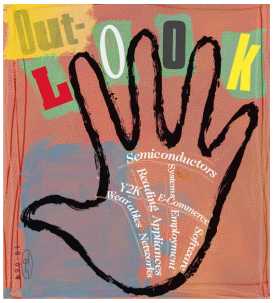

Semiconductor Research Corporation

Taking Moore's Law Into the Next Century

"Cooperation" is not a word associated with the competitive world of semiconductor manufacture. Yet one organization not only gets all the major players to the table, it also helps set the course for the chip industry's future.

Scott Hamilton
Computer

Far from Silicon Valley, amid the green and rolling hills of North Carolina, is an organization that propels the multibillion dollar semiconductor industry. It's one of the few organizations to get fierce competitors like Intel, Motorola, Texas Instruments, and IBM to the same table, let alone cooperate. These companies and many others contribute fees of about $37 million annually to fund future research to keep the engines of semiconductor production churning.

In short, the Semiconductor Research Corp. is bent on taking Moore's law into the next century.

PRECOMPETITIVE RESEARCH

The SRC was formed in 1982 by the Semiconductor Industry Association with the unique charter of coordinating academic research to meet industry's needs. Its research aims at technology that will support Moore's law in the next five to eight years. SRC pools funds from its 30 member companies to sponsor "precompetitive" research and makes that research available to all members.

SRC is unusual for an industry consortium in the depth to which it manages research projects (see the "Managing R&D" sidebar). Without SRC, the US semiconductor industry would have a much less coordinated strategy for initiating and managing precompetitive R&D.

Together with its sister organization, Sematech, SRC helps chart and promote the continued progress of the semiconductor industry. (See the "Spectrum of Semiconductor Research" sidebar.) Their main guide for doing so is the *National Technology Roadmap for Semiconductors*.[1]

CHARTING MOORE'S LAW: THE SIA ROADMAP

Moore's law posits that microprocessor performance (as defined by the number of transistors on a chip) will double every 18 months—a 58 percent compounded annual growth rate.[1] Historically, the semiconductor industry has kept pace by

- continuously shrinking feature size to increase the number of transistors on a chip, and thus
- increasing the speed of the circuits.[1]

To ensure that the industry focuses on the technical challenges to continued productivity, the SIA publishes its biannual *Roadmap*. This document analyzes industry progress in process and design technologies, charting research directions necessary to meet specific goals. Overall, this *Roadmap* calls for the exponential scaling of feature sizes shown in Table 1.[1,2] In developing the 2012 targets, the SIA constrains projections so interim targets are believed to be attainable. These targets assume that technology barriers will be overcome in a timely manner.

Beyond 2006, physical barriers ultimately include atomic properties that will come to the fore with aggressive device shrinkage, as well as capacitance, crosstalk, soft errors, and other effects I discuss later.

More immediate concerns include the approaching limitations of lithographic techniques. Current techniques use deep-ultraviolet (DUV) light with a 193-nm wavelength. DUV exposure tools will reach their physical limits as feature sizes approach 100 nm.[1]

Potential technology barriers—or *grand challenges*, as the *Roadmap* calls them—include *the ability to continue affordable scaling*, which directly involves two

0018-9162/99/$10.00 © 1999 IEEE

Table 1. Technology roadmap for semiconductors.

	Year						
Characteristic	**1997**	**1999**	**2001**	**2003**	**2006**	**2009**	**2012**
Process technology (nm)	250	180	150	130	100	70	50
No. of logic transistors (millions)	11	21	40	76	200	520	1,400
Across chip clock speed (MHz)	750	1,200	1,400	1,600	2,000	2,500	3,000
Die area (mm^2)	300	340	385	430	520	620	750
Wiring levels	6	6-7	7	7	7-8	8-9	9

areas addressed in this article: interconnect and design and test.

