

Revolutionising multicore architecture design

Professor Radu Marculescu has recently been formally recognised for his groundbreaking work in redesigning the on-chip communication for embedded multicore systems. Here, he talks about his ambitious approach to research and the potential applications of his studies to computing

To start, how are you working towards the overall aims of your project?

In short, my work is about thinking outside the box, taking risks, and connecting science and engineering. We are living in one of the most creative periods in the history of engineering – computer engineering in particular – so we cannot rely on traditional methods. My project is living proof of how far we can go if imagination is the only limit. When one hears about quantum mechanics, energy levels, one thinks of atoms, electrons and things of that nature. However, this project demonstrates that this view is too narrow. Indeed, we have shown that quantum-like behaviour can be observed at a macro-scale in modern multiprocessor computing platforms and, more importantly, that it can be exploited to design better computing systems.

Our adventure in 'quantum-land' demonstrates that this is not the exclusive territory of physicists, but engineers can also partake in this endeavour. Despite the specific focus of this project, our hope is that our results can impact other research areas by improving the level of understanding of networking concepts in order to design and control complex systems.

Has your interest in this field changed over time?

My entire professional career has evolved around two focal points: low power and network-based design of computing systems. Over the years, I realised that these may be the two most powerful design paradigms for future computing systems. Placing the network at the very centre of my research was, and still is, a major priority for me. As time went by, I became interested in a variety of networks occurring in integrated, biological and social systems. Behind this research lie deep concepts and disciplines such as statistical physics, quantum physics, criticality phenomena, stochastic communications and the 'small worlds' idea from social networking. Fortunately, Carnegie Mellon University is

a great place intellectually; I have always been surrounded by wonderful students and collaborators who have helped me continuously shape my own ideas, deepen my understanding and, more importantly, bring some of these ideas to life. I cannot think of a greater blessing for a researcher, so I owe each and every one of them a great deal for helping me to re-invent myself every 10 years or so.

In what way will your latest results be useful to computing? What are the potential applications?

There are many. Firstly, my hope is that this research on multicore systems will bring some new blood into computer engineering from a modelling and optimisation perspective. Multicore systems will become central to the entire computing infrastructure of the future. Designing such systems will require a new science involving more than mastering very-large-scale integration and standard computer system design. We will need to use concepts from statistical physics, critical phenomena, and small world networks and so on.


In a small way, our project has already shown how such systems can be used in multicore platform design, but they can also be exploited in software, data centres and cloud computing. It truly is science and engineering at their best.

Are there any particular hurdles to technological development?

The greatest challenge is gaining a comprehensive understanding of systems. In current technologies, we can still build things that we cannot fully understand or explain in precise terms. In fact, the gap between what we can build in silicon and what we can rigorously explain or formally verify is huge. This needs to change in future.

Where do you see this field of technology progressing? Do you have any projects in the pipeline?

At the moment, we are conducting several interesting projects in the areas of embedded systems, biological systems and social networks. This subsumes a large intellectual diameter encompassing everything from power management and software for multicore systems to pacemaker optimisation and drug delivery systems using bacteria-based micro-robots, Twitter and social network analysis. However, it is much harder to make an impact in these areas. Real experimentation is more complex and important than in standard computing systems, so the entire dynamic of the research process changes and things can take a lot longer to pan out. It is a fantastic learning experience, a self-discovery process, and this makes everything interesting and worthwhile and, again, it is about having great students around who can provide motivation and new ideas when things seem to stall.



The chip is the network

Pioneering research at **Carnegie Mellon University** in Pittsburgh, USA, is challenging the dominant paradigms of on-chip communication. The researcher's Network-on-Chip technologies are already being applied commercially and there is optimism that eventually they will fully replace traditional approaches

AS BUSINESS, HEALTHCARE, governance, personal organisation and socialising become increasingly reliant on computer technology, the concept of the network has established its ubiquity. In response, a new field of research has sprung up, setting as its target the exploration and examination of the possibilities of improved network design. Within this field, a growing number of researchers are seeking to discover and explain the properties which affect a network's structure and behaviour, and the mathematical models which can help to predict and guide network dynamics and implementation in future multicore platforms. At the forefront of this field is Dr Radu Marculescu, Professor of Electrical and Computer Engineering at Carnegie Mellon University (CMU) in Pittsburgh, USA. Along with an outstanding team of support staff and students, Marculescu is making great conceptual and practical strides in the quest for more powerful approaches to Network-on-Chip based communication.

