

Perspectives

NSF Workshop on EDA: Past, Present, and Future (Part 1)

Robert Brayton

University of California, Berkeley

Jason Cong

University of California, Los Angeles

■ **EDA HAS BEEN AN IMMENSELY** successful field, helping to manage the exponential increase in our capability to implement ICs that incorporate billions of transistors. At the same time, EDA has fostered and applied theories of computation and modeling, successfully combining them with practice. EDA has completely transformed the way that electronics engineers design and manufacture ICs. It was one of the earliest fields in computer science to engage in interdisciplinary collaboration, in which computer scientists and engineers successfully worked with electrical engineers, physicists, chemists, applied mathematics and optimization experts, and application domain specialists.

In this two-part article, we review the July 2009 National Science Foundation (NSF) Workshop on EDA. The workshop had two objectives: first, to reflect on EDA's success and determine whether EDA practices and methodology can influence other fields of computer science and be applied to other application domains; and second, to review the progress made under the National Design Initiative established in 2006 and evaluate what new directions and topics should be added to it. We present an overview of EDA, its funding history, a discussion of major challenges and related emerging technologies, and a look at how EDA experience might help in developing these technologies, along with associated educational aspects and challenges. We also consider EDA's relation with computer science theorists and how to revive this collaboration. Finally, we present recommendations on how to promote EDA and help it meet the serious challenges it faces in the future. The recommendations are divided into three parts: promoting research, supporting educational

programs, and encouraging enhanced collaboration with industry. Part 1 of this article, which appears here, discusses the workshop objective, EDA definition and history, and EDA funding sources. Part 2 will discuss the foundational areas of EDA, key EDA challenges, emerging areas that may benefit from the EDA technologies, educational perspective, theory and EDA, and, finally, recommendations to NSF.

Background, motivation, and objectives

EDA of VLSI circuits and systems has made a significant impact on the development of information technology (IT) in computer science and engineering. EDA's impact is particularly evident in the successful scaling of Moore's law in the past 40 years, which in turn created the high-performance and cost-efficient IT infrastructure that has transformed our life and society. The EDA field's success is inspiring for several reasons.

First, EDA has successfully managed the exponential increase in design complexity—from the first microprocessor (Intel 4004), with 2,250 transistors, to the latest multicore processor, with over a billion transistors.

Second, EDA was one of the first fields in computer science and engineering (CS&E) to achieve remarkable success in applying the concepts and techniques of computational modeling, computational thinking, and computational discovery to an application domain (electronic circuit design).

Third, EDA has transformed the way that electronics engineers design and manufacture ICs. Every circuit designed today starts with a computational model (specified in an executable programming

language) at a high level of abstraction. The circuit design then goes through a sequence of synthesis and optimization transformations, followed by rigorous digital simulation and prototyping, as well as formal and semiformal verification, before the circuit is finally manufactured via advanced lithographical and chemical processes.

Finally, EDA was one of the earliest CS&E fields to engage in interdisciplinary collaboration. Computer scientists and EDA engineers have successfully collaborated with electrical engineers to derive various levels of circuit models. They have collaborated with physicists and chemists to derive manufacturing models. They have worked with theoretical computer scientists to conduct various kinds of complexity analysis. They have collaborated with applied mathematics and optimization experts to improvise highly scalable simulation and synthesis algorithms and with application domain specialists to develop intellectual property (IP) libraries.

Currently, the EDA field is facing serious challenges in its own domain. For example, nonrecurring engineering (NRE) costs associated with VLSI circuit design are skyrocketing, with estimates of more than \$30 million per ASIC design. The rapid increase in the number of transistors available on a single chip leads to SoC integration, with associated complex interactions—for example, between software and hardware designs and between digital and analog designs. Moreover, the applications spurred by semiconductor technology are increasing rapidly. They range from very high-performance microprocessors and signal processors to a broad array of low-power portable devices, to networks of chips for micro-sensing, microcommunication, and microactuation, driven by very low per-unit cost and extremely low operating power. Consequently, designers must create chips that function properly in conventional digital and mixed-signal operation and must also comprehend sensors that respond to signals from many physical domains, such as pressure, temperature, chemical, and optical. Design problems are further compounded by the introduction of many physical phenomena determining the performance of severely scaled semiconductor devices. For example, transistors' power and performance characteristics are becoming statistical in nature. The probability of soft or permanent errors is much higher in the new generation of CMOS devices at 32 nm or below or in

emerging non-CMOS devices. These present unprecedented challenges to the EDA field.

