The Tides of EDA

Alberto Sangiovanni-Vincentelli
University of California at Berkeley

Alberto bases this article on remarks from his invited keynote speech at the 40th Design Automation Conference. In that speech, he proposed a bold initiative in electronic design automation, EDATech, which would drive the industry much like SEMATECH propelled semiconductor manufacturing.

—Rajesh Gupta, Editor in Chief

IT WAS A GREAT HONOR to celebrate the 40th anniversary of the Design Automation Conference with a keynote lecture intended to place in perspective the most relevant research results presented at DAC in all these years and to identify trends and challenges for the future of electronic design automation (EDA). Going through the DAC proceedings, hunting for relevant papers, was indeed a formidable task. While doing so, I wanted to find regular patterns that those of us in EDA could use to understand the history of our field and its destiny.

Fishing into my background as a former student of an Italian classical Lyceum, I found unexpected help from a fellow countryman: Giovan Battista Vico, the first philosopher who analyzed history and its patterns from a philosophical point of view. He wrote his masterpiece, *Scientia Nova*, in 1650. Vico's fundamental contribution was that history repeats itself with a regular spiral-like pattern ("corsi e ricorsi storici").

Vico identified three phases in mankind's history: the age of gods, the age of heroes, and the age of men. The age of gods is characterized by knowledge that comes to people from the use of their senses. In this respect, events and natural phenomena are inexplicable (for the most part) and attributed to "external" entities, like the ancient gods.

The age of heroes is characterized by the use of imagination that lets people supersede the sensory information to find the first abstract interpretations of reality. It is the age of creativity, the foundation of great human achievements.

The following period, the age of men, is characterized by reason—rational analysis that dissects events, during which people fear novelty and creativity as jumps into the dark because no analysis can guarantee any initiative's success. Vico identified the age of men as the beginning of decay in society. Yet he found that after the decadence of this period, mankind would again loop through the three stages, returning to the next age of gods.

Ages of EDA

Quite surprisingly, I found that I could identify similar patterns in the history of EDA. To determine the time span of the ages of EDA, I used attendance data from DAC, shown in Figure 1.

In this diagram, we see an initial period, from 1964 to 1978, where attendance was fairly low. Next comes a period—from 1979 to 1992—of great prosperity with a very sharp increase in participation. Finally, from 1993 to today, there is relative stagnation and a marked decrease in the last few years. It seems natural to associate these three periods with the ages of gods, heroes, and men as outlined by Vico!

Another interesting set of statistical data relates to paper submissions and their origin, as shown in Figure 2.

In this diagram, the age of gods corresponds to a prevalence of industry papers; the age of heroes, to a marked increase in academic papers; and the age of men, to a predominance of academic papers.

Age of gods (1964 to 1978)

In this period, industry pioneers laid the foundations

0740-7475/03/$17.00 © 2003 IEEE Copublished by the IEEE CS and the IEEE CASS

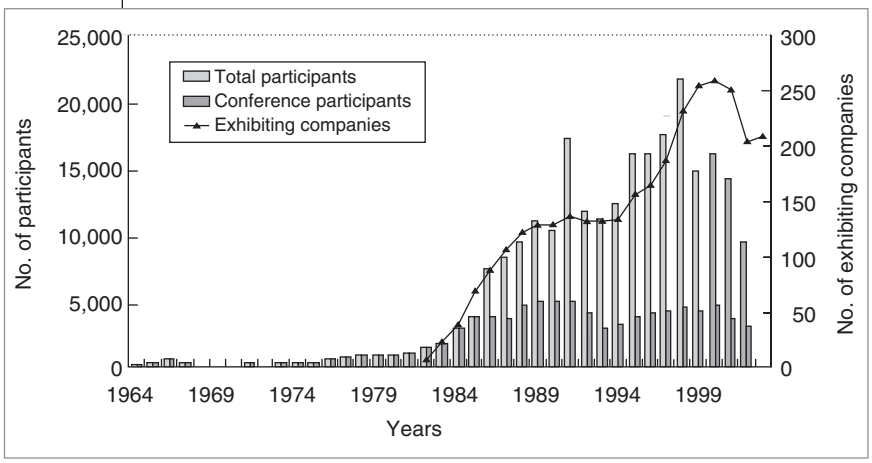

Figure 1. DAC attendance from 1964 to 2003.

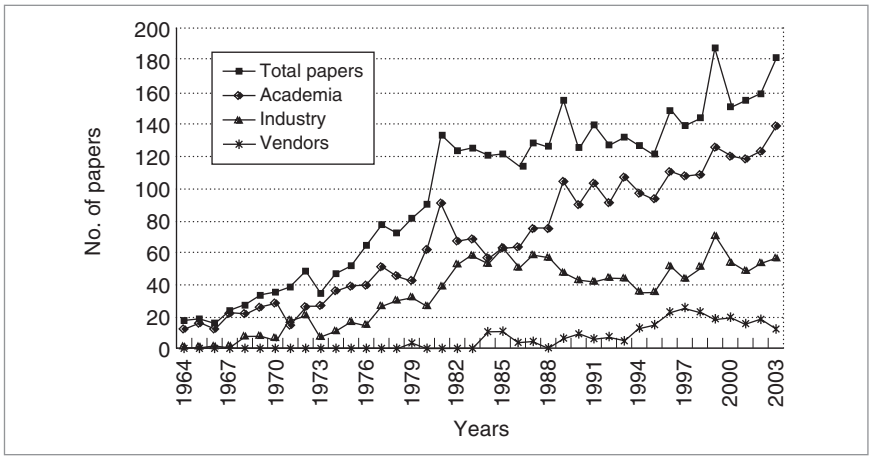

Figure 2. DAC paper statistics.

of EDA. While browsing through the first few years' proceedings, I found seminal papers that continue to have a strong impact today. I classified them by relevant topics, and I clustered fundamental contributions into five areas: circuit simulation; logic simulation and testing; MOS timing simulation; wire routing; and regular arrays.

Circuit simulation. Circuit simulation has been an important DAC topic, especially as it relates to IC design. The great success of circuit simulation in IC designs has been a dominant driver in the birth of EDA as an industry. IBM researchers were the most prominent force in these years: Frank Branin was a pioneer in determining the architecture for circuit simulation. In my opinion, Brayton, Hachtel, and colleagues at IBM T.J. Watson Research Center in Yorktown made fundamental, revolutionary contributions to this field: They introduced the algorithms behind circuit simulation as we know it today, from sparse matrices to backward differentiation formulas. They enabled the development of two important programs: IBM's Astap (advanced statistical analysis program) and UC Berkeley's Spice (simulation program for integrated circuits emphasis). The early contributions of Ron Rohrer and Don Pederson were essential to making these programs the workhorses of circuit simulation for years and years.

Logic simulation and testing. Logic simulation had already been in use for some time when DAC started. Computer companies, such as IBM and CDC, relied on this technology to debug their logic designs. The contributions to this field are countless, and the field's founders also had a fundamental role in developing automatic test pattern generation and fault simulation. The role of Ulrich and Hayes at the University of Michigan; Breuer at the University of Southern California; and Szygenda at the University of Texas at Austin, to name a few, was invaluable. The invention of a methodology for level-sensitive scan design by Tom Williams and colleagues, and the consequent development of the D-algorithm by Paul Roth of IBM, changed the way in which this industry designed computer hardware and generated test patterns.

MOS timing simulation. Exploiting the quasi-unidirectional characteristic of a MOS transistor to speed up circuit simulation first appeared in the work of Gummel and colleagues of Bell Labs in 1975. Gummel's intuition was to approximate the solution of the ordinary differential equations describing the circuits with a fast relaxation-based heuristic algorithm. His insight was to recognize that when analyzing a digital circuit's timing, its accurate waveforms are unimportant as long as the switching events could be correctly placed in time. This insight is still the basis of today's fast circuit simulation.

Wire routing. Most of the techniques that we use today in our tools are based on the results obtained in this period. The famous Lee maze router developed in 1961

(before the first DAC), is still the basis for most of the routers in use today. The Hightower line extension algorithm dates back to 1969, and the idea of channel routing came from the work by Hashimoto and Stevens in 1971.