INTERCONNECTS

The interconnect problem has to do with the signal propagation delay over interconnect lines as gates continue to shrink (becoming faster) and the chip area continues to grow. As Doug Matzke observed, "the percentage of the die that can be reached in a few clock cycles is decreasing at an alarming rate" with each new processor generation.[3] The *Roadmap* goes so far as to say, "Interconnect has been represented as the technology thrust with the largest potential technology gaps."[1]

Interconnect also adds a new dimension to design complexity. As interconnects also shrink and come closer together, previously negligible physical effects like crosstalk become significant. Interconnects are a problem, but are they a physical barrier to continued process shrink? Not according to science area directors and program managers at SRC.

The next couple of process generations will contain some fairly revolutionary changes. Companies such as IBM, Texas Instruments, and Motorola are already introducing copper interconnect to replace aluminum. Copper has an advantage over aluminum in terms of its electromigration properties (the self-diffusion of metal along the interconnect caused by the current flow through the metal), so designers can increase current densities and lower the height of the interconnect cross-section, enabling further device shrinkage and reducing the lateral, or adjacent line, capacitance.

Companies are also looking at low-κ dielectric (insulating) materials, and the process tools to implement copper *and* low-κ material are already in place. Interlevel dielectrics with lower permittivity (low-κ) reduce parasitic capacitance and crosstalk between levels of interconnect.

Interconnect myth and reality

According to James Hutchby, director of SRC's Interconnect Sciences, interconnects will likely present a physical barrier in some applications, particularly global interconnect. Pessimists maintain that even new copper and low-κ dielectric process technologies are not panaceas or solutions. They will merely delay the onset of serious barriers by one or two generations. At the generations defined by perhaps a 100-nm feature size—and certainly at 70 nm—pessimists say interconnect design will face serious physical barriers.

The reality is more complex, according to Hutchby. To understand, though, you must consider local and global interconnect components.

Local interconnects. A local interconnect system serves a particular module or macro-function on a chip—processor, logic, cache, data management, and so forth—and can contain lines as much as a few millimeters long. As technology scales from generation to generation, Hutchby contends that propagation delay for local interconnects will scale down as fast as the transistor or gate propagation delay. So the local interconnect delay will not be the factor limiting data rates.

Global interconnects. Global interconnects, on the other hand, are one or two centimeters long (length equal to a die's linear dimension) and connect major

Managing R&D

SRC research is designed to provide results that will be of value to its members. It is therefore more hands-on in terms of managing research projects. Member-based technical advisory boards review all contracts at the task level and recommend which projects to sponsor or discontinue. The SRC and its members review research projects annually, providing detailed feedback as to the research directions required by industry. In addition, the SRC's research customization program lets member companies sponsor targeted research—selected researchers and tasks—using a portion of their member dues, thus enabling members to further influence academic research directions.

At the same time, according to Lawrence Arledge, program manager for Design Sciences on assignment from Texas Instruments, there exists a fruitful tension—rather than antagonism—between the technical advisory boards and faculty researchers. Although subject to review, researchers also can lobby for research they want to pursue. It's up to researchers to convince the SRC and its members that a proposed project is one they can use.

Spectrum of Semiconductor Research

The Semiconductor Industry Association was founded in 1977 primarily as a trade association representing the industry in tax and trade issues. Recognizing the continuing need for advanced research in the face of reduced government funding, the SIA has formed a number of wholly owned subsidiaries that pursue R&D on behalf of its members:

- *Semiconductor Research Corporation (SRC)* for applied research in microelectronic design and process technologies,
- *Sematech* in the tool arena, and
- *Microelectronic Advanced Research Corporation (Marco)* for revolutionary solutions to technical problems.

The research continuum represented by these subsidiaries is shown in Figure A.

Sematech

Sematech is a SIA subsidiary formed in 1987 that is much more closely aligned with suppliers and member companies, and focuses much more strongly on current industry needs. Sematech was for a long time a dues-paying associate member of SRC, and the two organizations have a long-term research management relationship dating back to Sematech's founding. However, in 1996 Sematech decided to no longer accept federal funding and has opted instead to become an SRC strategic partner, meaning the strategic planning processes are coupled and SRC actively transfers technology to Sematech.

Marco

At the request of the SIA, SRC formed a wholly owned subsidiary called Marco to operate the Focus Center Research Program. The FCRP explores revolutionary solutions to impending physical barriers with projects targeted to impact industry 10 to 15 years down the road.

