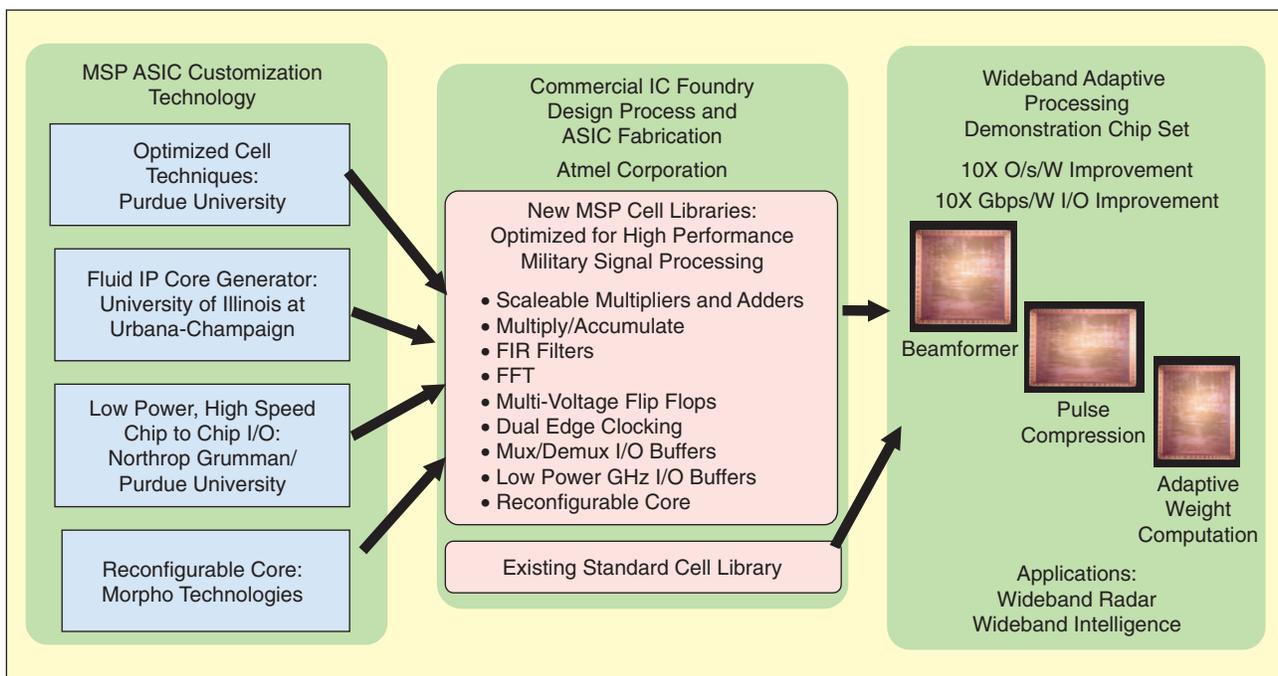
back if implemented as optimized MSP cells. The resulting cell set ranged from very large functions, such as finite-impulse response (FIR) filters and fast Fourier transforms (FFTs), down to low-level functions, such as improved flip flops that support multivoltage operation and lower power chip-to-chip I/O. The set of MSP cells addressed on the project consists of the following:

- ◆ scaleable bit-width core arithmetic cells: multipliers, adders, and multiply/accumulate
- ◆ large-scale DSP functions: FIR filters, polyphase filters, and FFT core
- ◆ low-power enabling cells: multivoltage level converting, and dual-edge transition flip flops
- ◆ chip-to-chip communication: mux/demux I/O cell, low-power gigahertz I/O buffers
- ◆ high-performance, low-power reconfigurable core.

The MSP project is developing this set of high-performance cells by incorporating optimization techniques that include a fluid IP core generator, reduced-complexity algorithms, advanced I/O circuits, and reconfigurable cores.

Fluid IP Core Generator: Power/Speed Optimized ASIC Cells Without Synthesis

The University of Illinois at Urbana-Champaign developed a fluid IP core generator that provides custom quality layouts in a short (comparable to synthesis) design cycle. The fluid IP



1. MSP integrates complementary ASIC design techniques to meet high-performance DoD system throughput and power requirements.

core generator accomplishes this by sacrificing the generality of a synthesis methodology in order to obtain significant improvements in power, speed, and area efficiencies. This is achieved by having the core generator target the design of specific DSP subsystems such as FIR filters, FFTs, polyphase filters, and forward error-correction blocks. The fluid core generator employs algorithmic and architectural considerations to optimize the circuit fabric. The resulting layouts are based on finely tuned transistor sizes, hence the term “fluid.” More

