

# Innovative processes

# bring systems on chips to life

Full system-on-chip design supplements digital cores with analog circuitry, memories, and, sometimes, RF blocks. Unique silicon-process technology lets these vastly different functions successfully coexist on a silicon chip.

*JIM LIPMAN,  
TECHNICAL EDITOR*

True cell-based system-on-chip (SOC) technology involves more than just placing huge amounts of digital logic on silicon. To implement full electronic systems on chips, you also have to be able to design or place other types of blocks—precision analog, special memories, and, for some designs, RF—on that silicon (**Figure 1**). Good EDA tools and design methodologies are necessary for successful SOC designs. However, the essence of true SOC design is silicon processing that supports multiple design technologies on a single chip.

The processing technologies you need to design technology blocks on silicon require process foundries to modify high-speed digital-logic processes. These modifications—often, add-on process modules—increase the cost of the baseline digital process. Along with adding silicon cost, this approach may lower performance for modules, such as analog blocks, designed with the added processing than you can get with a separate chip fabricated in a process optimized for that block type. However, the cost, power, and speed advantages of SOC designs often outweigh alternative multichip-based systems with individual chips optimized for digital, analog, and memory functions.

#### SOC process support

Most ASIC companies approach SOC process support with modular processing. The companies start with a high-speed digi-



Photo courtesy Samsung



## SYSTEMS ON CHIPS

tal-CMOS process and add process masking and other steps to implement technologies on the digital process. These technologies include precision analog; high-density SRAM, DRAM, and flash memory; and BiCMOS for RF or high-current-drive applications. Of these technologies, analog is the most important add-on for SOC-based systems. Unfortunately, adding analog capability to a digital CMOS also creates process- and circuit-design problems (see sidebar "Tackling analog-induced problems in SOC designs"). A few examples illustrate how semiconductor companies achieve full SOC capability and the impact of adding process modules on chip cost and performance.

Lucent Technologies starts with a 0.25-mm (0.18-mm-effective-channel-length), five-layer-metal digital-CMOS process. The company complements this baseline 300-MHz process with process-enhancement modules that you can use alone or in combination with: precision analog with low-threshold transistors and a precision metal-to-metal capacitor for voltage linearity and good component matching; BiCMOS with a 35-GHz bipolar transistor for RF and linear applications; high-density, six-transistor SRAM with fast access and a maximum 4-Mbit core; 5V flash memory with a maximum 8-Mbit core; and laser-programmable system chip using Chip Express' ([www.chipexpress.com](http://www.chipexpress.com)) laser-programmable-gate-array technology with cell-based blocks (Figure 2).

## @a glance

- c True system-on-chip (SOC) design supplements high-speed digital cores with analog blocks and special memory blocks.
- c The process changes needed to implement SOC are available but add to silicon-wafer cost.
- c Full SOC-process capability lags behind leading-edge digital-only capability by one or two process generations.
- c Because ASIC companies implement SOC processes differently, take the time to fully understand the capabilities of your SOC vendor.

Circuitry designed with Lucent's process modules offers good performance. For example, the analog module supports sigma/delta ADCs with resolution as high as 18 bits, and the BiCMOS module lets you design 1.9-GHz wireless RF blocks. However, you must pay for additional wafer processing. To the basic 17-mask, 0.25-mm logic process, each module adds two (for SRAM) to five (for flash memory) masking steps. Each masking step increases wafer-processing cost by 2 to 5%. This addition translates to approximately 7 and 18% higher processing cost for high-density SRAM blocks and flash-memory cores, respectively. To fill the DRAM gap,

Lucent is acquiring DRAM technology from a technical partner and should have it available by year-end.

Samsung also has a 0.25-mm, five-layer-metal CMOS process with variations for implementing technology blocks with high-speed digital logic. The company starts with a logic process and adds process steps for analog, DRAM, SRAM, and flash-memory design capability. Samsung offers combinations of high-speed logic and high-density SRAM, DRAM, or flash. The basic process (STD120) has one polysilicon layer; merged logic and flash memory (MFL120) adds two poly layers; and merged logic and DRAM (MDL120) adds four poly layers.

