

Challenges in Computer Architecture Evaluation

pp. 30-36

Kevin Skadron, Margaret Martonosi, David I. August, Mark D. Hill, David J. Lilja, and Vijay S. Pai

Reasoning about today's tremendously complex computer systems is difficult and developing them is expensive. Detailed software simulations are thus essential for evaluating computer architecture ideas. Industry uses simulation extensively during processor and system design as the easiest and least expensive way to explore design options.

Unfortunately, constructing accurate models of modern computer systems is becoming harder and more time-consuming, while the effort required to develop high-fidelity simulation tools typically yields few academic rewards. Without funding and promising prospects for academic recognition, research and development in these areas will likely languish.

Software Reliability from the Customer View

pp. 37-42

Alan P. Wood

The relationship between software defects and failures is not one-to-one. Some defects remain undiscovered and never cause a failure, but a single defect can cause many failures.

Many researchers have offered solutions to this problem, but their approaches typically reflect only the developer view of software reliability: how to predict and prevent the underlying defects. At Hewlett-Packard's Non-Stop Enterprise Division, researchers augment their defect-prevention activities with analyses of what their customers experience: software failures. NED provides customers with a regularly updated suite of hardware and software products for business-critical applications, then tracks the failure rate for each version so that it can measure the customer experience with that update.

The Future of Nanocomputing

pp. 44-53

George Bourianoff

Projections that show the end of CMOS scaling will occur around 2016 have spurred interest in emerging alternative technologies, particularly nanotechnologies, that promise to extend Moore's law beyond 2016.

The semiconductor industry has already entered the nanotechnology world: In 2000, it introduced the 130-nm node with a 70-nm gate-length feature size, followed in 2002 by the 90-nm node featuring a critical dimension of 50 nm. Industry leaders see new scalable technologies emerging from the novel alternative architectures and devices being proposed today that will take us through multiple processor generations for another 30 years or so.

Reviewing the lessons learned in the semiconductor industry over the past few decades helps us understand the emerging technologies and suggests some criteria for bringing current research efforts into the realm of high-volume manufacturing.

Programmable Stream Processors

pp. 54-62

Ujval J. Kapasi, Scott Rixner, William J. Dally, Bruceek Khailany, Jung Ho Ahn, Peter Mattson, and John D. Owens

The demand for flexibility in media processing motivates the use of programmable processors. However, very large-scale integration constraints limit the performance of traditional programmable architectures. In modern VLSI technology, computation is relatively cheap—thousands of arithmetic logic units operating at multigigahertz rates can fit on a modestly sized 1 square centimeter die. Yet delivering instructions and data to those ALUs is prohibitively expensive.

The Imagine media processor validates the hypothesis that careful management of bandwidth and parallelism, from the programming language to the hardware, results in both high performance and high performance per unit of power.

High-Level Language Abstraction for Reconfigurable Computing

pp. 63-69

Walid A. Najjar, Wim Böhm, Bruce A. Draper, Jeff Hammes, Robert Rinker, J. Ross Beveridge, Monica Chawathe, and Charles Ross

Reconfigurable computing systems typically consist of an array of configurable computing elements. The computational granularity of these elements ranges from simple gates to complete arithmetic logic units, with or without registers. A rich programmable interconnect completes the array.

Performance evaluation of Simple-Assignment C, a high-level, algorithmic language for one-step compilation to host code and field-programmable-gate-array configuration codes, has just begun, with the authors porting the system to a more complex board that contains three FPGAs.

Penny Tag Technologies for Removable Data Storage

pp. 70-76

Fred Thomas

The authentication and identification technologies typically used in retail and manufacturing have also found applications in removable data storage. For example, in addition to authenticating parts and preventing counterfeiting in various industries, cartridge identification systems also can protect a drive from damage that inserting a foreign object might cause.

Retail and removable data storage applications differ primarily in their cost constraints: The authentication and identification of removable data storage cartridges must be automated at a very low cost. Penny tag technologies and their associated low-cost automated detection systems for removable data storage cartridges provide innovative identification and authentication methods that have a high overall security-to-cost quotient.