POTRA: A Framework for Building Power Models for Next Generation Multicore Architectures

[Tutorial Overview]

Ramon Bertran
Barcelona Supercomputing Center
C. Jordi Girona 1-3
08034 Barcelona, Spain
ramon.bertran@bsc.es

Marc Gonzàlez
Universitat Politècnica de Catalunya
C. Jordi Girona 1-3
08034 Barcelona, Spain
marc@ac.upc.edu

Xavier Martorell
Barcelona Supercomputing Center
C. Jordi Girona 1-3
08034 Barcelona, Spain
xavier.martorell@bsc.es

Nacho Navarro
Barcelona Supercomputing Center
C. Jordi Girona 1-3
08034 Barcelona, Spain
nacho.navarro@bsc.es

Eduard Ayguadé
Barcelona Supercomputing Center
C. Jordi Girona 1-3
08034 Barcelona, Spain
eduard.ayguade@bsc.es

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Modeling Techniques; C.4 [**Performance of Systems**]: Measurement Techniques

General Terms

Measurement, Experimentation

Keywords

Power estimation, Performance counters, DVFS

1. INTRODUCTION

Controlling the trade-off between performance and power is a complex task that currently is addressed at different levels by hardware, firmware and system software. The effectiveness of a power management policy depends on many parameters that require an actual understanding of the power and performance levels observed for a workload to effectively turn the appropriate knobs. One key aspect of any power management solution is that of relying on meaningful data that bring this understanding and explain the power, temperature and performance levels being generated within the architecture.

Last generation multi-core chips have exposed the limitations for sensor distribution within the chip and the unfeasibility of having per-core measurements. Besides, heuristics exclusively based on sensor distribution do not allow for predicting the effect on performance of any decision at the power management level. Alternative methods are required to get power, temperature and performance estimates: for instance, the IBM POWER7 includes power models based on hardware counters collected on a per-core basis that are used by the embedded firmware to perform the power management.

SIGMETRICS'12, June 11–15, 2012, London, England, UK.
ACM 978-1-4503-1097-0/12/06.

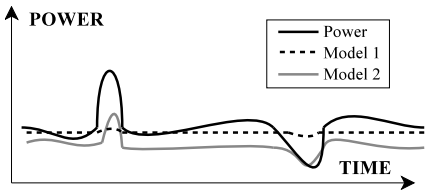

Figure 1: Power model examples. Both models show similar accuracy, but Model 2 is more responsive than Model 1.

The methods based on Performance Monitoring Counters (PMCs) have been shown to be a good solution to estimate power consumption. As a result, their applicability has been demonstrated on several fields such as power management and application profiling. PMC-based power models are used to perform live predictions of power behavior in order to guide power aware policies [2, 3, 4]. Moreover, they are also used in research for quickly exploring new approaches since they allow to profile real systems and full executions of applications, avoiding the need of performing long-time and limited simulations [5]. In the end, they have been crucial in the process of addressing power issues.

Modeling techniques to be embedded in a power management framework demand three primary features: first, simplicity and systematic generation; second, a per architectural component decomposition of the overall estimate, so that the power management decisions can be accompanied by a real understanding of the observed performance levels; and finally, accuracy and responsiveness (see Figure 1).

2. TUTORIAL OVERVIEW

The main topics covered in this tutorial are:

- A systematic methodology for producing PMC-based power models on CMP architectures [4]. The methodology ensures by definition the decomposability and

the generality of the models, while at the same time it produces accurate, responsive and CMP-aware power models.

- A evaluation metric for power models: the *responsiveness*, the ability to track dynamic power variations. This metric is based on power phase detection accuracy and it allows to evaluate the suitability of the models to guide policies built on top of them.
- A case study is presented for an Intel® Core™ 2 Duo. Single and multiple core models are presented for several DVFS states and an empirical evaluation of the produced models is performed using the SPEC-cpu2006, NAS and LMBENCH benchmark suites.