A merged-logic/DRAM process requires around 35% more process steps and is one-third more expensive than the baseline logic process. You can get as much as 128 Mbits of DRAM on a chip with I/Os, or words (in ASIC terminology), as wide as 1024 bits. Samsung's flash-memory capability requires 30% more masks and associated process steps. The technology is available for Samsung's 0.35-mm process now, and the company is migrating it to a 0.25-mm process. Analog core designs need one or two additional masking steps—one for metal-to-metal capacitors and a second for poly-to-poly capacitors. Samsung estimates that using its 0.25-mm process provides as much as 1.5 Mbits of SRAM and 0.5 million gates of logic or 64 Mbits of DRAM and 1 mil-

## REPRESENTATIVE SYSTEM-ON-CHIP COMPANIES

For more information on companies discussed in this article offering system-on-chip products, circle the appropriate numbers on the Information Retrieval Service card or use EDN's Express Request service. When you contact any of the following manufacturers directly, please let them know you read about their products in EDN.

Artisan Components  
Sunnyvale, CA  
1-408-734-5600  
fax 1-408-734-5050  
[www.artisan.com](http://www.artisan.com)  
Circle No. 347

LSI Logic  
Milpitas, CA  
1-800-574-4286  
fax 1-408-433-8989  
[www.lsillogic.com](http://www.lsillogic.com)  
Circle No. 348

Lucent Technologies  
Allentown, PA  
1-800-372-2447  
fax 1-610-712-4106  
[www.lucent.com](http://www.lucent.com)  
Circle No. 349

NEC Electronics  
Santa Clara, CA  
1-408-588-6000  
fax 1-408-588-6130  
[www.nec.com](http://www.nec.com)  
Circle No. 350

Samsung Semiconductor  
San Jose, CA  
1-800-423-7364  
fax 1-408-544-4980  
[www.sec.samsung.com](http://www.sec.samsung.com)  
Circle No. 351

Silicon Access  
San Jose, CA  
1-408-441-7390  
fax 1-408-441-7370  
[www.siliconaccess.com](http://www.siliconaccess.com)  
Circle No. 352

### VOTE . . .

Please also use the Information Retrieval Service card to rate this article (circle one):  
High Interest 586  
Medium Interest 587  
Low Interest 588

### Super Circle Number

For more information on the products available from all of the vendors listed in this box, circle one number on the reader service card. **Circle No. 353**



## SYSTEMS ON CHIPS

lion gates of logic on a chip. Using the 0.25- $\mu\text{m}$  process also provides a maximum flash-memory core size of 32 Mbits. Samsung also estimates an upper frequency limit of 0.25 to 0.5 GHz for mixed-signal blocks on its process, allowing you to implement some lower frequency RF functions on chip.

Taking the same approach as Lucent and Samsung, NEC supplements its 0.25- $\mu\text{m}$  digital-CMOS process with additional processing for analog, flash-memory, BiCMOS, and DRAM support. For high-speed analog, NEC uses two additional masks to get precision capacitors and polysilicon resistors. Flash memory requires two more masking steps. Adding bipolar transistors to create a BiCMOS process for RF cores requires four more masks. The BiCMOS process lets you design I/Os with speeds as high as 2.4 GHz. NEC's DRAM-process module appears more complicated than Samsung's. NEC doubles the number of masks used for pure logic chips to add DRAM to the logic process, resulting in a 60 to 80% cost increase.

You can use LSI Logic's 0.18- $\mu\text{m}$  (0.13- $\mu\text{m}$ -effective-channel-length) process, the G12, for designing RF blocks on a chip. Transistor  $f_t$ 's greater than 70 GHz for the n-channel transis-

