Energy-Secure System Architectures (ESSA): A Workshop Report

Pradip Bose
IBM T. J. Watson Research Center

Saibal Mukhopadhyay
Georgia Institute of Technology

Modern microprocessor chips have multiple processing engines (or cores) that are architected to solve a variety of problems in individual and cooperative execution modes. In the current regime of commercial designs, we already see double-digit core counts; and if one considers the degree of hardware multithreading supported in each core, the number of hardware threads that can be supported in concurrent execution add up to many dozens or scores. For example, IBM's prior generation POWER8 processor chip already supported up to 96 hardware threads via its 12 cores, each of which can execute in up to an eight-way simultaneously multithreaded mode. As explained in recent ISSCC technology trend data, while the core count growth has been steady, the clock frequency has saturated around the 4-GHz mark—mainly limited by power density (or temperature) constraints. Effective parallelization of application codes, supported by many-core/many-thread hardware engines, is the established trend in current computing. Since 96-thread POWER8 server chips have already been in the market for a few years, it is not unrealistic to expect around 50 cores and perhaps ∼200 hardware threads supported in a couple of generations. Of course, due to area pressures, one can expect to see leaner (simpler) cores with only modest single-thread performance growth.

Ever since on-chip and system-level power management architectures have become routine in the industry, concerns about reliable operation and associated security vulnerabilities have been present in the minds of both the designer and researcher community.

This technology- and market-driven trend toward throughput-oriented (scale-out) designs implies a major challenge in terms of chip-level power and/or thermal management—in a regime where balanced performance growth (single-thread versus throughput) at affordable power becomes a steeper challenge over time. And, at the full system (i.e., server, rack, or data center) level, the challenge can be even greater. At whatever scale one is interested in managing such metrics (i.e., power or temperature, or even related ones, like system reliability),

Digital Object Identifier 10.1109/MM.2019.2921508
Date of current version 23 July 2019.

on-chip and system-level power management control architectures will need to be carefully architected. These must ensure the right tradeoffs to be applied at runtime to make sure that workload-dependent performance is maximized (at least to the extent that customer service-level agreements are met) while adhering to system-imposed power consumption limits.

Ever since on-chip and system-level power management architectures have become routine in the industry,[1-3] concerns about reliable operation and associated security vulnerabilities have been present in the minds of both the designer and researcher community. What if the sense-and-actuate feedback control system(s) implemented in such a design had latent bugs, wherein a corner case workload (launched maliciously or inadvertently) could disrupt the intended functionality and cause the system to fail? Could the chip or system incur irreparable physical damage? At a minimum, could a power virus attack result in significant performance degradation for other (regular) customer workloads—effectively signaling a denial of service attack? In fact, even before power management control systems were in vogue, research papers[4] had demonstrated physical attacks, where thermally induced memory bit-flips could enable Trojan software to take over the full system, evading immediate detection. Later on, the Charlie Miller hacks of Apple laptop battery control loops created quite a stir[5]; and data-center level power (energy) attack vectors were demonstrated.[6]

In light of the above motivational background, it is not surprising that researchers from IBM Research (led by Pradip Bose) initiated IBM internal research on a variety of power attacks and their mitigation around the year 2010. This research was supported by a DARPA seedling grant in 2011–2012, and this also helped that research team launch the ESSA workshop series in the beginning of 2011 [in conjunction with the International Symposium on Computer Architecture (ISCA)]. An early visionary paper on the ESSA theme was presented by Bose in 2012.[14] This workshop series was interrupted for a few years, before being resurrected recently in conjunction with the Hardware Security-Focused Conference (HOST 2019). The workshop was jointly organized by Dr. Pradip Bose from IBM Research and Prof. Saibal Mukhopadhyay from Georgia Institute of Technology. The key new feature of ESSA 2019 was to bring circuit-level power-management experts within the ESSA community to explore a cross-layer approach to energy-secure processor design. In the next section, we provide a summary description of ESSA 2019, which was held on May 9–10 at Tysons Corner, VA, USA.