Regular arrays. With the advent of large-scale ICs, the appeal of regular layout patterns for reducing design time was strong. Researchers developed gate arrays (called *master slices* at IBM) and standard cells (*master images* at IBM) as alternatives to custom layouts. In this period, IBM and Bell Labs made extensive use of integrated tools using these design styles for the automatic layout of circuits. I remember the beginning of this idea, when a gate array IC had four gates and a Schlumberger researcher discussed the use of automated tools to optimize the circuit's utilization. Look at the distance we've covered since then: in 40 years, from four to 400 million gates on a single chip.

Business side

Observers considered computer-aided design to be of strategic value to the system industry, and in particular, to the computer industry. IBM and the other large companies considered it important enough to warrant a sizable investment in funding and resources. Internal CAD groups were powerful engineering organizations. Looking at the DAC proceedings, I also noticed the strong presence of the Japanese computer and communication system industry. These contributors wrote papers describing unified approaches to system design using tools, thus showing the strategic value of CAD technology worldwide.

During this period, entrepreneurs founded the first-generation CAD companies: Applicon in 1969, Calma in 1970, and Computervision in 1972. They all supported pretty much the same design activity: artwork editing on customized workstations. Their business models focused on workstation sales; they considered the software as an add-on. Of these first-generation CAD companies, none is alive today. I can identify several causes for their demise:

- The architectures of their products consisted of customized hardware with complex software written mostly in assembly code. It was difficult for them to keep abreast of technology advances because of the huge investments needed to design novel workstations. This problem was further aggravated by severe software porting problems and a limited sales volume to support the investments.
- Their products had limited loyalty because customers perceived their value added as low.
- These companies had a limited understanding of market evolutions and customer needs.

These factors contributed to a complete lack of innovation and eventual obsolescence.

I cannot conclude the age of gods without mentioning Pat Pistilli's particular role: He nurtured DAC from a handful of attendees to over 5,000 at the technical sessions. Our celebration of DAC's 40th anniversary is a tribute to his vision.

Age of heroes (1979 to 1993)

It is surprising to see that between 1979 and 1993 the EDA field exploded in all its aspects. To be at DAC during this period was quite an experience. The vibrancy and enthusiasm that permeated the presentation rooms and the exhibits was a clear sign of the community's healthy growth. DAC became the source of many essential contributions in all aspects of EDA from physical verification to layout synthesis, from logic synthesis to formal verification, and from system-level design to hardware acceleration. The technical community expanded to reach areas of expertise in nonlinear and combinatorial optimization, control, artificial intelligence, and logic. In this quest for new methods and tools, the community also explored avenues that did not yield the promised results. There were years when papers using expert systems and neural networks dominated the conference; many panels discussed the potential impact of these methods. Yet today, little is left to show of their impact, and I am glad to report that DAC presented no papers in these areas in 2003!

During these years, entrepreneurs founded the most successful EDA companies. The most prominent research groups hired several PhDs in EDA; industry exerted a strong pressure on students to enter this area. The main contributions from this age cluster into several distinct topics.

Verification and testing. This field encompassed two main lines of work. One focused on making circuit simulation orders of magnitude faster than with Spice; the other focused on formal techniques to prove that the circuit would perform correctly. In the first domain, the work on relaxation-based techniques and mixed-mode simulation by, among others, Richard Newton, Albert Ruehli, and myself around 1980, established the basis for the fast MOS simulators in use today. Other work by

Bryant focused on the simulation of digital circuits with a two-order-of-magnitude speedup over Spice and accuracy between the one offered by circuit simulation and the one offered by logic simulation.

Interconnects had an increasing impact as geometries scaled down. The interconnect delay model introduced by Penfield and Rubinstein in 1981 saw extensive use. Pileggi and Rohrer at CMU tackled interconnect simulation, developing their results on asymptotic waveform evaluation (AWE) in 1988.

Formal techniques first aimed at answering the question of whether two different networks of gates computed the same Boolean function. The first use of formal verification was again at IBM by Bahnsen during the age of gods, but this work never appeared in the DAC proceedings. The seminal work by Bryant on binary decision diagrams (BDDs) in 1986 revolutionized the field by introducing a canonical form for Boolean functions and very fast manipulation algorithms. The work by Coudert and Madre on finite-state machine equivalence using BDDs, and the work by Ed Clarke, Ken McMillan, Dave Dill, and Bob Kurshan on model checking in the early 1990s, elevated formal verification to higher levels of abstraction. Model checking tackled the problem of verifying whether a sequential system represented by an FSM would satisfy a property described as a logic proposition defined on the states and transitions, or as another FSM.

In the testing area, the famous Podem-X program based on the D-algorithm emerged at IBM under the direction of Prabhu Goel in 1981.

Layout. At the representation and basic manipulation level, artwork editing saw a fundamental change. In the previous period, designers used different data repositories for each of the design steps and performed tedious translation work from one data format to another. In the early 1980s, Newton's work at UC Berkeley with Squid, Oct, and VEM; and Ousterhout's work with Magic revolutionized the field by showing that it was possible to have a unified database and graphical user interface. I cannot overemphasize the impact of this work. The OpenAccess database has borrowed from Newton's seminal work; some of the data representations for advanced physical design are based on the idea of corner stitching introduced by Ousterhout.

Another revolutionary idea arose from work on silicon compilation and layout languages. Several research groups explored this area at around the same time. These groups were at MIT (Battali), Caltech (Mead's group with the Bristle Block silicon compiler), and Bell Labs (Gummel; Buric and colleagues with the L layout language). The idea was to inject IC design with some parts of the computer science culture. Although the approach was intellectually elegant and powerful, little of it remains in present tools.

Two physicists, Kirkpatrick and Gelatt, used their knowledge of physics to develop simulated annealing, introducing it at IBM in 1980 to solve a placement problem for gate array layout. Once publicized, this approach gave rise to a great deal of research on its efficient implementation and theoretical properties. Customizations of the algorithm took place for standard-cell and macrocell layout as well as for global routing. Most of the major companies, from Intel to DEC, and from Motorola to TI, used the TimberWolf system written by Sechen at UC Berkeley.

These years also witnessed significant interest from theoretical computer scientists about studying the combinatorial aspects of layout design. The work of Rivest and Pinter (MIT) on routing, and of Karp (UC Berkeley) on placement and routing, were examples of this involvement. This was a clear sign of EDA's success in attracting other communities to contribute. During this time, the work of Kuh at UC Berkeley yielded integrated layout systems for macrocell design and point tools for placement and routing; these systems had a great impact on the field.

Logic synthesis. A seminal paper by Darringer, Joyner, and Trevillyan introduced *logic synthesis* in 1979. These researchers used peephole rule-based optimization to generate efficient gate-level representations of a design. Immediately following this work and somewhat independently, another approach to Boolean optimization began at IBM in 1979. Brayton and Hachtel, in collaboration with Newton and me at UC Berkeley, developed the two-level logic optimizer Espresso, and two multilevel logic optimizers, the Yorktown Silicon Compiler and the Multilevel Interactive Synthesis (MIS) system. This work, supported by the Defense Advanced Research Projects Agency (DARPA), spanned more than 10 years. In MIS' early development, it incorporated a technology-independent phase that manipulated and optimized Boolean functions, followed by a technology-mapping step that mapped the optimized Boolean functions to a library of gates. All the major companies—Intel, ST, TI, Motorola, Honeywell, DEC, and Philips—soon adopted this version of MIS.

In the second phase, the common approach was to

use rule-based techniques like those in the Socrates system, developed at General Electric by Aart DeGeus and colleagues. Kurt Keutzer later showed how to use the compiler work by Aho and Ullman to obtain a highly efficient technology mapper. The idea was to formulate technology mapping as a tree-covering problem to be solved with dynamic programming. This idea still finds a home in most of today's logic synthesis systems. The work in Japan at Fujitsu, NTT, and NEC was outstanding in producing working logic synthesis systems based primarily on the original work of Darringer at IBM.

