

System-Level Policies to Reduce Power Consumption in Fault-Tolerant Embedded Systems

Mohsen Ansari*

*Department of Computer Engineering, Sharif University of Technology, Tehran, Iran

Abstract—Multicore platforms provide great opportunity for implementation of fault-tolerance techniques to achieve high reliability in real-time embedded systems. In spite of the huge potential for fault-tolerant techniques in multicore platforms, due to the Thermal Design Power (TDP) constraint, designers of fault-tolerant embedded systems face a challenge in deciding how to use them. TDP is considered as the highest sustainable power that a chip can dissipate without triggering any performance throttling mechanisms, e.g., Dynamic Thermal Management (DTM). If a chip temperature violates its TDP, it automatically restarts or significantly reduces its efficiency to prevent permanent damage. The continually increase in the degree of integration along with using fault-tolerant techniques can increase power consumption and rise the peak power which can lead to violate the TDP constraint. In this research, we aim at considering effects of fault-tolerant techniques on the power consumption in real-time embedded systems.

Index Terms—Power Consumption, Real-time Embedded Systems, Reliability, Mapping, Scheduling, Fault Tolerance.

I. INTRODUCTION

TECHNOLOGY scaling allows integrating multiple cores onto a single chip to make the mainstream for advanced embedded systems [1][2][3][4][5][6][7][8]. However, technology scaling results in aggravating the reliability problem of on-chip systems because of increasing transient fault rate due to lower voltages and shrinking transistor dimensions that lead to smaller critical charges [9][10][11][12][13][14]. Multicore systems provide a great opportunity to implement reliability mechanisms such as redundant multithreading (RMT) and process level redundancy. Task replication (e.g., RMT) is a well-established technique to achieve high reliability against transient faults. However, replicated executions may increase power consumption beyond the chip Thermal Design Power (TDP) constraint. TDP is considered as the highest sustainable power that a chip can dissipate before being forced to exploit a performance throttling mechanism, e.g., Dynamic Thermal Management (DTM). If a chip violates its TDP, it automatically restarts or significantly reduces its performance to prevent a permanent damage. Therefore, due to the unwanted system restarts, DTM techniques may not be applicable for the systems that require satisfying strict timing constraints. In this dissertation, at first, we illustrate how fault-tolerance techniques may increase peak power consumption and consequently may result in a chip TDP violation. Then, we propose a peak-power-aware scheduler that manages peak power consumption for different task models on multicore

embedded systems. The scheduler schedules different applications on different types of cores in two types of multicore platforms without violating timing constraints. The proposed methods aim at removing overlaps of the peak power of concurrently executing tasks to keep the power consumption below the chip TDP. To do this, considering the tasks' power profiles, at first, we partition the tasks into parts with different peak power values. Then, based on the task models, we propose policies for scheduling the partitioned tasks. In this dissertation, we have proposed five policies to schedule applications of embedded systems such that the power consumption is kept below the power constraints. These policies are: (i) the Peak-Power-Aware Longest Task First (PPA-LTF) policy for scheduling the task graph model, (ii) the Maximum-Peak-Power-First (MPPF) and Maximum-Peak-Power-Last (MPPL) policies for scheduling the frame-based task model, (iii) the Peak-Power-Aware Earliest-Deadline-First (PPA-EDF) and Peak-Power-Aware Earliest-Deadline-Late (PPA-EDL) policies for scheduling the periodic tasks. In summary, our proposed schemes try to spread out the parts of tasks that consume high power over the entire interval before the deadline with the aim of keeping the total peak power below the chip TDP.