SUMMARY PROCEEDINGS OF ESSA-2019

The initial years of workshop offerings around the ESSA theme resulted in a successful spawning off of a few academic research projects—which was the underlying objective. Perhaps the best-known research that has been reported in the literature since that time is the CLKSCREW attack modality[12] published by Simha Sethumadhavan's group at Columbia University. This work not only demonstrates the use of software code segments to disrupt the power management controls in a processor, it also shows how side-channel attacks can be orchestrated around this basic attack paradigm. In general, the scope of "energy attacks" has expanded to side channel attacks that exploit energy leaks of various types. As such, in ESSA 2019 (https://www.essa-workshop.org/), the technical scope of the workshop was expanded to broadly cover: "the range of research being pursued within industry and academia in order to ensure robust and secure functionality while meeting the energy-related constraints of the *green computing* era." The technical program of the workshop (https://www.essa-workshop.org/#program) consisted of one keynote, three visionary invited talks, three ESSA-relevant special invited talks, four contributed regular papers, and one panel session.

Keynote

The workshop began with an informative keynote address by John Marsh, who represented Linton Salmon, Program Manager of the ongoing DARPA program called System Security Integrated Through Hardware and Firmware (SSITH).

This was a valuable readout of the most promising research projects that are currently being pursued in the mitigation of software-assisted hardware attacks. The list of currently active SSITH program performers and the key innovations of their approach, as quoted from Marsh's talk, is shown in Table 1.

Table 1. Summary description of SSITH projects.

Prime	Technical area	Point of contact	Technical approach
Charles Draper Lab	H/W Architecture	Arun Thomas arun@draper.com	Every word has metadata + every instruction is checked based on flexible security micro-Policies defined in software (DSL); compartmentalization;PPASS workbench
Lockheed Martin	H/W Architecture	Jim Eiche james.eiche@lmco.com	Combination of efficient tagging,fenced,regions, Protection,domains,per-Thread keying,and memory encryption;security hardware in parallel with CPU; No source changes (Binary analysis)
MIT	H/W Architecture	Adam Chlipala adamc@csail.mit.edu	Hardware security compiler with end-to-end formal verification; generic support for tagging policies; compartmentalized secure enclaves.
SRI International	H/W Architecture	Robert Watson Robert.watson@cl.cam.ac.uk	Tags on every word of data and every instruction implement bounds checking and permissions; encapsulation; formal methods to verify security;security architecture extended to DMA engines
UC San Diego	H/W Architecture	Dean M Tullsen tullsen@cs.used.edu	Anti-fragility approach learns from attacks;machine learning based on Hardware Performance Registers; efficient X86 implementation leveraging micro-ops.
University of Michigan	H/W Architecture	Todd Austin austin@umich.edu	High entropy (hard to hack) Plus rapid churning (no time to exploit) to mediate “undefined semantics”; tagging; encryption;relocation of memory areas.
Galois	Test and Metrics	Joe Kiniry kiniry@galois.com	Automated system security metrics framework and tools; objective analysis; used formal methods; GFE IP development and support: RICS-V baseline for evaluation.

Visionary Talks

The first visionary talk was delivered by Dr. Vivek De from Intel Labs. In his visionary talk titled: “Attack-Resistant Energy-Efficient SoC Design,” Dr. De provided an in-depth perspective on the fundamental issues and tradeoffs in the design of power-performance-area (PPA) efficient SoCs that can also be architected to be resilient to malicious attacks. Figure 1 depicts the basic features of resilient platforms, as presented in Dr. De’s talk. The software, firmware, and hardware layers of abstraction in the design stack that need to be co-designed to factor in targeted resiliency features, while meeting critical PPA metrics were covered. In order to bake in the “attack resistance” shield into the general framework of resilient platforms, the architecture of the secure roots of trust (including the use of secure and variation-tolerant PUF/TRNG) was described. The second part of the talk showed examples from recent literature where existing power management circuit techniques are leveraged and re-purposed to improve resistance to power and electromagnetic emission based side-channel attack. The talk showed