Logic synthesis was a great achievement of our community, and as such, it was the source of many papers and attracted significant interest from those in other connected fields. In a keynote address at the International Test Conference in 1985, I asserted that test pattern generation and logic synthesis were indeed two faces of the same coin. The work by Keutzer, Devadas, Malik, McGeer, and Saldanha in collaboration with Brayton, Newton, and me demonstrated several results in redundancy removal and delay testing using logic synthesis techniques. New testing algorithms from V. Agrawal, Tim Cheng, and others arose from the cross-pollination of the two fields.

In the mature period of logic synthesis, Coudert and Madre were able to considerably accelerate logic synthesis algorithms based on conjunctive normal form (CNF) representations of Boolean functions by using BDDs.

Hardware description languages. The technical work characterized as logic synthesis would be better classified as logic optimization because the algorithms change a digital circuit's Boolean representation into an optimized equivalent representation. Synthesis implies a bridge between two layers of abstractions. Hardware description languages were born to more efficiently and compactly represent digital circuits than Boolean functions. The real synthesis job became mapping an HDL description into a netlist of gates. Unfortunately, HDL development began independently of the work on logic optimization. This implied that logic synthesis algorithms could not tackle all the constructs of HDLs such as Verilog (proposed by Moorby and colleagues) and VHDL. It thus became necessary to restrict the use of these languages to their so-called synthesizable subset. Although HDLs were indeed a great advance in terms of introducing verification early in the design cycle to reduce design time, the need for subsetting showed that HDLs had problems on the semantic side.

The HDL battle was very interesting during these years: Verilog was a proprietary language (Gateway Design was selling a Verilog simulator and licensed the language) while VHDL was an open standard supported by the DARPA within the VHSIC program. We practitioners had countless debates about the superiority of one language over the other at DAC. When Verilog became public, there were mostly cosmetic differences between the two, even though most people expressed a strong attachment to one or the other, according to their personal taste. Joe Costello very eloquently argued that the adoption of two standards for a single task is in general a bad idea: "... Adoption of VHDL was one of the biggest mistakes in the history of design automation, causing users and EDA vendors to waste hundreds of millions of dollars. ..."

Hardware acceleration. All EDA approaches require a massive amount of compute time to execute complex algorithms on very large data sets. In a great expansion period for the technology, customized hardware to speed up the execution of EDA algorithms was extremely appealing. As usual, IBM (which I believe is the single institution with the most impact on this field) proposed a special-purpose architecture for logic simulation, the Yorktown Simulation Engine (YSE), by Pfister and colleagues. The great advantages in performance drove strong industry activity; entrepreneurs formed several new companies to serve the hardware acceleration market. The idea of hardware acceleration extended to other EDA fields, including wire routing (via Ravi Nair and colleagues at IBM). However, it did not have enough appeal in other areas to repeat the success of YSE.

In parallel with this work, researchers pursued the alternative idea of using general-purpose parallel computers to achieve similar performance advantages but at a lower development cost. In the late 1980s, the computer design community showed significant interest in parallel architectures. Thinking Machines, using a massively parallel architecture designed by Danny Hillis at MIT, generated excitement in research and industrial communities, including EDA. The company formed around this machine disseminated this approach to computing and attracted some of the very best minds in the field. For example, Nobel laureate Richard Feynman designed and implemented the routing algorithms for the communication among processors; he spent many hours on Thinking Machines' premises. Impromptu debates about algorithms and applications were common; the most prominent scientists of the time

gave seminars. Researchers developed algorithms for this architecture, including those for circuit and logic simulation, and for placement and routing. However, the lack of understanding that final users needed complete solutions, rather than powerful hardware, limited the industrial use of these machines to research laboratories; this situation eventually led to the demise of the company.

During the same time period, others in EDA actively pursued and used parallel architectures such as the N-cube, Sequent, and Intel hypercube. Up to now, however, no one has sold a commercial tool based on these machines. I believe that parallel computing with various degrees of heterogeneity is still an untapped source of important results for EDA and other engineering applications. Before parallel computing is of widespread use, however, researchers must solve the fundamental problem of software support.

High-level design. High- or system-level design is a bridge to the future. We all agree that raising the level of abstraction is essential to increasing design productivity by orders of magnitude. I am indeed very passionate about this field, and I believe our future rides on the success of design methodologies and tools in this area. The foundational work (by Thomas, Parker and Gajski) started in the 1980s with high-level synthesis. This work started almost in parallel with logic synthesis, and researchers developed several commercial tools. Despite these facts, the design community has not widely accepted this approach; much work remains to be done.

The basic question is what made logic synthesis successful, and what made the adoption of high-level synthesis so difficult? I believe that the original work on high-level synthesis was too general; designers had to explore too many alternatives, and the tools had a difficult time beating humans at this game. However, when system-level design focuses on constrained architectures such as DSPs and those based on microprocessors, high-level synthesis has had a degree of success in industry. The IMEC work by DeMan and Rabaey on the Cathedral system is one example of success in this narrower domain. Hardware-software codesign approaches embedded in systems—such as Vulcan at Stanford, Flex by Paulin, Cosyma at Braunschweig University, and Polis at UC Berkeley—are other examples.

During this period, Ed Lee developed Ptolemy and Harel developed Statecharts for design capture and verification at the algorithmic level; this work influences present approaches to embedded-system design. In software design, Berry and Benveniste at INRIA, and Caspi and colleagues at Verimag have proposed synchronous languages (Esterel, Signal, and Lustre).

Relevance of EDA research to the scientific community

During this period, EDA research was in the spotlight and equally pursued in electrical engineering and in computer science. The relevance of this work becomes apparent in the many awards that EDA researchers collected and by the following astonishing statistic: According to the CiteSeer database, the three most cited computer science papers come from EDA. They are, in order,

- "Optimization by Simulated Annealing," Kirkpatrick et al., 1983;
- "Graph-Based Manipulation of Boolean Functions," Bryant, 1986; and
- "Statecharts: a Visual Formalism for Complex Systems," Harel, 1987.

You could argue that this data does not show scientific excellence but rather that our community is more generous in quoting other people's work than other, more-traditional computer science areas. But I consider frequent citation a positive aspect, anyhow.

Business side

Second- and third-generation EDA companies formed during this period. Second-generation companies—Daisy, Mentor, and Valid—arose in the 1980 to 1981 time frame to serve the digital design market with schematic data capture and simulation on workstations. Daisy and Valid would build their own workstations, in line with the traditional approach of the first-generation companies. In contrast, Mentor sold Apollo workstations with an original equipment manufacturer agreement. Hardware sales were a substantial part of all three companies' revenues.

In 1982, entrepreneurs founded SDA and ECAD, which merged in 1987 to create Cadence. These companies were the first example of software-only companies. Because I was involved in SDA's creation, I would like to share a tidbit that many of you might not know. ECAD and SDA were actually supposed to be one company from the very beginning, but Paul Huang had already completed Dracula, his physical-design verification system, and wanted to go to market quickly, before SDA's side of the equation was ready. SDA prepared to go public in 1987—on the day that investors still

call Black Monday. Stock market conditions prevented any IPOs for a few quarters afterward. Thus, Costello, Jim Solomon, Paul Huang, and Glen Antle felt that the best strategy was to make SDA public through its merger with ECAD. They called the joint company Cadence.

During this period, others founded Silicon Compilers and Silicon Design Labs, based on the concepts of silicon compilation and symbolic layout. ViewLogic was pursuing the EDA market from the PC angle, aiming at low-cost solutions for digital design. Entrepreneurs formed Gateway to commercialize Verilog and its associated simulator.

In 1987, Optimal Solutions Inc. (OSI) was incorporated in North Carolina. Not many people know that this was the original name for Synopsys! By that time, it was clear that selling workstation hardware was not an appealing business model, given the difference in margins. Despite this developing consensus, as late as 1990, visible figures in the EDA community commented negatively on the business model of "software only" companies, saying that they knew of no long-term successful EDA company that did not sell workstation hardware. Specialized hardware acceleration companies also started during that period: Quickturn, PiE Design, and IKOS, among others. Today, there are no more independent hardware acceleration companies; mergers and acquisitions have absorbed most.

