

Energy-Efficient Cluster Computing via Accurate Workload Characterization

S. Huang and W. Feng
Department of Computer Science
Virginia Tech
{huangs, feng}@cs.vt.edu

Abstract

This paper presents an eco-friendly daemon that reduces power and energy consumption while better maintaining high performance via an accurate workload characterization that infers “processor stall cycles due to off-chip activities.” The eco-friendly daemon is an interval-based, run-time algorithm that uses the workload characterization to dynamically adjust a processor’s frequency and voltage to reduce power and energy consumption with little impact on application performance.

Using the NAS Parallel Benchmarks as our workload, we then evaluate our eco-friendly daemon on a cluster computer. The results indicate that our workload characterization allows the power-aware daemon to more tightly control performance (5% loss instead of 11%) while delivering substantial energy savings (11% instead of 8%).

1. Introduction

According to a recent IBM report [17], the annual budget for power and cooling is fast approaching the annual budget for new server spending. This is arguably why Google located one of their new data centers in rural Oregon on the Columbia River — to take advantage of the cheap hydroelectric power generated by the nearby Grand Cooley Dam [19]. Rapidly increasing utility bills, coupled with how heat from excessive power affects reliability [5], motivates the need for power awareness in cluster computers, whether in a supercomputing center or a large-scale data center.

One way to address this growing problem is to improve the power and energy efficiency of cluster computers at different levels of abstraction: hardware [15], [4], systems integration [5], systems software [8], [7], middleware [13], and applications software [18]. In this paper, we use a *systems-software approach* that leverages *accurate workload characterization* via a *unique*

synthesis of hardware performance counters in order to determine *when and how* to use dynamic voltage and frequency scaling (DVFS) to improve power efficiency while strictly maintaining performance. Because the power consumption of a processor is proportional to its clock frequency and the square of its voltage supply, we use DVFS, available on virtually all modern processors, e.g., SpeedStep on Intel and PowerNow! on AMD, to set the voltage and frequency of the processor so as to reduce power consumption.

In a DVFS-enabled processor, a low (or high) power-consuming mode corresponds to processor that runs at a low (or high) frequency and voltage. Thus, a DVFS algorithm should reduce the frequency and voltage of a processor only when the processor is not needed to do useful work, e.g., waiting for the completion of a large block of I/O accesses. Application performance during such periods of off-chip access is insensitive to processor performance. Thus, we can reduce the processor voltage and frequency during such periods to reduce power consumption while maintaining application performance.

However, given that the time to scale voltage and frequency takes $O(10,000,000)$ clock cycles, sophisticated use of DVFS is needed if power and energy savings is to be realized within a performance bound. Enabling such use requires accurate workload characterization, one of the main contributions of this work. Then, the challenge is to make use of this workload characterization to *autonomically* scale the frequency and voltage.

Thus, this paper presents a novel methodology for workload characterization on a per-node basis that is then used to enable intelligent power-aware computing across computational nodes in a cluster, data center, or grid. For the purposes of this paper, however, our focus will be on the former.

We refer to our power-aware, eco-friendly algorithm as *eco* and its implementation as *ecod*. The *ecod* daemon manages application performance and power consumption in real time based on an accurate mea-

surement of CPU stall cycles due to off-chip activities and does *not* require application-specific information a priori. The paper will show that `ecod` limits performance impact to only 5.1% (if our performance-bound knob is set to 5%) with less than 3.5% variance, both better than the current state of the art, while saving up to 50% in processor energy and up to 19% in overall system energy.

The remainder of the paper is organized as follows. Section 2 discusses related work on power-aware algorithms and its workload characterization. Section 3 presents our novel workload characterization based on CPU stall cycles due to off-chip activities. In Section 4, we present our power-aware, run-time algorithm called *eco*, which is implemented as a daemon. Sections 5 and 6 present our experimental setup and results, respectively, followed by a conclusion in Section 7.

2. Related Work

The past few years has seen significant research in power-aware cluster computing, which can be categorized into two types [7]: off-line, trace-based scheduling [1], [6], [16] and on-line, profile-based scheduling [8], [11], [7], [3]. For brevity, the related work below focuses on the latter, which is the more challenging problem.