SIA formed Marco to continue the long-term research that used to be the province of industrial research laboratories. The FCRP initiative will add another $50 to $60 million on top of SRC-funded research. So instead of $40 to $50 million, the industry will eventually invest $90 to $100 million across grand-challenge problems, such as design, interconnect, lithography, and so on.

Marco also has a slightly different constituency than SRC. Although funded primarily by the SIA (50 percent) and members of SEMI/Sematech (25 percent), it receives a heavy dose of government funding (25 percent).

Focus centers are large, distributed university research centers devoted to a particular research problem, each looking at the 8 to 15 year timeframe. Marco intends focus centers to create the vision of where the semiconductor industry wants to be in 10 to 15 years, and industry will have an arms-length relationship with the focus center universities. Two research areas initially identified are design and interconnect, with UC Berkeley and a dozen other universities and Georgia Tech and half-dozen universities spearheading these efforts.

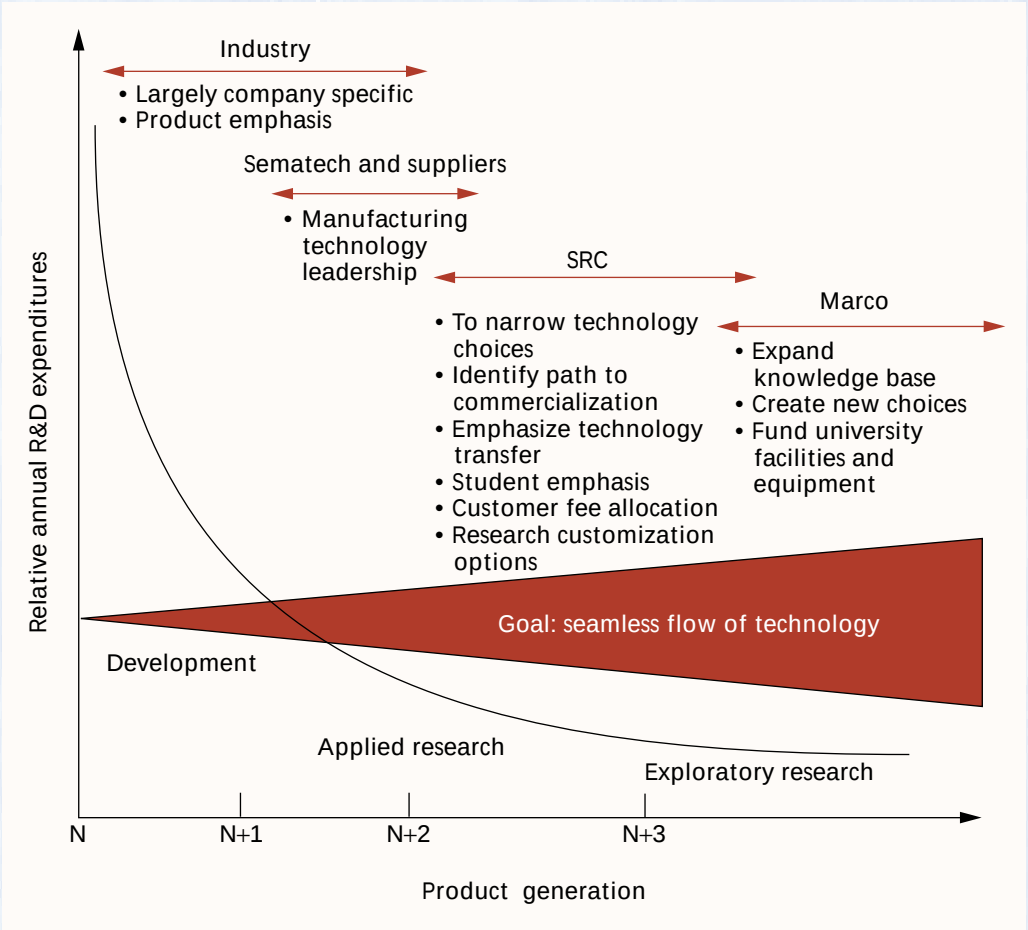

Figure A. Each of several organizations focuses on research at different points along the adoption curve.

