Volume 8, Issue 2
Winter '08

CECS eNEWS

Center for Embedded Computer Systems, University of California, Irvine

Highlights

- Rainer Doemer wins NSF CAREER Award
- Minyoung Kim wins two at PhD forum

Inside this Issue:

Rainer Doemer wins NSF CAREER Award

Dr. Rainer Doemer, a CECS affiliate and assistant professor of Electrical Engineering and Computer Science at the Henry Samueli School of Engineering, has been honored by the National Science Foundation with a Faculty Early Career Development (CAREER) award. The award, along with a $400,000 grant, was in recognition of his research in modeling embedded computer systems, more specifically, "Result-Oriented System-Level Modeling for Efficient Design of Embedded Systems." The award will support a project which addresses the creation and optimization of the system model itself for effective use in existing design processes rather than the traditional method of focusing largely on simulation and synthesis from a given model. Furthermore, the results of his project will be directly applicable to established system design flows in the industry as well as fit into existing and new courses in computer engineering. Doemer's research utilizes a new model of computation, named ConcurrenC, which refines the generic capabilities of common C-based system-level description languages. The creation of the system model is automated by using computer-aided re-coding, then optimized using Result-

Oriented Modeling (ROM). In contrast to the traditional Transaction-Level Modeling (TLM), ROM offers a higher simulation speed and accuracy, which allows the system designer to quickly evaluate his design decisions. Doemer's project also explores TLM of computation, which is an

Continued on page 4

Minyoung Kim wins two awards at the RTSS-2007 PhD forum

Minyoung Kim, a CECS graduate student pursuing a PhD in Computer Science, received two awards at the 2007 IEEE Real Time Systems Symposium (RTSS) PhD Forum held in Tucson, AZ; "Best System Architecture Award" and "Best Overall Idea Award". The RTSS PhD Forum, sponsored by the National Science Foundation, is designed to discuss innovative research challenges and application ideas in deeply embedded real-time computing systems, encourage student involvement in new research directions, and reward the most innovative student ideas in an exciting emerging research field.

In this regard, it is notable that Minyoung Kim received two awards out of 30 PhD participants, who were in turn selected from a field of international applicants for this prestigious forum. Kim's research focuses on power-aware distributed embedded systems, formal methods and multimedia systems.

She is co-advised by CS Professors Nikil Dutt and Nalini Venkatasubramanian.

PROJECT PROFILE

Efficient Modeling of Complex Embedded Systems

- Gunar Schirner

Modern embedded systems are getting increasingly more complex. Current manufacturing capabilities allow integrating hundreds of millions of transistors onto a single chip. Multiple processors together with complex custom hardware accelerators and custom sensors can be squeezed onto a single die, creating a complete embedded system within a chip: a Multi-Processor System-on-Chip (MPSoC). With the enormous complexity and the ever reducing time-to-market, their design however, is extremely challenging. Early in the process, a suitable platform needs to be defined, choosing the HW/SW split, selecting processors, hardware accelerators, and communication hierarchy. Additionally, the extreme time pressure demands developing and analyzing software applications running on top the designed platform right at the same time.

Describing a system ahead of time using abstract, executable models tremendously simplifies the process. Such abstract models provide a functional and performance evaluation already early in the process. They abstract away much of the implementation detail and offer high simulation speed. Hence, such models lend themselves to investigate what-if-scenarios during the design space exploration. However, the question remains: How much detail can be abstracted away while still obtaining sufficiently accurate simulation results?

Gunar Schirner, who successfully defended his Ph.D. recently, and his advisor, Prof. Rainer Dömer, investigate into analysis and optimization of abstract models for communication and computation.

Transaction Level Models (TLM) abstract system communication to whole transactions, hiding low level details about pins, wires and waveforms. Although TLMs have been generally accepted as one solution to tackle MPSoC design complexity, their abstraction levels are not clearly defined in the current literature and the TLM trade-off has not been examined in detail.