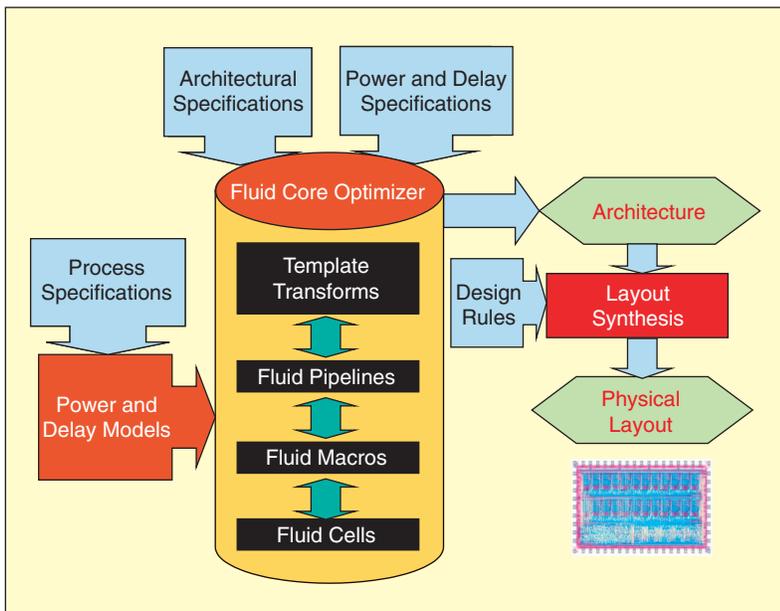
Advanced DoD surveillance sensors are continually pushing the need for ever-increasing electronics performance.

tor encapsulates cross-domain optimizations, such as those that span algorithmic, architectural, and circuit domains, and makes them available to the average designer. For example, a large number of algorithmic and architectural optimization

but suffer from similar problems as synthesized designs. Hard cores, in the form of completed layouts, are predictable but are not portable across technology generations.

The Illinois fluid IP core generator builds upon a core-based design methodology between these extremes. The core genera-

tor encapsulates cross-domain optimizations, such as those that span algorithmic, architectural, and circuit domains, and makes them available to the average designer. For example, a large number of algorithmic and architectural optimization techniques exist for signal processing and communication systems, such as pipelining, strength-reduction, folding/unfolding, parallel processing, and DÉCOR [1], [2]. These techniques, when combined with circuit-level optimizations, provide significant improvements in power and delay.



2. Fluid IP core generator produces optimized physical layouts from architecture, power, delay, and process specifications.

Fluid IP Core Generator Architecture

Figure 2 illustrates the structure of the fluid IP core generator. The IP core generator has two major components: a fluid core optimizer and a layout synthesizer. It accepts as inputs power and delay models, a library of template transforms, algorithmic specifications, and power and delay specifications, and generates optimized physical layouts.

Fluid Core Optimizer

The fluid core optimizer begins the design process with a data flow graph (DFG) based on the desired type of filter or DSP function. The DFG represents computational units, such as multipliers and adders, as nodes and the flow

importantly, the core generator bypasses logic synthesis completely, providing designs with predictable quality. The MSP project initially focused on the design and demonstration of a fluid IP core generator for a complex data FIR filter; extensions to other functions were also performed.

IC design methodologies in common use today can be broadly classified into full custom design, logic synthesis based, and the more recent core-based methodology. Custom designs offer the best energy, delay, and area benefits but require a long and expensive design cycle, making it economically unfeasible for the small production quantities typical of military designs. Synthesis-based methodologies reduce the design cycle time and cost by employing an automated logic synthesis step followed by automatic placement and routing of standard cells but suffer from unpredictable quality and sacrifice area, power, and throughput. Soft cores, usually in the form of VHDL code, are portable from one technology generation to the next

of data between the computational units as edges. Architecture-specific information, such as word sizes and delays, is embedded in the graph. Using the computation-efficient delay models developed on the project, the delay of each node in the DFG can be computed. A shortest-path algorithm is then used to determine the critical path of the DFG. This can be compared with the desired throughput specified by the user.

After the application of transform templates, the fluid core optimizer reduces transistor sizes in the fluid cells to minimize power consumption. The fluid cell library consists of one parameterized cell for every logic cell in the library, and each cell can be instantiated at virtually any size desired. This provides a tremendous benefit over traditional standard cell libraries, which consist of cells with a limited number of drive strengths. The fluid cell library also has cells with multiple scale factors, each corresponding to a different group of transistors within the cell. While a buffer has a single scale factor for the entire cell, a full adder has four scale factors that separately

affect the p-channel metal-oxide semiconductor (PMOS) and n-channel MOS (NMOS) transistors in the carry-out and sum paths. By carefully choosing the groups of transistors to be scaled, as well as the relative scaling of transistors within groups, a cell with one or two scale factors can be as effective as customized sizing of each transistor. Figure 3 shows a full adder cell instantiated at two different sizes. The scale factor α_C controls scaling in the first stage, which generates the carry output, and α_S controls scaling in the second stage, which generates the sum output. Thus, α_C affects the delay through both the carry and sum paths, while α_S affects only the sum delay. It is infeasible to characterize each cell for every possible instantiated size. Therefore, a large part of the effort involved modeling development to quickly and accurately describe the delay and power of the cells in terms of the scaling parameters α_C and α_S .

The fluid cells are used to reduce the overall power consumption of structures. The existing core generator uses ripple-carry adders and Baugh-Wooley multipliers due to their regularity in layout. These adder and multiplier architectures have many paths shorter than the critical path. The shorter paths can be exploited to reduce power by employing the scaling factors to shrink the transistors in cells off the critical path.