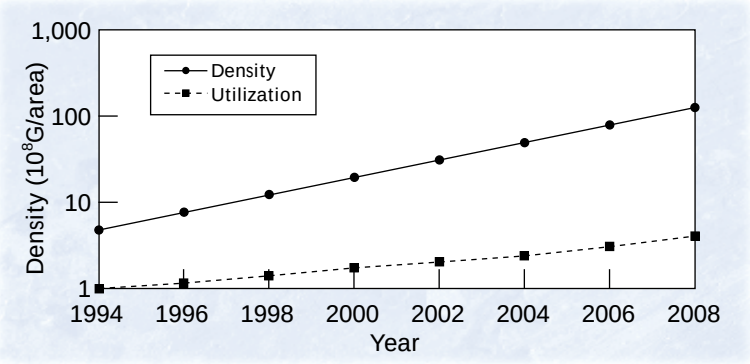

Figure 1. Technology growth is outstripping design productivity.

functional areas on the chip—logic with memory, for example. More of a delay problem exists as feature size scales down and chip size scales up. Jason Cong's simulations of NTRS data at the 70-nm level suggest that delays on local interconnect will decrease by more than 50 percent, whereas delays on unoptimized global interconnect will increase by 150 percent (from 2 ns to 3.5 ns).[2]

But do major, showstopping barriers arise at 100 nm? Maybe not, says Hutchby, as there are some immediate workarounds. First, designers are using buffer amplifiers as repeaters to reduce the signal propagation delay in global interconnect lines. Current optimization techniques such as optimal wire sizing, buffer insertion, and simultaneous device and buffer sizing will continue to be applicable, although as feature size continues to shrink, the interconnect itself becomes complex circuitry in its own right.[2] Second, designers are exploring and even implementing architectural innovations that relocate functional areas to shorten global interconnects.

Furthermore, as feature size dips below 100 nm to 70 or 50 nm, alternative means of providing global interconnect arise. Designers can perhaps translate the schemes they use to communicate clock and signals *among* chips to communicate clock and signal *within* a chip; similarly, a certain portion of on-chip global interconnect could be offloaded to the package.[4] New communication protocols permitting asynchronous or differential signaling might also obviate the delay problem.[2] Researchers are exploring 3D architectures, which allow active interconnect systems realized in two and three dimensions.[1,5] Finally, fundamentally new technologies like optical interconnect or on-chip RF connections might also emerge.[2]

Crosstalk

Delay isn't the only problem designers are running into: Crosstalk is increasingly becoming a problem, says Justin Harlow, director of Integrated Circuits and Systems. If you look at a cross-section of earlier-generation interconnect wires, they're much flatter, and their edge-to-edge or adjacent wire capacitance is very small compared to the interconnect's vertical capacitance.

Today, interconnects are becoming taller and narrower, and are also placed much closer together. As a result, the majority of capacitance is lateral or between wires so there is a lot of coupling that never used to be a problem. These effects are quite subtle and can cause logic errors. For instance, you can have a signal pulse traveling down one interconnect wire at the same time a neighboring wire is supposed to be quiet. The capacitance from the active line couples the signal to the supposedly inactive one, possibly causing the latter to reset a latch. So suddenly the system possibly implemented an incorrect function.

Thus, interconnect design must change because designers and design tools have not previously had to focus on designing for these distributed effects. Delay and crosstalk have always existed, but the degree of capacitive coupling between adjacent wires has increased substantially in the last three process generations.

DESIGN CHALLENGES

Aside from challenges posed by physical laws, the semiconductor industry must also find solutions to problems in the processes for design and test themselves.

Designing for complexity

Traditionally, the industry has kept pace with Moore's law by throwing more resources at the problem. Figure 1 shows the industry's historical 58 percent annual growth in technology progress versus a 21 percent annual growth in design productivity.[6] What the graph doesn't show is that this productivity growth results largely from the rapid growth of design teams—now commonly 250 to 300 people—whose size is rapidly becoming unmanageable. The *Roadmap* signals a "consensus among the semiconductor design technology working group members… that a crisis is approaching."[1] The industry will need new design methodologies and tools.