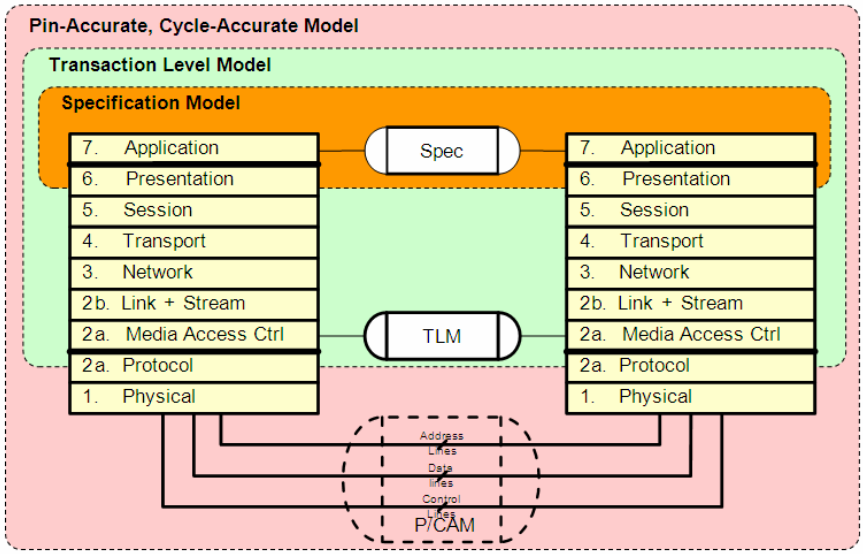

Figure 1 Abstraction layers of communication.

Possible abstraction levels can be derived from the ISO/OSI reference model, which describes layers of a traditional communication stack. Figure 1 shows two communicating application stacks captured at three different levels of abstraction: a Specification Model that captures the application, a Transaction Level Model for exploration and evaluation, and a Pin-Accurate and Cycle-Accurate Model for detailed timing evaluation. The models differ in the number of layers they implement.

How many layers are needed to produce fast and accurate results? To analyze the potential, we investigate TLMs based on the lowest three layers. These layers also determine the granularity with which the model handles arbitration and data transfers. The granularity ranges from fine grained modeling of individual bus cycles (Pin-Accurate Cycle-Accurate Model, the BFM), to a coarse grain representation of user-transactions (data blocks of an arbitrary size, like messages) in the TLM.

Our analysis, based on three common bus architectures, shows that an abstraction based on a decreasing (coarsening) granularity yields at least an order of magnitude speedup per granularity level. However, TLM abstraction results in a serious loss in accuracy, which defines the TLM trade-off. In general, a model is either fast or accurate. Our fast (coarse grain) TLM models with up to 100MBytes/s simulation bandwidth show an error of up to 47%. Accurate models, on the other hand, are slow. Our fine grained BFMs simulate with less than 0.2MByte/s bandwidth.

To escape the trade-off posed by traditional abstract model-

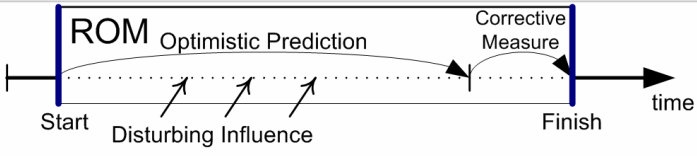

Figure 2 Generic ROM concept.

Continued on page 3

A Virtual Target for Next Generation Hardware Accelerated Multi-Core Systems

- Tony Givargis

Increasingly, the software subsystem of today's embedded devices has become the key design focus. In recognition of increased demand for features and time-to-market pressures, embedded system developers turn to software to improve design efficiency. Moreover, the competitive nature of the market has made innovation in product design and the need for product differentiation a requirement for success. In the future, such innovations will be mainly software-driven. For instance, the automotive industry expects 70% of future innovations to be enabled by software.

On the hardware side, two clear design trends have emerged. One trend is a shift from a single high performance core to multi-core platforms. A multi-core platform integrates multiple (possibly less complex) processors in a single physical package. The second trend in hardware is the use of FPGA technologies to provide flexible computational support for highly demanding applications whenever the use of general purpose microprocessors cannot satisfy performance requirements. Naturally, the fusion of multi-core platforms with dense FPGA fabric in a single physical package is likely to become the preferred execution platform of the increasingly complex embedded software.

