

Memristors and More

Lieven Eeckhout
Ghent University

Welcome to *IEEE Micro*'s special issue on memristor-based computing. Memristors were a theoretical concept (proposed by circuit theorist Leon Chua in 1971) for over three decades until HP Labs built a physical device in

2008. Memristors have gained substantial interest since then because of their nonvolatility; a memristor remembers its past, even when the supply voltage is turned off. Moreover, memristors facilitate co-mingling of storage and computation. A number of memristor applications have been proposed and evaluated, including resistive memory, memristor-based logic, and even neuromorphic computing.

Over the past decade, numerous research programs have focused on how to leverage memristors to enhance microprocessor design. This special issue contains four peer-reviewed articles and two Expert Opinion articles on the topic. I want to thank the guest editors, Lizy K. John and Earl Swartzlander from the University of Texas at Austin, for having done such a wonderful job in compiling this special issue. I refer you to their guest editorial for an introduction to the theme.

In addition, this issue contains four more articles. Two articles cover a topic that *IEEE Micro* has highlighted very recently, namely automotive computing (January/February 2018). Wang and co-authors from the Technical University of Munich and the Sun Yat-sen University propose a heterogeneous electronic control unit (ECU) for automated driving. Multiprocessor SoCs (MPSoCs), GPUs, and FPGAs can be plugged into the modular ECU to support co-execution of safety-critical and non-safety-critical applications through appropriate partitioning mechanisms. Kornaros, Tomoutzoglou, and Coppola (from the Technological Educational Institute of Crete and STMicroelectronics) focus on security and privacy issues in cyber-enabled automotive systems. They combine on-chip network physical isolation with cryptographic techniques at the system level to provide authentication and confidentiality.

The last two articles cover real processor designs. Maruyama and co-authors from Fujitsu present the SPARC64 XII 12-core server processor. They describe how this processor improves performance by a factor 2.3 to 2.9 over the previous-generation SPARC64 X+ through a variety of innovations. Shin and colleagues from the Korea Advanced Institute of Science and Technology (KAIST) describe DNPU, a heterogeneous multi-core processor architecture optimized for convolutional neural networks (CNNs) and recurrent neural networks (RNNs). DNPU's high energy-efficiency enables embedded deep-neural-network processing in mobile devices.

With that, I would like to wish a happy reading, as always.

ABOUT THE AUTHOR

Lieven Eeckhout is a professor in the Department of Electronics and Information Systems at Ghent University. Contact him at lieven.eeckhout@ugent.be.