Lim et al. [11] design an MPI run-time system that dynamically reduces CPU performance during communication phases in MPI programs. Curtis-Maury et al. [3] present a framework for autonomic power-performance adaptation of multi-threaded programs using thread throttling. However, these two works have limited application in that they are designed only for MPI and OpenMP applications, respectively. For power-aware research using general workload characterization, Choi and Pedram [2], Hsu and Feng [8] (β algorithm), and Ge et al. [7] possess the current state of the art for general computing systems.

Choi and Pedram propose a DVFS approach based on the ratio of off-chip access to on-chip computation time that is targeted to embedded systems. It uses the number of instructions and external memory accesses to compute the ratio of off-chip computation time to on-chip computation time. However, this workload characterization is *CPU-frequency dependent* and cannot characterize an application. Why? Theoretically, the off-chip access time is constant no matter what CPU frequency is used while on-chip computation time will decrease as CPU frequency increases. Hence, the ratio of off-chip access to on-chip computation time depends on the CPU frequency. Moreover, Choi's work

only considers memory access and ignores thread synchronization in exploring energy-saving opportunities.

The β algorithm [8] of Hsu and Feng assumes that CPU boundedness is indirectly reflected via the MIPS (millions of instructions per second) rate. Since the MIPS rate only approximately reflects CPU boundedness and is dependent on CPU frequency, it cannot accurately characterize application workload nor can it effectively bound performance loss. Another (arguable) drawback is that the β algorithm takes the entire history of workload into consideration when making DVFS decisions. While appropriate for some applications, it is not for many other applications.

CPU MISER [7] relies on retired instruction and cache-access statistics to provide information about the on-chip workload. It also uses constant values to approximate two important variables in its workload characterization. As such, this approach only accurately characterizes workload on average. Moreover, the workload characterization of CPU MISER is CPU frequency dependent.

The Linux on-demand governor is widely provided in the `CPUFreq` subsystem of a recent Linux kernel. It dynamically changes CPU frequency depending on CPU utilization [14]. Because CPU utilization is misleading in terms of characterizing a program's workload, the on-demand governor cannot efficiently deliver both power savings while controlling performance loss.

3. Workload Characterization

From a power-aware perspective, different applications create different opportunities for energy savings. Execution phases with memory-intensive activities have been an attractive target for DVFS algorithms because the time for a memory access is independent of how fast the processor is running. When frequent memory or I/O accesses dominate a program's execution time, they limit how fast the program can finish executing. It is this *memory wall* that provides an opportunity to reduce power and energy consumption while maintaining performance. In cluster computing and grid environments, there are further opportunities for power and energy savings, e.g., during network or I/O synchronization, e.g.,

Below we derive a parameter λ to characterize application workloads in order to assist in simultaneously optimizing performance and power. We then present our methodology for measuring λ using CPU stall cycles due to off-chip activities.

3.1. Deriving a Workload Model

Let $T(f)$ denote the time to execute a program at a CPU frequency f . The total number of clock cycles to execute a program is then given by $T(f) \cdot f$. Alternatively, as noted in Eq. (2), this value can be realized as the number of clock cycles to execute on-chip activities, C_{on} , plus the number of clock cycles to execute off-chip activities, C_{off} , where the execution time of the former is frequency-sensitive and the latter is frequency insensitive. That is, C_{on} is the number of CPU cycles whose execution is affected by frequency variation while C_{off} is the number of CPU cycles whose execution is *not* affected by frequency variation.

$$C_{on} + C_{off} = T(f) \cdot f \quad (1)$$

We define T_{off} to represent the execution time that is CPU frequency-insensitive.

$$T(f) = C_{on} \cdot \frac{1}{f} + T_{off} \quad (2)$$

When a program runs at maximum frequency f_{max} ,

$$T(f_{max}) = C_{on} \cdot \frac{1}{f_{max}} + T_{off} \quad (3)$$

Note that T_{off} in Eq. (3) is the same as in Eq. (2) when executing the same amount of program instructions since T_{off} is not affected by the change of CPU frequency f .