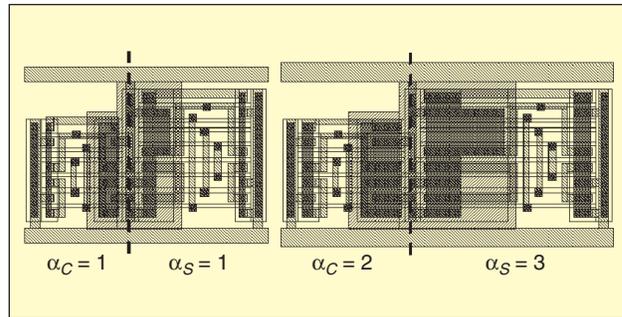
Core Generator Layout Synthesis

The layout synthesizer accepts architectural parameters and transistor scaling factors from the core optimizer to synthesize a layout. As the architectural template is known, the synthesizer starts off with a good initial placement and routing. From the power estimates of the filter and taps, the synthesizer calculates the number of power and ground lines and their dimensions. At the end, the layout synthesizer routes the clock, which is buffered as determined by the amount of load that the buffer is expected to face and routed in the opposite direction to that of the data in order to combat the effect of skew.

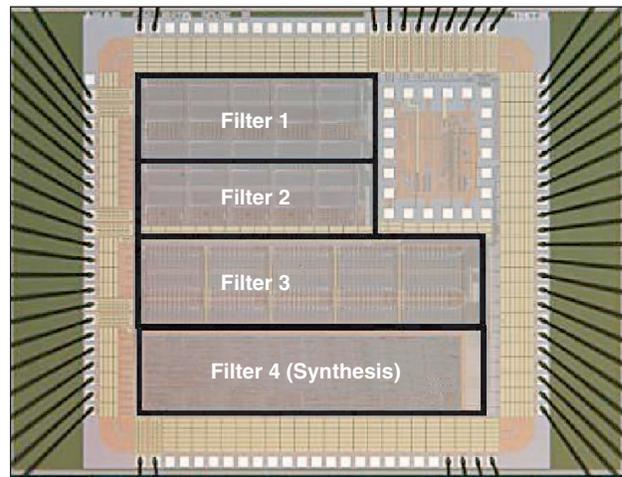
Core Generator Filter Test Case Results

A test chip was designed and fabricated with three core-generated filters and a synthesized filter for comparison. All filters were five-tap, complex FIR filters, with 8-b input, 12-b coefficient, and 23-b accumulator precision. The filters were developed for the Taiwan Semiconductor Manufacturing Corporation (TSMC) 0.25- μm , 2.5-V complementary MOS (CMOS) process. Figure 4 shows the layout of the test chip with the three core-generated filters and one synthesized filter.

The first core-generated filter (Filter 1) was a 20-MHz design using a fixed cell size with a scale factor of one, which is comparable to a standard cell with a drive of $1\times$. The second core-generated filter (Filter 2) is identical, but employs fluid cells. The third core-generated filter (Filter 3) was targeted for 150-MHz through-



3. Full adder cell instantiated with different scaling values for carry and sum.



4. Test chip for fluid IP core generator.

put with fluid cells. Here, the core generator employed fine-grained pipelining of the taps to meet the throughput requirements. These examples show the ability of the core generator to design high-performance, low-power cores for a wide range of design specifications. The synthesized filter (Filter 4) was targeted for 20-MHz area/power critical operation as a benchmark. The design was synthesized with a 43-cell standard cell library, consisting of logic gates at multiple drive strengths, several flip-flops, and two full adder cells; low-speed, low-power static logic, and high-speed, high-power pass-transistor logic. A VHDL description of the filter was pipelined at the tap level. Synopsys Design Compiler was

Table 1. Comparison between core-generated and synthesized filters.

Filter Number	Design Methodology	Speed MHz	Power Consumption	Area mm ²	Power-Delay Product pW-s
1	Core generated slow filter (fixed cells)	20	0.95 mW at 0.98 V	1.12	47.5
2	Core generated slow filter (fluid cells)	20	0.90 mW at 0.98 V	1.11	45.0
3	Core generated fast filter	160	109.8 mW at 2.5 V	1.47	686.3
4	Synthesized slow filter	20	6.47 mW at 1.5 V	1.59	323.5



5. DCM computation reductions for a FIR filter for different word bit widths and first and second orders (m) of differential coefficients.

employed for synthesis and Silicon Ensemble for timing-driven placement and routing.

Power and performance results for the FIR filters are shown in Table 1. In all cases, the core generator produced lower power filters with less area than synthesis. These results show that the core generated slow filters offer a $7\times$ improvement in energy-efficiency over the corresponding synthesized filter, with a 30% reduction in area. In the fast-filter case, the core-generated filter is shown to provide $2.5\times$ improvement over the synthesized filter along with an area reduction of 45%.

Reduced Complexity Algorithms: Complex DSP Results With Less Computation

Recasting mathematical functions into different forms can provide the same results as the original form, but require fewer arithmetic operations. Purdue University has used this approach to develop several methodologies for the design of low-energy consumption, high-performance DSP functions, such as filter banks, FFTs, and DCTs, using low-complexity design techniques at different levels of design abstraction. Two examples of these complexity reduction techniques—differential coefficients and shared multipliers—were the focus of the MSP project.

