

In This Issue

- Chair's Message
- Article: Concurrent Memory, A Call to Action
- New DATC Chair Elected
- Article: State of the Union in MSE
- Preview of DATC-sponsored EDPS Workshop
- Contribution opportunities

Editorial team

Gary Smith, Editorial Chair

Juan-Antonio Carballo, DATC Chair

Featured Links

tab.computer.org/datc

Contact DATC

jantonio@ieee.org

Message from the Chair

As the year 2009 starts, we are delighted to announce that activity and change has not only not declined at DATC but actually grown. We are happy to report that we have elected a new Chair, Dr. David Kung, and we are receiving an increasing number of contributions for our newsletter, and an increasing number of requests from conferences worldwide that are interested in our support, sponsorship, and brand association.

Please see more details below in our online newsletter. We look forward to a 2008 year full of even more growth and excitement.

Article: Concurrent Memory, A Call to Action

By *Gary Smith*

Background

This century brought on the greatest change in the computer field since the first electronic computers that were developed during the late 1940s. At that time the computer science community settled on the Von Neumann compute model as the mainstream Model of Computation. This was the sequential single processor architecture found in most computers today.

At the beginning of this century the semiconductor vendors, who produce the microprocessors that run today's computers, ran into a power problem. As we began reaching processor frequencies of over a gigahertz not only were power budgets being exceeded but the processors themselves were literally burning up. Now the beauty of the Von Neumann model was that performance was directly proportional to frequency and the semiconductor vendors had been providing faster processors every two years since the early 1970s. Since we could no longer do that the semiconductor vendors moved to a multi-core architecture. In doing so they solved their increased performance problem but abandoned the Von Neumann architecture. What that does is put into jeopardy all the companies that have made a living based on the Von Neumann compute model; Microsoft, Intel and all of the computer manufacturers.

Actually building the computer is not a problem. Semiconductor vendors can place thousands of processors on an IC today. The problem is, "Can you program it?" In the hay-day of the Mainstream computer Gene Amdahl discovered that communications problems with algorithms that contained sequential code limited the number of processors you could use to roughly four. That became known as Amdahl's Law. Now there are companies that have found tricks that allow them to use sixteen processors but that's as far as we been able to go. Of course algorithms that are basically parallel, the so called Embarrassingly Parallel Problems, can use hundreds of processors but so far we haven't found very many Embarrassingly Parallel Problems.

What's Needed

Actually pretty much everything; the software development infrastructure needs a complete overhaul. That being said we need four different areas developed.

1. The first is obvious: we need a new set of debuggers, profilers, and software virtual prototypes, etc. Some of these are being worked on by the Embedded Software Development Tools companies and some by the EDA vendors. In fact it wouldn't be a surprise if those two markets merged.
2. The second gets into languages and compilers. We've pretty much given up hope for the magic concurrent software compiler. Now the question is do we throw away C. Keep in mind that if we stick to C it will be a much different language than the C we use today. It is possible though, Fortran took this route and by the 1995 release became one of the better concurrent languages out there.
3. The next problem is what to do with the OS? Especially in a many core system it makes sense to use a kernel on the actual cores and then have a master OS that manages the scheduling. In addition it has a new task, handling the power management system. This sounds like Virtualization, however to be efficient Virtualization has to be a fairly thin layer of software. A lot of research is still needed.
4. That brings us to the topic of this paper, the Concurrent Memory.

Concurrent Memory

The concurrent memory is a memory that can handle multiple parallel programs. There is a lot more to it than that but to keep it simple we won't get into the details. We thought we had the answer with the Transactional Memory. Its only drawback was lack of speed. Last year researchers found ways to increase the speed. Unfortunately they found that it was a power hog; so back to the drawing board.

The problem we are facing is the same one that Intel and AMD are facing in the microprocessor market. The main demand for Concurrent Memory will be as an IP embedded into a SoC, not a standalone memory you put on a PC Board. In fact the view many researchers have is a stacked die (Concurrent Memory on top of a SoC) that is optimized for each SoC design. This does not fit into the mass market, standard product business model followed by the Memory Vendors. Just like Intel or AMD do not have a business model that fits into the embedded processor as IP business model that ARM has perfected. That means that the major Memory Vendors have no interest in developing the Concurrent Memory.