To quantify the performance loss, we define a parameter δ that indicates the performance bound in employing DVFS,

$$\frac{T(f) - T(f_{max})}{T(f_{max})} < \delta \quad (4)$$

Substituting $T(f)$ and $T(f_{max})$ from Eq. (2) and (3), respectively, into Eq. (4), we get

$$\frac{C_{on}}{C_{on} + T_{off} \cdot f_{max}} \cdot \frac{f_{max} - f}{f} < \delta$$

The equation can be reformulated as

$$\lambda \cdot \frac{f_{max} - f}{f} < \delta \quad (5)$$

where

$$\lambda = \frac{C_{on}}{C_{on} + T_{off} \cdot f_{max}} \quad (6)$$

The workload characterization, denoted by λ in Eq. (6), can be reformulated as

$$\lambda = \frac{C_{on}}{C_{on} + C_{off} \cdot \frac{f_{max}}{f}} \quad (7)$$

By combining Eq. (1) and (7), we eliminate the direct

dependence on C_{on} , thus resulting in

$$\lambda = \frac{f^2 T(f) - f C_{off}}{f^2 T(f) - f C_{off} + f_{max} C_{off}} \quad (8)$$

where $0 \leq \lambda \leq 1$.

From Eq. (6), the workload characterization λ is a parameter that is independent of the CPU frequency that the application is running at. λ only depends on the application itself. Eq. (7) shows that λ characterizes the percentage of on-chip cycles out of the total CPU cycles when running at frequency f_{max} . In Eq. (7), when λ equals to 1, C_{off} is 0, which means that the program spent all its time on on-chip activities. When λ equals 0, C_{on} must be 0, which means the program spent all its time on off-chip activities. Eq. (8) provides a method to quantify the behavior of applications even if they are not running on frequency f_{max} .

3.2. Methodology for Measuring CPU Off-Chip Stall Cycles

In this section, we present our methodology for measuring C_{off} . In order to achieve the desired accuracy, we obtain the CPU stall cycles due to off-chip activities from two aspects: on-chip (C_{off}^{on}) and off-chip (C_{off}^{off}).

3.2.1. Measuring from the On-Chip Perspective.

$$C_{off}^{on} = SC_{total} - SC_{on} \simeq SC_{total} - SC_{branch} - SC_{reorder}$$

where C_{off}^{on} is the on-chip measurement of CPU stall cycles due to off-chip activities. For our platform, we measure SC_{total} using the CPU's decoder/dispatch stall cycles and measure SC_{on} using the sum of the CPU's decoder stall cycles due to branch misprediction (SC_{branch}) and full reorder buffer ($SC_{reorder}$). Why choose these two events? They dominate CPU stall cycles due to on-chip activities and hardly overlap with each other. There are also other stall cycles contributors, e.g. segment load, serialization, and so on. However, our empirical results show that CPU stall cycles contributed by these events are small; thus, we ignore them in our estimation.

3.2.2. Measuring from the Off-Chip Perspective.

$$C_{off}^{off} = N_{mem} \cdot \tau_{mem} \cdot f + T_{io} \cdot f + T_{idle} \cdot f$$

where C_{off}^{off} is the off-chip measurement of CPU stall cycles due to off-chip activities. N_{mem} is the number of off-chip memory accesses; τ_{mem} is the memory-access latency; T_{io} is the CPU stall time for waiting on I/O completion; and T_{idle} is the CPU idle time. We use L2 cache misses to emulate the number of off-chip

memory accesses and use LMBench [12] to measure the memory-access latency τ_{mem} . T_{io} and T_{idle} can be obtained through `/proc/stat` on Linux systems.

3.2.3. Synthetic Measurement. We obtain our final measurement by taking the minimum of on-chip and off-chip measurement of CPU stall cycles due to off-chip activities, i.e.,

$$C_{off} = \min(C_{off}^{on}, C_{off}^{off})$$

Why take the minimum? Both measurements overestimate the number of CPU stall cycles. On the one hand, for on-chip measurement, there is no such hardware event that can measure CPU stall cycles due to off-chip activities directly. In order to estimate it, we choose a set of events out of many that can cause on-chip CPU stalls, e.g. branch abortion, serialization, full reorder buffer [10]. Moreover, most of the events involve both on-chip activities and off-chip activities. Therefore, an event cannot be simply treated as an event due to on-chip activities or off-chip activities. To exacerbate the problem, the events sometimes overlap with each other.