MODERN MAINSTREAM TECHNOLOGY

Currently, complementary metal-oxide-semiconductor (CMOS) technology is dominant almost everywhere, including applications such as digital logic, microprocessors and memories. "Interestingly, what made CMOS the dominant technology was its low power consumption," Marculescu explains. "However, power dissipation has now become one of the main concerns due to the move towards nanoscale devices and the general interest in portable electronics." Significant shifts in the field of network research are expected in coming years

making Marculescu's fusion of science and engineering, and its application in investigating non-traditional approaches to system-level design, more relevant than ever.

THREE DESIGN PARADIGMS

In order to approach their work in the most comprehensive and rigorous way possible, Marculescu and his fellow researchers have approached network design through the lens of each of the three significant paradigms that prevail over the field: from deterministic, probabilistic and statistical physics. The first of these concerns itself with the natural representation of networks as graphs and analysing deterministic approaches to Network-on-Chip (NoC) design; the second assesses network design through a probabilistic framework aimed at representing and optimising networks; and the third approach considers the network as a thermodynamic system, seeking to characterise network traffic using statistical physics. The formal structure of their work has allowed Marculescu and his students to identify the key limitations to the solutions being proposed by each approach, as well as highlighting a number of open-ended questions around which to base future research.

INTEGRATION OF COMPONENTS

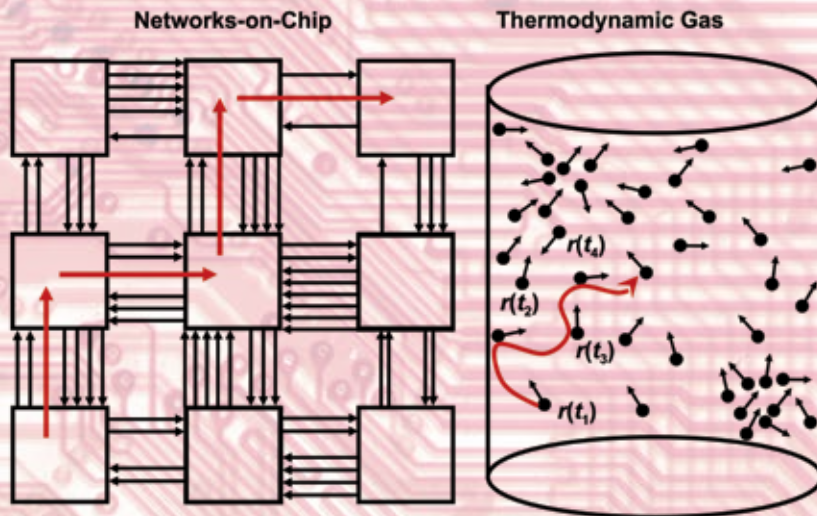
As part of their work, Marculescu and his colleagues are seeking to develop greater capabilities in System-on-Chip (SoC) technology. Recent advances in CMOS have enabled researchers to integrate tens or

even hundreds of individually programmable elements onto the same SoC, along with a significant amount of dedicated memory. These elements can range from traditional central processing units and digital signal processing cores to highly sophisticated video and audio processing technology. "Since the bus-based communication architectures are characterised by poor scalability, low performance and high power (or energy) consumption, the NoC approach has emerged as the only truly scalable communication architecture for SoC systems," Marculescu elaborates.

NoC architectures are considered superior in terms of their ability to replace global wires with a regular communication structure, and this presents a huge opportunity for the researchers at CMU to explore better scalability, power consumption and predictable behaviour in networks, in the face of a variety of workloads and parameter variations.

IMPROVING MATHEMATICS EDUCATION

In order for the full capabilities of the systems to be utilised, Marculescu believes that both current researchers and the next generation of experts must be highly proficient mathematicians: "Personally, I find the mathematics we teach students at engineering schools is pretty limited, especially with regards to real applications in modelling and optimising of modern computing systems". In traditional mathematical teaching the practical applications and potential for utilisation in the design of computer platforms are often not



INTELLIGENCE

THE CHIP IS THE NETWORK: RETHINKING THE THEORETICAL FOUNDATIONS OF MULTICORE ARCHITECTURE DESIGN

OBJECTIVES

This project aims at introducing a new statistical-physics approach for performance analysis and optimisation in multiprocessor SoCs. More precisely, it introduces a completely new mathematical framework by building on an analogy between a thermodynamic gas and the information flow in a communication network. The new framework can enable the development of more efficient multiprocessor SoCs which have a dramatic impact via applications, ranging from entertainment and gaming to security and bio-engineering. More broadly, the results can significantly improve the level of understanding of networking concepts needed to design and control complex systems.