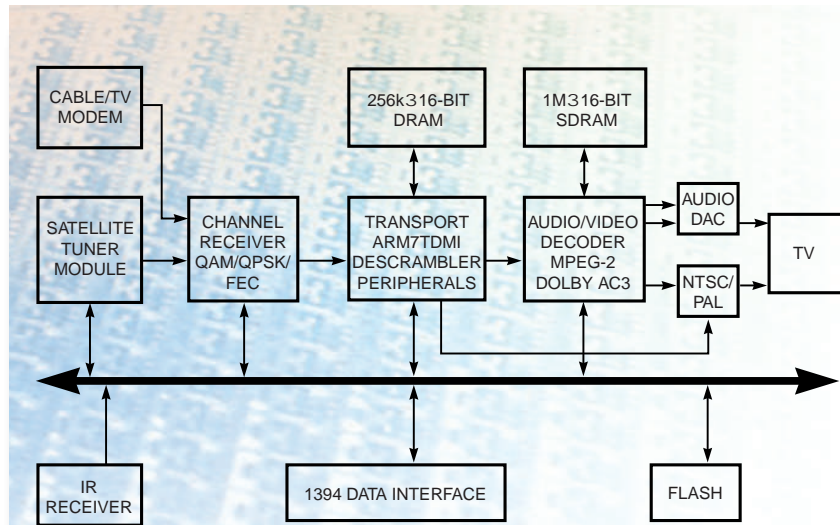


FIGURE 1

This set-top decoder box includes a 32-bit RISC core (ARM7TDMI), high-speed digital logic, mixed-signal blocks, flash memory, and DRAM. Implementing such a system on chip requires a process that can support all these technologies (courtesy Samsung).

tor and greater than 40 GHz for the p-channel transistor, low-K dielectric to reduce layer-to-layer capacitance, and specialized core-isolation techniques

with more than 110-dB isolation let LSI realize RF cores on CMOS SOC designs. LSI also has added a one-transistor DRAM technology to its 0.25- $\mu\text{m}$  G11

## TACKLING ANALOG-INDUCED PROBLEMS IN SOC DESIGNS

Adding analog-block design to a high-performance digital CMOS chip creates problems for both process engineers and SOC designers. These problems involve creating precision-analog passive components and making an "analog-friendly" SOC environment (Reference A).

Certain types of analog circuits require precision resistors. For example, some DACs need linear resistors for current-to-voltage conversion, whereas active filters require resistors with accurate values and high linearity. To get these resistors on chip, a process must either offer access to resistors masked from the process' transistor-gate layer (polysilicon or silicide) or add a polysilicon layer dedicated to resistor use. Switched-capacitor blocks, such as sigma-delta modulators and continuous-time filters, require floating capacitors. A digital process with added analog capability can produce these capacitors either through an additional polysilicon layer or by forming capacitors between two or more metal-interconnect layers.

For protecting analog blocks from digitally induced noise, an SOC process must isolate analog and digital blocks. Processes achieve this isolation by introducing extra distance between the blocks, which wastes silicon area; placing an

oxide barrier between them; or surrounding them with heavily doped isolation rings. Analog blocks are more susceptible than digital cores to thermal problems. Good analog mixed-signal design techniques include balanced analog designs and the isolation of analog blocks from high-heat-producing digital blocks, such as clock generators or output buffers. The use of low-K dielectrics in a mixed-signal process can reduce heat generation and layer-to-layer crosstalk. A move from aluminum to copper metallization, still in the early stages for ASIC vendors, results in lower total-chip power dissipation and smaller on-chip temperature rise, reducing analog-component drift (Reference B).

### References

- A. Liu, E. C Wong, Q Shami, S Mohapatra, R Landy, P Sheldon, and G Woodward, "Complete Mixed-Signal Building Blocks for Single-Chip GSM Baseband Processing," 1998 Custom Integrated Circuits Conference Proceedings, May 11 to 14, 1998, pg 101.
- B. Lipman, Jim, "How copper-topped silicon creates faster, longer-lasting chips," *EDN*, June 4, 1998, pg 48.



## SYSTEMS ON CHIPS

process and will carry the technology to the G12 process. For analog blocks, the G12 uses poly resistors, a technique similar to NEC's analog-resistor implementation.

### Specialized RAMs go on chip

Most ASIC libraries include some SRAM capability, either with fixed blocks or by using an SRAM compiler (Reference 1). ASIC companies either internally develop the technology or get it from the a third-party embedded-memory vendor through a purchase or a technology agreement. Embedded-SRAM vendor Artisan offers process-specific memories tuned to a customer's process that represent the high end of third-party embedded SRAMs. For 0.25-mm processes, Artisan's HS300 single- and dual-port SRAMs operate as fast as 300-MHz (worst case) with a 64-kbit SRAM rated at 500 MHz under typical conditions. The recently announced single- and dual-port HS500 family for 0.18-mm processes operates as fast as 500 MHz (worst case) and 850 MHz (typical). Artisan also offers the LP133 low-power SRAM family targeting 0.25-mm processes; a 64-kbit, single-port device dissipates 0.2 mW/MHz at 1.8V. These embedded SRAMs are expensive; expect to pay around \$500,000 for each SRAM generator (one target process) and supporting materials.

