



Emerging Computing Challenges in the Interaction of Hardware and Software

Patrick Schaumont, Worcester Polytechnic Institute

This installment of Computer's series highlighting the work published in IEEE Computer Society journals comes from IEEE Transactions on Emerging Topics in Computing.

will require joint innovation at every abstraction level of the joint hardware/software integration stack. That observation is well known. However, an emerging characteristic is that software becomes essential for novel hardware platforms to even be usable, and both hardware and software designers have started

Until very recently, hardware and software have coexisted with maximum reciprocal professional respect, but often limited awareness of each other beyond a minimum acceptable level. Year after year, hardware performances have experienced substantial growth thanks to increased integration density and shrinking feature sizes. In parallel, software's growth has mainly been in sophistication and utility, often by riding the tailcoats of hardware performance improvements. Unless a fundamental technology switch occurs that further postpones the limits of technology integration, those free performance upgrades will come to an end. In the future, faster applications

strong collaborations, well beyond the domain of embedded systems.

One of the best illustrations of the profound changes to the hardware/software stack is the impact of machine learning and its applications to computer architecture and programming. The need for efficient machine learning has simultaneously affected silicon, architecture, and applications. This domain has led to novel chip designs, accelerator architectures, and application mapping techniques. For example, several top IEEE and related semiconductor industry-leading conferences on high-performance microprocessors and related chips (ACM/IEEE Design Automation Conference, Design, Automation and Test in Europe, the IEEE International Solid-State Circuits Conference, HoTChips, and so on) have repeatedly hosted tutorials and dedicated sessions on machine learning topics.



The efficient processing of complex signals and understanding of relationships in data sources remains a topic of intense interest for *IEEE Transactions on Emerging Topics in Computing (TETC)*. The upcoming special section “Emerging Trends and Advances in Graph-Based Methods and Applications,” edited by Conte et al. examines the use of graph-based data structures in this field (<https://www.computer.org/digital-library/journals/ec/cfp-graph-based-methods>). The special section covers graph-based solutions’ scalability, computing adaptability, and effectiveness in various application domains.

Yet, dwindling hardware performance is just one of the causes of the tightening link between hardware and software. Two other important aspects have made software an indispensable hardware partner: the real risk of diminishing reliability of hardware at advanced technology nodes and the real-versus-perceived inability of hardware architectures to deliver sufficiently robust and long-term security guarantees.


A dramatic illustration of the risks of decreasing hardware reliability is perhaps the recent observation of silent data corruption in large-scale deployments, such as by Google or Meta.¹ Modern hardware complexities are such that manufacturing-level testing alone can no longer catch all errors that may develop in the field. A software application may develop a computing error at any moment during a chip’s lifetime. Because error rates are still low, hardware redundancy is not a cost-effective solution to deal with these faults. Yet, silent data corruption is highly unacceptable even at a consumer level, and a risk of unpredictable consequences. The detection and mitigation of such errors is an emerging topic of open but urgent research where robust and fault-resilient software plays an important role. Among other approaches, one promising development is approximate

computing, a paradigm that recognizes the need for error-resilient computing given inaccurate data. *IEEE TETC* will host the upcoming special section “Approximate Data Processing: Computing, Storage, and Applications,” edited by Chen et al. (<https://www.computer.org/digital-library/journals/ec/cfp-approximate-data-processing>). The special section will address recent advances in error-resilient and approximate systems, emphasizing memory storage and data processing.

On top of the challenge of the risk of reduced reliability, applications in the cloud also require strong security. In the cloud scenario, it is common for different users to share the same server or become cotenants of the same machine. A secure cloud environment means that, at the very least, tenants operate in isolation and that the cloud service provider does not have unfettered access to the tenants’ data. The tenants can be kept strongly isolated by hardware using a secure enclave such as Intel’s Software Guard Extensions (SGX) or AMD’s Secure Encrypted Virtualization (SEV). However, the reality is that such isolation often occurs at the microarchitecture level. The tenants still share the same physical processor and memory chips on the same server boards in the analog domain. Research has repeatedly shown that this physical sharing among tenants leads to potential vulnerabilities in side-channel attacks and faults attacks with lofty names such as MeltDown, ZombieLoad, CopyCat, and PlunderVolt. The paradigm is now slightly changing. One can practically argue that it is not so much the question of if side-channel leakage can be exploited but rather when side-channel leakage will be converted into an exploit.

This hardware security challenge explains an increased interest in a computing paradigm that

makes side-channel leakage worthless to an adversary. Homomorphic encryption enables tenants to perform arbitrary computations on encrypted data. Functional encryption enables tenants to obtain specific functions over encrypted data. There is extensive research on algorithmic solutions that allows homomorphic encryption, but even the best schemes today are still much more complex than equivalent nonhomomorphic solutions. The upcoming *IEEE TETC* special section “Advances in Emerging Privacy-Preserving Computing,” edited by Han and Susilo, will collect recent innovations in privacy-friendly computing (<https://www.computer.org/digital-library/journals/ec/cfp-privacy-preserving-computing>).

With these special sections and an open call for results on emerging topics in computing, *IEEE TETC* seeks to highlight the interaction of hardware and software to address resiliency, machine learning, and privacy-friendly computing. *IEEE TETC* also encourages regular and special issue submissions on other novel emerging topics such as emerging hardware for computing, including quantum computing, and novel solutions for high-performance computing, hardware security, and computational networking. 

REFERENCE

1. H. D. Dixit et al., “Silent data corruptions at scale,” 2021, *arXiv:2102.11245*.

PATRICK SCHAUMONT is a professor of computer engineering at Worcester Polytechnic Institute, Worcester, Massachusetts, 01609, USA, and the associate editor in chief for Special Sections at *IEEE Transactions on Emerging Topics in Computing*. Contact him at pschaumont@wpi.edu.