Even though there was a strong incentive for IC and system companies to adopt these commercial solutions, the strategic value of EDA kept internal investment high. In particular, Bell Labs and IBM were pulling ahead of the competition in terms of tools and environments.

In the next few years, the second-generation business model—hardware plus software—proved unsustainable because of the dominance of general-purpose workstations. Daisy and Valid died a slow death through acquisition, while Mentor reinvented itself to sustain competitiveness and remain economically viable. Silicon Compilers and Silicon Design Labs also disappeared through a series of mergers and acquisitions. In its basic form, silicon compilation did not pay enough attention to the final result's performance and area. In addition, layout languages and symbolic layout systems did not find acceptance in a community used to representing designs with images and geometries.

Age of men (1993 to 2002)

In my opinion, 1993 was the beginning of a new phase in our community. Technical innovation began slowing down. The vendor community became mature from Wall Street's point of view, leading to more attention to the bottom line and less risk taking. I quote the 1995 DAC keynote address of my colleague and friend, Richard Newton: "If there is a single point I wish to make here today, it is that as a discipline, both in industry and in academia, we are just not taking enough risks today."

This period coincides with the explosion of the Web and its applications. The emerging Web was commandeering the best energies and minds in electrical engineering and computer science; it was also capturing venture capital funding. This situation naturally yielded a lower rate of EDA innovation. At the same time, the semiconductor sector continued to drive technology along the lines of Moore's law, increasing the technical challenges to EDA. System on chip (SoC) became a reality.

SoC means many things to different people. I found that in Japan and Korea, SoC meant the integration of memory and microprocessors; elsewhere, anything that uses large numbers of transistors would qualify as SoC. In my opinion, SoC is about integrating different design styles into a coherent whole. Interdisciplinary approaches are necessary to solve the complex problems posed by advances in electronics.

Relevant contributions. The jury is still out as to the fundamental contributions of this period. In line with this period's characterization, I do not wish to take a position here by labeling various research approaches as successes or failures (it is simply too early to tell), but simply comment on some important topics. I will leave the task of identifying key contributions to the speaker delivering the keynote 10 years from now!

Physical verification has attracted much attention as geometries march toward the submicron range. Self-test has emerged as the only solution to rising costs and requirements for test equipment. Researchers have studied asynchronous design methods and the associated synthesis problem as potential solutions to performance problems arising from power consumption constraints and the unpredictability of wire delays. Designers interested in pushing the envelope still grapple with the deep-submicron paradigm change, seeking its limits and whether the gains are worth the approach's risks. As chips incorporate increasing functionality, analog design has become the bottleneck. In SoCs, the name of the game has become finding the best matches between analog and digital components, rather than optimizing the analog part's performance to death. Because of an analog circuit's dependency on many second-order effects, analog design has been more of

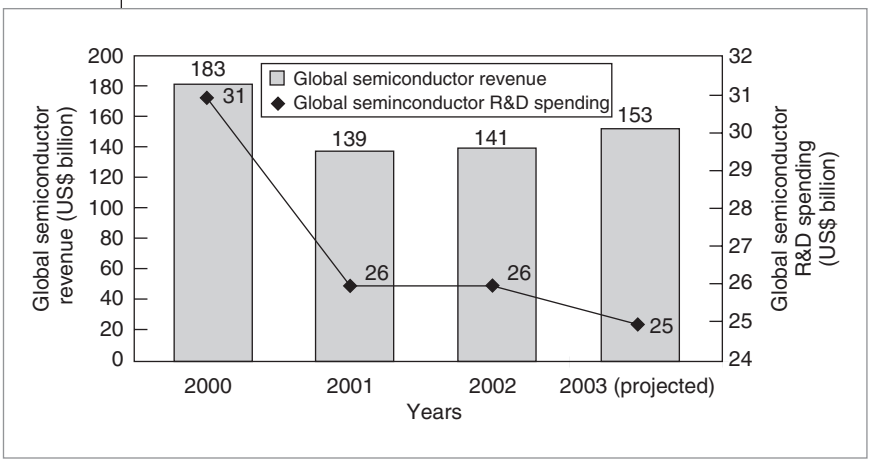

Figure 3. Market conditions.

an art than a science. I believe we must bring science into the picture to make this activity repeatable and much faster than it is today.

When wiring delays became relevant in chip design, the separation of concerns that enabled the layering of logic synthesis and layout began having problems achieving *design closure*. That is, circuits designed at the logic level to satisfy timing constraints had difficulty satisfying those constraints after final layout. Obviously, failing to achieve design closure generates unacceptable time-to-market delays and costs. Today's proposed solution is to merge layout and logic synthesis, at least in terms of gate sizing, into one single optimization loop.

Embedded system design is moving toward increasingly software-rich solutions. This creates a strong interest in hardware-software codesign as a means to speed up the design cycle by parallelizing hardware and software development, and by reducing integration time. Hardware-software cosynthesis also means using a high-level functional model to derive detailed, optimized software and hardware implementations. Note the difference between synthesis and compilation here: I say *synthesis* when I use a mathematical representation of the original design that contains limited bias toward a particular implementation style. *Compilation* implies the translation from a programming language to assembly or machine code. In this case, the mathematical abstraction is the same.

Thinking in this context, I was surprised to see that as early as 1967, DAC presented papers on software design. Are we going back to the future, then? The basic difference is in the type of software that was of interest. In the beginning, DAC software papers spanned activities from buildings to structures, from electronic circuits to "standard" software (database and airline reservation software). Now the focus is mostly on embedded software.

Business side

I already alluded to the stress on the EDA industry created by the Internet and high-tech financial frenzy. At that time, I often received two calls a day from recruiters, asking desperately if I had students interested in working at one of the established EDA companies, because there were not enough people to keep EDA going in the vendor space. Established vendors had unwanted terminations in the range of 20% of their workforce, losing them to Internet and EDA startups. In 1999, there were about 80 startups in EDA. During this period, companies such as Avanti, Ambit, Magma, Monterey, Get-to-Chip, Verisity, and Verplex started, trying to challenge the dominant players. However, if we look at the years after the Internet bubble burst, we see a different landscape: very few IPOs—if any—and acquisitions as the only exit strategy; not a pretty picture.

The future of EDA

The period from 2000 to 2003 has been very challenging to the high-tech industry at large. As shown in Figure 3, overall market conditions are gloomy to say the least; global semiconductor R&D spending is decreasing, and consequently the total available market for EDA companies is shrinking as well (a 3% decline in 2002, according to EDAC). Staying alive and prospering is a major challenge.

At the same time, we have to innovate and invest in new technology. EDA must adapt to a changed business condition and structure. We are witnessing a substantial change in the client-vendor relationship. Partnerships are increasingly important as semiconductor and system companies rationalize their investments in EDA technology; all of this in view of a fundamental change in semiconductor technology that requires investments of a size never seen before, while at the same time exposing the limitations of the present design methodology and tools. It is thus no wonder that many companies have delayed introducing their 90-nanometer technology node. The cost of ownership for ASIC design is increasing rapidly because of nonrecurring engineering (NRE) and mask costs. Companies avoid design starts, choosing instead to work with standard solutions and