We are of the opinion that today's research effort should focus on innovations and technologies that enable the designer to rapidly build software applications at the highest possible abstraction level with minimal knowledge of low-level details. Our efforts aim to explore and develop a parametric virtual platform virtualizing a multi-core system augmented with FPGA fabric. Such a virtual platform is intended to provide an abstract and universal underpinning for design and verification of complex embedded software.

It is an established fact that the virtual machine revolution of the past decade has had a profound impact on the design of portable, secure, and reliable software enhancing the human experience when using the Internet, personal computers, PDA devices, and so on. Similarly, we envision the virtualization of the next generation compute platforms as an important step towards the design of tomorrow's ubiquitous, high performance, high confidence, and smart embedded devices.

Project Profile:: Cont'd from page 2

ing, our research introduces a novel modeling technique: Result Oriented Modeling (ROM). ROM is a modeling approach similar to TLM that hides internal states of bus communication. Unlike traditional TLMs that incrementally model the bus state, ROM avoids modeling of intermediate bus states.

Figure 2 illustrates the generic concept. ROM uses an optimistic prediction approach. Right at the beginning of a transfer, it predicts the end result. ROM constructs a bus schedule taking pending transactions into account, and determines the earliest possible finish time for the requested transaction. While waiting for the predicted time, ROM records any disturbing influence (i.e. preemptions by higher priority transfers). If a preemption has occurred, ROM re-calculates the bus schedule at the end of the predicted time, adjusts the prediction as a corrective measure, and waits again to achieve 100% accuracy.

Our experimental results, based on two complementary busses: the on-chip AMBA AHB and off-chip CAN, reveal ROM's tremendous benefits. Figure 3 summarizes the results by showing the inaccuracy in simulated timing (y-axis) in dependency of the achievable simulation speed (x-axis).

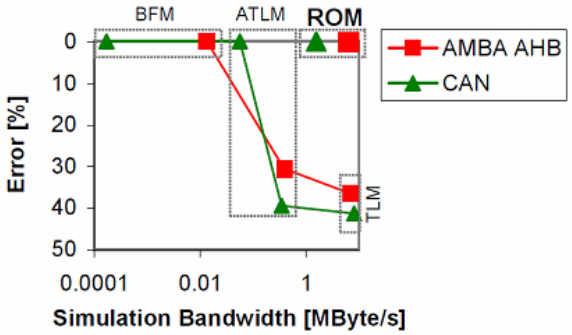

Figure 3 ROM beats the TLM Trade-Off

Traditional models (BFM, ATLM and TLM) are either accurate but slow with less than 0.4 MBytes/s, or they are fast (i.e. the TLMs) but exhibit an error of more than 35%. ROM, on the other hand, escapes the TLM trade-off and delivers both highest speed (reaching near TLM performance) and 100% accuracy at the same time. Abstract models based on the ROM technique allow rapid design space exploration with high fidelity. They also enable utilizing abstraction benefits in real-time applications, where 100% timing accuracy is required.

To conclude, Transaction Level Modeling has become a main approach to tackle the complexity challenges in today's system level design. This research work provides guidelines for the model designer when developing abstract communication models. It guides in categorization and systematically analyzing communication models. It introduces ROM, a novel technique for communication modeling, which escapes the traditional TLM trade-off. The analysis of both communication and computation models aid the system designer to navigate the TLM trade-off effectively and choose the most suitable model for a given application.

New Student Profiles

WeiWei Chen

My name is Weiwei Chen. I am now a first year Ph.D. student from the Department of Electrical Engineering and Computer Science at UCI. I got my masters degree in computer engineering from Shanghai Jiao Tong University in 2007, and my bachelors degree in computer science and engineering from the same university in 2004. I was born in Shanghai, China, and lived there until I came to the U.S.

In September 2007, I joined CECS as a Ph.D. student and currently work with Professor Rainer Doemer in the laboratory for Embedded Computer System (LECS). My research work focuses on methodologies of system level modeling of embedded computer systems. The development complexity of embedded system design is growing rapidly due to its great computational and high performance demands. System level design becomes more and more important in order to meet these competitive requirements. System modeling methodology provides the proper abstraction for each level of system level design in order to create and optimize the system model for effective use in design process. Our research work will utilize a new model of computation which will refine the generic capabilities of common C-based system-level description languages. It will be applied in a new automatic system-level design flow with our recoding tool and modeling methodologies. Thus we could build a robust and efficient embedded computer system in a short time. I am very proud to be able to do my Ph.D. work with so many intelligent and outstanding researchers and scientists here in CECS, UCI.