That's why we need the call to action. The Concurrent Memory is an integral part of the answer for parallel computing. Without it everything grinds to a halt, including Moore's Law. It's a big problem and it must be solved.

New Chair Elected

We are delighted to announce the next Chair of IEEE DATC: Dr. David Kung. Please see below for his impressive background, and for his view on how DATC will thrive even more under his leadership – his position statement.

David Kung's Bio

Dr. Kung manages a department of design automation researchers and charts the future direction of design tools research for IBM. He joined the Advanced Simulation Group in IBM Research in Yorktown Heights, N.Y., in 1986 to work on a massively parallel hardware simulation engine. Subsequently, Dr. Kung

joined the Logic Synthesis group and was a contributor to IBM's BooleDozer logic synthesis system. In 1999, he became the manager of the Logic Synthesis group and led project management and technical development of Placement Driven Synthesis (PDS), IBM's physical synthesis software. He received a Bachelor of Arts degree from the University of California at Berkeley, a Master of Arts degree from Harvard University and a Ph.D. from Stanford University.

David Kung's Position Statement for DATC Chair-Elect

The hardware and semiconductor industry will be facing considerable headwind in the coming months, and as a result the design automation industry will encounter great uncertainties. However, I would submit that these challenging times are also the most interesting ones and present a golden opportunity for design automation to establish a central and indispensable position in the design and technology development process. Design Automation is largely about productivity and enablement; how to do more with less is precisely the right medicine to encourage the hardware industry to take calculated risk to invest and lift the industry out of the current recession. My agenda for the DATC will be to provide a clear design automation roadmap for the near and long term with a keen eye on expansion into new areas, focus on a holistic view of design tools and the design process for maximal productivity gain, and drive a collaborative model to facilitate research in an adversarial economic milieu.

The high level themes of my agenda are integration and collaboration. To understand what to integrate and what to collaborate on, we need a clear vision of the future. In a time when the DA research budget across industry and academia is reduced, the DATC should step up to the challenge and provide a long term DA strategy. I support the roadmap effort of the base technology, designs, and systems, but propose to explore emerging and new areas for potential expansion such as design automation for biotech and pharmaceutical applications, Design for Environment tools, DA techniques to facilitate the synergy between chip and system architecture and the software stack. I propose to set up a task force to explore whether and how design automation should expand into these adjacent spaces.

During my tenure as design automation strategist at IBM, the main de-railers to productivity I have witnessed are compartmentalization and the lack of a global view. The value is not in treating each component in the design process as separate entities and optimize each one within its individual silo. There is tremendous gain in developing an efficient design methodology by streamlining, modularization and convergence. Therefore, I support the work of the EDPS sub-committee to raise the awareness of design tools researchers with respect to how their tool fit in the overall methodology. Moreover, with each technology advancement, different components within the design process interact with each other more closely. For example, the performance many-core architecture is highly sensitive to the physical floor-plan, cell layout impacts manufacturability and yield, many-core chips are hard to program, and so on. I propose to have the sub-committees to focus on the interaction and possible integration of components that have tight coupling with each other.

Funding for DA research, in both academia and industry will be increasingly challenged. We must go beyond traditional funding models and seek creative ways to extract value from collaborations. A win-win strategy for academia and industry is to create long-term and extensive internship programs so that students can get support for working on exploratory but industry related work and industry can benefit from harvesting top talents and reduce overhead and stretch investment dollars at the same time. Faculty members can also

leverage such programs to form close collaborations with industry. The case for this type of "collaboratories" is even more compelling if government money is available to sweeten the pot (e.g. the Taiwan government currently funds such internship programs). The main hurdles to overcome are IP ownership issues and access of proprietary data. I propose to set up a sub-committee on collaboration which draws on membership from leading academic and industrial organizations to focus on providing recommendation for these issues and set up several of these prototype collaborations by end of 2009.

Conferences and workshops play a vital role in facilitating collaboration among attendees, disseminating research, and spurring innovation. The economic downturn has threatened the well-being of some of the DA conferences. We must continue to work with CEDA to revitalize and influence existing key conferences through sponsorship and oversight, and seed new workshops to explore new frontiers.

I am passionate about the future of Design Automation and I believe that I have an actionable plan to work with CEDA to drive the agenda. Thank you for your consideration.