On the other hand, off-chip measurement is also not accurate enough. Both off-chip memory accesses and memory latency are hard to determine precisely. The L2 cache misses measured by the hardware counter usually include misses due to speculative execution. Additionally, due to CPU prefetching and block transfer, some L2 cache misses will be combined and transferred together. Thus, the actual number of memory accesses will be smaller than the measured value.

Two facts lead us to combine on-chip and off-chip measurements. For CPU-bound applications, L2 cache misses are smaller and the opportunity for combining and overlapping cache misses is small. Thus, off-chip measurement works better for CPU-bound applications. For non-CPU-bound applications, however, CPU stall cycles due to off-chip activities dominate the total CPU stall cycles. Therefore, on-chip measurement fits non-CPU-bound applications well.

4. *eco* Algorithm

Here we present our workload-aware, eco-friendly algorithm called *eco*. The algorithm consists of multiple components: (1) the high-level algorithm itself that periodically determines whether to scale the frequency and voltage, (2) workload prediction to enable the decision of what to scale the frequency (and voltage) to, and (3) once a frequency is determined, how to schedule and emulate the frequency (and voltage) if the platform does not explicitly support the frequency.

4.1. Overview of Algorithm

The *eco* algorithm is an interval-based, run-time algorithm, whose execution time is divided into intervals that span the running time of an application program. Within each interval, the algorithm performs the following:

1. *Characterizes the workload for the current interval*, as noted in Section 3. As stated before, frequent memory and I/O access, network process synchronization, as well as CPU idling constitute the three main opportunities for power-aware computing. However, these three opportunities vary from application to application and change from time to time. In short, the *eco* algorithm quantifies the application behavior at run time for each interval.

2. *Predicts the workload characterization for the next interval*. The *eco* algorithm predicts the workload for the next interval based on that of previous intervals. It uses the average of a λ window of previous intervals to predict the workload, since we observe that workload tends to be constant for short periods of time.

3. *Schedules the frequency for the next interval*. The *eco* algorithm schedules the CPU frequency based on the predicted workload characterization in order to maintain the performance bound while saving as much energy as possible. However, we must address two problems with frequency scheduling in real systems: (1) CPUs only support discrete frequencies, and (2) CPU frequencies have a lower and upper bound.

4.2. Workload Prediction

Though workloads may vary from application to application, the workloads can still be predictable at some level. For example, we set a window size of L and use the average across the window to predict the λ in current interval. The window size cannot be too large, otherwise the DVFS scheduler will not be reactive to workload variation. The window size cannot be too small either as it risks significant prediction error. Empirically, we set the window size to be 3 by default in our implementation of *eco* algorithm — *ecod*, short for *EcoDaemon*.

4.3. Frequency Scheduling and Emulation

Assuming that $\bar{\lambda}$ is the predicted workload characterization for the current interval, then based on Eq. (5), the ideal frequency for the current interval is

$$f^* = \frac{\bar{\lambda} \cdot f_{max}}{\bar{\lambda} + \delta} \quad (9)$$

However, the available frequencies in a real system are limited. Thus, as in [8], f^* needs to be calculated as

$$f^* = \max(f_{min}, \frac{\bar{\lambda} \cdot f_{max}}{\bar{\lambda} + \delta})$$

Finally, the calculated frequency f^* may not be directly supported on a real system. So, we apply the method proposed in [8] to emulate the calculated frequency f^* .

4.4. The *eco* Algorithm

Synthesizing the steps shown above, we design our *eco* algorithm. Figure 1 presents the pseudocode for the *eco* algorithm. Steps 1 and 2 encompass workload characterization. Step 3 is workload prediction, and Steps 4 and 5 deal with frequency scheduling and emulation.

