

Energy-Efficient Real-Time Heterogeneous Server Clusters

TR-05-127 *

Cosmin Rusu, Alexandre Ferreira, Claudio Scordino[†], Aaron Watson,
Rami Melhem and Daniel Mossé
Department of Computer Science, University of Pittsburgh

Abstract

With increasing costs of energy consumption and cooling, power management in server clusters has become an increasingly important design issue. Current clusters for real-time applications are designed to handle peak loads, where all servers are fully utilized. In practice, peak load conditions rarely happen and clusters are most of the time underutilized. This creates the opportunity for using slower frequencies, and thus smaller energy consumption, with little or no impact on the Quality of Service (QoS), for example, performance and timeliness.

In this work we present a cluster-wide QoS-aware technique that dynamically reconfigures the cluster to reduce energy consumption during periods of reduced load. Moreover, we also investigate the effects of local QoS-aware power management using Dynamic Voltage Scaling (DVS). Since most real-world clusters consist of machines of different kind (in terms of both performance and energy consumption) we focus on heterogeneous clusters.

For validation, we describe and evaluate an implementation of the proposed scheme using the Apache Web-server in a small realistic cluster. Our experimental results show that using our scheme it is possible to save up to 45% of the total energy consumed by the servers, maintaining average response time within the specified deadlines and number of dropped requests within the required amount.

1 Introduction

Until recently, performance had been the main concern in server farms, but energy consumption has also become a main concern in such systems. Due to the importance of customer care service in commercial installations and the importance of timely responses for embedded server clusters, current clusters are typically designed to handle peak loads. However, peak load conditions rarely happen in practice, and clusters are most of the time underutilized. In fact, their loads often vary significantly depending on the time of the day or other external factors, therefore the average processor use of such systems may be even less than 50% with respect to their peak capacity [BEK⁺02].

Clusters with high peak power need complex and expensive cooling infrastructures to ensure the proper operation of the servers. With power densities increasing due to increasing performance demands and tighter packing, proper cooling becomes even more challenging: fans driving the cooling system may consume up to 50% of the total system power in some commercial servers [LRR⁺03, MSS⁺04], and manufacturers are facing the problem of building powerful systems without introducing additional techniques such as liquid cooling. Electricity cost is a significant fraction of the operation cost of data centers [BR04]. For example, a Google 10kW rack consumes

*Supported by NFS through the Secure-CITI project (ANI-0325353) and the Power-Autonomous Networks project (ANI-0121658).

[†]Claudio Scordino is a PhD student at the University of Pisa, visiting the University of Pittsburgh.

about 10MWh a month (including cooling), which is at least 10% of the operation cost [BDH03], with this fraction likely to increase in the future.

These issues are even more critical in embedded clusters [XZR⁺05], typically untethered devices, in which peak power has an important impact on the size of the system, while energy consumption determines the device lifetime. Examples include satellite systems or other mobile devices with multiple computing platforms, such as the Mars Rover and robotics platforms.

Power management (PM) mechanisms can be divided into two categories: *cluster-wide* and *local* [BR04]. Cluster-wide mechanisms involve global decisions, such as turning on and off cluster machines, according to the load. Local techniques put unused (or underutilized) resources in low-power states, for example self-refresh, standby and off modes for DRAM chips, Dynamic Voltage Scaling (DVS) and low-power states for the CPU, disk shutdown, etc. A PM mechanism (local or cluster-wide) is *QoS-aware* if it reduces the power consumption while guaranteeing a certain amount of Quality of Service (QoS), such as average response times or percentage of deadlines met.

To the best of our knowledge, this is the first Power Management scheme that is *simultaneously* (a) cluster-wide (i.e., turning on and off machines), (b) **designed** for heterogeneity, (c) QoS-aware **and** power-aware at the local servers (i.e., deadline-aware), (d) measurement-based (instead of theoretical modeling), (e) implementation-oriented, and (f) reconfiguration decisions are made online.