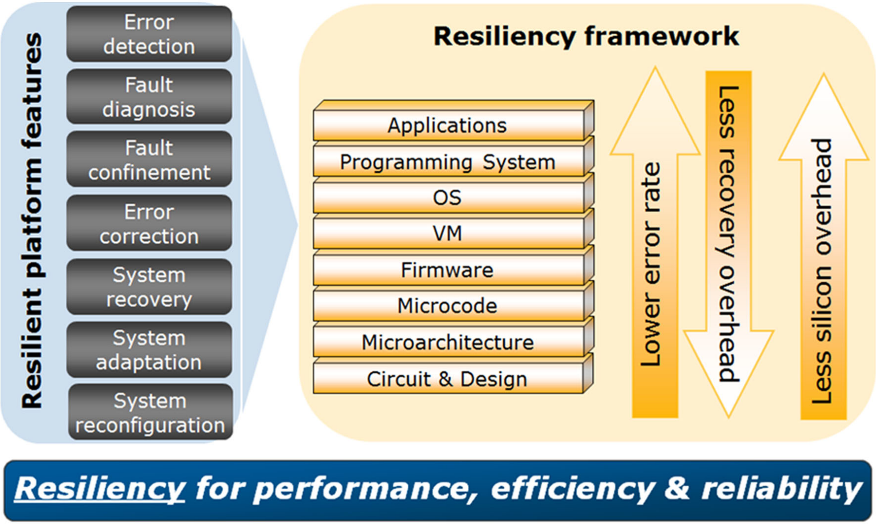

Figure 1. Resilient platforms.

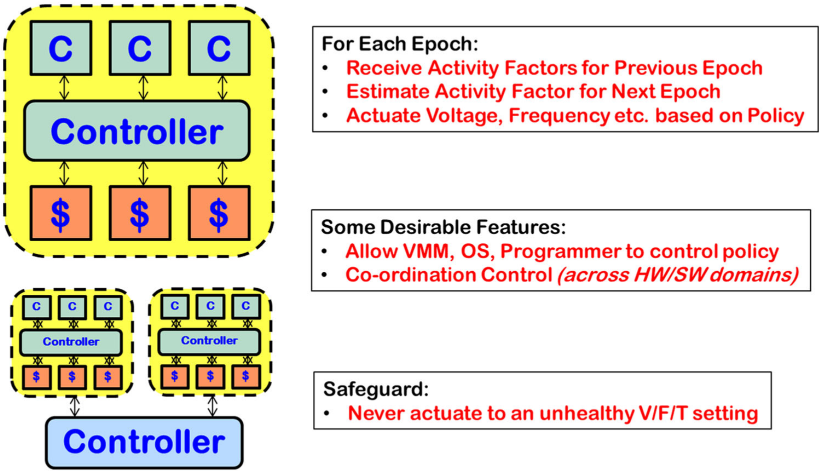

For Each Epoch:
- Receive Activity Factors for Previous Epoch
- Estimate Activity Factor for Next Epoch
- Actuate Voltage, Frequency etc. based on Policy

Some Desirable Features:
- Allow VMM, OS, Programmer to control policy
- Co-ordination Control *(across HW/SW domains)*

Safeguard:
- Never actuate to an unhealthy V/F/T setting

Figure 2. Abstract view of energy management.

that circuit techniques can strongly impact the energy-security tradeoff in a SoC.

The second visionary talk was delivered by Prof. Simha Sethumadhavan from Columbia University. In his talk, titled: "Software vectored fault attacks," Sethumadhavan recalled the CLKSCREW attack work[7,8] that his group had pioneered, and then painted a picture of the work that is emerging beyond that groundbreaking prior work. The CLKSCREW research exposed the vulnerabilities in classical DVFS-based power management in embedded systems (see Figure 2), and has since resulted in solutions to fix such gaps in security in a class of popular commercial mobile platforms driven by the Android OS. Sethumadhavan presented a wish list of future energy-secure architectures, with a focus on hardware security features that would need to be augmented in order to mitigate the threat of energy-sourced side channel intrusions and physical attacks.

Prof. Onur Mutlu's (ETH Zurich) visionary talk, titled "Using commodity memory devices to support fundamental security primitives," was an in-depth journey into the fundamentals of memory system architectures, and the associated security-related vulnerabilities. Initially, Mutlu spoke about the solution approaches to reduce data movement related power consumption through architectural innovations. Subsequently, a significant part of the talk was on ways to use memory devices to generate true random numbers with low latency and high throughput. Generation and evaluation of physically unclonable functions (PUFs) was another aspect that was covered. These capabilities were pointed up to be crucial elements in system-level security mechanisms. Another item covered was quick destruction of in-memory data (for DRAMs). The security-centric aspects of the presentation focused mainly on the speaker's work published in HPCA 2018,[15] HPCA 2019[16] and arxiv 2019.[17]

Special Invited Talks

Prof. Mingoo Seok (Columbia University) presented a paper (M. Seok, A. Tang, Z. Jiang, S. Sethumadhavan, "Blacklist core: machine-learning-based power management tampering,") where machine-learning-based dynamic operating performance point blacklisting is used for mitigating software based power-management tampering. Using CLKSCREW attack as an example, Seok argued that static guard-banding could mitigate such attacks but it incurs performance degradation, power efficiency loss, and long testing time. Instead, the talk introduced a *detection-then-mitigation* approach, which uses a neural-net model and detects a malicious command to put the system on an unsafe operating performance point. If detected, it then mitigates the attack by ignoring the command. The algorithm and hardware realization of the technique shows the ability to detect and mitigate CLKSCREW attempts at a reasonably small amount of overhead in power, delay, and area. The talk illustrated that co-design of circuit and algorithm is necessary to optimally tradeoff the energy and security behavior of an SoC.