Jesse Dannenbring

I am currently working under Professor Dutt on a SystemC implementation of the H.264 Decoder. I was previously at the University of Washington in Seattle, and am enjoying the change of weather! My transition into graduate classes at UC Irvine has been a smooth one, and I have had some great teachers this quarter.

Visitor Profile

Yongjin Ahn is visiting CECS for one year, starting December 2007.He finished his graduate studies in EECS at Seoul National University, Korea under the supervision of Prof. Kiyoung Choi who also was a CECS visitor, and received a Ph.D. degree in 2007. He then worked in Prof. Choi's lab as a postdoctor for 7 months and developed a static analysis based system level design tool starting from a process network model.

Yongjin is interested in research in system level design methodology for MPSoC design. In particular, he has been working on a model of computation, application to architecture mapping and design space exploration. "I am very pleased to work in Prof. Gajski's lab, one of the best known labs for system-level design methodology. He has been developing an efficient system-level design tool," says Yongjin. He adds, "I am glad to join the design team and am very thankful to the team for the knowledge I gain from them. In the team, everyone is very kind and energetic. I hope my experience can help the team in developing the tool."

Yongjin is very excited about his visit to Irvine. He says, "this is my first visit in Irvine. Irvine is a very safe and beautiful city. I am happy to be able to work in such a good city. I am trying to learn American culture and excited in many interesting things."

Rainer Doemer wins award

:: Cont'd from front page

emerging research area promising to build better products in a shorter time.

The CAREER award is the NSF's most prestigious award to new faculty members. The CAREER program recognizes and supports the early career-development of those teacher-scholars who are most likely to become academic leaders of the 21st century. CAREER awardees are selected on the basis of creative, career-development plans that effectively integrate research and education within the context of the mission of their institution. The CAREER award plans should build a firm foundation for a lifetime of integrated contributions to research and education.

CECS extends congratulations to Professor Rainer Doemer on receiving this prestigious NSF CAREER award and the accompanying recognition.

PUBLICATIONS

The following papers were published by CECS affiliates during Winter 2008 quarter

