SPECIAL ISSUE ON MOBILE SYSTEMS

Calin Cascaval
Qualcomm Research
Silicon Valley

Mobile devices are changing the landscape of computing and how people perceive and interact with each other and with the physical world. Smartphones and tablets are now ubiquitous, permanently available, and always connected. Increasingly, we rely on mobile connectivity for most of our daily tasks, whether shopping, meeting with friends, taking pictures, or checking in for work. To continue to provide the capabilities for this level of usage, mobile systems need to meet several challenges.

First, users expect a smooth experience from their devices and increasingly demand their privacy to be protected. Mobile systems are used by everyone, and consumers are ever more demanding, requiring high availability and fast response time. They expect mobile devices to continue to function in the face of adversarial attacks, such as malware, and protect the privacy of the user data and the networks accessed by these devices. Understanding the behavior of applications, isolating the execution of malware, and providing different personalities for one device are real problems that consumers take for granted, and which are yet to be solved.

Second, increases in the computational capabilities are stressing the limits of power, energy, and thermal capabilities of mobile devices. These limits are imposed by the constrained packaging and the batteries available on devices. Since these are unlikely to change, solving these challenges requires new approaches to designing both hardware and software optimizations for mobile computing. Full-system heterogeneous computing is a promising direction, attractive from a hardware perspective, but it has a huge impact on application development because it requires very specialized skills to code.

IEEE Micro is fortunate to present in this issue a palette of articles that address many of these challenges. The articles run the gamut from designing specialized accelerator cores, to discovering vulnerabilities in Android applications, from a full system analysis of web browsers with respect to performance and energy, to virtualizing the system in order to enable enterprise computing, and to bringing new memory technologies to improve application startup time. The cross-section of these articles barely scratches the surface of mobile computing, but it serves well to highlight the complexity of this computing paradigm.

Accelerator cores

We begin this issue with an article by Mircea Horea Ionică and David Gregg, "The Movidius Myriad Architecture's Potential for Scientific Computing," which looks at saving power using a specialized core. They analyze the energy efficiency of the Movidius Myriad architecture on dense matrix multiplication, and demonstrate how specialization provides a significant performance/watt advantage. Single-precision matrix multiply is a key kernel for mobile computing because it is used in emerging applications such as scene detection, face recognition, and several other computer vision algorithms, as well as machine learning and game physics simulations. The authors also demonstrate that the mobile energy efficiency designs are extremely valuable for high-performance and scientific computing, because large systems have a significant power and energy problem.

Published by the IEEE Computer Society

User experience and battery life

The next two articles look at technologies and system design to improve user experience.

In "Accelerating Application Start-up with Nonvolatile Memory in Android Systems," Hyojong Kim et al. study the memory behavior of Android devices using LMBench and AndroBench. They propose using nonvolatile memory as a secondary memory between DRAM and Flash to store frequently used shared libraries and applications to improve start-up time.

"The Role of the CPU in Energy-Efficient Mobile Web Browsing" by Yuhao Zhu, Matthew Halpern, and Vijay Janapa Reddi analyzes system behavior while web browsing. The authors use page load time as a metric of browsing performance and evaluate it with respect to network latency and CPU frequency. A key observation is that as network latencies improve, by moving to faster networks such as LTE and Wi-Fi, the CPU plays an increasingly important role. However, even so, the most benefits with respect to power and energy are obtained when the optimization takes into account full system behavior—for example, adapting the behavior of different components (CPU and network in their case) to match each other.

Security and privacy

The next two articles fall in the realm of security and privacy. As devices become ubiquitous, they become an attractive platform to either attack directly or use to penetrate enterprise networks. "A Virtualization Solution for BYOD with Dynamic Platform Context Switching" by YaoZu Dong et al. presents a virtualization approach to enable a device to have different "personalities" with negligible performance degradation. Their vNative design allows a foreground virtual machine access to all the device resources, while a set of background virtual machines are suspended and ready to run immediately when triggered by certain factors, such as connecting to an enterprise network.

"VulHunter: Toward Discovering Vulnerabilities in Android Applications" by Chenxiong Qian et al. proposes a static-analysis framework to detect vulnerabilities in Android applications. The authors' framework, VulHunter, constructs an application property graph using required permissions, intents, and other Android application properties. Vulnerabilities are modeled as property graph traversals that reach method calls that should be infeasible given the application permissions. Applying this analysis on 577 applications from GooglePlay, the authors identify 375 applications that violate their stated permissions.

I hope that this collection of articles provides a picture of the challenges and breadth of mobile computing and will encourage new research in this exciting area. The impact of mobile computing is undeniable, as billions of people are using mobile devices every day, and in many ways are becoming dependent on them. Moreover, there are areas on the globe where the only form of computing penetrating is mobile. The challenges discussed in this special issue are especially relevant in those situations. MICRO

Calin Cascaval is a senior director at Qualcomm Research Silicon Valley, where he is leading projects in parallel and heterogeneous computing, including the MARE runtime system and programming model for heterogeneous computing; parallel libraries and applications, such as the Zoomm parallel browser; and compiler technologies, such as the MCJS JavaScript engine. Cascaval has a PhD in computer science from the University of Illinois at Urbana-Champaign. He is a member of IEEE and the ACM. Contact him at cascaval@acm.org.