To address these challenges, the NSF and Semiconductor Research Corporation (SRC) held a joint workshop in October 2006 to study the future directions of design automation (DA). They recommended that “research in design technology and tools be increased through a National Design Initiative which focuses on three research areas:

1. The development of a powerful new, physically aware, system design science and methodologies to increase design productivity by one order of magnitude over current techniques for integrated systems containing billions of elements.
2. The creation of robust optimization methodologies that provide guaranteed performance of integrated systems composed of devices whose characteristics are highly variable, that operate in several different physical domains, and that have uncertain reliability.
3. A revamped, systematic, and greatly improved interface to manufacturing to support the design of high-yield systems that obtain maximum utilization of the technology.”

Such an effort was deemed to be critical “to maintain U.S. leadership in design for integrated nano- and microsystems.” The second objective of this (2009) workshop was to review the progress made under the National Design Initiative proposed in the 2006 NSF/SRC workshop and evaluate what new directions and topics should be added to that initiative.

EDA defined: Evolution and history

Wikipedia defines EDA as “the category of tools for designing and producing electronic systems ranging from printed circuit boards (PCBs) to integrated circuits. This is sometimes referred to as ECAD (electronic computer-aided design) or just CAD.” The workshop attendees felt that this definition was somewhat narrow, focusing mainly on the use of EDA technologies. We agreed that a definition should equally emphasize the following three aspects of EDA:

- EDA consists of a collection of methodologies, algorithms, and tools that assist and automate the design, verification, and testing of electronic systems.

- EDA embodies a general methodology that seeks to successively refine a high-level description to a low-level detailed physical implementation for designs ranging from ICs (including SoCs), to PCBs and electronic systems.
- EDA involves modeling, synthesis, and verification at every level of abstraction.

The second and third aspects of EDA in this definition can be applied easily to application fields other than electronic system design.

Although the history of DA algorithms, such as Kernighan and Lin's bipartitioning heuristics and Hall's r -dimensional quadratic placement procedures, predates the VLSI era, it was not until the early 1980s that the mystique surrounding VLSI chip design and fabrication was unveiled by Mead and Conway. Their work led to the cultivation of VLSI education in universities and to the flourishing of research in VLSI system design and EDA, which, in turn, created automation tools for logic synthesis, layout generation, circuit simulation, design verification, reliability modeling, chip testing, debugging, and yield analysis. It is beyond the scope of this article to exhaustively chronicle the evolution of EDA research during the past three decades.

One way to characterize the advancement of VLSI automation tools is to demarcate them into different eras depending on their chief optimization criterion—namely, area, timing, power, reliability, and nanoscale issues. In the early 1980s, silicon area was at a premium, and the underlying theoretical foundation of layout tools was rooted in graph theory (min-cut, quadrisectioning), convex optimization (geometric programming), stochastic methods (simulated annealing), evolutionary biology (genetic algorithms), physical laws (force-directed relaxation methods), circuit theory (energy minimization in resistive networks), and spectral methods (eigenvalue-based partitioning). In the mid-1980s, stringent timing constraints motivated the development of theoretical underpinnings for interconnect delay models by using linear algebra and matrix solvers (Krylov subspace method; Arnoldi, Padé, and Padé via Lanczos algorithms), numerical analysis (differential quadrature method, finite difference), frequency domain solvers (method of characteristics, FDFD), and other theoretical techniques.