Integrated design. The industry is just now emerging from the 1980s version of system design, which abstracted away the circuit-level details. Currently, there is so much interdependence among interconnects, layout, architecture, and the circuitry's performance and reliability that designers can't abstract away these issues anymore. Because everything is now layout dependent, designers must predict what the electrical performance will be on the chip and how subsequent optimizations will affect this performance.

Chip design has become so complex that designers need more education, experience, and exposure to a broad range of fields—device physics, wafer processing, analog effects, digital systems—to understand how all these aspects come together. Successful deep submicron designs will take a different breed of designer—a combined computer architect and physicist—and we are not producing them in sufficient numbers, according to Harlow. (See the "The SRC, Education, and Technology Transfer" sidebar.)

Design tools. For the same reasons, designers need smarter tools that comprehend distributed effects like crosstalk. Early logic synthesis tools took a logic-level description and translated it into a gate-level description. When fed into physical design tools, that description would generate a circuit that came within about five percent of the area or speed of a manual design. These partitioned tools enabled a real increase in productivity. By partitioning tool sets, the industry was able to advance the tools and enhance design productivity and throughput.

Now, however, as industry needs to capture that last few percent, it's almost as if these tools must come back together. The next generation of tools must bridge the traditional partitioning to implement techniques that advanced designers produce. These include

- *synthesis tools* that take into account placement, layout, and timing information;
- *system-level verification tools* to demonstrate the correctness of an entire chip design; and
- *electrical simulation and modeling tools* that account for increased noise and coupling in today's designs, and at the same time abstract those models so they can be fed back into higher levels of design.

Above all, these new tools must permit forward and backward interaction among design levels and tools. New tools must pass information from, say, the synthesis tool to the layout tool—for example, that two interconnects may affect each other, things that can be determined statically but also dynamically.

Synthesis tools—whether at the high, behavioral, or logic level—must also handle such issues as power constraints, performance, test insertion, test constraints, reliability, and so on. A major advance, according to William Joyner, director of SRC's Computer-Aided Design and Test Sciences on assignment from IBM, would be a straight-through, noniterative design flow where each layer of tools would allow tools later in the flow enough flexibility to adjust designs to meet performance and other requirements.

Formal verification of higher, system-level design uses mathematical techniques to prove the correctness of circuit functionality. Formal methods are now transitioning from academia to industry. The verification challenge becomes even greater as intellectual property reuse becomes an important design strategy, enabling designers to build increasingly sophisticated circuits without continuing to throw human resources at the problem. Of course, the growing emphasis on analog and mixed signal design raises some interesting challenges for formal verification given that today's formal methods tools are digitally focused.

There has already been some good academic research on streamlined models for predicting interconnect performance that might fit well in the overall design strategy. Systems have become so complex and contain so many variables that industry needs streamlined models to channel the design effort in the right direction.

SYSTEM DESIGN: THE FINAL FRONTIER

SRC is expanding its role in system design. Architecture goes beyond traditional electronic design to the building of complete systems. Its members are looking at design for application-specific ICs, and RF, analog, and MEM (microelectromechanical) systems, where there is a dire need for design tools. Although traditional computer engineering has effectively explored various processor architectures, closely

The SRC, Education, and Technology Transfer

SRC's charter is to provide precompetitive research to its members. This means providing both research results *and* technical talent, and SRC takes the latter responsibility very seriously. Since its inception, more than 2,000 PhD students have participated in SRC projects, more than half have gone on to join member companies, and a quarter have joined nonmember companies.

By coupling SRC and university research, students are trained to think critically about problems—where the state of art is, where the state of knowledge is. Every research project has one industry mentor to work with the student and professor, and companies give students access to company resources such as advanced testing equipment and fabs, as well as internships. Researchers see great benefit to their students in helping them focus their research and helping them prepare their students for industry jobs. The biannual Techcon (university focused) conference exposes the students' work to members and, according to William Joyner, is quite impressive, with the quality of the presentations every bit as good as at major technical conferences. The SRC member Web site includes human resources pages where participating students and member companies exchange information.

