

New approaches to DVFS in mobile MPSoC for power-, thermal-efficiency and reliability

Somdip Dey* Amit Singh*,
Klaus McDonald-Maier*,

** School of Computer Science and Electronic Engineering, University of Essex, UK.*

ABSTRACT

Modern mobile devices come equipped with heterogeneous multi-processor system-on-chip (MP-SoC) consisting of different types of cores such as CPU and GPU to cater for performance requirement of different types of applications. These MPSoCs also supports dynamic voltage and frequency scaling (DVFS) to reduce dynamic power consumption, which retrospectively affects the peak temperature of the multi-processors on the chip as well. In this poster presentation, we introduce three novel approaches to DVFS in mobile MPSoCs such that the executing applications can meet the performance requirement while consuming less power, reducing peak temperature and improving thermal side-channel security and reliability in terms of lifespan of the device.

KEYWORDS: agent system; reinforcement learning; machine learning; CPU; GPU; power optimization; thermal optimization; MPSoCs; user behaviour; user interaction; smartphone; mobile; convolutional neural network (CNN); program source code; thermal side-channel attack, reliability

1 Introduction

Recently we could see an increase in embedded computing systems [D⁺], which employs heterogeneous multi-processor systems-on-chips (MPSoCs) consisting of different types of cores such as CPU and GPU. These architectures provide opportunities to exploit distinct features of CPU and GPU cores to meet performance and power consumption requirements. Furthermore, the cores in these SoCs support dynamic voltage and frequency scaling (DVFS) [D⁺19b, D⁺19c, D⁺19a, S⁺20], which can be used to reduce dynamic power consumption ($P \propto V^2 f$) [I⁺19, D⁺19d]. This helps to reduce the power consumption by executing the workload over extra time at a lower voltage and frequency, which could be accounted for reduced power consumption.

One of the most popular heterogeneous MPSoCs is the Samsung Exynos 5422 SoC [exy] used in Samsung Galaxy S5 smartphone, which incorporates 4 ARM Cortex-15 (big) CPU, 4 ARM Cortex-7 (LITTLE) CPU and 6 ARM Mali-T628 GPU cores, are excellent playground for demanding applications. Since many modern mobile applications are mixed workload,

¹Corresponding E-mail: somdip.dey@essex.ac.uk

Abstract for poster presentation: ACACES 2020

both compute intensive and memory intensive, [D⁺19e] such applications are really good candidates for utilizing the full computing capacity of MPSoCs such of Samsung Exynos 5422. But making such applications more power efficient and reliable while catering for the performance requirements is a big challenge.

Since most modern mobile and embedded systems, which are running on heterogeneous MPSoC, operate under limited power constraints [S⁺20], some of the biggest challenges to execute mobile applications in such systems are as follows:

- Develop methodologies to execute applications on such heterogeneous multi-processors in an power efficient manner.
- The applications should operate at low temperature on the device in order to increase the lifespan of the device itself and hence leading to reliable operations of such systems.
- These methodologies should also cater for performance requirements while being power efficient and reliable at the same time.

To resolve the aforementioned challenges we propose 3 novel methodologies as follows:

- **SoCodeCNN**: Automated feature extraction from program source-code such that proper computing resources could be allocated to the program is very difficult given the current state of technology. In this work, we introduce a novel human-inspired approach, named SoCodeCNN (program **Source Code** for Visual **CNN** classification) [D⁺19e], to automatically convert program source-codes to visual images, understandable by machines. The images are then utilized for automated classification of the program/application into three classes: *Compute, Memory & Mixed* by visual convolutional neural network (CNN) [D⁺18, D⁺19f, D⁺20b] based algorithm. After the source-code of executing application is classified by the CNN, appropriate DVFS is applied to CPUs of the Exynos 5422 MPSoC to cater for performance of the particular type of application, i.e. Compute, Memory & Mixed, and to optimize power consumption on the device.
- **Next**: Mobile device user's usage behaviour changes throughout the day and the desirable Quality of Service (QoS) could thus change for each session. On the other hand, performance of every application in an Android mobile device is based on the generated frame rate of the touch screen of the device. In this work, we introduce a QoS aware agent, named Next [D⁺20c], to monitor the mobile user's usage behaviour to find the target frame rate, which satisfies the user's desired QoS, and applies reinforcement learning based DVFS on the CPU-GPU of the MPSoC to satisfy the frame rate requirement while achieving power saving and reducing peak temperature of the device.
- **DATE**: Side channel attacks remain a challenge to information flow control and security in embedded/mobile devices. One such important security flaw could be exploited through temperature side-channel attacks, where heat dissipation and propagation from the processing cores are observed over time in order to deduce security flaws. In this work, we introduce DATE, **D**efense **A**gainst **T**emperature side-channel attacks [D⁺20a], a novel approach of reducing spatial and temporal thermal gradient, which makes the system more secure against temperature side-channel attacks, and at the same time increases the reliability of the device in terms of lifespan.