Prof. Swaroop Ghosh (Penn State) spoke on: "Security of persistent memories." Excellent properties, such as zero leakage, high-density, scalability, and high endurance, of emerging non-volatile memories (NVMs) make them an attractive candidate for energy management in SoCs. However, Ghosh pointed out that although NVMs can reap energy and performance benefits they may face new security issues that were not perceived before. His talk discussed several potential vulnerabilities of NVMs, and how they can be exploited to compromise data integrity (e.g., tampering and row-hammering) and data privacy. The talk also explored circuit and system level methods for sensing and inhibiting attacks on NVMs. The talk reminded the audience that new technologies introduced for energy management can also lead to new security challenges.

Prof. Vijay Janapa Reddi (Harvard) spoke on: "Closing the performance, power and reliability gap in autonomous aerial machines." Reddi has been working on the topic area of "aerial computing," and this particular presentation addressed the fundamental issues of power-performance efficiency and resilience in designing the embedded processor engines that power autonomous drones. The connection between reliability and security, as also articulated in Sethumadhavan's visionary talk, was re-examined briefly in Reddi's presentation.

Contributed Regular Papers

The regular technical presentations consisted of the following contributed papers:

- K. Khatamifard, L. Wang, S. Kose, A. Das, U. Karpuzcu, "A novel class of covert channels enabled by power budget sharing."
- A. Krishnan, P. Schaumont, "Hardware support for secure intermittent architectures."
- I. Tochukwu and A. Ismail, "Holistic hardware security assessment framework: a microarchitectural framework."
- D. Trilla, C. Hernandez, J. Abella, F. Cazorla, "Four birds with one stone: on the use of time randomized processors and probabilistic analysis to address timing, reliability, energy and security in critical embedded autonomous systems."

It is well known that runtime power management is in charge of the optimal distribution of the power budget—a very critical shared resource—among system components. The paper by Tochukwu *et al.* argued that any system-wide shared resource can give rise to covert communication, if not properly managed, and power budget, unfortunately, does not represent an exception. The paper presented a proof-of-concept demonstration of covert communication exploiting shared power budget and discussed the potential design space for countermeasures. The paper argued that a secure power management infrastructure must be aware of the potential threats associated with sharing power across multiple entities.

The paper by Krishnan *et al.* stressed the need for a secure power transition mechanism to convert the active system state into a protected nonvolatile form and back in energy harvesting based IoT edge platforms. The paper observed that secure checkpointings are necessary, but are expensive to compute and require hardware-accelerated cryptography and isolated secure non-volatile storage. The paper defined an energy-harvester subsystem interface that drives the optimized execution of a secure communication protocol such that wasted energy is eliminated and that run-time performance is improved.

The paper by Tochukwu *et al.* presented the challenging but critical need to enable a holistic hardware security evaluation from the microarchitectural point of view. The paper introduced an important step toward this direction by proposing a framework that categorizes threat models based on the microarchitectural components being targeted and provides a generic security metric that can be used to assess the vulnerability of components, as well as the system as a whole.

Finally, the paper by Trilla *et al.* discussed that as complexity and time-criticality of operations being performed in the autonomous vehicles continue to increase, this creates conflicting requirements in designing such processors. On one hand, a simple and predictable design of the processors facilitate verification of functional and nonfunctional metrics; but on the other hand, using high-performance and complex processor designs with some degrees of obfuscation can deliver high computing performance and security. The paper argued that time-randomized processors (TRP), an alternative to traditional (deterministic) designs, can address these conflicting requirements. TRP facilitates timing analysis via the use of statistical/probabilistic techniques, while also show capabilities to effectively tackle the challenges of reliability, security, and energy consumption. The paper reviewed the TRP opportunities and show that they are a natural fit to fulfill the requirements of autonomous critical systems. The paper showed that disruptive ideas in the macro- and microarchitecture may be necessary to design future energy-secure autonomous systems.