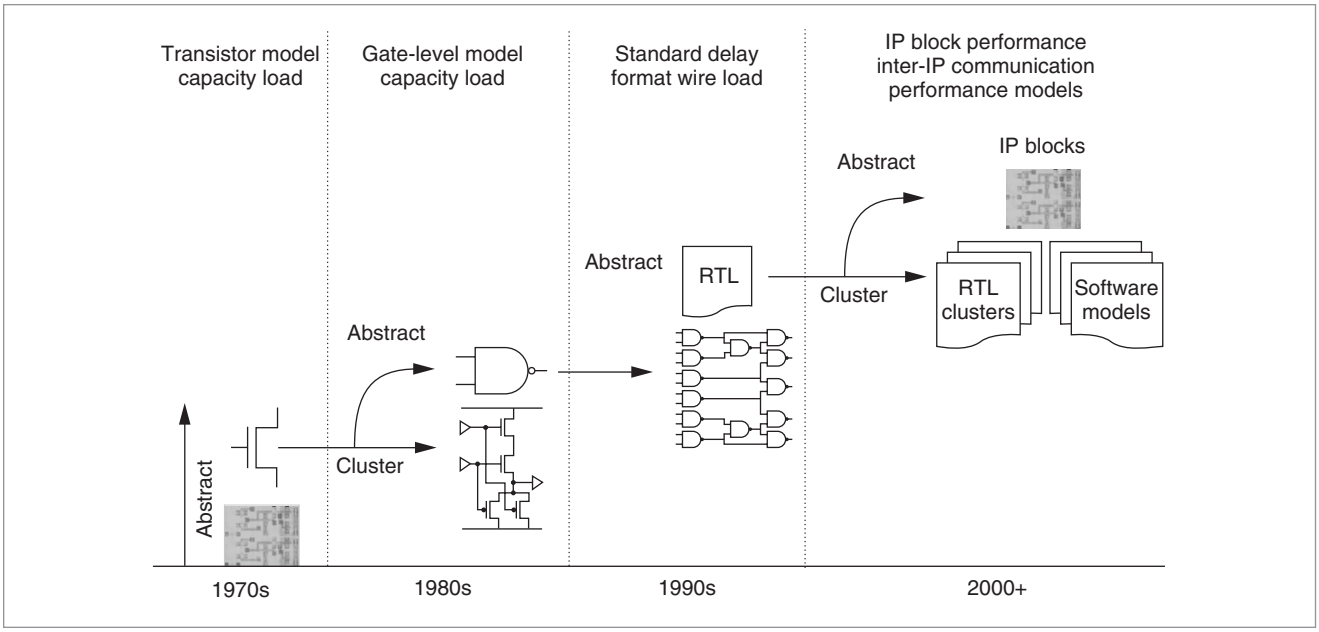

Figure 4. Raising the level of abstraction.

customization by software. If we extrapolate the data, it has only one message: The traditional EDA market, centered on ASICs, is evaporating. We need to re-think the entire design process. System-level design concerns are now dominating the definition of new platforms for future electronic systems. This trend is not very visible today, but it will become increasingly important until it dominates the design process.

There is no other choice for the EDA community than to look for other areas of application. The semiconductor industry, EDA's main customer, is looking for the next killer applications for its products, after PCs and cell phones.

Societal-scale applications

Given my affiliation with a university known for its involvement with social issues, I firmly believe that the next drivers for the high-tech industry will involve the global interests of society. A consensus is forming that electronics has yet to penetrate application domains of great interest. Such potential applications are the focus at the Center for Information Technology Research in the Interest of Society, a very broad University of California program centered in Berkeley and sponsored by the State of California and industry. The central role in this research belongs to devices such as the Smart Dust, developed by Kris Pister and his colleagues at UC Berkeley. Smart Dust combines a wireless communication sensor and information elaboration node in a very rich network.

If we take for granted that these applications will dominate the future landscape of electronics, what must EDA do to support them? As already mentioned, the design style of choice should favor reuse in all its forms and, given the constant increase in NRE and mask costs, make software even more pervasive than it is today.

Ad hoc communication protocols will also play a substantial role in the design process. In the history of design methods, changes in design productivity were always associated with raising the level of abstraction in design capture. In 1971, the year I graduated, the highest level of IC abstraction was a transistor schematic; 10 years later it was the gate. By 1990, HDL was pervasive and design capture was at the register transfer level. Figure 4 shows these transitions in level of abstraction.

In the future, EDA must work with blocks of much coarser granularity than today to provide the required productivity increase. We must bring system-level issues into chip design. The emergence of novel languages for the description of systems such as SystemC (pioneered by Gupta, Liao, and colleagues) and System Verilog are clear indications of this trend. However, they fall short in addressing the system design problem at higher levels of abstraction, mostly due to their lack of a clear, unambiguous synthesis semantics.

Support for the design chain

I have been talking about these issues for more than 15 years, but I am still passionate about them. The trend

SEMATECH Evolving: A New Model for Global Industrial R&D Coordination

Kenneth S. Flamm
University of Texas, Austin

The late 1970s were a period of radical change in a global semiconductor industry previously dominated by US producers. Japan launched a series of government-industry semiconductor R&D consortia, the so-called VLSI projects. Most observers perceived these efforts to have greatly advanced the technological and manufacturing competence of Japanese semiconductor producers.

In 1987, the US Department of Defense's Science Board issued a report noting a rapid deterioration in the relative position of US semiconductor manufacturers, characterizing this as a national security issue. Responding, the US government decided to have the DoD pay half of the cost of a joint industry consortium dubbed SEMATECH (for semiconductor manufacturing technology), which had a total budget of $200 million annually. The objective of improving US semiconductor manufacturing technology might have been fairly clear, but the means of doing so sparked considerable debate. In its first few years of existence, SEMATECH's organizational focus shifted about, and it was not always wholly effective. One constant was that it was restricted to US companies—the organization turned away Japanese producer NEC, which had a US production plant, when NEC sought SEMATECH membership in 1988.

SEMATECH refocused its structure and research direction in the early 1990s. Even in earlier years, it had placed increasing emphasis on projects aimed at improving the equipment and materials that US semiconductor makers procured from suppliers. In 1992, under new CEO William Spencer, SEMATECH carried out an internal reorganization and explicitly defined a new long-range strategy (dubbed SEMATECH II). As part of this plan, members targeted a significant reduction in the elapsed time between introductions of new technology nodes into manufacturing plants. A crucial element in this strategy was the institutionalization and acceptance within the US semiconductor industry of a so-called roadmap process—a systematic attempt by the major players in both the US IC industry and its materials and equipment suppliers to

- jointly work out the details of likely new technologies required for manufacturing next-generation chips,
- coordinate the required timing for their introduction, and
- intensify R&D efforts on technologies that were likely "showstoppers" and required further work if the overall schedule was to succeed.

SEMATECH oversaw the publication of the first such *National Technology Roadmap for Semiconductors* in 1992. The next one, issued in 1994, still had new technology nodes introduced at the historical pace of approximately every three years. But the effort to step up the pace succeeded: The 250-nm technology node came online a year earlier than predicted by the 1994 roadmap. The *1997 National Technology Roadmap* called for maintaining the two-year interval, rather than returning to the historical three-year pattern, for the next technology node (180 nm) and those to follow.

This acceleration in the rate of manufacturing technology improvement within what had become a globalized semiconductor industry clearly was assisted by factors beyond the walls of the SEMATECH consortium. Competitive pressures intensified around the world; the quickening pace of new technology deployment was a logical economic response. However, the open discussion of industry-wide R&D needs and explicit coordination of R&D efforts across companies through an industry-wide program were significant new developments.

The industry-wide embrace of an accelerated, two-year rhythm for technology introductions coincided with a major structural change within SEMATECH. In 1995, the consortium decided to join with foreign producers on a project to quicken the deployment of materials and equipment designed for use with 300-mm (12-inch) silicon wafers. The US government terminated its funding for SEMATECH in 1996 by mutual consent. A new International SEMATECH formed in 1998 to house the increasing number of projects involving foreign chip producers. Finally, in 1999, the original SEMATECH reorganized itself as International SEMATECH.

Today, SEMATECH's activities little resemble the classical vision of an industrial research laboratory. As an organization, it is mainly concerned with coordination and standards:

- bringing materials and equipment suppliers together with its members to work on technology projects largely executed outside its walls,
- serving as executive agent for the industry roadmap, and

continued on p. 69

continued from p. 68

- uniting a broad array of firms to organize industry standards for tools, software, and metrics for manufacturing.