Differential-Coefficients Method

The differential coefficients method (DCM) is a novel algorithm-level technique for realization of low-power FIR/IIR filters with a large number of taps (of the order of hundreds) [4]. DCM relies on reducing computations to reduce power. The algorithms for the differential coefficients use various orders of differences between the coefficients in conjunction with stored precomputed results, rather than the coefficients themselves, to compute the canonical form convolution. These algorithms result in lesser computations per convolution as compared to direct form computation. However, they require more storage and storage accesses and, hence, more energy for storage operations. The net energy savings using DCM is dependent on various parameters such as the order of differences used, energy dissipated in a storage access, and the word widths used for the digitized input data and coefficients. DCM can also lead

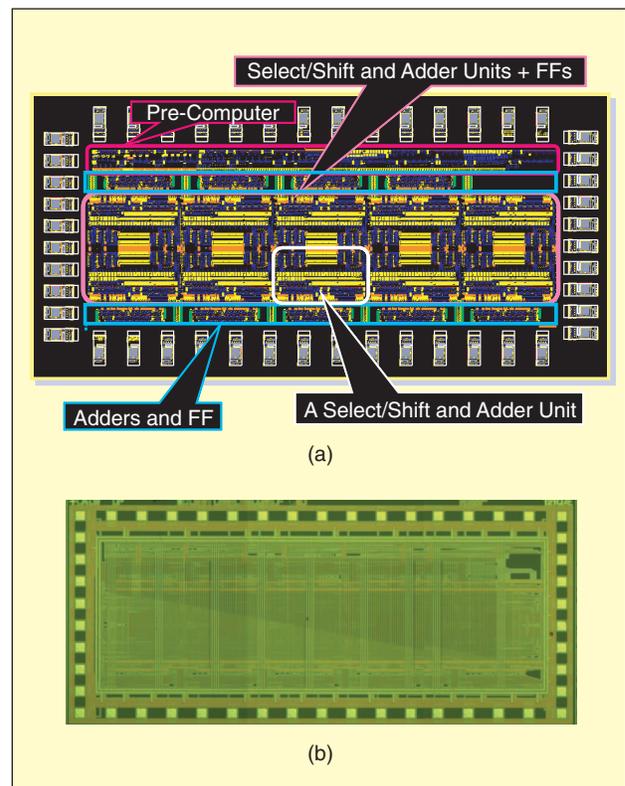
to a reduction in the time needed for computing each convolution and may provide an added advantage of higher speed computation. Analogous to the savings in energy, the speed enhancement obtained is dependent on the order of differences used and various other parameters. DCM does not rely on using a specific binary encoding of the coefficients or input data for the results and is applicable to a wide range of FIR filters.

This technique is useful if, and only if, the differences between coefficients are small compared to the coefficients themselves. In the multiplication for computing a product term using this algorithm, the technique trades a long multiplier for a short one and some additional overhead. If the energy savings in multiplication is greater than the net energy dissipated due to the overheads, there is a net energy savings as compared to using the direct form. The algorithm can be generalized to use with higher order differences (differences of differences) for even higher performance, as shown in Figure 5.

Shared Multiplication Technique for Digital Filters

Computational complexity can also be reduced by sharing multiplications for vector-scaling operations, such as FIR filtering. The computation shared multiplication (CSHM) scheme considerably reduces redundant computation by decomposing the vectors in a manner that results in maximal computation sharing, resulting in a faster and potentially low-power implementation [3].

Consider a product of a given vector, representing coefficients, with the scalar x . In this operation, the input x is multiplied by all



6. CSHM FIR test chip.

the coefficients simultaneously. One can carefully select a set of small bit sequences so that the same multiplication result can be obtained by only add and shift operations. For instance, $(1011) \cdot x$ can be decomposed as $(0011) \cdot x + 2^3(0001) \cdot x$. If both $(0011) \cdot x$ and x are available, the entire multiplication process is reduced to a few add and shift operations. These chosen basic bit sequences are referred to as *alphabets*. Also, an *alphabet set* is a set of *alphabets* that spans all the coefficients in the vector.

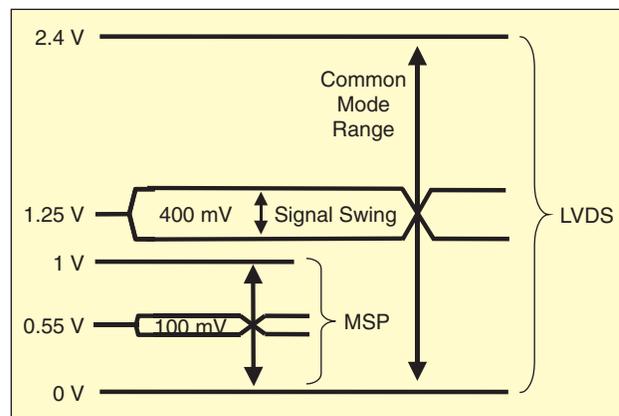
Depending on the selection of the alphabet set, the number of required add and shift operations changes. As the number of coefficients increases, there can be many choices for *alphabet sets* on the coefficients, and each *alphabet set* gives rise to a different combination of add and shift operations. Obviously, an alphabet set should cover all the coefficients in the coefficient vector. In addition, there are two other desirable characteristics of a “good” alphabet set. First, the total number of add operations should be minimized. Because the add operations lie on the critical path, reduction of the number of add operations improves overall performance. Second, the number of alphabets in an alphabet set should be minimized. [5] By slightly modifying the coefficients, the number of alphabets can be reduced. In the above example, the coefficient is 01101011. If we change the coefficient to 01101100 by adding 1, the coefficient can be computed as $2^5(11) + 2^2(11)$. In this case, the alphabet set is reduced to {11}. Reducing the number of alphabets gives rise to a lower complexity vector scalar, which leads to high performance as well as low-power design.