Our scheme is realistic because most clusters have one or more front-ends, are composed of different kind of machines, and need both local and cluster-wide QoS-Aware PM scheme. While the methodology and the algorithms proposed apply to any kind of cluster, we show their use in a web server context. Our measurements show a reduction of energy consumption equal to 17% using only the local PM, 39% using the On/Off scheme, and 45% using both schemes. With respect to delays, the local PM added 0.5ms, while On/Off added about 4ms; in all cases, the average delay was quite small with respect to deadlines.

The remainder of the paper is organized as follows. We first present related work in Section 2. The cluster model is given in Section 3. The cluster-wide PM scheme is explained in Section 4, while the local real-time DVS scheme is presented in Section 5. Both schemes are then evaluated in Section 6. In Section 7 we state our conclusions.

2 Related Work

With energy consumption emerging as a key aspect of cluster computing, much recent research has focused on PM in server farms. A first characterization of power consumption and workload in real-world web servers was made in [BEK⁺02]. DVS was proposed as the main technique to reduce energy consumption in such systems. Software peak power control techniques were investigated in [FRKR05]. However, these studies considered only power consumption of processor and main memory in single-processor systems.

The problem of *cluster configuration* (i.e., turning on and off cluster machines) for homogeneous clusters was first addressed in [PBCH01]. An offline algorithm determines the number of servers needed for a given load. Cluster reconfiguration is then performed online by a process running on a server, using thresholds to prevent too frequent reconfigurations, even though there is no explicit QoS consideration. The authors have extended their work to heterogeneous clusters in [HDC⁺03]. Models have been added for throughput and power consumption estimation. Reconfiguration decisions are made online based on the precomputed information and the predicted load. The authors also proposed to add request types to improve load estimation in [HDC⁺05].

Our work differs from the above previous studies in the following ways: we consider QoS directly; individual servers are both power-aware and QoS-aware; we rely on offline measurements instead of using models; reconfiguration decisions (i.e., number of active servers and load distribution) are not expensive and are performed online; and reconfiguration thresholds are based on the time needed to boot/shutdown a server.

One of the first attempts to combine cluster-wide and local PM techniques [EKR02] proposed five different

policies combining DVS and cluster configuration. However, the theory behind this work relies on (a) homogeneous clusters, and cannot be easily extended to heterogeneous machines; and (b) the often-incorrect assumption that power is a cubic function of the CPU frequency. Another work proposed to use the cluster load (instead of the average CPU frequency) as the criteria for turning on/off machines [XZR⁺05]. However, this study assumed homogeneous clusters as well. To the best of our knowledge, this work is the first attempt to combine cluster-wide and local techniques in the context of heterogeneous clusters.

In real-time computing, dynamic voltage (and frequency) scaling has been explored to reduce energy consumption. DVS schemes typically focus on minimizing CPU energy consumption while meeting a performance requirement [YDS95, HQPS98]. DVS work for aperiodic tasks in single processors includes: offline and online algorithms assuming worst-case execution times [YDS95, SL04], automatic DVS for Linux with distinction between background and interactive jobs [FM02], and use of knowledge about the distribution of job lengths for voltage scaling decisions [LS01, RXMM04]. However, these techniques typically aim at reducing the energy consumed only by the CPU [LS01, SR03, SL04, SB05] and do not take into account other devices (such as memory, power supplies, or disk) that contribute with an important fraction to the total energy consumed by the system. In our model, instead, servers can put their resources in low-power states, and no assumption is made about their local PM schemes.

Most related to our *local scheme* is Sharma et al.'s work on adaptive algorithms for DVS for a QoS-enabled web server [STA⁺03]. Their scheme uses a theoretical utilization bound derived in [AL01] to guarantee the QoS of web requests. However, they take into account only local PM, assuming that a good load balancing algorithm is used at the front-end. In that sense, our works are complementary, since we describe how to achieve such load balancing. Moreover, they only focus on a single server, not in a cluster as we do.