Companies in Japan viewed SEMATECH as a major success. The SEMATECH model (ironically, a US reaction to the Japanese VLSI consortia of the 1970s) became the inspiration for a new generation of Japanese semiconductor R&D consortia in the mid-1990s. Japan's semiconductor industry formed its own R&D consortium, Selete, with a single non-Japanese member, Korean producer Samsung. Today, two transnational R&D organizations coexist within the international semiconductor industry: Selete, headquartered in Japan; and International SEMATECH, headquartered in the US. The 1997 roadmap became the last *national* technology roadmap, replaced by the *International Technology Roadmap*, sponsored and coordinated through these two global R&D consortia, and semiconductor industry associations in the US, Europe, Japan, Korea, and Taiwan.

Economic studies show that the move to a two-year cycle coincided with accelerating declines in quality-adjusted semiconductor prices in the late 1990s. Faster semiconductor price declines had a large impact on price declines for computer and communications equipment. These in turn had a major impact on aggregate economic growth and productivity improvement in recent years. The two-year cycle for the introduction of new technology nodes remains a feature of recent roadmaps, which continue to call for a reversion to the slower-paced three-year cycle in later years. Calls for a slower cycle have mainly gone unanswered.

Before there was a roadmap, semiconductor companies organized their technology planning around something approximating Gordon Moore's prediction of a doubling of transistors per IC every 18 months. As Moore's law continued to be approximately true, companies organized technical plans around this timetable. They didn't do so because that schedule necessarily maximized their profit were everyone else not to innovate on the same timetable. Rather, they did so because they believed that all their competitors would introduce new products and technology on the Moore's law schedule, and they, too, had to stick to this plan to stay competitive. This changed in the 1990s, when SEMATECH sponsored the roadmap coordination mechanism in pursuing its goal of technology acceleration. By explicitly coordinating an increasingly complex array of decentralized pieces of technology, requiring simultaneous improvement to create a new generation of manufacturing systems, the roadmap appears to have succeeded in altering the tempo of innovation. In fact, the industry's unsuccessful (to date) efforts to get off the "technology treadmill" and return to an older, slower pace of technological change by the end of this decade might indicate that the acceleration genie, once unleashed, is not so easily put back into its bottle. On the one hand, an individual company gains no competitive advantage if it slows innovation to a level matched by the rest of the industry. On the other hand, it has a lot to lose if it slows down and the rest of the industry continues at the original, faster pace.

Economists are largely accustomed to thinking of the speed of technological change as something that is exogenous, dropping in gracefully from outside their models. One moral of the SEMATECH story is that the pace of technological change might have an internal policy component as important as its external scientific foundations. This might be particularly true where an industry must precisely coordinate many complex items of technology secured from a broad variety of sources to create economically viable new technology platforms. Vague and diffuse factors like expectations and even political coalitions might then play an important role.

Kenneth S. Flamm holds the Dean Rusk Chair in International Affairs in the LBJ School of Public Affairs at the University of Texas, Austin. He is an economist and expert on international trade and investment in high-tech industries. Contact him at kflamm@mail.utexas.edu.

has been clear over this time span: The electronics industry sector has been segmenting at a rapid pace. System companies have retrenched, returning to their core competencies of product specification and market analysis, and shifting to others the task of delivering the engineering and the system components. For example, companies like Ericsson and Nokia are increasingly less involved in chip design. Consequently, semiconductor companies must do more for their strategic customers. Some of the engineering responsibilities have transferred over. At the same time, semiconductor companies are increasingly relying on intellectual property provided by specialized companies, such as ARM for processor cores and Artist for libraries. Some manufacturing has transferred to companies like UMC, IBM, and TSMC. Aside from Intel, these are the leaders

in bringing new technology to market.

In PCBs, specialized manufacturing companies such as Solectron and Flextronics are taking in the lion's share of manufacturing. Supporting this movement requires viewing design as a highly integrated activity across company boundaries, a task that is far from easy. Competencies in engineering disciplines as distant as mechanical engineering and electronics, RF, and microelectromechanical systems, must become part of an integrated environment. Such an environment must provide an economic analysis of alternative solutions together with a set of tools capable of exposing the tradeoffs in available designs. EDA must support design representations with rigorous semantics, which in turn supports clean hand-offs between design teams and more robust design verification methods. Databases that can handle design and manufacturing data will have to emerge. We call this emerging field design-chain support, a great opportunity for EDA at large to increase its relevance both in terms of value added and economic opportunities. To respond to this challenge, we must think about design as a science instead of a set of technicalities.

The concept of platform as an "opaque" layer of abstraction that exposes the critical parameters of the layers below to the layer above is, in my opinion, an essential part of a design methodology for design-chain support. I think of platform-based design as a sequence of steps that takes a system from concept to reality. Logic synthesis is an example of the paradigm applied from the RTL to the logic-gate layer of abstraction. In this view, the library of gates that we limit ourselves to for implementation is the lower-level platform, and the RTL description is the top-level platform. Selecting the best implementation platform means using the gates in the library to "cover" the logic functions in the RTL description.

In contrast, consider using a com-

The SEMATECH Experience

William J. Spencer, SEMATECH International

The US semiconductor industry lost its market leadership to Japan in 1985. This capped a six-year decline as the industry recovered from a market downturn in the late 1970s.

At the same time, US semiconductor equipment manufacturers were losing market share to the same competitor at about 5% per year. Although equipment makers had a large market share at the start of this decline in both industry segments, they were destined to lose market leadership by the late 1980s. Earlier, the entire consumer electronics market had left the US, so many had already written off the semiconductor market as having the same eventual fate.

Two groups decided to do something about this situation. The Semiconductor Industry Association under the leadership of Charles Sporck, CEO of National Semiconductor, examined opportunities for reversing the market loss and the technology lag that had led to it. The US Department of Defense under the leadership of Norman Augustine, CEO of Martin Marietta, commissioned a study by the Defense Science Board, measuring the loss in terms of whether it would create a national security issue.

The findings in these studies prompted the semiconductor industry and several government organizations to negotiate a proposed solution: The government and the industry would each provide $100 million per year to form a consortium to focus on manufacturing technology. The name for the consortium, SEMATECH, comes from the phrase semiconductor manufacturing technology.

The money would support a consortium of originally 14 members to improve the US position in semiconductor manufacturing. After a nationwide review of proposed sites in 1987, the new organization set up home in Austin, Texas. The technical staff came from member companies that assigned engineers to SEMATECH, usually for a period of two years; SEMATECH also hired a few of its own full-time employees. The member companies' engineers and their families moved to Austin for the length of the assignment. After a yearlong search, Bob Noyce of Intel agreed to become SEMATECH's first CEO. He joined the new organization in Austin; that was in July 1988.

SEMATECH struggled with its mission for two years. The prolonged discussion led to the resignation of three of the original companies and the selection of a new chief operating officer, before the organization began focusing on improving current semiconductor equipment and developing new systems. The funding came principally from SEMATECH, with the development typically performed in the equipment manufacturers' facil-

continued on p. 71

continued from p. 70

ity. A team of SEMATECH engineers oversaw the R&D activity. These programs required approval from member company committees that met regularly in Austin and at the equipment company sites. Program participants regularly reported their results at committee meetings.

There was no work on specific products or even manufacturing techniques for specific products; SEMATECH considered this the competitive area for each member company. The results of the equipment improvement projects were available to any company, member or not, US or not. The only restriction was that the member companies' requirements for new or improved equipment be met before selling equipment to nonmembers. The sale to all semiconductor manufacturers was essential if US equipment manufacturers were to compete in the global semiconductor market.

The US share of both the semiconductor market and the equipment market had changed by 1993, and both industries returned to leadership positions, a situation that continues today. The question continually arises as to the role of SEMATECH in that turnaround. Some argue that other factors—the opening of the Japanese semiconductor market, changes in the cost of capital in the US and Japan, the resurgence of Intel—had some impact as well. Nearly every major TV broadcaster, newspaper, and magazine in the world descended on Austin to determine what had happened. Perhaps *The Economist* best reported the situation as "the US has had a major recovery in the semiconductor industry. Many give some credit to SEMATECH; no one can prove they didn't have an effect."

Whatever SEMATECH's role in the turnaround, the experiment generated similar activities in other economic areas. It has changed the way industries undertake R&D in developing new technology, manufacturing processes, and cooperative efforts.

