



Academic Insights

CMU project redefines network design

By Richard Goering

11/30/07

The System Level Design Group at Carnegie Mellon University (CMU) is thinking big about networks. The project is not only seeking to redefine the way multiprocessor ICs are wired through network-on-chip (NoC) technology, but is also developing ambient networks of wireless video sensors.

Directed by [Radu Marculescu](#), associate professor of electrical and computer engineering at CMU, the [System Level Design Group](#) has two projects. One is Silicon Networks (SlicNets), where CMU has published some key papers in NoC technology and built an FPGA-based prototype. The other is Secure Camera-based Ambient Networks (SCAN), where CMU is researching and developing wireless ad-hoc video networks.

"Two of the most important concepts that will drive design in the future are low power and network design," Marculescu said. "I'm interested in networks in a variety of ways, and these are two incarnations — network on chip, and networks of wireless sensors."

In a broader sense, Marculescu said, the project focuses on "creative engineering" rather than just solving immediate problems. "I'm not interested in solving problems the industry begs about now, but rather pointing out a problem the industry will hit three to five years from now," he said. "I'm a true dreamer and a true academic person. I'm interested in redesigning CAD tools and methodologies for the new century."


As for the System Level Design Group, Marculescu said, "pretty much anything goes as long as it's based on good science, and is related to this general space of new design platforms and methodologies." Behind the research lie such concepts and disciplines as statistical physics, quantum physics, criticality phenomena, stochastic communications, and the "[small worlds](#)" idea from social networking.

"I want to change the perception that if you want to do VLSI, you just have to learn VLSI," Marculescu said. "I believe you have to learn a lot of things aside from VLSI, from statistical physics to criticality phenomena, because this is good engineering." Marculescu said he wants to "change the mentality" of computer engineers to think outside of their discipline, so as to get "stronger and more interesting results."

Because he believes there should be fundamental science behind engineering, Marculescu is not a pure experimentalist. "I want to make this [research] formal with a solid theoretical foundation. I'm not just simulating these things to see if they work." Nevertheless, the SlicNet project did build a Xilinx Virtex-II prototype of an MPEG2 encoder to verify its NoC architecture.

Marculescu said his original dream was a three-phase research project that would include not only NoCs and networks with "ambient intelligence," but also systems biology. The systems biology has been postponed, but Marculescu hopes to have some activity in this area next year. "What's common about the three application domains is the math and the optimization techniques, and the way you think about a system, rather than the way in which you solve a concrete problem," he said.

The System Level Design Group started in 2000, and today involves [Marculescu and several students](#). Funding comes from the National Science Foundation ([NSF](#)), Semiconductor Research Corp. ([SRC](#)), and the 17-institution Gigascale Systems Research Center ([GSRC](#)), where Marculescu contributes to the [core design technology](#) theme and alternative themes. Companies including Intel and Xilinx have offered support and feedback.



The CMU System Level Design Group includes (left to right) Jung-Chun Kao, Paul Bogdan, Radu Marculescu, Chen-Ling Chou, and Umit Ogras.

The CMU project has given SRC member companies new ideas about the design of on-chip communications, said David Yeh, director of integrated circuit and systems sciences at SRC. "The ideas developed in the NoC work can be used in multiple applications, including wireless video sensors, and this helps tie the results to a real-world application," he said.

The System Level Design Group is so named, Marculescu said, because it's aimed at algorithms and design methodologies used to design and optimize integrated systems, be they on a chip or off chip. "I'm interested in probabilistic approaches that can be used to design the systems of the future, and I'm interested in the modeling analysis and optimization of these systems," he said.

Silicon Networks

The purpose of the [SlicNets](#) project is to "provide a design methodology for future communications fabrics for multiprocessor SoCs [systems on chip]," according to Marculescu. In specific terms, this means network on chip, which seeks to replace fixed busses with packet-based networks for on-chip communications.

Today, NoCs are fairly well known, and there are commercial implementations from companies including [Arteris](#) and [Silistix](#). But in 2001 when Marculescu began researching NoCs, it was a new area. "I gave a talk at IBM in 2002 and they almost threw chalk at me because they believed in busses," Marculescu recalled. "Today there is a much better picture than in 2002. For me, it's not a question of whether network on chip will happen, it's when."

Putting numbers to claims, Marculescu and student Umit Ogras, jointly with two collaborators from Seoul National University, co-authored a 2007 paper entitled "[On-chip communication architecture exploration: a quantitative evaluation of point-to-point, bus, and network-on-chip approaches](#)." Providing direct measurements using an FPGA prototype and actual video clips, it compares area, performance, and power consumption of point-to-point, bus, and NoC implementations of an MPEG-2 encoder.

The paper concludes that the NoC architecture scales well in terms of area, performance, power consumption, and design effort. The point-to-point architecture scales poorly except for performance, and the bus-based architecture scales poorly except for area, the paper claims.

Marculescu said that the System Level Design Group has looked at three problems with respect to NoCs. One is designing the fabric, which involves optimizing the topology. Another is choosing a routing protocol. A third is optimization for power, performance, and fault tolerance through application mapping.

