SRC Student Symposium 2006 Technical Paper Form

Presentation Title:	Phase Detection and Prediction on Real Systems for Workload-Adaptive Power Management
Presenter, University:	Canturk Isci, Princeton University
Presenter's Expected Graduation Date:	June 2007
Author(s), Organization(s):	Canturk Isci, Princeton University Margaret Martonosi, Princeton University
SRC Research Task ID:	1121.001 (Adaptively-Controlled Execution for Power and Performance)
Task Leader:	Margaret Martonosi

Abstract Text:

Phase Detection and Prediction on Real Systems for Workload-Adaptive Power Management

As power dissipation continues to be one of the primary design constraints of computer systems, runtime, dynamic management and adaptation techniques emerge as essential components of modern systems at various abstractions, from circuits to systems. These techniques allow system hardware and software to recognize and respond on-the-fly to changing patterns in workload behavior to achieve desired power-performance trade-offs. These workload patterns, commonly referred to as *application phases*, exhibit certain repetitive behavior, following from the procedure calls and loops in applications. Therefore, it is important to develop methods to identify and predict repetitive phases, to proactively apply feasible dynamic management responses.

While most of the prior phase analysis work is focused on performance characterization studies and simulation-based approaches, real-system techniques that account for both power and performance phases are especially important for two related reasons. First, to be able to deploy such phase analysis on a real machine, we need techniques that are resilient to variations introduced by the operating system, which pose additional challenges to detecting available recurrent behavior. Second, large-scale system level solutions can only be implemented and tested practically and realistically under such experiments. On the other hand, most of the prior dynamic management approaches on real systems base their actions on recently observed application behavior, thus being prone to misconfigurations due to sudden phase changes. These approaches can also benefit from tracking dynamic phase information to observe patterns for better optimized adaptations to changing workload behavior.

Our work examines the phase behavior of applications running on real systems and describes techniques to detect and predict repetitive application behavior. This information can then be used to guide workload adaptive dynamic management techniques such as dynamic voltage and frequency scaling. In our real-system experiments, we mainly rely on hardware performance monitoring counters (PMCs) to track application behavior and perform multimeter and data acquisition hardware measurements to evaluate power consumption.

Phase Detection: Workload execution can be characterized with many different features ranging from dynamic control flow, such as traversed basic blocks, to performance metrics as provided by PMCs. Our previous experiments show that specific subsets of PMCs provide a close proxy to actual processor power consumption and characterizations based on these can reveal accurate representations of application phases. However, our observations demonstrate that identifying the repetitive behavior directly from this information is very challenging under real-system variability, as both the timing and metric behavior changes at each occurrence of a specific phase. These variations are reflected in the observed phases as various transformations. We classify these as *time shifts, glitches, gradients, mutations* and *time dilations.* Time shifts represent the lag between two observed similar patterns. Glitches and gradients are artifacts of sampling that lead to many spurious phase transitions. Mutations result from metric variability and cause similar behavior to be characterized as different phases. Time dilations are due to small time variations that alter the extent of observed phases. To extract available recurrence information under these effects, we propose a *transition-guided* phase detection framework that relies on phase change information to identify repetitive execution. We develop effective methods to improve detection of phase transition signatures, such as *glitch/gradient filtering* and *near-neighbor blurring*. Under this transition-guided detection framework, we achieve 2-8X better detection capabilities with respect to the initial phase characterizations, under real-system variability. Our best detection scheme can identify experimented phase signatures with less than 5% false alarms, demonstrating the effectiveness of our detection framework.

Phase Prediction and Application to Dynamic Responses: The primary motivation for our research is to assess the potential of phases for guiding dynamic power management techniques. Currently, most of the system-level management methods are applied *reactively*, based on the recent behavior an application exhibits. However, phase information can provide more insight to application behavior, enabling these actions to be exercised proactively. To evaluate the opportunities with such an approach, we develop an experimental on-the-fly phase prediction framework with application to dynamic voltage and frequency scaling. We define application phases corresponding to different operating modes of the system and propose a configurable runtime phase predictor that can seamlessly monitor and predict workload behavior. This, in turn, provides hints to possible adaptive control opportunities. We experiment with various types of predictors and achieve the highest accuracy with a table based predictor derived from a global branch predictor architecture. Our proposed *global phase history predictor* predicts upcoming phases based on the observed phase pattern history. In comparison to the original reactive response mechanisms, this phase based prediction methodology provides significant improvements for workloads with varying behavior. While for a widely stationary benchmark such as `twolf` both approaches achieve close to 100% accuracies, for a highly variable application such as `applu`, the reactive approach can choose the correct operating mode for only 47% of the time. By applying phase based prediction to the same application, we can improve this to 92%. On average, for benchmarks with significant variations, our predictive methods can improve the accuracy of chosen operation modes by more than 20%, enabling better adaptation of dynamic management to workload execution patterns. In our current work, we are extending our runtime phase predictor framework to integrate with different system power and thermal management applications. Currently, we are investigating the benefits of phase prediction guided dynamic voltage and frequency scaling, where the operating modes are chosen based on predicted next phases. Our future research plans include also applying phase history and prediction information to detect and avoid thermal emergencies via thermal-aware scheduling.

Overall, our presented research shows a roadmap to effective on-the-fly phase detection and prediction on real-systems for application to workload-adaptive dynamic management techniques. With the increasing focus on adaptive and autonomous system management, such methods for predicting and reacting to dynamically varying application behavior become critical for power-performance efficiency. We believe our described research offers practical techniques that can serve as useful guides for dynamic management applications in current and emerging power aware systems.