When SEMATECH voluntarily ceased accepting government funding in 1994, the organization and its members thanked the government for its help in a difficult time. They expressed the opinion that the turnaround could not have been accomplished without government help. Although the government supplied $100 million per year until 1994, it left the management and direction of SEMATECH to the members. This proved to be a unique and successful approach.

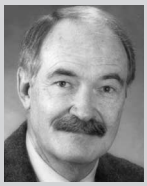

William J. Spencer is retired and chairman emeritus of International SEMATECH. He served as SEMATECH's CEO and chairman from 1990 to 1997. Contact him at spencerucb@yahoo.com.

mon semantic domain to represent both the top and the bottom layer of abstraction in the stack. In this case, we can always formulate the selection of the best lower-level platform as a covering problem. In my view, our community's focus has been the "intermediate" level, which corresponds to the ASIC design methodology. We have not paid attention to the top level—the system-level layer of abstraction—and the bottom layer, the link between circuit design and manufacturing. These two layers are the best opportunities to add value to the design process.

Embedded-system design

At the system level, we should look closely at embedded-software design as a great opportunity to innovate. For the past six years, keynotes at DAC pointed out the great importance of software for electronics, even for the semiconductor industry. There is consensus about the need to change the way we design software in general and embedded software in particular. Several horror stories trace the failure of very expensive systems, such as the Mars Lander and the Ariane rocket, to software bugs. In addition, Fabio Romeo, in a DAC keynote panel in 2001, presented an interesting statistic about embedded software for the automotive industry. In his data, productivity for embedded software ranged from 6 to 10 lines of code per day, depending on the specific application. These companies record productivity measures from the start of the project to the end of testing. However, even after extensive testing, this software has an average 3,000 errors per million lines of code. After visiting several companies in different industrial sectors, I find this situation to be quite common. The design methodology for embedded software must become more of a science and treat software as an implementation choice, not an isolated aspect of the product.

EDATech: Driving Future Innovations in Semicondutor Research

Fred Shlapak, Motorola Semiconductor Products Sector

Semiconductor process technology has been advancing at a tremendous pace over the last four decades. Silicon productivity is improving every 18 months, while design productivity continues to lag significantly. Much has been written about the design challenges we faced as the industry entered the 0.13-μm era. They spanned both digital and analog domains. Issues included power management, functional verification, leakage, complexity management for designs of more than 150 million transistors, and mixed-signal and digital design for below 0.13 μm. The list goes on. These types of formidable challenges have been addressed before: No one really doubts that semiconductor technologists will find solutions in time to meet or beat the *ITRS* predictions.

The evidence validating this confidence and optimism stems from the fact that the semiconductor industry continues to spend billions of dollars on research and development through effective industry-wide partnerships. Consequently, even though semiconductor technology costs are increasing, the barriers of entry in manufacturing have been lowered on a global basis where "patient" money has created sizable, competitive manufacturing in the Asian marketplace. In my opinion, this trend will continue, as Asia grows and prospers. (Thinking otherwise is pure denial.) Semiconductor manufacturers will take advantage of this in their business models because of economies of scale in 12-inch production and lower wafer cost in smaller wafer dimensions.

To make rapid advances in process technologies and avoid cost duplications, some manufacturers have formed partnerships. One example of this is the Philips, Motorola, and STMicroelectronics partnership, which will develop technology down to the 32-nm node. More cooperations like this will occur over time.

However, design productivity and design technology effectiveness at higher levels of abstractions remain dismal. The reason for this is lack of a concerted industry effort by both the producers and consumers of design technology. Given the EDA industry's size, this is a daunting task—perhaps no less than those facing semiconductor equipment manufacturers.

Should the EDA industry support a set of standards? The answer is *yes*. There are many technical areas where such concerted effort in silicon systems engineering will be valuable. These include a global strategy for design tools, design for manufacturability, statistical design methods, low-power design, and system level validation.

With the complexities of current and future design requirements, I believe that the semiconductor industry and its customers would benefit by creating an industry EDA consortium that will set industry-wide standards and provide a source of innovation to feed the changing IC system design industry—be it fabless or "fablite." The US currently has the lead in the know-how to build complex system chips; our challenge is to build an engineering ecosystem that turns this into a long-term technology and economic advantage.

Fred Shlapak is past president and CEO of Motorola's Semiconductor Products Sector.

Since embedded software correctness often relates to its timing behavior, we must somehow link behavior to implementation platforms; this implies a different paradigm from the classical one, which carefully hides the computing engine's details. We must think of embedded-system design holistically, rather than focusing only on embedded software to solve the problem! To do so, the educational system must change, broadening the background of engineering students and training them to consider embedded systems as a whole. We need ways to design hardware and software concurrently but not independently. Software developers will have to deal with parameters that characterize the hardware's behavior. Doing so will permit them to predict physical quantities associated with the software's execution on the implementation platform; these quantities include timing, memory occupation, and power consumption. In turn, hardware designers must know what is important for the application software to work correctly.

Design for manufacturing

The decrease in feature size jeopardizes the separation of concerns between circuit design and process development. Circuit designers can no longer ignore the impact of detailed physical effects on the manufacturing process. As minimum feature sizes shrank below the wavelength of light, mask making became complex to account for the light diffraction patterns. Estimates point to an intolerable 50% in design re-spins for 0.13-μm designs; this, in part, has delayed the widespread adoption of the 0.09-μm technology node. The situation tempts us to combine layers of abstraction, so that designers account for manufacturing in circuit designs, and manufacturing is conscious of the needs of design.

However, as I argued earlier, separation of concerns is essential for design

productivity and lower error levels. Hence, platform-based design principles should be relevant at this level as well. We must identify the important parameters that manufacturing should "export" to design so that circuit designers are fully aware of the implications of their choices. Designers should also propagate the constraints that they need satisfied to the manufacturing level to ensure that the design will work correctly. Design for manufacturing will be an important battlefield for the semiconductor industry as well as for the EDA industry.

Business side

I have outlined the grand challenges for EDA in two spaces. I believe the EDA industry must invest the right amount of attention to enlarge the business boundaries where it operates to embrace these spaces. However, it must do so when the overall business situation is difficult. Innovation is an expensive and risky proposition that sometimes clashes with the status of medium-sized public companies. The consequence is what we have seen over the years: a wave of mergers and acquisitions carried out by the major EDA companies.

As Robert Stern, a Smith-Barney analyst said a few years ago, "Buying companies is a legitimate way of doing research and development in this industry." Although I agree with Stern, it is also important to find additional mechanisms to innovate. The innovation from startups in EDA has been incremental, except for a few (albeit important) cases; successful startups found better ways to do things we already knew how to do. In addition, in this economic climate, larger companies tend to outperform the market in terms of both earnings per share and capitalization. Venture capital investment is in a four-year decline; the IPO window closed at the end of 2001. From then until today, not a single EDA company has gone public. In 2002, the market share of small companies was only 7.8% of the total EDA market and losing ground. The stock price of the large EDA companies aligned with the Nasdaq market, but the small-cap companies performed much worse than the Nasdaq index.

Thus, although the startup concept is still valid and necessary to foster an important part of the innovation landscape, we need new ways to support large-scale innovation. We must think about ways to make innovation "safe" inside established companies. We must support a model where the financial and commercial strength of an established company combine with the vibrant environment of a startup. People who have ideas must be able to innovate, free from the day-to-day drudgery that, in general, is typical of any corporate environment. They must have incentives with the associated accountability that make the motivation for success high. Yet additional bolder efforts are needed.

Partnering and the EDATech concept

I see partnering between vendors and customers as an important way to alleviate the innovation problem. It has been very effective in the age of heroes. I cannot overemphasize the effects of partnerships in the success of Cadence and Synopsys, for example. Partnering for innovation should include academia, forming the "virtuous triangle," which was first presented by DeMan and touted for many years. However, we have not seen many system engineers—representing the application domains—at DAC. Neither have I seen process engineers. We need a concerted effort to bring these communities to DAC to make it the forum for EDA's new age. However, this is clearly not enough. Some initiatives such as the Marco Focus Research Centers bring these three constituencies together to find new ways of designing circuits, but I believe we must go further.

