# SPECIAL SERIES ON HARSH CHIPS


Augusto Vega
Alper Buyuktosunoglu
Pradip Bose
IBM T.J. Watson Research Center

Chip power consumption is one of the most challenging and transforming issues that the semiconductor industry has encountered in the past decade, and its sustained growth has resulted in various concerns, especially when it comes to chip reliability. It translates into thermal issues that could harm the chip. It can also determine (that is, limit) battery life in the mobile arena. And even more critically, it can have serious consequences for proper chip operation in terms of adopting low-power techniques such as near-threshold voltage computing. For example, chips become more susceptible to soft errors at lower voltages. This scene becomes even more disturbing when we add an extra variable: a hostile (or harsh) surrounding environment. Harsh environmental conditions exacerbate already problematic chip power and thermal issues, and can jeopardize the operation of any conventional (that is, non-hardened) processor. For example, NASA's *Curiosity* rover (exploring Mars in recent months) uses two RAD750 radiation-hardened microprocessors in its on-board computer.[1] The engineers who designed this microprocessor had to consider the high radiation levels and wide temperature fluctuations in the red planet to ensure mission success. Similarly, on-board computer design for unmanned aerial vehicles (that is, drones) poses hurdles due to the tight power budget to perform mission-critical tasks, which must be successfully accomplished under a wide range of environmental conditions. In this context, we can describe a harsh-environment-capable chip (or harsh chip, for short) as an embedded computing system specially designed and manufactured to guarantee certain levels of real-time computing performance and reliability under well-defined, tight constraints. These constraints include limited power budget and battery life, maximum tolerable levels of radiation, electromagnetic interference and ambient temperature, strict fault tolerance, and security requirements.

Interplanetary rovers and drones are just two examples of harsh-environment-capable embedded systems. Embedded computing has become pervasive and mobile and, as a result, many of the day-to-day devices that we use and rely on are subject to similar constraints—in some cases, with critical consequences when they are not met. For example, cars, trucks, and even motorcycles are becoming smarter, and in some cases autonomous (driverless).[2,3] Highly reliable, low-power embedded chips for harsh environments are the key enablers for autonomous and semiautonomous vehicles, and it is not difficult to imagine the safety consequences if the on-board control systems fail or are improperly operated.

The adoption of harsh chips that can operate properly even under extreme conditions is experiencing an unprecedented growth, in tune with the revolutions related to mobile systems and the Internet of Things, emergence of autonomous and semiautonomous transport systems (such as connected and driverless cars), and highly automated factories and the robotics boom. The *IEEE Micro* special series on harsh chips will feature several articles that discuss the most critical aspects of new-generation harsh-environment-capable embedded processors. In addition, we have successfully organized two editions of the Workshop on Highly-Reliable Power-Efficient Embedded Designs (HARSH), which have attracted the attention of researchers from academia, industry, and government research labs.

In this issue, we present the first two articles of the series; additional articles will be published in subsequent issues of *IEEE Micro*. The first article, "Timing Verification of Fault-Tolerant Chips for Safety-Critical



Published by the IEEE Computer Society

Applications in Harsh Environments," examines the time predictability aspect of critical real-time embedded systems (CRTES). Mladen Slijepcevic et al. focus on the conditions that are required in faulty CRTES architectures to satisfy worst-case execution time (WCET) bounds. They adopt a measurement-based probabilistic timing analysis approach to address the timing-related unpredictability issues of complex multicore architectures. More specifically, the authors present a processor design with trustworthy and tight WCET estimates for safety-critical applications running on high-performance hardware that are subject to hard and soft errors.

In the second article, "R3TOS-Based Autonomous Fault-Tolerant Systems," Xabier Iturbe et al. exploit hardware reconfigurability to enable autonomous fault-tolerant systems (AFTS) for use in harsh environments. The authors present the Reliable Reconfigurable Real-Time Operating System (R3TOS), their solution for building adaptive electronics using commercial-off-the-shelf reconfigurable field-programmable gate arrays. They demonstrate the benefits of R3TOS through a traction controller prototype of a real-world railway traction system, where it proves to be able to recover from most of the errors it encounters without requiring any human intervention.

We would like to thank all the authors who submitted their papers to this special series. We also express our deep gratitude to all the reviewers for their valuable time. The reviews helped the authors to considerably improve the technical quality of the accepted papers. We sincerely hope that you enjoy this special series and find its contents informative and useful! MICRO

## References

1. W. Harwood, "Slow, but Rugged, Curiosity's computer was Built for Mars," *CNET*, 10 Aug. 2012; www.cnet.com/news/slow-but-rugged-curiositys-computer-was-built-for-mars.
2. T. Fishman, "Flattened Butterfly: A Cost-Efficient Topology for High-Radix Networks," *Deloitte Univ. Press*, 1 Dec. 2012; http://dupress.com/articles/digital-age-transportation.
3. E. Ackerman, "Google Wants Option to Test Autonomous Motorcycles and Trucks in California," *IEEE Spectrum*, 20 Aug. 2014; http://spectrum.ieee.org/cars-that-think/transportation/self-driving/google-autonomous-motorcycles-and-trucks-in-california.

**Augusto Vega** is a research staff member in the Reliability and Power-Aware Microarchitecture department at the IBM T.J. Watson Research Center. His research interests include high-performance, reliable, and power-efficient computer architectures; distributed and parallel computing; and performance-analysis tools and techniques. Vega has a PhD in computer architecture from the Polytechnic University of Catalonia.

**Alper Buyuktosunoglu** is a research staff member in the Reliability and Power-Aware Microarchitecture department at the IBM T.J. Watson Research Center. His research interests include high-performance, reliable, and power-efficient computer architectures. Buyuktosunoglu has a PhD in electrical and computer engineering from the University of Rochester.

**Pradip Bose** is a distinguished research staff member in the Reliability and Power-Aware Microarchitecture department at the IBM T.J. Watson Research Center. His research interests include high-performance, reliable, and power-efficient computer architectures. Bose has a PhD in electrical and computer engineering from the University of Illinois at Urbana-Champaign.

Direct questions and comments about this issue to Augusto Vega at ajvega@us.ibm.com, Alper Buyuktosunoglu at alperb@us.ibm.com, and Pradip Bose at pbose@us.ibm.com.