A ten-tap programmable FIR filter was designed and fabricated in TSMC 0.25- μm technology based on the architectural and circuit-level techniques. Separate power supplies for the core and I/O pads allowed exact measurement of the power consumption of the core of the chip. The CSHM layout and test chip is shown in Figure 6.

Table 2 shows the resulting characteristics of the CSHM test chip versus FIR filters using a Wallace tree multiplier (WTM) and a carry save array multiplier (CSAM) for comparison. (WTM and CSAM are the two most widely used multipliers. Generally, WTM has better performance than CSAM due to the tree-like structure of partial-sum adders. However, WTM has the disadvantage of having very irregular interconnect.) As shown in the table and based on the simulation results, the FIR filter using CSHM has a 19 and 43% performance improvement over the FIR filter using WTM and CSAM, respectively. In terms of power consumption, the CSHM scheme has a 17 and 20% improvement with respect to the FIR filter based on WTM and CSAM, providing power-delay product improvements of 1.7 and 2.6 \times , respectively.

New I/O Techniques Improve Chip-to-Chip Gb/s/W

DoD radar sensor systems typically have high data input rates and chip-to-chip communications requirements. In many cases



7. High-speed I/O uses a reduced swing for lower power and higher data rate.

the chip I/O accounts for half of the power consumed by the chip. MSP-class ASICs typically have high communication requirements with chip-to-chip data rates of 50–100 Gb/s. Northrop Grumman developed an MSP I/O receiver and driver to provide a greater than 10 \times -Gb/s/W improvement over the existing low-power standard, low-voltage differential signaling (LVDS). The I/O cell implements a lower swing (100 mV) version of the LVDS standard, as shown in Figure 7, and provides higher speed: up to 4 Gb/s double data rate (DDR).

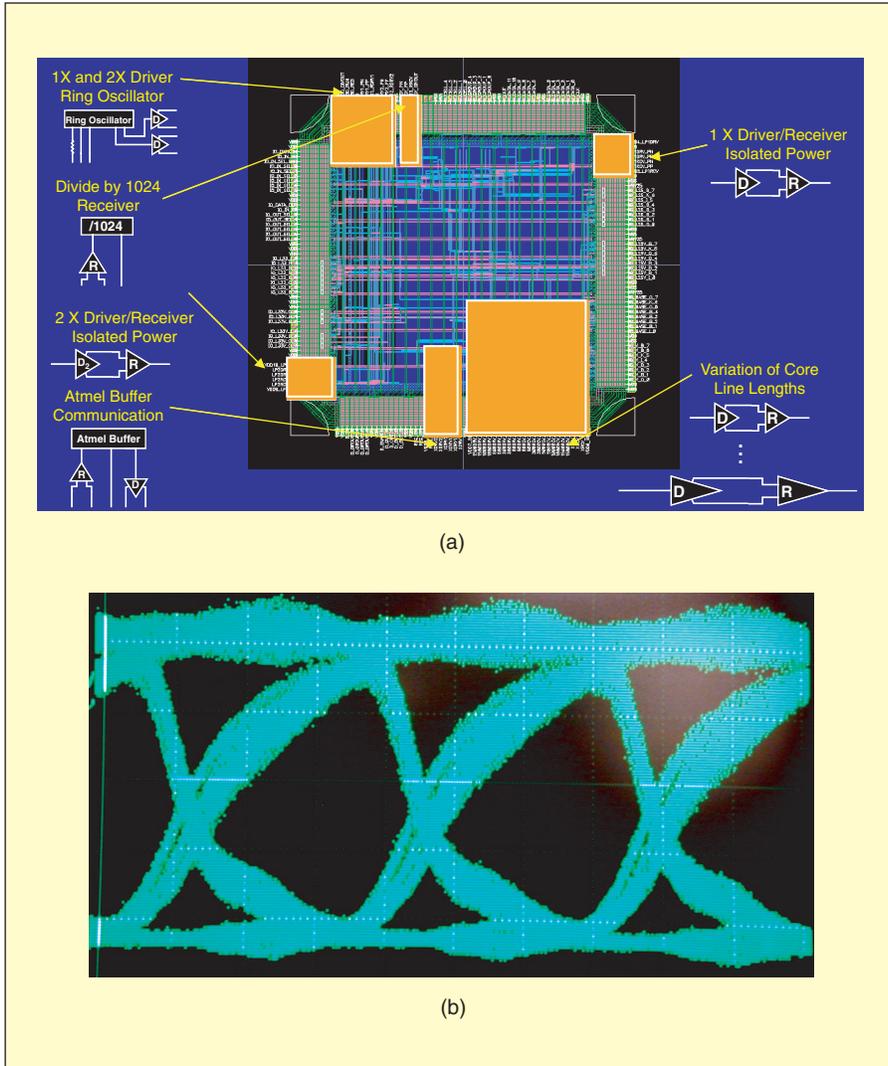
An I/O test ASIC was designed for 2-GHz operation and contained drivers, receivers, and internal test support structures. The I/O ASIC was tested with a 12-GHz bit error rate tester (BERT) using pseudo random bit stream (PRBS) data sets. Bit error rate tests were performed at frequencies up to 2.2 GHz, demonstrating better than 10⁻¹² bit error rates. Figure 8 shows the test structures and eye patterns for the test chip.