IMPACT OF THE ESSA THEME WORKSHOPS

In this section, we will briefly examine the ongoing impact of the ESSA theme workshop

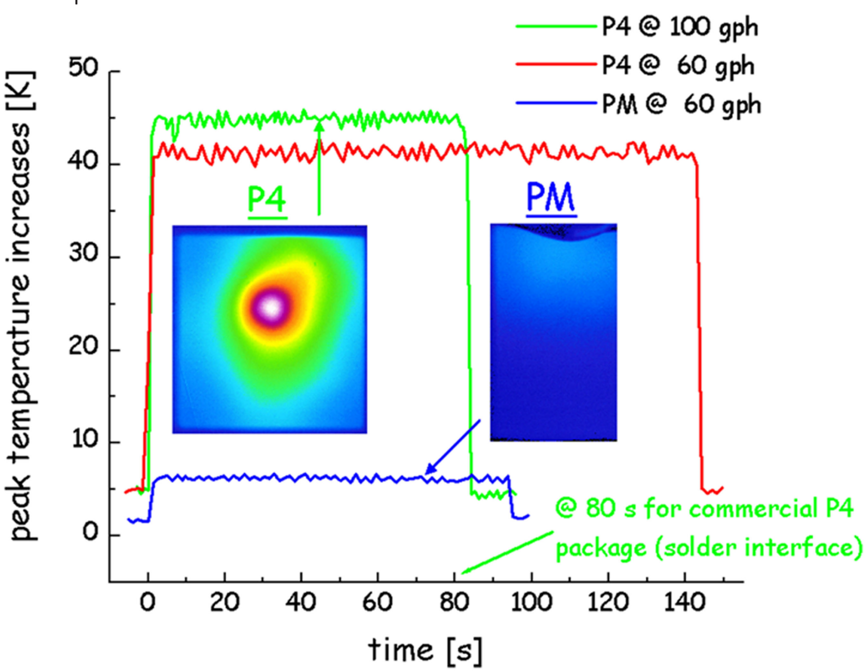

Figure 3. Peak temperature driven throttle points and performance effects for Intel Pentium IV (P4) and Pentium M (PM) class processors. Experiments done using gcc workload within SPEC95 and different cooling solutions: 100 gallons per hour (gph) is the highest fluidic cooling rate used, resembling the commercial P4 package; 60 gph is a reduced cooling rate to represent a cheaper packaging solution. (Experimental data: courtesy Hendrik Hamann, IBM Research).

series in the context of identifying new security-related vulnerabilities and devising mitigation solutions thereof. So far, in terms of core security domain impact, the CLKSCREW work from Columbia[7,8] is the leading example of new innovation and impact related to the ESSA theme. This work has been acknowledged by the Android Security Team: (https://source.android.com/security/overview/acknowledgements): "Adrian Tang of Columbia University (CLKSCREW paper), CVE 2017-8252." Qualcomm also acknowledged and reported the fixing of bugs associated with the above-quoted CVE (common vulnerabilities and exposures) item.

Within IBM, an in-depth research study (led by Pradip Bose) was conducted, under the sponsorship of a small DARPA seedling grant (during 2011–2012), to assess the threat level imposed by maliciously launched power/thermal viruses. For commercial IBM high-end processor systems (e.g., POWER7 at that time), the study was not able to demonstrate any performance or functional degradation through user-level application software access alone. (Supervisory mode access to OS and/or power management firmware code is, of course, a different issue.) One of the saving factors, as assessed, was that the IBM POWER processors up to the POWER7 generation were not subject to performance throttling under even the highest power workloads conceivable. In other words, for such high-end server-class processors, the heat sink and cooling solution were over-designed to make sure that the worst-case applications would not result in exceeding the power limit that could be handled by the packaging-cum-cooling solution. Note that power or electromigration (EM)-based side channel attack vulnerabilities were not within the scope of this study; only physical damage and denial of service type attacks were under consideration at that time.

