GreenLight Technical Report

2011

Power and Thermal Management

1.A.1. Tajana Simunic Rosing: Power/Thermal Management

The goals of our work are to develop strategies to analyze the power, thermal and workload dynamics, and to design methods to reduce power consumption while mitigating the temperature induced problems in datacenter environments. Thus far we have implemented and tested a number of reactive and proactive thermal management techniques evaluated the benefits of optimizing temperature while minimizing energy versus only focusing of lowering the system power/energy consumption. We are in process of extending Xen, a virtual machine monitor for IA-32 IA-64 and PowerPC 970 architectures, to enable power and thermal management within the virtualization technology.

Energy management in virtualized environments [TODAES'11]
We designed vGreen, a multi-tiered software system for energy efficient virtual machine management in a clustered virtualized environment. The system leverages the use of novel hierarchical metrics that work across the different abstractions in a virtualized environment, and capture power and performance characteristics of both the virtual and physical machines. These characteristics are then used to implement policies for scheduling and power management of virtual machines across the cluster. We show through real life implementation of the system on a state of the art testbed of server machines that vGreen can improve both average performance and system level energy savings by close to 40% across benchmarks with varying characteristics.

Power management in general purpose systems under tight performance constraints [ISLPED'11, ISLPED'10]
We developed a new DVFS algorithm for enterprise systems that elevates performance as a first order control parameter and manages frequency and voltage as a function of performance requirements. The algorithm uses closed-loop control based on state space model and as such is able to guarantee meeting tight performance constraints. We implement our algorithm on Intel Westmere platform in Linux and demonstrate its ability to meet stringent performance targets while minimizing power dissipation. Our algorithm reduces the standard deviation from target performance by more than 90% over state of the art policies while minimizing average power by 17%. Runtime characteristics of individual threads (such as IPC, cache usage, etc.) are a critical factor in making efficient scheduling decisions in modern chip-multiprocessor systems. They provide key insights into how threads interact when they share processor resources, and affect the overall system power and performance efficiency. We propose and implement mechanisms and policies for a commercial OS scheduler and load balancer which incorporates thread characteristics, and show that it results in improvements of up to 30% in performance per watt.

Power estimation and prediction in virtualized environments [DAC'10]
We designed a system for online power prediction in virtualized environments. It is based on Gaussian mixture models that use architectural metrics of the physical and virtual machines (VM) collected dynamically by our system to predict both the physical machine and per VM level power consumption. A real implementation of our system shows that it can achieve average prediction error of less than 10%, outperforming state of the art regression based approaches at negligible runtime overhead.

Thermal management in heterogeneous MPSoCs [ASPDAC'10]
We designed PASTEMP, a solution for Package Aware Scheduling for Thermal and Energy management using Multi- Parametric programming in heterogeneous embedded multiprocessor SoCs (MPSoCs). Based on the current thermal state of the system and current performance requirements of the workload, PASTEMP finds thermally safe and energy efficient voltage/frequency configurations for the cores on a MPSoC. The tasks are assigned to the cores

depending on their performance demand and the current voltage/frequency of the core. The voltage/frequency settings of the cores are chosen through an optimization process which is based on the instantaneous thermal model we introduce to decouple the effect of package temperature from the temperature changes caused by the power consumption of the cores. To be able to find the best voltage/frequency settings at runtime, we use multi-parametric programming to separate the optimization into offline and online phases. According to our experimental results, compared to similar DTM techniques, PASTEMP results in up to 23% energy saving and 26% throughput improvement and reduces the deadline misses to more than a half while meeting all thermal constraints.

Proactive thermal management of memory subsystems [ISLPED'10R]
Energy management of memory subsystem is challenging due to performance and thermal constraints. Big energy gains can be obtained by clustering memory accesses, however this also leads to a higher need for cooling due to larger temperatures in active areas of memory. Our solution to memory thermal management problem is based on proactive thermal management that intelligently allocates workload pages to few memory units and powers down rest of the memory. Our experimental results show that this approach improves energy savings by 43% and reduces performance overhead by 85% with respect to the state of the art polices.

References:

[ASPDAC'10]
S. Sharifi, A. Coskun, T. Simunic Rosing, "Hybrid Dynamic Energy and Thermal Management in Heterogeneous Embedded Multiprocessor SoCs," ASP-DAC, 2010.

[ICCD'10]
S. Sharifi, T. Simunic Rosing, "Package-Aware Scheduling of Embedded Workloads for Temperature and Energy Management on Heterogeneous MPSoCs", ICCD'10.

[DAC'10]
G. Dhiman, K. Mihic, T. Simunic Rosing, "A System for Online Power Prediction in Virtualized Environments Using Gaussian Mixture Models", DAC'10.

[ISLPED'10]
G. Dhiman, V. Kontorinis, D. Tullsen, T. Simunic Rosing, E. Saxe, J.Chew, "Dynamic Workload Characterization for Power Efficient Scheduling on CMP Systems," ISLPED'10.

[ISLPED'10R]
R. Ayoub, K. Indokuri, T. Simunic Rosing, "Energy Efficient Proactive Thermal Management in Memory Subsystems," ISLPED'10.

[ISLPED'11]
R. Ayoub, T. Simunic Rosing, "Power minimization under Tight Performance Constraints in General Purpose Systems," submitted to ISLPED'11.

[TODAES'11]
G. Dhiman, G. Marchetti, T. Simunic Rosing, "vGreen: A System for Energy Efficient Management of Virtual Machines," TODAES'11.