In the early 1990s, energy optimization requirements led to the adaptation and extension of finite-element methods and gridded techniques to solve

large second-order partial differential equations associated with the heat equation. EDA researchers realized the need to integrate full-chip thermal analysis with chip layout packages by applying fast Fourier and discrete cosine transformation techniques and so to accelerate the multilayer Green's function solvers over the quadruple integral spaces. In the late 1990s, higher reliability and quality assurance demands led to the development of technology-centric EDA tools that tackled lithographic limitations, process variations, and anomalous operating conditions. EDA tools adopted formal design verification and validation methods to ensure system design correctness by using model checking, SAT solver, static analysis, and other mainstream computer science theories. Multiscale/multilevel optimization techniques also received much attention in the late 1990s, especially in physical design, to cope with the rapid increase in design complexity.

By 2000, the VLSI industry had made a quantum leap by shattering the barriers of 100-nm and rapidly inching toward 32-nm nodes. Consequently, new-generation EDA design tools for nanoscale CMOS chips started adopting computational quantum physics (density function theory, nonequilibrium Green's function, Wigner's function) to tackle nanoscale issues such as leakage currents through high- k gate dielectrics, as well as to develop multiscale modeling for the beyond-Moore's-law technologies such as carbon nanotubes, graphene and tunneling FETs, quantum dots, single-electron transistors, and molecular devices.

In parallel with the progress in modeling and optimization techniques used for EDA problems, another significant development in responding to design complexity was the rise of design abstraction and the use of accurate estimation methods. Mead and Conway led the way with their book that opened up VLSI to the masses with its simplified design rules. Accurate estimation methods, such as model order reduction, delay modeling, high-level models, and others, have allowed the design flow to be divided into separate concerns. Design abstraction has risen from polygon drawing to schematic entry, then to RTL specification using hardware description languages, and most recently to behavior specifications—for example, in C, C++, SystemC, and Matlab code.

Funding for EDA research

The NSF's annual funding of academic EDA research has fluctuated between US\$8 million and

US\$12 million in recent years. For the past 15 years, most of this funding came from the Computer & Information Science & Engineering (CISE) Directorate under its core DA program, while additional awards were provided by the Electrical Communications and Cyber Engineering (ECCS) Division in the Engineering Directorate. (Funding for the CISE Directorate and the ECCS Division of the Engineering Directorate in FY2009 was \$574M and \$125M, respectively. Total NSF funding in FY2009 was \$5,183M.) To a lesser extent, awards were also provided by the interdisciplinary Mathematics of Computer Sciences program jointly run by CISE and the Division of Mathematical Sciences.

Because the main NSF mission is to foster groundbreaking discoveries, EDA awards are multifarious, covering the broad spectrum of EDA research that ranges from high-level synthesis of mixed-signal systems to multiscale simulation of emerging technologies. Special research initiatives like the Information Technology Research (ITR) and National Nanotechnology Initiative (NNI) have provided additional funding in the past to boost EDA research activities. It's fair to estimate that ITR added approximately US\$3 million per year to the DA area during its five-year term from FY1999 to FY2004. The CISE Directorate also funded approximately \$2 million per year of design-automation-related research from its share of the NNI until about FY2005.

The CISE Emerging Models and Technologies (EMT) program (recently disbanded) also stimulated EDA research in disruptive nano, bio, and quantum technologies. Typical EDA projects aimed at developing photonics clock distribution networks, synthesis of biomolecular computing systems, DNA computing using self-organized Wang tiles, multiscale modeling of nano and molecular systems, and synthesis of quantum logic circuits. It's not clear how the nearly \$2 million that EMT spent on EDA research will be distributed in the reclustered Computer and Communication Foundations (CCF) Division.

The Division of Mathematical Sciences has often funded proposals with mathematical underpinnings for EDA research: for example, Markovian modeling to study convergence of simulated annealing algorithms, convex programming and polyhedral methods, algorithmic semi-algebraic geometry and topology, and nonlinear programming under density inequalities. The ECCS Division funded several multiscale modeling and simulation projects for emerging

nanotechnologies and quantum engineering systems.