One should note, however, that well before engaging in this POWER7-based study, the IBM researchers had studied the performance degradation characteristics of Intel's Pentium IV series processors. Figure 3 shows the experimental characterization data that compares the temperature-driven performance differences between Intel's Pentium IV (abbreviated here as P4) and Pentium M (PM) class processors. In this laboratory experiment (which was conducted back in 2005), the packaging lid of the processor was taken out and replaced by a cooling solution provided by a controllable (special) heat-transparent fluid flow. The thermal imaging, measurement, and calibration were conducted using a special infra-red camera setup.[9] As Figure 3 shows, the "normal" execution time of the full gcc workload on a P4 using a commercial grade cooling solution is 80 s. Whereas, if the cooling mimics a low-cost packaging solution, the execution time degrades to about 140 s due to the throttling-based dynamic temperature management built into these processors. In contrast, the lower power PM processor runs much cooler, without incurring any throttling and executes the same workload in about 90 s, even with the low-cost packaging solution. The experiments demonstrated the possibility that throttling-based performance degradation (at a very significant level) could be instigated for a given processor-package system product by launching a high-power (possibly synthetic) virus workload.

The literature on side-channel attack mechanisms that exploit the power and/or EM monitoring has advanced quite a lot since the inception of the first ESSA workshop in 2011. This is evidenced even from some of the technical and visionary talks presented at ESSA 2019. The threat imposed by sharing of a common power budget in a multicore chip setting has been described in work by Sasaki *et al.*[11] and the paper by Khatamifard *et al.* presented at ESSA 2019 shows the consequence of this threat model in explicit terms.

Future work must connect unreliable control loops explicitly to vulnerabilities from a mainstream security research viewpoint. In other words, the guarded power management principle must be tested as a mitigation technique against CLKSCREW-inspired attacks.

While at the system scale, ESSA themed workshop has studied the potential security threats introduced by the power management solutions, there have been significant progress in recent years in understanding the energy-security tradeoff at the circuit level. In particular, recent research threads have emerged in designing energy-efficient security engines, as well as exploring on-chip power-management circuits for security. A specific example of this new direction has been in the domain of designing low-overhead techniques for improving power and EM-based side-channel-attack resistance of encryption engines. In particular, collaborative work between Georgia Tech and Intel labs has demonstrated a set of studies where on-chip integrated voltage regulators (IVRs) and adaptive clocking circuits, introduced mostly for power management, have been leveraged to improve SCA resistance. A paper presented by Kar *et al.*, at the 2017 International Solid State Circuit Conference (ISSCC) demonstrated the promise of the using inductive IVRs for inhibiting power attack on AES engines. A second article from the group, authored by Singh *et al.* and published in the IEEE Journal of Solid State Circuits (JSSC) in 2019, showed inductive IVR coupled with fine-grain dynamic voltage scaling and adaptive clocking can inhibit power and EM-based side channel attacks on AES engines. More recently, an article presented at ISSCC 2019 by Singh *et al.* showed that on-chip low-dropout regulators, coupled with adaptive clocking and fine-grain DVFS provide security against power-/EM-based side-channel attacks against AES engines. There are many more examples in recent literature where circuit level studies are being performed, and techniques are being developed to enable a bottom-up approach to energy-secure hardware designs. The advancements at the circuit level security research showed the need for engaging circuit community within the ESSA theme, and ESSA 2019 took a positive step toward this goal. The success is evident from the talks presented at ESSA 2019 by Dr. De, Dr. Ghosh, and Dr. Seok, all of which have pointed out the need for circuit level research in this domain.

FUTURE DIRECTIONS

In this section, we provide our view of the future directions of research and development within the ESSA theme. One of the early research agenda items at IBM that fell out of the ESSA-theme was that of guarded power management.[10] In this solution approach, the baseline power management architecture is protected through a guard mechanism. The latter is a higher level monitor-and-control system that observes the operation of the baseline architecture through specialized activity (performance) counters. Anomalies detected in the observed counter-based signatures can serve to trigger mitigation actions. The latter could include adapting the hardware parameters of the baseline mechanism on-the-fly. The above-referenced work was pursued with robust power management in mind. Future work must connect unreliable control loops explicitly to vulnerabilities from a mainstream security research viewpoint. In other words, the guarded power management principle must be tested as a mitigation technique against CLKSCREW-inspired attacks.

The thrust of research in support of power reduction in the wake of GPU-centric high-performance compute nodes has led to techniques like adaptive voltage guard-band management (e.g., J. Leng *et al.*, MICRO 2015). In future accelerator-rich systems, the task of balancing power, performance, and reliability will have to be managed using systematic hardware-software management systems. Purely software-based scheduling heuristics will need to get supported by hardware-based monitors. As we progress toward many-core processor chips, old-style on-chip power control architectures (with a single, centralized management unit) will give way to scalable, distributed control, and management systems. An initial vision on so-called swarm power management architectures has been portrayed in recent invited papers (e.g., the one at DATE 2018[11]).