2 Conclusion

In this poster presentation, we introduced three novel methodologies: *SoCodeCNN* [D⁺19e], *Next* [D⁺20c] and *DATE* [D⁺20a] to execute applications on heterogeneous MPSoCs in a power-, thermal-efficient and reliable manner.

3 Code Availability

Please contact the corresponding author, Somdip Dey, to get the program code/implementation details of *SoCodeCNN*, *Next* and *DATE* methodologies.

References

- [D⁺] Somdip Dey et al. Energy efficiency and reliability of computer vision applications on heterogeneous multi-processor systems-on-chips (mpsocs). In *Adaptive Many-Core Architectures and Systems workshop, York, UK*.
- [D⁺18] Somdip Dey et al. Mat-cnn-sopc: Motionless analysis of traffic using convolutional neural networks on system-on-a-programmable-chip. In *2018 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*, pages 291–298. IEEE, 2018.
- [D⁺19a] Somdip Dey et al. Deadpool: Performance deadline based frequency pooling and thermal management agent in dvfs enabled mpsocs. In *2019 CSCloud/2019 EdgeCom*. IEEE, 2019.
- [D⁺19b] Somdip Dey et al. Edgecoolingmode: An agent based thermal management mechanism for dvfs enabled heterogeneous mpsocs. In *2019 VLSID*. IEEE, 2019.
- [D⁺19c] Somdip Dey et al. P-edgecoolingmode: An agent based performance aware thermal management unit for dvfs enabled heterogeneous mpsocs. *IET CDT*, 2019.
- [D⁺19d] Somdip Dey et al. Rewardprofiler: A reward based design space profiler on dvfs enabled mpsocs. In *2019 CSCloud/2019 EdgeCom*. IEEE, 2019.
- [D⁺19e] Somdip Dey et al. Socodecnn: Program source code for visual cnn classification using computer vision methodology. *IEEE Access*, 2019.
- [D⁺19f] Somdip Dey et al. Tmav: Temporal motionless analysis of video using cnn in mp-soc. *arXiv preprint arXiv:1902.05657*, 2019.
- [D⁺20a] Somdip Dey et al. Date: Defense against temperature side-channel attacks in dvfs enabled mpsocs. *IEEE Transactions on Computers*, 2020.
- [D⁺20b] Somdip Dey et al. Temporal motionless analysis of video using cnn in mp-soc. In *2020 31st IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP 2020)*. IEEE, 2020.
- [D⁺20c] Somdip Dey et al. User interaction aware reinforcement learning for power and thermal efficiency of cpu-gpu mobile mpsocs. In *2020 DATE*. IEEE, 2020.

Abstract for poster presentation: ACACES 2020

- [exy] Exynos 5 octa (5422). <https://www.samsung.com/exynos> Accessed: 2018-07-23.
- [I⁺19] Samuel Isuwa et al. Teem: Online thermal-and energy-efficiency management on cpu-gpu mpsoes. In *2019 DATE*, pages 438–443. IEEE, 2019.
- [S⁺20] Amit Kumar Singh et al. Dynamic energy and thermal management of multi-core mobile platforms: A survey. *IEEE Design & Test*, 2020.