KEY COLLABORATORS

Professor Paul Bogdan, Department of Electrical Engineering – Systems, University of Southern California, Los Angeles • **Professor Diana Marculescu**, Department of Electrical and Computer Engineering, Carnegie Mellon University, Pennsylvania • **Drs Umit Ogras** and **Mike Kishinevsky**, Intel Corp., USA

FUNDING

National Science Foundation – award no. CCF 0916752 • Synergistic Projects • Semiconductor Research Corporation 2011-HJ-2168 • Semiconductor Research Corporation 2012-TJ-2232 • National Science Foundation – award no. CNS 1128624

CONTACT

Professor Radu Marculescu
Principal Investigator

Electrical and Computer Engineering
Carnegie Mellon University
Hamerschlag Hall 2108
5000 Forbes Ave, Pittsburgh
Pennsylvania 15213-3890, USA

T +1 412 268 8710
E radum@ece.cmu.edu

www.ece.cmu.edu/~sld

RADU MARCULESCU is a professor in the Department of Electrical and Computer Engineering at Carnegie Mellon University, USA. He received his PhD in Electrical Engineering from the University of Southern California in 1998. Marculescu has received several 'best paper' awards in the area of design automation and embedded systems design. He has been involved in organising several international symposia, conferences, workshops, as well as being guest editor of special issues in archival journals and magazines. His research focuses on modelling and optimisation of embedded systems, cyber-physical systems, and biological systems. Marculescu is a Fellow of the Institute of Electrical and Electronics Engineers.

fully explained, leading it to become heavily theoretical and even dry. This form of teaching often relies on a number of simplifying assumptions which often fall down in practice; Marculescu's work in the field of multiprocessor SoC platforms has demonstrated the limitations of such an approach. This situation predicts that this situation could become even worse when the focus shifts to the development of thousand-core systems, data centres and cloud systems, and for this reason he is committed to aiding students to discover and appreciate the fundamental science which underpins engineering. This, it is hoped, will equip future generations of researchers with the necessary mathematical skills to tackle the challenges faced in systems design.

MOST INFLUENTIAL PAPER AWARD

Marculescu's research has widely been acknowledged as pioneering, breaking new ground in the development of more advanced and intelligent NoC. In January 2013, his work was formally recognised with the 10-Year Retrospective Most Influential Paper Award by the Asia and South Pacific Design Automation Conference. In part, this award was granted to Marculescu in recognition of his very early identification of a great problem facing network design. He and his fellow researchers began working on NoC technology in 2002 when it was still in its infancy, and considered highly experimental: "I remember giving talks at some major semiconductor companies, and they

almost threw chalk at me because they believed in busses," he recalls.

Today, the picture is very different, and there are already a number of commercial implementations available. The idea that NoC technology would catch on was seeded in 2003 paper in which Marculescu was the very first to address issues surrounding energy-aware mapping of IPs/cores onto a generic multicore platform, and to propose a solution for NoC architecture. Funded by the National Science Foundation and the Semiconductor Research Corporation, the project was able to come to full fruition, and the results it has yielded have gone on to have an enormous influence on follow-up work, both within academia and in industry.

'WILD AND CRAZY IDEAS'

Alongside his groundbreaking research projects, one of Marculescu's aims is to teach subjects at CMU that will inspire the next generation of researchers to become creative and experimental thinkers. This is typified by his flagship class entitled 'Wild and Crazy Ideas', a graduate-level class aimed at exploring concepts which are expected to dictate the design of future computational platforms, extending beyond the use of traditional silicon. "Specifically, this class discusses topics that may have a significant impact five or 10 years from now, examining the merits of such ideas and identifying the larger context to which they belong," Marculescu explains. As part of this process, students gain a thorough grounding in the science behind ideas such as stochastic modelling for nanotechnologies; critical phenomena in information networks; swarms and group behaviour; and emerging multi-scale behaviour software. The module reflects Marculescu's passion for long-term, high-risk research, and through it he hopes to create a new generation of engineers and researchers who can share his commitment to thinking outside the box to create technologies with a major social contribution.

Marculescu's research has widely been acknowledged as pioneering, breaking new ground in the development of more advanced and intelligent Networks-on-Chip