Table 2. Comparison of CSHM-based FIR filter with convention FIR filters.

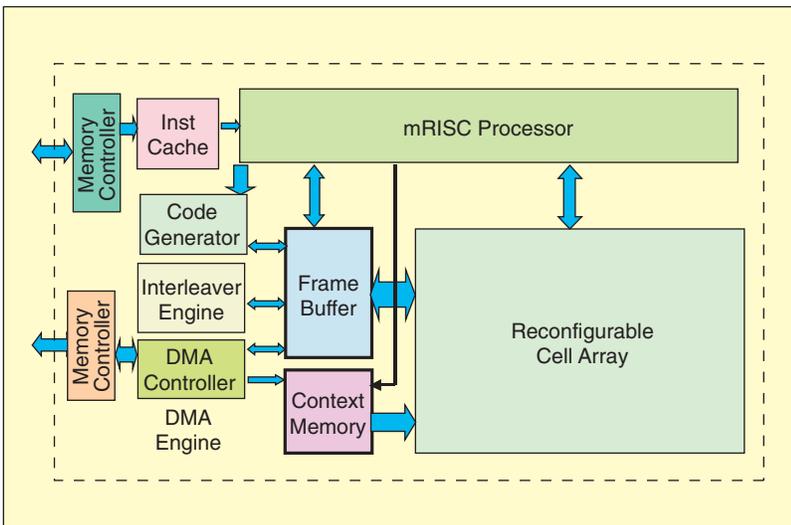
	FIR using CSHM (Measurement)	FIR using WTM (Simulation)	FIR using CSAM (Simulation)
Minimum Cycle Time (ns)	5.7 (7)*	7.0	10
Power Consumption at 100 MHz (mW)	238.75	344.3	357.1
Power-Delay Product (pW-s)	1,362	2,408	3,571

*Simulation showed 5.7 ns clock cycle, but test instruments were limited to 7 ns.

The power savings over current LVDS I/O designs was significant. The comparisons were made against the existing LVDS design in 0.18- μm technology with a common supply voltage of 1.8 V. The existing LVDS cell has a maximum speed of 600 MHz, equivalent to a DDR rate of 1.2 Gb/s. The MSP I/O cells operated at over 2 GHz (4 Gb/s DDR). The power consumption for a 2-GHz link was 11 mW and resulted in a 13 \times -Gb/s/W improvement over existing LVDS. A



8. I/O Test ASIC with $<10^{-12}$ bit error rates and >4 Gb/s operation.



9. The MS1 core architecture

companion mux/demux design by Purdue was also designed and simulated, and simulation data showed operational speed capable of supporting the 4-Gb/s I/O rate.

Reconfigurable MSP Core Provides Flexible Performance Where Needed

A reconfigurable core cell was also defined as part of the MSP approach. This core was targeted at functions that need higher performance-per-watt than central processing unit (CPU) or field-programmable gate array (FPGA) devices can provide but more flexibility to respond to changes in modes and missions than is feasible to put in completely dedicated ASICs. The Morpho Technologies MS1 reconfigurable DSP (rDSP) architecture was selected for this core cell. Architectural optimizations to accommodate high-performance radar requirements and cell-level optimizations using the MSP core generator and low power cells were analyzed to define the final core design.

MS1 rDSP Architecture

rDSP solutions are provided based on a reconfigurable array-processing paradigm known as the MS1 rDSP. Figure 9 shows the MS1 architecture. As in conventional reconfigurable systems [6], the MS1 Core contains a reconfigurable block, called the RC array, and a 32-b reduced-instruction set computer (RISC) processor, called the mRISC. The RC array consists of reconfigurable cells (RCs) interconnected by a reconfigurable network. Both the functionality of the RCs and the network interconnections are determined by a configuration program, called Context. By writing the appropriate context, the developer can use the RC array to exploit the parallelism available in the application. The mRISC core controller determines the application's control flow, executes the sequential tasks of the application, and starts transfers to/from the off-core memory. The instructions for the mRISC are stored in the instruction memory. The MS1 Core also contains a context memory, a data buffer, and an I/O Controller. The RC is the

programmable element within the RC array. It performs general-purpose operations as well as word-level and bit-level DSP functions. Input operands can be either internal to the RC, from other RCs, or from the data buffer. Due to the flexible interconnect structure and reconfiguration capabilities, the RC array can operate in several modes

whereby groups of cells can perform similar operations. The size and composition of these groups, as well as the mode of operation, can be changed dynamically by switching contexts during runtime. The operation of the RC array is coordinated by the mRISC core controller, which selects and switches contexts at runtime.

Reconfigurable MSP (RMSP) Core Optimization

To develop the optimized RMSP core, several techniques were introduced for performance enhancement of the MS1. Briefly, these techniques can be categorized to two approaches:

- ◆ optimization of the RC components while maintaining the overall architecture
- ◆ enhancements that required architectural changes in RCs and other MS1 components.