One of the project's early discoveries, Marculescu said, is that the way buffers are typically sized for video SoCs is wrong. A [2004 paper](#) presents a new approach to traffic modeling and synthesis for MPEG-2 video applications, and discusses the impact on buffer space allocation. The paper claims that an approach to traffic modeling based on "self-similar" or long-range dependent stochastic processes will help designers discover an optimal buffer length distribution. It also presents a method of generating synthetic traces to speed buffer simulation.

In another research direction, Marculescu proposed stochastic communications as a fundamental paradigm shift for on-chip communications networks. This concept was first proposed in 2003, and research has been ongoing. A [2007 paper](#) describes stochastic communications as a new approach for fault-tolerant NoCs.

With stochastic communications, intellectual property (IP) blocks communicate using a probabilistic broadcast scheme, similar to "[randomized gossip](#)" protocols used in databases or sensor networks. According to the paper, a stochastic communications scheme can tolerate a large number of data upsets, packet losses, and synchronization failures, while providing high levels of performance.

SRC's Yeh identified on-chip stochastic communications as the most interesting research result from the System Level Design Group. "The results show that this non-traditional approach to communication within chips has real-world, fault-tolerant benefits," he said.

"Prof. Marculescu has done a great job in developing these ideas."

Another research direction applies the social networking concept of "small worlds" to NoCs. A [2006 paper](#) entitled "It's a small world after all" shows how NoC performance can be optimized by the insertion of a small number of long-range links. The paper claims these links significantly increase the critical traffic workload level at which the network transitions from a free to a congested state, and thus reduce packet latency.

"In short, we made the cores on a chip socialize, even if they are remote," Marculescu said. "We made them look as if they were close to each other, similar to human communities."

The System Level Design Group is also doing some basic tutorial work with respect to NoCs. A [2005 paper](#) from the International Conference on Hardware-Software Codesign and System Synthesis ([CODES+ISSS](#)) outlines some key problems in NoC research, and proposes a unified representation for NoC applications and architectures. Problems described include topology synthesis, channel width, buffer sizing, floorplanning, routing, switching, scheduling, and IP mapping.

To help facilitate its research, the SlicNet project has developed NoCmap, an energy-aware mapping tool for NoC architectures, and Worm_Sim, a cycle-accurate simulator for NoC optimization. "The network simulators that were out there were not useful to us," Marculescu explained. These tools aren't open source but are available to GSRC members.

The project also built the MPEG-2 FPGA prototype. According to Marculescu, this NoC implementation achieves a 46 frames/second encoding rate for 352x288 pixel frames when using a single motion estimation module. Power consumption is about 2 Watts at 100 MHz. The design uses 10,442 slices from the target Virtex-II Xilinx FPGA.

Marculescu said the prototype uses a mesh based structure, and employs a preferential buffer assignment. "It's a completely non-democratic approach," he said. "We give more buffer space to those routes that need it the most, and take from others that don't need it. This is an optimization that's done at the router level." A second optimization puts in long-range links. The CMU NoC approach, Marculescu said, involves a "negligible" area penalty of 5 to 10 percent and does not increase power consumption. Meanwhile, he said, it provides a "huge savings" in latency. The prototype is not meant to represent a fully optimized design, he said, but as a proof of concept for an "exotic" NoC implementation.

As of today, Marculescu said he's working on four network design problems. One is voltage and frequency island control, coordination and power management. The second involves improvements to the stochastic communications approach. The third is network design based on computational physics, including quantum effects. The fourth is energy modeling and power management for wireless sensor networks.

SCANning a video network

This fourth research effort is embodied in the [SCAN](#) project, which aims to develop secure, low-power, high-performance video networks. Research here is aimed at secure and anonymous wireless routing, distributed power management, and resource-constrained video processing techniques.

Marculescu said that SCAN is another way to explore two key research interests — networks and low power. "Wireless is a huge research area, but there is a much smaller number of people working on video sensors," he said. "It is extremely demanding. For bandwidth, you have real time constraints, and when you have a limited power budget it becomes a significantly more difficult problem."

A paper given at the Design Automation and Test Conference in Europe (DATE 2007) describes [distributed power management techniques for wireless network video systems](#). It proposes two coordinated power management policies for video sensor networks that claim to be scalable as the network grows. The goal is to extend the system lifetime of wireless networks operating on limited energy resources.

The System Level Design group built a small prototype involving four nodes that monitored a parking lot at CMU. It was able to send images of cars coming in and out over a wireless network. Unfortunately, Marculescu said, the cameras were fragile and several broke.

The group also launched a research effort into electronic textiles that embed wireless sensors into wearable garments. It's not an active area of research today, but [several papers](#) have been published on this topic.

"I'm interested in long-term and high-risk research," Marculescu said. "I don't expect an immediate impact. I hope that in the end, we'll be able to create a new breed of engineers and researchers. That is the most important impact of this work."