Hardware:

n frequencies f_1, \dots, f_n

Parameters:

I : time-interval size

δ : performance bound

L : prediction window size

Algorithm:

Initialize the λ window

Repeat

1. Measure CPU stall cycles due to off-chip activities for current interval C_{off}
2. Compute coefficient λ for current interval

$$\lambda = \frac{f^2 I - f C_{off}}{f^2 I - f C_{off} + f_{max} C_{off}}$$

3. Predict the workload for next interval for all λ in window $[0, L]$

$$\bar{\lambda} = \text{average}(\lambda)$$

4. Compute the desired frequency f^*

$$f^* = \max(f_{min}, \frac{\bar{\lambda} \cdot f_{max}}{\bar{\lambda} + \delta})$$

5. Schedule next interval I at f^*

Figure 1. Pseudocode for *eco* algorithm

5. Experimental Set-Up

Here we detail the experimental set-up for evaluating our *eco* algorithm, including hardware and software platform, power and energy measurement, and *ecod* implementation.

5.1. Experimental Platform

The hardware platform in our experiment includes a four-node cluster for computing and an additional node for recording the power and energy consumption. Each compute node contains two dual-core AMD Opteron 2218 processors and 4-GB main memory. Each CPU core includes one 128-KB split instruction and data L1 cache. Two cores on the same die share one 1 MB of L2 cache. Each processor supports *six* power/performance modes. Finally, the nodes are interconnected with Gigabit Ethernet.

We run Red Hat Linux (kernel version 2.6.18) on each compute node. The Linux kernel `CPUFreq` subsystem is used for controlling DVFS and `PERFCTR` for hardware counter monitoring. With respect to the benchmarks, we use the latest NAS Parallel Benchmarks (NPB3.2-MPI). We use `mpich2` (version 1.0.6) to run the benchmarks.

5.2. Energy Measurement and Processing

We use the “Watts Up? PRO ES” power meter to measure the total system energy for each node. Energy values are recorded immediately before and after the benchmark runs. The difference of the two energy values is the energy consumed by the system when the benchmark ran. Since DVFS scheduling only affects the power consumption of CPU, it is (arguably) misleading to evaluate our *eco* algorithm based on the energy consumption of total system. So, in addition to reporting the total system energy, we also evaluate the effect of the *eco* algorithm on CPU energy by applying a CPU power model used in [8] to isolate the CPU energy from the total system energy.

5.3. The *ecod* Implementation

Figure 2 illustrates the software architecture of our *ecod* implementation. We implement *ecod* as a lightweight daemon that monitors all the cores in a node and schedules appropriate frequencies for them. When *ecod* starts up, it reads the configuration file and dynamically detects processor settings, e.g. available frequencies, number of cores, etc. In each sampling interval, the master daemon fetches hardware-event information from the “Hardware Event Monitor Module.” Then, workload prediction and performance rectification are performed. The former is discussed in Section 4.2 while the latter is a mechanism to compensate for the performance loss due to the misprediction of λ . Due to space limitations, we refer to technical report [9] for more details. In the end, the

daemon dispatches the desired frequency to “DVFS Scheduler Module,” which then takes care of frequency scheduling of the cores.

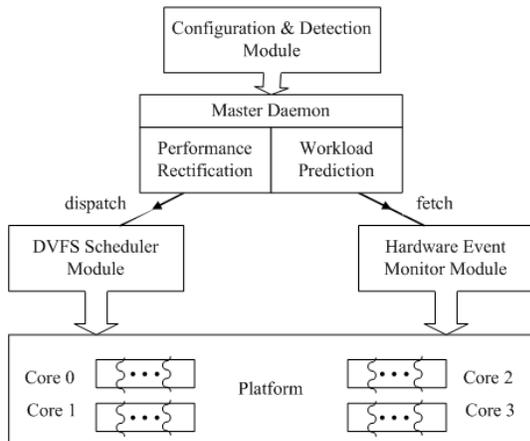


Figure 2. Software architecture of `eod`

5.4. Parameters and Sensitivity Analysis

`eod` is configurable and tunable. The user-configurable parameters are *sampling interval*, *performance bound*, and *prediction window size*. Below are the tradeoffs of these user-configurable parameters.

Sampling Interval. As sampling intervals increase in length, the precision of workload characterization and its prediction will worsen, resulting in performance that cannot be tightly controlled. Conversely, when the sampling intervals get too short, the overhead of sampling the workload and scheduling the frequency is not as easily amortized.