The original MS1 core was optimized for low-cost commercial applications and had relatively narrow (16 b) data paths and limited register files within the RC array. The DoD radar problem requires higher precision data paths and has large data set sizes requiring more intermediate storage. A number of architectural changes were identified to add these capabilities. Also, two implementation-related MSP enhancements, fluid IP cores from the Illinois core generator and the low-power circuit techniques from Purdue, were evaluated. When implemented in a multicore ASIC, these improvements can provide a 17× improvement in O/s/W over the existing design when applied to radar applications.

MSP CELLS BECOME PART OF AN ASIC FOUNDRY'S LIBRARY

The MSP program also demonstrated a path to commercial silicon production. The approach for application of the optimized MSP cells is to incorporate them into a commercial ASIC foundry standard cell library, so they can be used in a conventional design flow along with the foundry's existing standard cells in the library. Techniques for automatically generating front-end and back-end design tool libraries were explored, and only a minimal tailoring effort was required for the integration of MSP cell designs into the library. Initial cells were transferred to the commercial foundry team member, Atmel, to demonstrate the approach. This approach is divided into two areas—standard cell components and compiled macrocell functional blocks.

- ◆ *Standard Cell Components:* Purdue's research produced candidate cells for improved power or speed performance

Emerging DoD sensors need more sensor channels, wider bandwidth, and adaptive cancellation of interference, jamming, and clutter.

at the basic logic functional level. An example of this type cell is the edge triggered flip flop that clocks data on both rising and falling edges, thereby reducing the clock distribution power on an integrated circuit. Cells of this type are designed to match the standard cell requirements of Atmel's 0.18- μ CMOS design kits. Cells are character-

ized with HSPICE for functional behavior, timing, and power over process, temperature, and voltage corners.

- ◆ *Macrocell Blocks:* Macrocell blocks, such as a compiled FFT or FIR filter, are created by compiling a design of leaf cells with variable transistor sizing for power optimization. The leaf cells are created using Atmel's 0.18- μ CMOS design rules and technology files. Components of this type are created with the core generator software each time they are used. This software has been delivered to the government for use in Macrocell block generation.

