

HARPA

Harnessing Performance Variability

HARPA

Harnessing Performance Variability

Project ref. FP7-612069

Call ref. FP7-ICT-2013-10

Activity ICT-10-3.4

HARPA Press Release

Dimitrios Soudris
ICCS, GR

Report Number:	D6.6
Version:	v1.0
Date:	November 14, 2013

Revisions List

Date	Version	Author(s)	Description
November 4, 2013	v0.1	Dimitrios Soudris	First Draft
November 11, 2013	v0.2	All partners	Revised Version
November 14, 2013	v1.0	Dimitrios Soudris	Final Version

Executive Summary

This deliverable (issued at M2) presents the results of the activity carried out in Task 6.1 (Dissemination) under the leadership of ICCS. The deliverable consists of the first press release of HARPA project, which will appear in the 37th issue of [HiPEAC Newsletter](#). The press release describes the main objectives and the technical approach of HARPA project.

Updated version of the press release will be delivered by end of project.


HiPEAC *info* 37

COMPILATION **ARCHITECTURE**

APPEARS QUARTERLY
JANUARY 2014

NETWORK OF EXCELLENCE ON
HIGH PERFORMANCE AND EMBEDDED
ARCHITECTURE AND COMPILE

WELCOME TO
THE HIPEAC 2014
CONFERENCE,
JANUARY 20-22, 2014,
VIENNA, AUSTRIA



WWW.HIPEAC.NET

FP7 STREP HARPA PROJECT

Project name: HARPA (Harnessing Performance Variability)
Coordinator: Prof. William Fornaciari, Politecnico di Milano
Partners: Politecnico di Milano, Italy
 IMEC, Belgium)
 ICCS/NTUA, Greece
 UCY, Cyprus
 IT4I, Czech Republic
 THALES, France
 HENESIS, Italy
Start date: September 2013
Duration: 36 months
Website: www.harpa-project.eu



GOAL AND CHALLENGES

The overall goal of the HARPA 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 goal will be achieved by utilizing both proactive techniques (in the absence of hard failures) and reactive techniques (in the presence of hard failures).

The term "performance dependability guarantee" refers to time-criticality in ES (i.e., meeting deadlines), and, in the case of HPC, a predefined bound on the performance deviation from the nominal

specifications. The promise is to achieve this reliability guarantee in both domains within a reasonable energy overhead (e.g. less than 10% average). A significant improvement is hence achieved compared to the SoTA, which currently provides guarantees but with at least 50% overhead. In addition, we will provide 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 for 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) the 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 requirements (timing/throughput) 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 the HPC domain we will use Disaster and Flood Management Simulation, while from the ES domain we will use applications for Radio frequency spectrum sensing, Face Detection, Object Recognition, and Real-time data monitoring for wearable human motion acquisition system.



Project Coordinator, Prof. William Fornaciari, Politecnico di Milano




Figure 1a: HARPA High-Level View




Figure 1b: HARPA Engine