Focus	Title, Author, Publication
Fourier Descriptors	Weisheng Duan, Falko Kuester, Jean-Luc Gaudiot, and Omar Hammami, Automatic Object and Image Alignment using Fourier Descriptors, Image and Vision Computing, in press (2008)
Data Distribution	Jung-Yup Kang, Sandeep Gupta, and Jean-Luc Gaudiot, "An Efficient Data-Distribution Mechanism in a PIM (Processor-In-Memory) Architecture Applied to Motion Estimation", IEEE Transactions on Computers, Vol. 57, No. 3, March 2008
Energy Awareness	Minyoung Kim, Sudarshan Banerjee, Nikil Dutt, Nalini Venkatasubramanian, "Energy-aware Co-synthesis of Real-time Multimedia Applications on MPSoCs Using Heterogeneous Scheduling Policies", ACM Transactions on Embedded Computing Systems (TECS). 7(2): Article 9, 2008.
Cross Layer System Adaptation	Minyoung Kim, Mark-Oliver Stehr, Carolyn Talcott, Nikil Dutt, Nalini Venkatasubramanian, "Constraint Refinement for Online Verifiable Cross-Layer System Adaptation", IEEE/ACM Design Automation and Test in Europe (DATE '08), Mar. 2008, Munich, Germany.
Cross Layer System Adaptation	Minyoung Kim, Nikil Dutt, Nalini Venkatasubramanian, Carolyn Talcott, "xTune: Online Verifiable Cross-Layer Adaptation for Distributed Real-Time Embedded Systems", IEEE International Real-Time Systems Symposium (RTSS'07) Ph.D. Forum, December 2007, Tucson, AZ, USA. Also published as SIGBED Review, Volume 5, Number 1, January 2008 Special Issue on the RTSS Forum on Deeply Embedded Real-Time Computing
On Chip Communication	S. Pasricha, Y. Park, S. Pasricha, Y. Park, F. Kurdahi, N. Dutt, "Incorporating PVT Variations in System-level Power Exploration of On-Chip Communication Architectures", IEEE VLSI Design Conference (VLSID 2008), Bangalore, India, January 2008
On Chip Communication	S. Pasricha, N. Dutt, "ORB: An On-chip Optical Ring Bus Communication Architecture for Multi-Processor Systems-on-Chip", IEEE Asia & South Pacific Design Automation Conference (ASPDAC 2008), Seoul, Korea, January 2008
Bus based Communication	S. Pasricha, N. Dutt, M. Ben-Romdhane, "Fast Exploration of Bus-based Communication Architectures at the CCATB Abstraction", ACM Transactions on Embedded Computing Systems (TECS), Vol. 7, No. 2, February 2008
Microprocessor Design	Won Woo Ro and Jean-Luc Gaudiot, A Complexity-Effective Microprocessor Design with Decoupled Dispatch Queues, Parallel Computing, in press (2008)
Result Oriented Modeling	Gunar Schirner and Rainer Dömer, "Introducing Preemptive Scheduling in Abstract RTOS Models using Result Oriented Modeling", In Proceedings of Design Automation and Test in Europe (DATE), Munich, Germany, March 2008.
Software for MPSoCs	Gunar Schirner, Andreas Gerstlauer, Rainer Dömer, "Automatic Generation of Hardware dependent Software for MPSoCs from Abstract System Specifications", Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC), Seoul, Korea, January 2008.
Power Profile Correlation	Love Singhal, Sejong Oh and Eli Bozorgzadeh, "Statistical Power Profile Correlation for Realistic Thermal Estimation", IEEE Asia and South Pacific Design Automation Conference (ASPDAC), Seoul, Korea, January 2008
FPGA	S. Sirowy, G. Stitt, and F. Vahid. "C is for Circuits: Capturing FPGA Circuits as Sequential Code for Portability", International Symposium on FPGAs, 2008.
Thread Accelerators	G. Stitt and F. Vahid, "Thread Warping: A Framework for Dynamic Synthesis of Thread Accelerators", International Conference on Hardware/Software Codesign and System Synthesis (CODES/ISSS), 2007, pp. 93-98

Continued on page 6

***CECS** Mission Statement:*

To conduct leading-edge interdisciplinary research in embedded systems emphasizing automotive, communications, and medical applications, and to promote technology and knowledge transfer for the benefit of the individual and society.

Primary Contact
Melanie Kilian
Center for Embedded Computer Systems
University of California, Irvine
Email: mbkilian@uci.edu

Publications :: Cont'd from page 5

Focus	Title, Author, Publication
Simulation-Guided Model Checking	G. Madl and N. Dutt, "Real-time Analysis of Resource-Constrained Distributed Systems by Simulation-Guided Model Checking," ACM SIGBED Review: Special Issue on the RTSS Forum on Deeply Embedded Real-Time Computing, Volume 5, Number 1, January 2008.
Exploration of SMT Processors	D. Kannan, A. Gupta, A. Shrivastava, N. Dutt, and F. Kurdahi, " PTSMT: A Tool for Cross-Level Power, Performance and Temperature Exploration of SMT Processors," Proceedings of the 2008 International Conference on VLSI Design, Hyderabad, India, January, 2008.
Cache Architectures	A. Shrivastava, I. Issenin , N. Dutt, "A Compiler-in-the-Loop Framework to Explore Horizontally Partitioned Cache Architectures," Proceedings of ASPDAC-2008, January 2008.
SoCs	N. Dutt, "Quo Vadis, BTSoCs (Billion Transistor SoCs)?" Panel Position Statement, Proceedings of ASPDAC-2008, January 2008.
Design Methodology	N. Dutt, "Design Methodology for Memory-aware NoC Exploration and Design," Special Session on The Memory Challenge in NOC based Systems, Proceedings of the 2008 Conference on Design, Automation and Test in Europe (DATE 2008), March 2008.
Routing of VLSI chips	A. Gupta, F. Kurdahi, N. Dutt, K. Khouri, M. Abadir., "Thermal Aware Global Routing of VLSI Chips for Enhanced Reliability," Proceedings of ISQED 2008, March 2008.