Performance Bound. The larger the performance bound (or percentage slowdown), the more energy that will be saved. However, once the frequency reaches the system’s lowest frequency, it cannot save any more energy.

Prediction Window Size. If the window size is large, the algorithm will depend on a larger amount of historical information, thus making more instantaneous workload prediction inaccurate. If the window size is small, the algorithm will be too sensitive to the workload variation.

In our experiments, we compare `eod` with the β -algorithm [8] and the Linux on-demand governor [14]. For `eod`, I is set to 1 second and δ is set to 5% and L is set to 3. For β algorithm, the performance constraint is set to 5%. As for Linux on-demand governor, we use the default configuration with a sampling rate of 560,000 ms and up threshold of 80%.

6. Experiments and Analysis

In this section, we first validate the workload characterization λ obtained by measuring the CPU stall cycles due to off-chip activities against an off-line approach [9]. Then, we evaluate the workload prediction method used in `eco` algorithm. Finally, we demonstrate the efficacy of `eod`, our power-aware daemon based on `eco`, on the NAS Parallel Benchmarks (NPB3.2-MPI) in a cluster environment.

6.1. Validation of Workload Characterization

Before evaluating `eco` on the NAS Parallel Benchmarks, we first validate our workload characterization (λ) on a representative set of 10 SPEC CPU2000 benchmarks: three CPU-bound, three memory-bound, and four in between. Specifically, by evaluating λ , we indirectly evaluate our approach to measure CPU stall cycles due to off-chip activities.

Figure 3 shows our evaluation of measured λ to that of an off-line approach [9], with the benchmarks arranged in such a way that the CPU-boundedness (i.e., Y-axis) of the benchmarks decrease going left to right. The error of the measured λ to off-line value is only 3.4% on average.

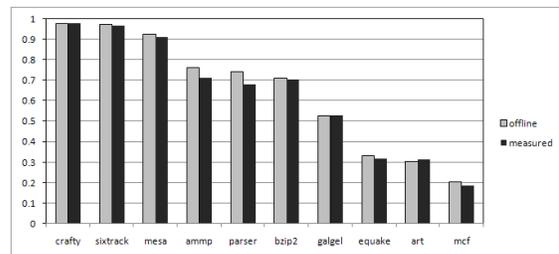


Figure 3. Validation of measured λ against that obtained offline

6.2. Evaluation of Workload Prediction

Here we use the workload characterization (λ) obtained by CPU stall cycles due to off-chip activities as a baseline in order to evaluate the effectiveness of our workload prediction method. Due to space constraints, we only show the *crafty* benchmark from SPEC CPU2000 to illustrate the predictive performance of our methodology.

Figure 4 shows a comparison between the measured λ and predicted λ for the *crafty* benchmark, where the y-axis denotes the workload characterization λ . Over the execution time of the benchmarks, the difference

between measured λ and predicted λ is within 2%. The figure also shows that the predicted λ changes more smoothly than the measured λ . This reflects the stability of our algorithm, which in turn, avoids significant DVFS scheduling overhead since the larger the frequency transition, the more overhead that is induced in DVFS scheduling [8].

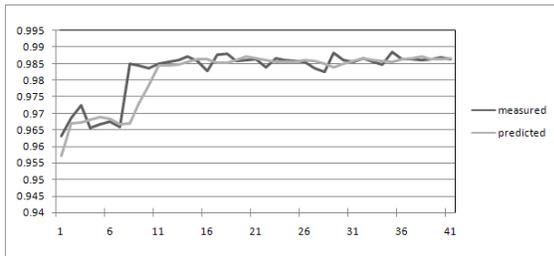


Figure 4. Measured λ versus predicted λ for *crafty* over execution time

6.3. Parallel Experiment

With the validation of our workload characterization and workload prediction, coupled with our sensitivity analysis, all on a per-node basis as shown above, we are now ready to evaluate our *eco* algorithm, implemented as an eco-friendly daemon that we call *ecod* in a cluster environment. In such an environment, we expect the performance of our eco-friendly daemon to be quite good given the additional opportunities for energy savings due to frequent memory and I/O access, network process synchronization, as well as CPU idling.

