

EC FP7 STREP HARPA PROJECT: Harnessing Performance Variability



FP7-ICT-2013-10-612069

Project Coordinator

Prof. William Fornaciari

Politecnico di Milano

william.fornaciari@polimi.it

Project website:

www.harpa-project.eu

Partners:

Politecnico di Milano (IT),

IMEC (BE),

ICCS/NTUA (GR)

UCY (CY),

IT4I (CZ),

THALES (FR),

HENESIS (IT)

Duration:

Sep. 2013 – Aug. 2016

Goal and Challenges

The overall goal of our project is to provide architectures (both Embedded Systems (ES) and High Performance Computing (HPC)-oriented) with efficient mechanisms to offer performance dependability guarantees in the presence of unreliable time-dependent variations and aging throughout the lifetime of the system. This will be done by utilizing both proactive (in the absence of hard failures) and reactive (in the presence of hard failures) techniques.

The term “performance dependability guarantee” refers to time-criticality in ES (i.e., meeting deadlines), and a predefined bound on the performance deviation from the nominal specifications in the case of HPC. The promise is to achieve this reliability guarantee in both domains with a

reasonable energy overhead (e.g. less than 10% average). A significant improvement is hence achieved compared to the SotA, which now provides guarantees at the payoff of at least 50% overhead. In addition, we will provide a better flexibility in the platform design while still achieving power savings of at least 20%. To the best of our knowledge, this is the first project to attempt a holistic approach of providing dependable performance guarantees on both ES and HPC systems. This is done while taking into account various non-functional factors, such as timing, reliability, power, and ageing effects. The HARPA project aims to address several scientific challenges in this direction:

(i) Shaving margins. Similar to the circuit technique Razor, but with different techniques at the microarchitecture and middleware, our aim is to introduce margin shaving concepts into aspects of a system that are typically over-provisioned for the worst case.

(ii) A more predictable system with real-time guarantees, where needed. The different monitors, knobs, and the HARPA engine will make the target system more predictable and proactively act on performance variability prior to hard failures.

(iii) Implementation of effective platform monitors and knobs. HARPA will select the appropriate monitors and knobs and their correct implementation to reduce efficiency and performance overheads.

Technical Approach: HARPA Engine Overview

Figure 1(a) below shows the main concepts of the HARPA architecture and the main components of an architecture that can provide performance-dependability guarantees. Note that this generic framework applies to both embedded systems and high-performance general-purpose systems. The main elements that distinguish a HARPA-enabled system are: (i) Monitors and knobs, (ii) User requirements and (iii) HARPA Engine. Conceptually, the HARPA Engine mainly consists of a feedback loop (Figure 1(b)), where the different metrics (performance, timing, power, temperature, errors and manifestations of time-dependent variations etc.) of the system are continuously

monitored. The HARPA engine actuates the knobs to bias the execution flow as desired, based on the state of the system and the performance (timing/throughput) requirements of the application.

The concepts that are to be developed within the HARPA context address equally both the HPC and ES domains. More specifically, from HPC domain we will use Disaster and Flood Management Simulation, while from ES domain Radio frequency spectrum sensing application, Face Detection Application, Object Recognition, and Real-time data monitoring for wearable human motion acquisition system.




Fig.1(a): HARPA High-Level View

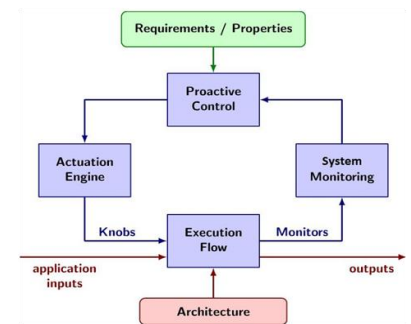


Fig. 1(b): HARPA Engine



Project Coordinator: Prof. William Fornaciari
Politecnico di Milano
william.fornaciari@polimi.it