Semiconductor Research Worldwide

Several organizations outside the US are also working on semiconductor research. Here are a few we found information on.

Japan

Three consortiums are developing next-generation technology:

- The Semiconductor Technology Academic Research Center (http://www.starc.or.jp) performs basic research in silicon semiconductors. Similar to the SRC, it organizes research at universities.
- Semiconductor Leading Edge Technologies Inc. (http://www. selete.co.jp) focuses on semiconductor manufacturing equipment and materials.
- The Association of Super-Advanced Electronics Technology (http://www.aset.or.jp) focuses on semiconductors and other complementary technology, like magnetic recording media and displays.

Europe

In Europe, Esprit sponsors several semiconductor-related projects. Medea, a program in Micro-Electronics Development for European Applications (http://www.medea.org) sponsors a project in CMOS-based technology platforms.

Taking the Roadmap International

Beginning this year, the SIA has begun soliciting input from semiconductor research organizations outside the US. Their goal is to make the next edition of the *Roadmap*—due out in 1999—international. Paolo Gargini, Director of Strategic Research at Intel, will direct that effort.

related issues haven't received the same attention. Local clocking and asynchronous design, for example, are gray areas where SRC can investigate technology alternatives. Moreover, design now depends on physical properties as we have seen, and SRC looks to foster meaningful interaction between process technologists and design architects.

SRC sees another promising area in 3D process architectures. 3D integration could lend itself to systems architecture, so designers can collocate different kinds of circuits—memory and logic, for example—using this third dimension. This would have a tremendous impact on design and provide whole new ways of organizing systems on a chip. This is the realm SRC has been exploring in systems architecture—and it intends to continue pushing out in those realms off the roadmap.

At the moment, chip designers aren't constrained by the number of transistors they have to play with—the industry is well on its way to a billion transistors on a chip. Now, in an odd twist of fate, designers appear limited by the very tools and processes that made those billion transistors a reality in the first place.

But most of the researchers I talked to relish the seemingly insurmountable challenges that physics lays before them. The industry has bested predictions of doom before, they say, and there's no reason it won't continue to do so—*if* it remains focused.

So that's why researchers in a corner of North Carolina quietly continue to train people, write road maps, and generate new ideas. They're focused on a mission: to take the semiconductor industry—and Moore's law—into the next century. ❖

Acknowledgments

I thank Lawrence Arledge, Ralph Cavin, Justin Harlow, James Hutchby, William Joyner, Richard LaScala, and Sonny Maynard for the time they devoted to this interview. I also thank Ken Sakamura and Dante Del Corso for their input.

References

1. *The National Technology Roadmap for Semiconductors*, Semiconductor Industry Association, San Jose, Calif., 1997; http://notes.sematech.org/ntrs/Rdmpmem.nsf.
2. J. Cong, "Challenges and Opportunities for Design Innovations in Nanometer Technologies," *SRC Tech. Report*, Research Triangle Park, N.C.; http://www.src.org/areas/design.dgw.
3. D. Matzke, "Will Physical Scalability Sabotage Performance Gains," *Computer*, Sept. 1997, pp. 37-40.
4. B. Atkins, "Interconnect Generations Beyond the Final Shrink Limit of 0.007µm," *SRC Tech. Report*, Research Triangle Park, N.C.; http://www.src.org/areas/design.dgw.
5. K.L. Wang and W.T. Lynch, "Opportunities and Challenges for Si-Nanoelectronics in the Next Decade, *SRC Tech. Report*, Research Triangle Park, N.C.; http://www.src.org/areas/design.dgw.
6. "Design Sciences TAB Physical Design Task Force Report," *SRC Tech. Report*, Research Triangle Park, N.C., 1997; http://www.src.org/areas/design.dgw.

Scott Hamilton is senior acquisitions editor for Computer.

Contact him at s.hamilton@computer.org.