To evaluate *ecod*, we use the NAS Parallel Benchmarks. We run the benchmarks with a Class C workload on 16 cores across four compute nodes, with each compute node containing four cores. Since the cores on the same die have a common power/performance mode, we schedule the core frequency according to the higher one on the same die in order to guarantee performance.

Figures 5 and 6 show the performance control and energy savings of *ecod* in comparison with the β algorithm and Linux on-demand governor, respectively. Table 1 summarizes the statistics on performance loss and energy savings. The performance loss averages 5.1%, which is better than the β algorithm (10.6%) and Linux on-demand governor (7.9%). The standard deviation of performance loss for *ecod* is also the best among the three algorithms.

The CPU energy savings are comparable between *ecod* (average of 31.5%), β -algorithm (average of

Table 1. Statistics on Parallel Experiment

	<i>ecod</i>	β	on-demand
Perf. Loss (Avg)	5.1%	10.6%	7.9%
Perf. Loss (Std Dev)	3.5%	10.3%	7.7%
Energy Savings (Avg)	31.5%	32.9%	28.6%

32.9%) and on-demand governor (average of 28.6%). Considering that *ecod* achieves the same energy saving by sacrificing far less performance, *ecod* clearly performs better than the β algorithm and Linux on-demand governor. Theoretically, if we set the performance bound δ of *eco* to the actually performance loss of β -algorithm (10.6%) and on-demand governor (7.9%) respectively, the energy saving of *eco* will be more than that of the later two.

Finally, with respect to overall energy savings, *ecod* performs better than the β algorithm and the Linux on-demand governor on average, as shown in Figure 7. *ecod* can achieve 11% energy savings on average across the NAS Parallel Benchmarks. Both β and the Linux on-demand governor save 8% of energy for the same benchmarks on average.

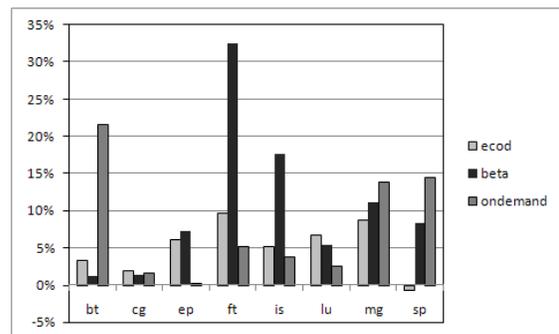


Figure 5. Performance loss on NAS parallel benchmarks

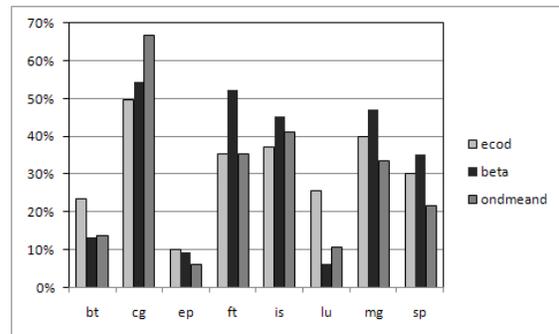


Figure 6. CPU energy savings on NAS parallel benchmarks

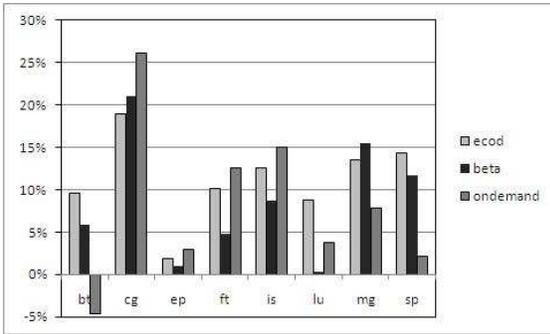


Figure 7. Overall energy savings on NAS parallel benchmarks

7. Conclusion

This paper presents a novel behavioral quantification of cluster workloads using CPU stall cycles due to off-chip activities. We leverage this quantification to create a power-aware, eco-friendly, run-time algorithm called *eco*. This algorithm dynamically monitors processor states and obtains the workload characterization at run time in order to guide the appropriate scaling of frequencies and voltages in a parallel computing environment. Results show that our implementation *ecod* achieves the best performance control over the β -adaptation algorithm and Linux on-demand governor while delivering an overall energy savings of 11%.