The NSF has also supported young researchers through its Career Award program, which funded about five new EDA projects per year from 1996 through 2001 from the DA program. Young DA researchers also received some funding from other programs we've mentioned, including the EMT program (for nano-related research), the CISE/CNS Division (for embedded systems research), and the Engineering Directorate (for analog and mixed-signal research). The total number of Career awards in such areas is estimated to be about 10 per year. Recently, this number has declined—for example, to five Career awards per year total in the DA areas. However, it should be noted that during the period 2003–2008, the DA Career awardees were winners of three prestigious Presidential Early Career Awards for Scientists and Engineers.

All of these funding opportunities contribute to the total US\$8 million to US\$12 million in funding that we estimate for NSF's DA program, which represents about 1.3% to 2% of the total available funding for the CISE Directorate and the ECCS Division of the Engineering Directorate—the two major funding sources for the DA program in the NSF.

Foundational aspects of EDA and its role in CISE

Although EDA traditionally has been supported at the NSF by CISE, some think it should be supported by the Engineering Directorate instead. However, CISE is a directorate that supports computer engineering (CE) as well as computer science (CS), and EDA has long been one of the pillars of CE, together with architecture and software. EDA needs strong support because, as both an engineering and theory discipline, its need for infrastructure support makes it more expensive.

EDA research rests at the interface between engineering and CS because it exemplifies the problem-solving application of CS to engineering. Although EDA research is necessarily governed by engineering requirements and limitations, the solutions to EDA research problems have long been achieved via math/CS core theories. For example, solutions to problems in physical design use graph theory; logic design uses Boolean algebra; simulation uses dynamical system theory; and verification uses numerous theories, including computation models and programming-language theory. These solutions also require the

development of efficient algorithms for solving optimization problems, numerical analysis of nonlinear systems, satisfiability problems, stochastic analysis, and so on. The development of such algorithms obviously requires researchers with strong backgrounds in computer and computational sciences. That EDA is a tightly connected combination of theory and practice can readily be seen by examining almost any paper published on EDA. Theoretical results needed to solve EDA problems might indeed solve open questions in math/CS theory through the application of an EDA perspective to the problem. For example, verification is solidly based on mathematics and CS theory; it is inherently multidisciplinary, involving researchers in verification, as well as domain experts knowledgeable about the kinds of systems to be verified (mixed-signal circuits, for example).

Moreover, researchers have recently demonstrated that this EDA perspective is valuable in solving problems in the sciences. For example, researchers are applying EDA ideas to the development of tools for solving problems in physics, chemistry, systems biology, and synthetic biology. For all these reasons, we believe that the correct home for EDA is within the CISE directorate at the NSF.

Limitations of industrial funding

Some might feel that because EDA problems are often of direct interest to industry, industry should be the primary financial supporter. Indeed, industry has been a good supporter of EDA, primarily through the SRC. However, these funds are quite limited now, and even in good economic times have been insufficient to cover all of EDA's needs. As we've described, EDA research requires theoretical studies that might not have immediate commercial application, thus making it difficult to obtain industrial sponsorship. Further, new research in applying EDA research to the sciences will lack support by industry. Industry will likely continue as a good supporter of near-term research problems, but we will continue to need NSF support for riskier, longer-term research with less-obvious commercial application.

Model checking and model reduction are examples of EDA technologies that, without initial long-term funding (e.g., by NSF and DARPA), probably would not have reached a level at which industry would be interested. This would have been the case even when industry was much more tolerant of "far

out" research. For example, model checking was developed from programming language theory over a period extending back to 1981. For the first 10 years, it was supported exclusively by NSF. Now major EDA companies such as Synopsys and Cadence are marketing model checkers.