Although some semiconductor companies combine DRAM processing with digital processes, start-up Silicon Access recently announced DRAMatic, a one-transistor, high-density embedded-DRAM core (Reference 2). DRAMatic cores come in a range of word widths and depths (Figure 3). The high-I/O-bit cores are interesting for SOC applications because they provide higher memory-access speed because they don't have to go off chip to interface with logic. Single-chip logic and DRAM also result in lower power dissipation than having the two functions on separate chips. Common logic and DRAM eliminate many of the pins needed for off-chip DRAM support. The resultant SOC's reduced pin count can substantially lower chip-package cost.

You can use DRAMatic for a merged DRAM/logic chip either by adding processing steps to a foundry's logic process

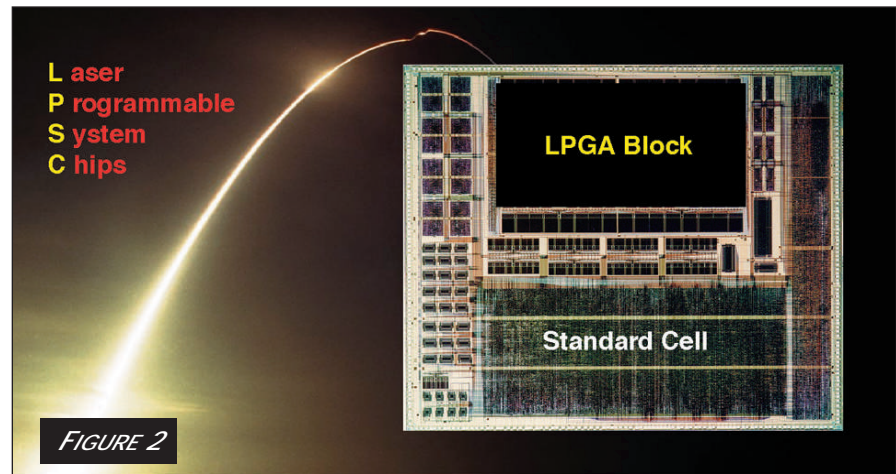
or by using a DRAM process and adding metal layers to support digital logic. Most SOC designs strive for maximum logic speed, so you may end up using the high-speed logic process with added DRAM processing because the baseline DRAM process reduces logic performance. The additional processing for DRAMatic adds approximately 10 mask layers and costs about 20 to 40% more than digital-logic-only processing at 0.25 mm. Silicon Access offers DRAMatic as a core for a targeted application or as a DRAM compiler. Core prices are approximately \$100,000 to \$150,000; the company has not yet set compiler prices.

ry available for these digital-core-based programmable-logic devices.

If programmable-logic vendors can overcome the technical challenges of placing analog or mixed-signal cores on FPGAs, the demand for such devices will probably remain low. On processes with similar dimensions, programmable chips are slower than cell-based chips. Because of this performance gap, cell-based designs remain the implementation vehicle for your SOCs for the foreseeable future.

### Choosing an SOC process

If you're ready to jump into a full-blown, mixed-signal SOC design, you



Laser-programmable system chips combine cell-based blocks with laser-programmable gate-array modules for fast-turnaround chip evaluation.

### Can SOC designs be FPGA-based?

Leading-edge programmable-logic chips can contain hundreds of thousands of logic gates, so you might wonder whether any FPGA- or CPLD-based SOC designs exist? The answer is "no" within the context of this article's definition of SOC: chips with large amounts of digital logic, analog cores, and different memory blocks. Many of the leading programmable-logic companies, including Actel ([www.actel.com](http://www.actel.com)), Altera ([www.altera.com](http://www.altera.com)), and Xilinx ([www.xilinx.com](http://www.xilinx.com)), have aggressive programs with third-party core providers for both soft and hard cores on FPGAs or CPLDs. However, these programs now include only digital cores. Furthermore, SRAM is the only memo-

have a lot of work to do before starting your chip. Along with selecting design tools, design methodology, cell libraries, and intellectual-property cores, you also have to determine which process to use to implement your design.