ACKNOWLEDGMENTS

The ESSA 2019 workshop,[13] the proceedings (and anticipated impact) of which were summarized in this article, would not have been possible without the help of an active and supportive program committee. The valuable contributions of our web and publicity chair, David Trilla, are to be noted in particular. We are grateful to the organizing committee of the HOST 2019 symposium for their support.

REFERENCES

1. S. Gunther and R. Singhal, “Next generation intel microarchitecture (nehalem) family: Architectural insights and power management,” presented at Intel Developer Forum, San Francisco, CA, USA, Mar. 2008.
2. M. Floyd *et al.*, “Introducing the energy management features of the POWER7 chip,” *IEEE Micro*, vol. 31, no. 2, pp. 60–75, Mar./Apr. 2011.
3. T. Webel *et al.*, “Robust power management in the IBM z13,” *IBM J. R&D*, vol. 59, no. 4/5, pp. 16-1–16-12, Jul./Sep. 2015.
4. S. Govindavajhala and A. W. Appel, “Using memory errors to attack a virtual machine,” in *Proc. IEEE Symp. Secur. Privacy*, 2003, pp. 154–165.
5. C. Miller, “Battery firmware hacking,” presented at BlackHat, August 2011. [Online]. Available: https://www.blackhat.com/html/bh-us-11/bh-us-11-briefings.html#Miller
6. Z. Wu, M. Xie, and H. Wang, “Energy attack on server systems,” in *Proc. 5th USENIX Workshop Offensive Technol.*, 2011, p. 8.
7. A. Tang, S. Sethumadhavan, and S. Stolfo, “CLKSCREW: Exposing the perils of security oblivious energy management,” in *Proc. USENIX Secur. Symp.*, 2017, pp. 1057–1074.
8. A. Tang, S. Sethumadhavan, and S. Stolfo, “Motivating security-aware energy management,” *IEEE Micro*, vol. 38, no. 3, pp. 98–106, May/Jun. 2018.
9. H. Hamann, A. Weger, J. Lacey, Z. Hu, and P. Bose, “Hotspot-limited microprocessors: direct temperature and power distribution measurements,” *IEEE J. Solid State Circ.*, vol. 42, no. 1, pp. 56–65, Jan. 2007.
10. N. Madan, A. Buyuktosunoglu, P. Bose, and M. Annavaram, “A case for guarded power gating for multi-core processors,” in *Proc. 17th Int. Symp. High Performance Comput. Arch.*, Feb. 2011, pp. 291–300.
11. H. Sasaki, A. Buyktosunoglu, A. Vega, and P. Bose, “Mitigating power contention: A scheduling based approach,” *Comput. Arch. Lett.*, vol. 16, no. 1, pp. 60–63, 2017.
12. A. Vega, A. Buyuktosunoglu, and P. Bose, “Energy-secure swarm power management,” in *Proc. Design Test Eur.*, 2018, pp. 1652–1657.
13. ESSA-2019 Workshop. [Online]. Available: https://www.essa-workshop.org/
14. P. Bose *et al.*, “Power management of multi-core chips: Challenges and pitfalls,” in *Proc. Design Test Eur.*, 2012, pp. 977–982.
15. J. Kim *et al.*, “DRaNGe: Using commodity DRAM devices to generate true random numbers with low latency and high throughput,” in *Int’l. Symp. High Perform. Comput. Arch. (HPCA)*, Feb. 2018.
16. J. Kim *et al.*, “The DRAM latency PUF: Quickly evaluating physical unclonable functions by exploiting the latency-reliability tradeoff in modern commodity DRAM devices,” in *Int’l. Symp. High Perform. Comput. Arch. (HPCA)*, Feb. 2018.
17. L. Orosa *et al.*, “Dataplant: in-DRAM security mechanisms for low-cost devices,” *arxiv*, 2019, [Online]. Available: https://arxiv.org/abs/1902.07344

Pradip Bose is with IBM T. J. Watson Research Center, New York. Contact him at: pbose@us.ibm.com.

Saibal Mukhopadhyay is with Georgia Institute of Technology. Contact him at: saibal@ece.gatech.edu.