REFERENCES

- [1] M. Ansari, S. Safari, A. Yeganeh-Khaksar, M. Salehi and A. Ejlali, "Peak Power Management to Meet Thermal Design Power in Fault-Tolerant Embedded Systems," in *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 1, pp. 161-173, 1 Jan. 2019.
- [2] M. Ansari, J. Saber-Latibari, M. Pasandideh and A. Ejlali, "Simultaneous Management of Peak-Power and Reliability in Heterogeneous Multicore Embedded Systems," in *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 3, pp. 623-633, 1 March 2020.
- [3] M. Ansari, M. Salehi, S. Safari, A. Ejlali, and M. Shafique, "Peak-Power-Aware Primary-Backup Technique for Efficient Fault-Tolerance in Multicore Embedded Systems," *IEEE Access*, vol. 8, pp. 142843-142857, 2020.
- [4] M. Ansari, A. Yeganeh-Khaksar, S. Safari, and A. Ejlali, "Peak-Power-Aware Energy Management for Periodic Real-Time Applications," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 4, pp. 779-788, April 2020.
- [5] Mohsen Ansari, M. Pasandideh, J. Saber-Latibari, and A. Ejlali, "Meeting Thermal Safe Power in Fault-Tolerant Heterogeneous Embedded Systems," *IEEE Embedded Systems Letters*, vol. 12, no. 1, pp. 29-32, March 2020.
- [6] S. Safari, M. Ansari, G. Ershadi, and S. Hessabi, "On the Scheduling of Energy-Aware Fault-Tolerant Mixed-Criticality Multicore Systems with Service Guarantee Exploration," *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 10, pp. 2338-2354, 1 Oct. 2019.
- [7] A. Yeganeh-Khaksar, M. Ansari, and A. Ejlali, "ReMap: Reliability Management of Peak-Power-Aware Real-Time Embedded Systems through Task Replication," *IEEE Transaction on Emerging in Computing*, 2020

- [8] Javad Saber-Latibari, Mohsen Ansari, Pourya Gohari-Nazari, Sina Yari-Karin, Amir Mahdi Hosseini Monazzah, and Alireza Ejlali, "READY: Reliability- and Deadline-Aware Power-Budgeting for Heterogeneous Multi-Core Systems," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2020.
- [9] A. Yeganeh-Khaksar, M. Ansari, S. Safari, S. Yari-Karin, and A. Ejlali, "Ring-DVFS: Reliability-Aware Reinforcement Learning-Based DVFS for Real-Time Embedded Systems," *IEEE Embedded Systems Letters*, 2020
- [10] M. Ansari, S. Safari, F. R. Poursafaei, M. Salehi, and A. Ejlali, "AdDQ: Low-Energy Hardware Replication for Real-Time Systems through Adaptive Dual Queue Scheduling," *The CSI Journal on Computer Science and Engineering (JCSE)*, vol. 15, no. 1, pp. 31-38, 2017.
- [11] S. Safari, M. Ansari, M. Salehi, and A. Ejlali, "Energy-Budget-Aware Reliability Management in Multi-Core Embedded Systems with Hybrid Energy Source," *The CSI Journal on Computer Science and Engineering (JCSE)*, vol. 15, no. 2, pp. 31-43, 2018.
- [12] F. R. Poursafaei et al., "Offline Replication and Online Energy Management for Hard Real-Time Multicore Systems," *Proc. of the 11th Int'l the CSI Symp. on Real-Time and Embedded Sys. and Tech. (RTEST)*, Tehran, Iran, 2015, pp. 1-7.
- [13] S. Yari-Karin, A. Sahraei, J. Saber-Latibari, M. Ansari, N. Rohbani, and A. Ejlali, "A Comparative Study of Joint Power and Reliability Management Techniques in Multicore Embedded Systems," *Proc. of the 3rd Int'l of the CSI Symposium on Real-Time and Embedded Systems and Technologies (RTEST)*, 2020.
- [14] Z. Shirmohammadi, M. Ansari, S. Kazemi, S. Safari, and S.-G. Miremadi, "PAM: A Packet Manipulation Mechanism for Mitigating the Crosstalk Faults in NoCs," *Proc. of the 13th IEEE Int'l Conference on Dependable, Autonomic and secure computing (DASC)*, Liverpool, England, October, 2015.



Mohsen Ansari received the M.Sc. degree in computer engineering from the Sharif University of Technology, Tehran, Iran, in 2016. He is currently working toward the Ph.D. degree in computer engineering at Sharif University, Tehran, Iran, from Sept. 2016 until now. He is now a visiting researcher in the Chair for Embedded Systems (CES), Karlsruhe Institute of Technology (KIT), Germany, from Oct. 2019 until now. Also, he is a group

leader of the Embedded Systems Research Laboratory (ESR-LAB) at the department of computer engineering, Sharif University of Technology. His research interests include the low-power design of embedded systems and multi-/many-core systems with a focus on dependability/reliability.