SRC's support of EDA

At the SRC's inception in 1982, design science was an initial focus area, and EDA research was a major part of it. Design science funding constituted about 25% of SRC's budget; the remainder was devoted to technology and manufacturing. Today's SRC organization has two design-related areas: computer-aided design and test, and ICs and systems. These areas have absorbed an increasing share of member-directed contributions and now comprise about 45% of SRC-funded research (SRC's total EDA-related funding is estimated to have been about \$5M for 2009). Interest in design comes not only from member CAD companies but also from integrated device manufacturers, fabless, and "fab-lite" companies, and even from equipment manufacturers and foundries. There are four reasons for this:

- It is getting harder to continue on the Moore's law density-performance curve using technology scaling alone. Improvements in design techniques and tools have provided a critical boost to this continued pace.
- The *International Technology Roadmap for Semiconductors (ITRS)* has increasingly emphasized design in recent years, with much more material in chapters on design and design drivers. The current *ITRS* emphasizes "the importance of software as an integral part of semiconductor products" and "software design productivity as a key driver of overall design productivity" (see <http://www.itrs.net/reports.html>).
- Productivity and time to market are key drivers of improvements in DA.
- To be competitive, most semiconductor companies must have a diverse set of applications (no longer just general-purpose microprocessors). These applications drive industry to show differentiation in design (not just technology) of both hardware and software.

SRC continues to be an important funder of DA research, supporting faculty in both computer

science and electrical engineering departments and serving as an industrial liaison.

Partnerships between NSF, SRC, DoD, and industry

The foundational efforts we've described, which have enabled optimized logic and physical design, formal verification, DFT, and design for manufacturability. These efforts have not relied, and cannot rely, on industry support alone. SRC plays a key role in combining industry resources to focus on precompetitive areas of mutual interest. One successful example is the Focus Center Research Program (FCRP), which is jointly funded by SRC member companies and DARPA. In 2008, there were five FCRP centers nationwide, with GSRC and C2S2 centers directly related to the DA program. The EDA-related funding from these two centers was estimated to be from \$4 million to \$5 million in 2008.

A similar joint funding model is also being explored by NSF. Recently, NSF partnered with industry in areas such as analog/mixed-signal design and multicore design and architecture to support work important to industry and to part of the CISE mission.

The total available funding of DA research is estimated at around \$20 million. It combines various programs from NSF, SRC, and FCRP, with about half provided by the NSF CISE program. However, other countries with competitive strength are making much more substantial investments in this area.

Comparison with international funding sources

During the workshop, discussions about funding-support levels for EDA overseas seemed to suggest that the US is underfunding this area.

Taiwan. EDA support in Taiwan comes from direct government funding and from university-industry-government partnerships. One of the significant government-funding programs related to EDA is the SoC program, funded since 2001 at US\$70 million annually. Research topics include ICs, embedded software, and DA. According to the estimate by the workshop attendees, a significant portion (say 50%) of this program supports EDA. Other government-funding programs also include EDA-related research programs, such as the nanotechnology program, funded since 2003 at US\$100 million annually and

the telecom program, funded since 1998 at US\$70 million annually (topics include wireless, broadband, and Internet telecommunications). These government programs include support of long-term visits by Taiwanese students and faculty for interaction with non-Taiwan companies.

In the university-industry-government program, the Taiwan government acts as the primary enabler by encouraging industrial matching and participation. Notably, government grants have only a 5% overhead and industry grants only a 20% overhead.

Europe. European funding for R&D occurs on various levels and through various channels, which are largely uncoordinated. Although no program is specifically dedicated to EDA, many European countries have national programs to support R&D in areas such as information and communication technology (ICT), nanotechnology, embedded systems, advanced computing, and software technology, which all contain EDA aspects. The European Community has budgeted for and supported IT, microelectronics, and similar technologies since the early 1980s via so-called framework programs.

Currently, the seventh such framework program (FP7) covers the period from 2007 to 2013. The program assigns EUR 9.050 billion to ICT and EUR 3.475 billion to nanosciences, nanotechnologies, materials, and new production technologies, distributed over the entire period (http://cordis.europa.eu/fp7/budget_en.html). The program attempts to foster cooperation between countries. Participants must organize themselves in project groups composed from institutions from at least three different countries.

The Eureka Consortium is another source, currently covering 38 countries. This supports European international projects with funding provided by national sources. The consortium raised funds on the order of EUR 3.1 billion in 2008. Key members are Austria, Belgium, Finland, France, Germany, Spain, Switzerland, the UK, and Israel (<http://www.eurekanetwork.org/>). For example, "Electronics, Microelectronics," supported 50 projects with a budget of EUR 105.26 million over 10 years.

