

Architectures for Silicon Nanoelectronics and Beyond

pp. 25-33

R. Iris Bahar, Dan Hammerstrom, Justin Harlow, William H. Joyner Jr., Clifford Lau, Diana Marculescu, Alex Orailoglu, and Massoud Pedram

A candidate to replace CMOS semiconductors, nanoelectronics could address some of the challenges facing the semiconductor industry. However, it also introduces new problems.

Molecular-scale computing will likely allow additional orders-of-magnitude improvements in device density and complexity, which raises three critical concerns: how to use these huge numbers of devices, how to modify design tools and methodologies to accommodate radical new ways of computing, and how to produce reliable, predictable systems from unreliable components with unpredictable behavior.

Designing for Software's Social Complexity

pp. 34-39

José Luiz Fiadeiro

The complexity that software and system designers and engineers complained about in 1994 has changed fundamentally. A decade ago, developers were concerned with physiological complexity—systems whose size and makeup required decomposition into smaller, more manageable and reusable parts that any skilled programmer could assemble. Yet services are inherently social, which means that complexity now has two sides: physiological and social.

Reengineering the Internet for Better Security

pp. 40-44

Manoj Parameswaran, Xia Zhao, Andrew B. Whinston, and Fang Fang

Ironically, what has made the Internet so successful—its open and decentralized structure—also sustains malicious online activity. Information on newly discovered vulnera-

bilities propagates quickly, and tools to launch ever more sophisticated attacks are readily accessible.

Most controls implemented by Internet service providers target inbound traffic. However, there is no similar economic incentive to control outbound traffic, as the potential damage is to other networks. This lack of accountability derives from the decentralized nature of datagram routing in the Internet and its decentralized organizational structure.

AQUA: An Amphibious Autonomous Robot

pp. 46-53

Gregory Dudek, Philippe Giguere, Chris Prahacs, Shane Saunderson, Junaed Sattar, Luz-Abril Torres-Mendez, Michael Jenkin, Andrew German, Andrew Hogue, Arlene Ripsman, James Zacher, Evangelos Milios, Hui Liu, Pifu Zhang, Martin Buehler, and Christina Georgiades

The aquatic environment is almost ideal for autonomous robot development. First, it provides a range of real tasks for autonomous systems to perform, including ongoing inspection of reef damage and renewal; tasks in the oil and gas industry; and aquaculture. Second, operating in the water requires robust solutions to mobility, sensing, navigation, and communication.

Experimental testing of the AQUA robot to date has concentrated on the independent evaluation of individual components. Over the next few years, the authors anticipate integrating the various sensor components within the robot itself.

The Strangest Thing About Software

pp. 54-60

Tim Menzies, David Owen, and Julian Richardson

The authors propose that software works because internally it is surprisingly simple. Recent research results from artificial intelligence offer much support for believing this. AI

has discovered certain previously unrecognized regularities that developers can use to quickly find solutions. A careful reading of the software engineering literature shows that these regularities also appear in conventional software. For problems with those regularities, random search can find much of what we can find via complex and costly methods.

Process Query Systems

pp. 62-70

George Cybenko and Vincent H. Berk

In the context of computer security, situational awareness typically means knowing which monitored systems are under attack and the nature of those attacks. For a physical sensor network within a building or extending across a geographic region, it might mean being cognizant of certain objects and activities. Extracting meaningful and actionable information from these events, however, remains a challenge. A new algorithmic and software paradigm, process query systems, addresses these event-processing challenges.

Steps Toward a Science of Service Systems

pp. 71-77

Jim Spohrer, Paul P. Maglio, John Bailey, and Daniel Grubl

The service sector accounts for more than 80 percent of the US gross domestic product and employs a growing share of the science and engineering workforce. Yet it's one of the least-studied areas of the economy. Some see economics, operations research, industrial engineering, or the science of complex systems as the appropriate starting point for a general services theory. Others contend that the pervasiveness of services creates a need for many specific disciplines. An interdisciplinary effort called Service Science, Management, and Engineering provides a solution that falls between those two approaches.