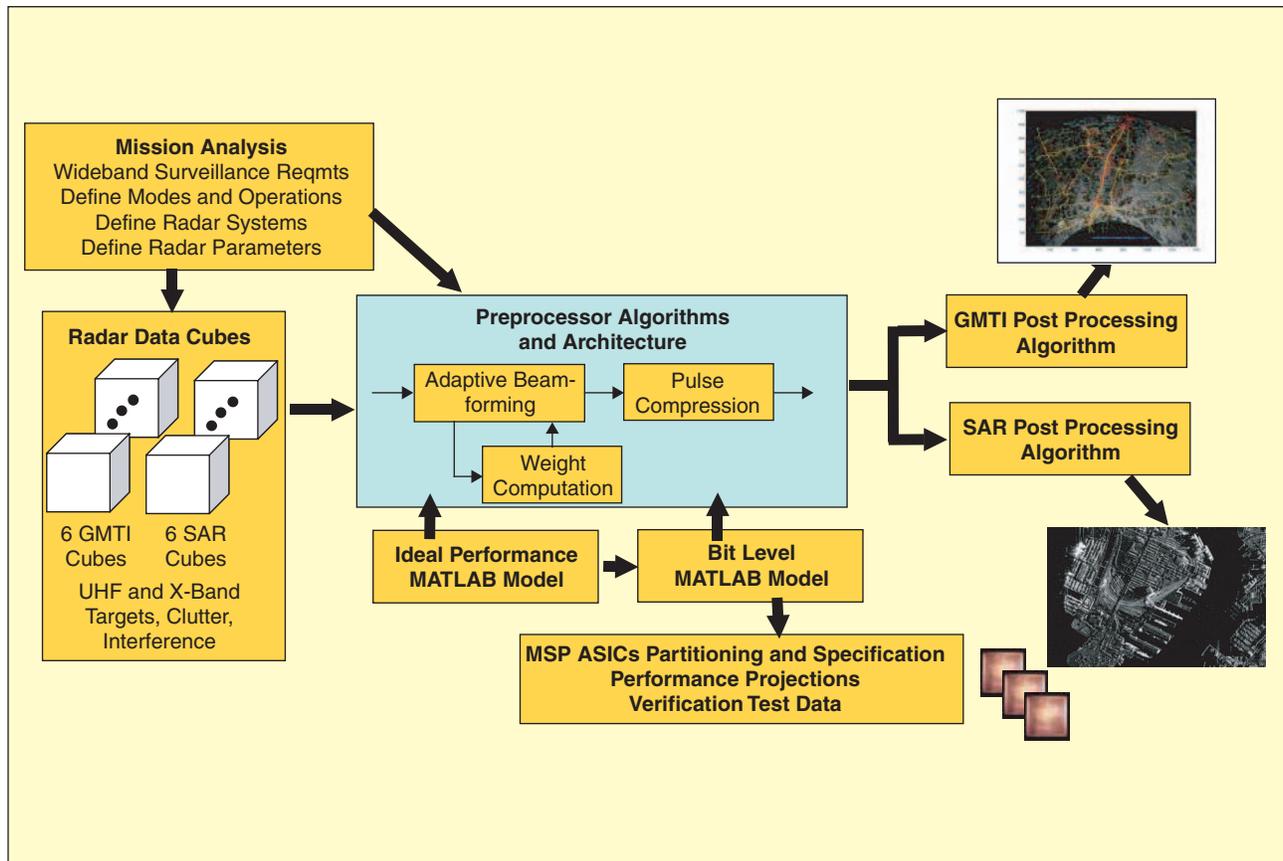
COMPLETE SYSTEM MODELING ENABLES PRECISE BIT-WIDTH TUNING

The culmination of the MSP ASIC technology development was the detailed definition of an MSP chipset and performance projections of a TO/s-throughput radar preprocessor. An end-to-end modeling methodology of the entire radar processing chain was developed and demonstrated as part of this ASIC definition process. Emerging DoD mission requirements were used to drive the processor architecture and MSP cell and device requirements. System configurations, modes, and performance parameters were derived from wideband surveillance programs, and the targeted system addresses synthetic aperture radar (SAR) and ground moving-target indicator (GMTI) surveillance modes in a dual-band radar configuration. A MATLAB model of the radar modes was built and used to derive and optimize the MSP processing architecture and MSP ASICs. The modeling activity consisted of four major parts: a data generator, preprocessor model, SAR post-processing model, and GMTI post-processing model. The end-to-end model with all these elements was integrated as shown in Figure 10.

The model was then exercised using the various data cubes to optimize the system performance. A key use of the model was to determine the required precision and word size needed in the MSP chip set computations. System performance parameters were determined as a function of bit length for different stages of the preprocessing. This allowed the radar performance targets to be met using only the minimum word sizes required, which provides the optimum O/s/W in the chip set.

MSP ASIC PROCESSOR PROVIDES REAL-TIME ADAPTIVE BEAMFORMING ON A BOARD

A three-chip set was defined to perform the wideband adaptive beamforming to support upcoming sensors. Two chips



10. An end-to-end system modeling approach was used to specify ASICs.

form the core of the beamforming function and use MSP cells from the core generator, reduced complexity algorithms, and low power flip flops and I/O to optimize O/s/W. A third ASIC was defined using the optimized reconfigurable MS1 core implemented with the MSP techniques to obtain a low-power, high-performance, on-the-fly reconfiguration capability. A beamformer module was defined for use with a real-time multichannel wideband radar test bed at Northrop Grumman. The MSP beamformer design has the capability to replace seven state-of-the-art custom ASIC boards with a single board, providing nearly a TO/s throughput at a 10× improvement in O/s/W.

ACKNOWLEDGMENTS

This work was supported by Dr. Robert Reuss of DARPA/MTO, under the Mission-Specific Processing Contract NBCHC010038. The authors thank the participating members of the team, most notably Ken Delson, Dave Fry, Rob Kober, Joe Matesic, Steve Shauck, and Jennifer Trotta of Northrop Grumman; Timothy M. Wilson, Jeffrey A. Geib, Eric J. Martina, Ming Zhang, Byonghyo Shim, and Brian Lam of the University of Illinois; Yongtao Wang, Hamid Mahmoodi-Meimand, Hunsoo Choo, Woopyo Jeong, Jong-sun Park, Lihyih Chiou, and Mark Johnson of Purdue University; and S. Safavi, M. H. Lee, A. Niktash of Morpho Technologies.

REFERENCES

- [1] R. Jain, P.T. Yang, B.Y. Chung, C. Chien, L.K. Tan, and T. Yoshino, "FIRGEN: A CAD system for automatic layout generation of high-performance FIR filters," *IEEE Trans. Signal Processing*, vol. 39, pp. 1655–1668, July 1991.
- [2] N.R. Shanbhag, "Algorithms transformation techniques for low-power wireless VLSI systems design," *Int. J. Wireless Inform. Networks*, vol. 5, no. 2, pp. 147–171, 1998.
- [3] J. Park, H. Choo, K. Muhammad, S. Choi, Y. Im, and K. Roy "Non-adaptive and adaptive filter implementation based on sharing multiplication," in *Proc. Int. Conf. Acoustics, Speech, and Signal Processing, 2000 (ICASSP 2000)*, pp. 460–463.
- [4] K. Muhammad and K. Roy, "A novel design methodology for high performance and low power digital filters," in *Proc. 1999 IEEE/ACM Int. Conf. on Computer-Aided Design*, pp. 80–83.
- [5] K. Muhammad and K. Roy, "Graph theoretic approach for design and synthesis of multiplierless FIR filters," in *1999 IEEE Int. Symp. System Synthesis*, pp. 94–99.
- [6] W.H. Mangione-Smith, B. Hutchings, D. Andrews, A. DeHon, C. Ebeling, R. Hartenstein, O. Mencer, J. Morris, K. Palem, V.K. Prasanna, H.A.E. Spaanenburg, "Seeking solutions in configurable computing," *IEEE Computer*, vol. 30, pp. 38–43, Dec. 1997.

Michael Lucas is with Northrop Grumman Electronic Systems in Baltimore, Maryland. Naresh Shanbhag is with the University of Illinois at Urbana-Champaign. Kaushik Roy is with Purdue University in Indiana. Fadi Kurdahi is with Morpho Technologies and UC Irvine. John Fagan with Atmel in Columbia, Maryland. E-mail: mr.lucas@ngc.com. CD ■