Article: State of the Union in Microelectronics Systems Education

By *Andrzej Rucinski*

This article is a message to both I-GEMS and MSE2009. Despite the fact that many of you participate in both Committees, I am obliged to explain what MSE2009 and I-GEMS are all about.

I - GEMS stands for IEEE Global Education in Microelectronic Systems initiative and is concerned with globalization of design issues. I-GEMS initial focus involves the development and proselytizing of the use of certified IP libraries for SoC, FPGA, and multi core designs. Dr. Ted Kochanski, IEEE Boston, is a primary driver behind this initiative.

The IEEE 2009 International Conference on Microelectronic Systems Education (MSE2009) is a leading conference, associated and collocated with DAC, concerned with curricula development in high tech areas. Dr. Mark Johnson, Purdue, led an excellent MSE2007 conference in San Diego two years ago.

Dr. Don Bouldin, U. Tennessee, is the spiritual mother and the father of both initiatives and I have the privilege of chairing both of them this year.

With MSE2009 approaching quickly, I would like to share with all of you some of the thoughts which might be timely and relevant. I did talk to some of you already about some of these issues, but I am eagerly soliciting your advice and help in making MSE2009 as excellent and outstanding as possible.

Like other challenges this country and the world are facing today, the microelectronic community is experiencing difficulties. Most specifically we are facing the challenge of living in the post VLSI era. As we all know, the VLSI industry and to some extent modern civilization owes its existence to the pioneering work of Mead and Conway in the late 1970's. DARPA and NSF reacted quickly by establishing MOSIS, and VLSI programs mushroomed throughout the United States and beyond. Designing VLSI has become a capstone of the computer engineering education. Students were able with confidence and pride to show to a potential employer their MOSIS-fabbed chip -- the essence of successful microelectronics system engineering..

However, a lot of the glamour associated with VLSI has disappeared. The

technology has moved forward and we are dominated by FPGA-based systems, and multicore is a catchword, with widespread acceptance just around the corner. MSE2009 is on top of these events and we are lucky to be able to have outstanding keynotes by, Drs. Dewilde, Patterson, and Vijaykumar.

I would like to take things one step further and consolidate our microelectronics educational efforts. I-GEMS and MSE Groups together can provide an optimistic and leading role in defining where we are heading with microelectronic system education. I am proposing to offer yet another keynote entitled "The State of the Union in Microelectronic Systems Education." The title was Don's suggestion. In my address I would like to present the heritage, current status, and the future of microelectronic systems education. For example, I would like to bring to the attention of our community, the values microelectronics is bringing to societies around the world. I am thinking about new venues in medical science, security, energy, supply chains and others. Dr. Kochanski will further address these new applications during the MSE2009 dinner. This new microelectronic domain based on embedded systems and ambient intelligences prompted IEEE to promote the development of a new scientific discipline called Globally Integrated Security Engineering (GISE) focused on the design and implementation of a Globally Integrated Secure Environment for commerce.

I am asking for your informal contribution to both of these areas. I would very much appreciate if you could share with me your thoughts which Ted and I would be happy to include and quote first in our presentations at MSE2009 and later to include them in papers to be submitted to a Special Issue of the IEEE Transaction on Education and a Special Issue on GISE, of the Proceedings of the IEEE.

I am looking forward to hearing from all of you. My best regards.

Preview of DATC-sponsored EDPS Workshop

By Gary Smith

As the General Chair of EDP 2009 I invite you to join us in Monterey. For those that don't know about it EDP is the IEEE workshop for Design Methodologists and has a high profile attendance. I think we have put together an interesting program this year and I encourage everyone to attend. Just a short drive down to the beautiful Monterey Resort Hotel on the beach.

This year, topics will include:

Are Threads Dead?

Challenges of RTL Handoff


Will We Miss the Bus?

Is the Analog Revolution Really Here?

Registering early will get you a discount and so will your IEEE membership. The temporarily unemployed can use the student discount.

Hope to see you there!

Contribution Opportunities



The IEEE DATC welcomes proposals for contributions to this newsletter. Contributions should shed light on non-obvious key EDA trends. Educational contributions in emerging areas such as ESL and DFM are especially welcome. The ideal length of a contribution is a half a page in the form of a short fact-based essay with data or references backing the stated position, but longer contributions may be considered. Publication of important graphics and data tables might be possible by request. Please send proposals in the form of a 2-paragraph abstract to the editors at jantonio@ieee.org.