Hardware Performance Counters for Detailed Runtime Power and Thermal Estimations: Experiences and Proposals

Canturk Isci, Gilberto Contreras, Margaret Martonosi
Dept. of Electrical Engineering
Princeton University
{canturk, gcontrer, martonosi@princeton.edu}

Hardware Performance Counters (HPCs) have seen increasing use over the past decade or more, both for performance debugging systems at hardware design time, as well as for software development and performance tuning efforts. Most recently, there has also been significant research effort dedicated to using HPCs for large-scale workload phase behavior characterization and high level adaptive management techniques. Although HPCs are primarily designed to track performance behavior, the looming problems of power-aware and thermal-aware computer systems design have led researchers to also consider their applicability as proxy estimators of system-level power behavior. Such approaches can give detailed real-system power estimates with negligible interference to application behavior. These estimates are then useful for a wide range of goals. First, for example, localized power estimations can be used for detecting thermal hotspots of a processor for dynamic thermal management. Second, variations in estimated power components can enable power-oriented workload phase analysis and prediction, which is of value for both profile-based and dynamic power/thermal adaptations. Finally, for systems with more stringent energy requirements, such as portable embedded devices (PDAs, smart phones and routers), power estimation through HPCs can serve as the cornerstone for a power-aware software infrastructure, in which software performance is dynamically modified at run-time to reflect current energy needs. Run-time environments such as Java Virtual Machines can take advantage of this infrastructure to perform power-aware thread scheduling and power-conscious garbage collection, for example.

In using HPCs for power-related studies, several research questions arise, including: How accurate are the power estimates generated? How does this accuracy vary across platforms? What sorts of system-level power/thermal adaptations do current HPC architectures enable? And how could HPCs be modified in future designs to better support this increasingly important area?

In this talk, we will address these questions in two stages. We will first present two extremes of our experiences with HPC-based power research. At the high-performance end, we will discuss an HPC-based power estimation framework we developed for Pentium4 processors. This work uses Pentium4 HPCs to estimate power for 22 processor sub-components identified from an annotated layout. We can also extend the acquired power estimates to a detailed thermal model, for more quantitative hotspot analysis and validation. On the low-power end, we show the estimation framework for the Intel PXA255 processor, an embedded Intel Xscale technology-based processor targeted at high-end PDAs and smart phones. Power consumption for this processor is estimated using a first-order linear equation that uses performance counters as weight factors for a predetermined set of power parameters. We extend this model to support power estimation in a DVFS environment by augmenting the set of power parameters to reflect various CPU voltage/frequency settings. External memory power consumption is estimated using a similar technique, which, along with CPU power estimation, provides a more detailed characterization of software power behavior.

In the second part of the talk, we will use these and other experiences to guide a set of proposals for how future HPC designs might better support power-related studies and dynamic power adaptation. While HPC "wishlists" for power and performance would certainly have much in common—including good documentation and better portability across platforms—there exist certain design directions that are specifically attractive to power estimation. First and foremost, as power and thermal estimations correspond directly to physical units located on die, counters that individually track accesses to each unit separately is extremely desirable. This set should provide high parallelism in concurrent counting to minimize the need for counter rotations. Such information, together with the documentation of maximum utilizations and maximum power per unit allows easier and more accurate activity based real-time power and thermal estimator implementations. Second, additional structures to track bitline activity and one/zero population counts are imperative for good power and thermal estimates. Third, depending on circuit implementation,

certain queues at the in-order processor front-end and out-of-order engine schedulers dissipate power in proportion to their occupancy. Specifically for these units, metrics that gauge the occupancy ratio lead to better power approximations. Furthermore, in leaner, lower-power embedded processors such as the Intel PXA255 processor, power consumption of support logic outside the core can significantly contribute to processor power consumption. Adding counter support for external core components (off-chip memory accesses, DMA unit activity, etc.) can increase the fidelity of HPC-based power estimation, providing greater opportunities for power behavior control and energy savings.

In summary, our work shows the potential of HPCs to characterize power and thermal behavior—even on very different frameworks and design constraints. We believe, performance counter based approaches will gain even more traction in coming years with their ability to provide feedback over very large timescales at real-time. Based on these experiences, we raise several future design proposals with a specific eye towards even better support of power analysis. We believe these ideas set the foundation for smarter, power-conscious and power-adaptive systems, and will enable architects and software designers to make even broader usage of HPCs in the near future.