CMOS digital-logic capabilities among ASIC vendors for a process with similar dimensions are similar. Most performance variations result from differences in library-cell architecture and topology and in core design. You may also find speed and power differences among processes, resulting from one process having more interconnect layers than or slightly different metal pitches from another process. More metal layers or a tighter metal pitch



## SYSTEMS ON CHIPS

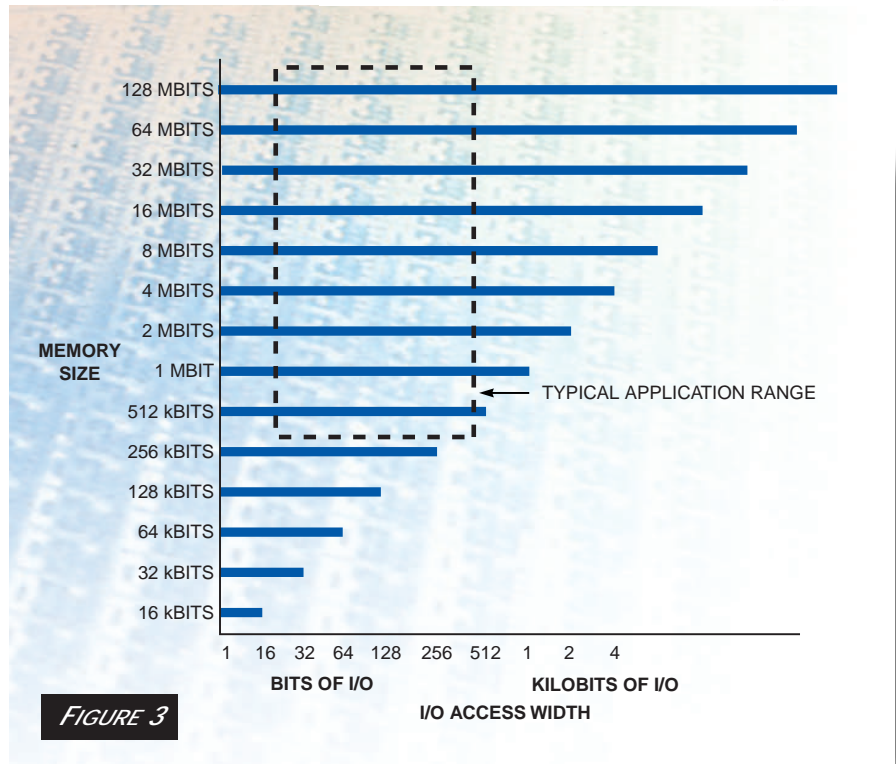


FIGURE 3

Silicon Access' DRAMatic core gives you a range of DRAM capability when you use it in a DRAM-modified, high-speed silicon process.

helps compress chip routing and can result in a higher performance device. Analog-block implementation, however, is a different story.

Although ASIC vendors all have similar digital capability in their comparably dimensioned processes, each has its own way of adding analog capability to a baseline digital process. Vendor differences in adding analog capability to processes require your making sure that a vendor's analog capability is sufficient to meet the analog requirements of your SOC design. This determination is important, particularly when you have to also consider how the process handles analog/digital isolation on chip.

Finally, determine your SOC's memory-technology requirements, and find a vendor with a process offering all needed technologies. Determine the vendor's memory compiler or block-size ranges and memory-configuration bounds, and then partition your design to meet these limits. Remember that maximum memory and maximum logic do not coexist on the same chip.

Verify that the ASIC vendor can meet your digital and memory needs on one piece of silicon. e

## References

1. Lipman, Jim, "Not just your basic ASIC libraries," *EDN*, June 5, 1997, pg 52.
2. Dipert, Brian, "Embedded memory: the all-purpose core," *EDN*, March 13, 1998, pg 34.



You can reach Technical Editor Jim Lipman at 1-925-606-1370, fax 1-925-606-1563, [ednlipman@mcimail.com](mailto:ednlipman@mcimail.com).

## VOTE

Please use the Information Retrieval Service card to rate this article (circle one):

High  
Interest  
586

Medium  
Interest  
587

Low  
Interest  
588