One of the latest Eureka initiatives is Catrene (Cluster for Application and Technology Research in Europe on NanoElectronics), which began in 2008 and is scheduled for completion in 2012. Catrene has a EUR 3 billion budget. As of August 2009, 143 partners from 13 European countries were participating. Partners are both industry (e.g., Infineon, NXP

Philips, STMicroelectronics, Airbus and Volkswagen, and Cadence) and universities. Currently, 13 projects are pending, and two explicitly address EDA (http://www.medeaplus.org/web/projects/project_list.php).

Complementary to the Catrene program is ENIAC-JRT, targeted at EUR 3 billion for the 2008–2013 time frame (http://www.medeaplus.org/web/downloads/clips_medeae/Elektronik%28Jan09%29.pdf).

The ENIAC-JRT predecessor, Medea+, covered the years 2001 to 2008 with a budget of EUR 4 billion. It supported 70 projects of which 15 focused on EDA. Of the total funding, 75% came from companies (http://www.medeaplus.org/web/downloads/medeaplus_brochure.pdf).

NSF workshop attendees learned about another program, the Nano-Tera initiative (<http://www.nano-tera.ch>), by Giovanni De Micheli, a professor at EPF Lausanne, Switzerland, who directs Nano-Tera and is a grantee of the NSF DA program. The nano-tera.ch program grant, secured for 2008–2011, will be delayed by 18 months. Ten projects were started in June 2009, with 10 more in early 2010. Funding of \$56 million (with an additional commitment of \$56M from each of six Swiss institutions) was inserted in the Swiss federal budget for 2008 to 2011. Money comes directly from the Swiss Ministry of Research to support collaborative research. Technically, six institutions formed a consortium to run Nano-Tera: EPFL (lead), ETHZ, CSEM, and Neuchatel, Basel, and Lugano universities. This program is independent of EU/FP7; links with FP7 may be developed, but for now this is a Swiss program. Finally, actual research is just starting with the hiring of pre- and post-doctoral students.

It is not clear what percentage of these funding opportunities is dedicated to DA-related research. However, if we simply take 2% of the funding (EUR 3,475 billion) dedicated for nanosciences, nanotechnologies, materials, and new production technologies as part of the FP7 program, it amounts to US\$105 million over five years, which is two to three times larger than DA program funding from NSF. (Recall that 2% is the share of DA programs in the CISE Directorate and the ECCS Division of the Engineering Directorate in NSF)

These comparisons raise serious concerns about the level of US investment in DA, a key IT area. For one example, the Taiwan government's funding for the DA program (estimated to be more than \$35M per year), with its GDP only 2.7% of that of the US (based on 2007 numbers), is 1.5 to 2 times higher than the total combined funding for DA in the US from the NSF, SRC, and FCRP. (It is probably not a coincidence that Taiwan is doing extremely well in the IC industry; for example, it has the largest share of the worldwide IC foundry market.) This underscores the need and urgency for the US government and industry to work together to strengthen the DA support needed to keep our competitive advantages in this area. ■

(To be continued: Part 2 will appear in the May/June 2010 issue of *Design & Test*.)

Robert Brayton is Cadence Distinguished Professor of Engineering at the University of California, Berkeley. His research interests include combinational and sequential logic synthesis for area, performance, and testability; formal design verification, and logical/physical synthesis for DSM designs. He has a PhD in mathematics from MIT. He is a member of the National Academy of Engineering, and a Fellow of the IEEE and the AAAS.

Jason Cong is Chancellor's Professor and director of the Center for Domain-Specific Computing at the University of California, Los Angeles. His research interests include synthesis of VLSI circuits and systems, programmable systems, novel computer architectures, and nano-systems. He has a PhD in computer science from the University of Illinois at Urbana-Champaign. He is a Fellow of the ACM and the IEEE.

■ Direct questions and comments about this article to Jason Cong, UCLA Computer Science Department, 4731J Boelter Hall, Los Angeles, CA 90095; cong@cs.ucla.edu.

cn Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.