I see a definite risk of the US losing leadership in software and electronic system design caused by a constant drain of resources toward Far East countries such as China and India. Although this global shift in resources might be seen as positive, reducing costs and increasing profitability, it might create a situation similar to the one that we have witnessed in the 1980s when IC manufacturing was migrating from the US at a fast rate (the share of IC manufacturing plummeted from 90% to 40% in a few years). Many argued then that IC manufacturing was a strategic resource and that the industry should mount a concerted effort to improve IC manufacturing in the US and to keep the equipment industry economically viable.

I believe that IC and electronic system design are even more strategic than manufacturing. Yet we are constantly losing jobs and technical ground as electronic system and software design migrate to the Far East. In a recent presentation (6 October 2003), Andy Grove of Intel independently painted an even bleaker picture of the state of software in an address at the Global Technology Summit in Washington.

As equipment manufacturers were essential for manufacturing, EDA is essential for design. Given the present economic conditions and industry structure, EDA is in a similar position as that of equipment companies in the 1980s. As software programming jobs migrate, it impoverishes the intellectual makeup of the workforce and considerably increases the likelihood of the emer-

In Support of an EDATech: Pooling Resources for Competitive Advantage

Ray Bingham, Cadence

Innovation is a grand challenge for the electronic industry in general and EDA in particular. Innovation in EDA comes in different flavors: from incremental evolution of methodologies and tools to provide better, faster, cheaper products to revolutionary propositions that change radically the way design is done. The major players in the EDA industry have reached a dimension that makes it difficult to be agile and change course rapidly. Hence, innovation is often brought about by improving existing tools and methodologies, integrating them more effectively, providing services to help in deploying them, and by acquiring smaller companies with proven technology. These smaller companies have been traditionally considered the agents of change and innovation. Although this role is certainly well suited for startups and acquisitions have been effective, I believe we cannot rely exclusively on acquisitions; we need to facilitate both evolutionary and revolutionary innovation *inside* our companies.

Innovation comes often from the dedicated work of researchers and technologists who devote their full, undivided attention to the technology they are trying to develop. In the day-to-day operation of EDA companies, integration of research innovation continues to be problematic. EDA developers find that supporting existing tools in terms of fixing bugs and adding features takes a significant amount of time and attention. Hence, I believe the EDA industry needs to implement a mechanism where people with brilliant, innovative ideas have a way of productizing them that, on one hand, provides access to the rich software base associated with existing products while at the same time shielding them from the development and support tasks associated with existing products and the installed base of customers. In addition, we must find a way of rewarding them if indeed their technology is successfully brought to market. The idea is to re-create the atmosphere and the motivation of a startup without its drawbacks. At Cadence, we have implemented an incubation model that follows this general idea, and we are confident this method of development will generate substantial innovation as advocated by Alberto in his article.

With the success of our internal program in mind, it is clear that we in EDA need a similar model available to the entire industry. We need to advance the industry to a stage where all players can be effective in innovating without losing time and financial resources to reimplement technology that does not bring value added to its products. Our first step in supporting this vision was to release our fourth-generation EDA database Open Access in an open source form to the industry. This donation of a more than $20million investment in infrastructure technology will help startups focus on their key technical deliverables. and, as a result, both reduce their risk in product development as well as provide a smooth integration path into potential customer design environments. In the "EDATech: Driving Future Innovations in Semiconductor Research" sidebar, Fred Shlapak describes an initiative he spearheaded to share the huge costs of developing manufacturing lines for digital processes in collaboration with STMicroelectronics and Philips. This approach is a clear indication that much can be accomplished if some (if not all) players of an industry take a strategic stance and share technology that will not give them a competitive advantage comparable to the expenses of development and deployment they will have to sustain if that technology were kept proprietary. I have mentioned two examples of "competitive-collaboration" that show how a high-technology industry such as ours can improve the effectiveness of bringing new ideas to market. The question now is what else can we do as a community and as a sizable industrial sector to improve our general position in the electronics industry landscape. In particular, how do we fuel the engine to produce radical innovation in the first place?

Research is clearly an important, if not the most important,

continued on p. 75

gence of strong competitors.

Arguably SEMATECH had a strong impact in keeping IC and equipment manufacturing afloat by sustaining innovation in hard times. Why not think of an *EDATech* as a form of SEMATECH for the EDA industry? There is not enough commitment to large-scale innovation in the EDA space. The EDA industry needs an infrastructure to support risky investment in revolutionary technology. In SEMATECH, a major government and corporate effort provided the foundations for next-generation manufacturing lines and equipment that were at risk of falling short of the necessary funding and investment. The effort could not have succeeded without a strong collaboration between the IC and equip-

continued from p. 74

potential source of innovation. Sustaining research is a strategic dimension that our industry cannot overemphasize. Yet carrying out research that can yield a revolutionary change in design and design methodologies necessitates the deployment of significant resources that are not productive in the short and even medium term. The risk is high because there is no guarantee that even the best research will be taken effectively to market to repay the investment. This is even more so than for other sectors of the electronic industry, because design methodologies are not exact science: There are significant subjective factors in evaluating whether a given method is better than others. This aspect brings additional uncertainty that, especially in a severe industry downturn, makes funding research challenging. To eliminate this risk at least in part, defining priorities and evaluating ideas jointly with competitors and customers are definitely important actions.

Globalization is forcing companies to move manufacturing and even design and software offshore to remain competitive. The intellectual substrate of our industry can suffer from this trend in the long run. The proposal of an EDATech must be supported by the EDA industry. It will provide a unique opportunity for EDA industry players and customers to team up in an open and collaborative framework. Of course, there will be problems to solve to find a strategy that is acceptable to all (such as defining the rules for IP and the transfer of results into actual products) but the rewards to the industry in being successful compel us to try.

Ray Bingham is president and CEO of Cadence Design Systems.

ment manufacturers that implemented the innovative ideas tested at SEMATECH.

I would like the industry and the government to think along the same lines for IC design and EDA. The road is not without obstacles: I do not have a precise recipe for implementing this program, but something must be done. To make an initiative like EDATech work, we also need the EDA industry to share technology, development efforts, and research in a common structure. It also must foster collaborations among designers and EDA professionals.

I have noted over the years a notable reduction in design expertise in the EDA community. This situation is of course creating difficulties, because EDA companies are no longer capable of "inventing" new design methodologies that offer designers substantial productivity gains. Today, IC company executives often lament that they run the risk of not being able to leverage billions of dollars of manufacturing investments because they lack the appropriate support from design methodologies, tools, and flows.

EDA IS A UNIQUE, wonderful field where research, innovation, and business have come together for many years, as demonstrated by its accomplishments over the past 40 years. Can the EDA community continue its quest for better design methodologies, increasing the productivity of electronic-system designers by orders of magnitude while also increasing quality? I believe so, but it is not an easy path. Many difficulties lay in front of us. I am calling for a deeper sense of urgency and for a stronger partnership among vendors, customers, and academia.

I would like to end with another quote by Vico (liberally translated using my long-forgotten Latin), which I believe will be a source of inspiration to us all. He characterized the age of heroes as "The holy furor for truth that lives in the eternal attempt to go beyond the limit, in the infinite possibility of self-realization and of overtaking ourselves to discover the power of the spirit and give a new push towards knowledge." Let us make an effort to live up to these words.

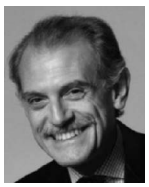

Alberto Sangiovanni-Vincentelli holds the Edgar L. and Harold H. Buttner Chair of Electrical Engineering and Computer Sciences at the University of California at Berkeley. His research interests include design tools and methodologies, large-scale systems, embedded controllers, and hybrid systems. Sangiovanni-Vincentelli has a DrIng in electrical engineering and computer science from Politecnico di Milano, Italy. He is a cofounder of Cadence and Synopsys, an IEEE Fellow, and a member of the National Academy of Engineering.