Acknowledgement

We would like to thank Chung-Hsing Hsu for his feedback on earlier versions of the manuscript.

This research was supported in part by NSF Grant 0848670.

References

- [1] K. W. Cameron, R. Ge, and X. Feng. High-performance, Power-aware Distributed Computing for Scientific Applications. In *IEEE Computer*, 2005.
- [2] R. Choi and M. Pedram. Fine-Grained Dynamic Voltage and Frequency Scaling for Precise Energy and Performance Trade-off Based on the Ratio of Off-Chip Access to On-Chip Computation Times. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 24(1), 2005.
- [3] M. Curtis-Maury, J. Dzierwa, C. D. Antonopoulos, and D. S. Nikolopoulos. Online Power-Performance Adaptation of Multithreaded Programs using Hardware Event-Based Prediction. In *International Conference on Supercomputing (ICS06)*, Queensland, Australia, 2006.
- [4] J. Ebergen, J. Gainsley, and P. Cunningham. Transistor Sizing: How to Control the Speed and Energy Consumption of a Circuit. 2004.
- [5] W. Feng and C. Hsu. Green Destiny and Its Evolving Parts. In *Innovative Supercomputer Architecture Award, International Supercomputer Conference*, Heidelberg, Germany, June 2004.
- [6] V. Freeh, D. Lowenthal, F. Pan, and N. Kappiah. Using Multiple Energy Gears in MPI Programs on a Power-Scalable Cluster. In *Principles and Practices of Parallel Programming (PPoPP)*, 2005.
- [7] R. Ge, X. Feng, W. Feng, and K. W. Cameron. CPU MISER: A Performance-Directed, Run-Time System for Power-Aware Clusters. In *International Conference on Parallel Processing (ICPP07)*, 2007.
- [8] C. Hsu and W. Feng. A Power-Aware Run-Time System for High-Performance Computing. In *ACM/IEEE Supercomputing 2005 (SC05)*, 2005.
- [9] S. Huang and W. Feng. A Workload-Aware, Eco-Friendly Daemon for Cluster Computing. *Technical Report, Computer Science, Virginia Tech*, 2008.
- [10] AMD Inc. BIOS and Kernel Developer's Guide for AMD Athlon 64 and AMD Opteron Processors. February 2006.
- [11] M. Y. Lim, V. W. Freeh, and D. K. Lowenthal. Adaptive, Transparent Frequency and Voltage Scaling of Communication Phases in MPI Programs. In *ACM/IEEE Supercomputing 2006 (SC06)*, 2006.
- [12] L. McVoy and C. Staelin. Imbench: Portable Tools for Performance Analysis. In *USENIX Annual Technical Conference*, 1996.
- [13] S. Mohapatra, R. Cornea, N. Dutt, A. Nicolau, and N. Venkatasubramanian. Integrated Power Management for Video Streaming to Mobile Handheld Devices. In *ACM Multimedia*, 2003.
- [14] V. Pallipadi and A. Starikovskiy. The On-Demand Governor – Past, Present, and Future. *Ottawa Linux Symposium*, 2006.
- [15] P. Penzes, M. Nustrom, and A. Martin. Transistor Sizing of Energy-Delay-Efficient Circuits. Technical Report. California Institute of Technology, 2002.
- [16] B. Roundtree, D. K. Lowenthal, S. H. Funk, V. W. Freeh, B. R. de Supinski, and M. Schulz. Adaptive, Transparent Frequency and Voltage Scaling of Communication Phases in MPI Programs. In *ACM/IEEE Supercomputing 2007 (SC07)*, 2007.
- [17] C. Scott. Energy Efficiency in the Data Center. 2007.
- [18] T. Tan, A. Raghunathan, and N. Jha. Software Architectural Transformations: A New Approach to Low-Energy Embedded Software. In *Design, Automation and Test in Europe*, 2003.
- [19] New York Times. Hiding in Plain Sight, Google Seeks More Power. 2006.