2.1 Tutorial organization

Initially we describe the basic options to generate a platform where to build a model (device for getting actual power measurements, sampling the device, getting the measurements) (*30 min*). Second, we describe how to breakdown a multicore architecture in components from where to obtain contributions on the overall power consumption (*45 min*). Third, we explain the set of counters commonly available in last generation multicore architectures and which ones better describe the activity of each architectural component (*45 min*). Fourth, we present how to linearly correlate the activity in each component with actual power. We describe a basic algorithm that uses stochastic methods to build power models (*30 min*). Finally, the tutorial describes how to validate a power model in terms of accuracy and responsiveness (*30 min*). More details of the tutorial can be found elsewhere [1].

3. ACKNOWLEDGMENTS

We thankfully acknowledge the support of the European Commission through the HiPEAC-2 Network of Excellence (FP7/ICT 217068), the support of the Spanish Ministry of Education (TIN2007-60625, andCSD2007-00050) and the Generalitat de Catalunya (2009-SGR-980).

4. REFERENCES

[1] POTRA Tutorial. `http://rbertran.site.ac.upc.edu/potra`.

[2] F. Bellosa. The benefits of event: driven energy accounting in power-sensitive systems. In *EW 9'00*, Kolding, Denmark, 2000.

[3] C. Isci *et al.* Runtime power monitoring in high-end processors: Methodology and empirical data. In *MICRO'03*, page 93, San Diego, CA, USA, 2003.

[4] R. Bertran *et al.* Decomposable and responsive power models for multicore processors using performance counters. In *ICS'10*, pages 147–158, Tsukuba, Ibaraki, Japan, June 2010. ACM.

[5] R. Bertran *et al.* Energy accounting for shared virtualized environments under dvfs using pmc-based power models. *Future Generation Computer Systems*, 28(2):457 – 468, 2012.

Biographies

Ramon Bertran received a bachelor degree in Computer Science and a master degree in Computer Architecture from Universitat Politècnica de Catalunya (UPC). Since 2005 he is a PhD student in the Computer Architecture Department at UPC. Currently, he is a resident student at Barcelona Supercomputing Center, and his research is focused on power and performance modelling of multicore architectures and energy aware computing.

Marc Gonzàlez received the Engineering degree in Computer Science in 1996 and the Computer Science PhD degree on December 2003. His research has been developed around parallel programming models and high performance computing, in particular, the OpenMP programming model. His current research is centered on this programming model and its use on heterogeneous architectures, based on accelerators and local memories. The main research activities are the design and implementation of software cache techniques for the IBM Cell BE processor as well as the implementation of runtime libraries and compiler support for the OpenMP programming model on this architecture.

Xavier Martorell received the M.S. and Ph.D. degrees in Computer Science from the Technical University of Catalunya (UPC) in 1991 and 1999, respectively. He has been an associate professor in the Computer Architecture Department at UPC since 2001, teaching on operating systems. His research interests cover the areas of paralellism, runtime systems, compilers and applications for high-performance multiprocessor systems. Since 2005 he is the manager of the team working on Parallel Programming Models at the Barcelona Supercomputing Center. He has participated in several european projects dealing with parallel environments (Nanos, Intone, POP, SARC, ACOTES). He is currently participating in the European HiPEAC2 Network of Excellence, and the ENCORE european project.

Nacho Navarro is Associate Professor at the Universitat Politècnica de Catalunya (UPC), Barcelona, Spain, since 1994, and Senior Researcher at the Barcelona Supercomputing Center (BSC). He holds a Ph.D. degree in Computer Science from UPC (1991), Spain. His current interests include: tools for evaluation of multi-core microprocessors, application-specific computer architectures, dynamic reconfigurable logic and resource management in heterogeneous environments and sensor networks. He is also doing research on the programmability and support of hardware accelerators like GPUs at the University of Illinois (IMPACT Research Group). Prof. Navarro is a member of the IEEE, the IEEE Computer Society, the ACM and USENIX.

Eduard Ayguadé received the Engineering degree in Telecommunications in 1986 and the Ph.D. degree in Computer Science in 1989, both from the Universitat Politècnica de Catalunya (UPC), Spain. Since 1987 he has been lecturing on computer organization and architecture and parallel programming models. Currently, and since 1997, he is full professor of the Computer Architecture Department at UPC. His research interests cover the areas of processor microarchitecture, multicore architectures and programming models and their architectural support. He has published more than 100 papers in these topics and participated in several research projects in the framework of the European Union and research collaborations with companies. He is associated director for research on computer sciences at the Barcelona Supercomputing Center (BSC-CNS).