3 Cluster Model

This section introduces the cluster model that we consider (see Figure 1). A front-end machine receives requests from clients and *redirects* them to a set of processing nodes, henceforth referred to as *servers*. The front-end is not a processing node and has three main functions: (a) accepting aperiodic requests from clients, (b) distributing the load to servers, and (c) cluster configuration (i.e., turning servers on/off) to reduce global energy consumption while keeping the overall performance within a prespecified QoS requirement. After receiving a request, the front-end communicates to the client to which server the request must be sent using *HTTP redirection* [CCSY99]. Then, the client sends its request directly to the server.

In our cluster scheme, each request is an aperiodic task (i.e., no assumptions are made about task arrival times) and is assigned a deadline. The specification of the QoS is system-wide and is, in our case, the percentage of deadlines met. The way to achieve the soft-real-time properties will be presented in detail in the next sections.

Each server in the heterogeneous cluster performs the same service (i.e., all servers can process all requests). No restriction is imposed regarding any aspect of their computation: process scheduling, CPU performance, memory speed/bandwidth, disk speed/bandwidth, power consumption, network bandwidth, etc. In addition, servers periodically inform the front-end about their current load, to aid the front-end in load distribution and cluster configuration decisions. After that a request has been processed by a server, the result is returned directly to the client, without the front-end as intermediary.

The aspects related to cluster configuration, PM and load distribution performed by the front-end will be presented in detail in the next section. Local PM is performed independently by each server, without front-end control, and will be presented in Section 5.




Figure 1. Cluster architecture

4 Front-end Power Management

Our proposed front-end follows a very general framework that is applicable to *any* heterogeneous cluster. To achieve this goal, we cannot impose any restriction on server characteristics. However, for ease of presentation, definitions and examples emphasize web server clusters.

4.1 Load Definition and Estimation

The front-end determines the number of active servers to meet the desired level of QoS while minimizing cluster energy consumption. The number of servers is computed (offline or online) as a function of the system load. Thus, defining load correctly is a crucial step. A measure of the load for clusters is the number of requests received per second, measured over some recent interval. Clearly, depending on the kind of service under consideration, other definitions of load may be more appropriate (such as the bandwidth for a file server).

At runtime, the front-end needs to correctly estimate (or observe) the load, in order to make PM decisions and to perform load distribution. The load estimation can be further improved by using some feedback from servers.

As observed in previous work [RXMM04, XZR⁺05, HDC⁺05], load estimation can be greatly improved by considering request types. The type of a request may be conveniently determined only by the header (e.g., the name of the requested file). Notice that the number of different types of request is a design issue. On one hand, different types may not be necessary (if the variability of the time to service a request is low). On the other hand, each request could be of a different type, leading to an improved estimation but also to a higher overhead (to measure all different types of requests and update statistics tables).

In the case of a web server there are two main types of requests, with different computational characteristics: static and dynamic pages. Static pages reside in server's memory and do not require much computation. Dynamic pages, instead, are created on-demand through the use of some external language (e.g., Perl or PHP). For this reason, dynamic pages typically require more computation than static ones.

Consider a generic server, and let A_{static} and $A_{dynamic}$ be the average execution times to serve a static and a dynamic page, respectively, at the maximum CPU speed. For example, for one server in our cluster we measured an average execution time $A_{static} = 438\mu s$ for static pages and $A_{dynamic} = 24.5ms$ for dynamic pages. On average, the time needed by the CPU to serve N_{static} static requests and $N_{dynamic}$ dynamic requests is thus $N_{static}A_{static} + N_{dynamic}A_{dynamic}$ seconds. If the number of requests is observed over a period of $monitor_period$ seconds, then the load of the machine serving the requests is

$$Load = \frac{N_{static}A_{static} + N_{dynamic}A_{dynamic}}{monitor_period} \quad (1)$$

Notice that this definition of load assumes a CPU-bound server. This is normal for most web servers because much of the data are already in memory [BEK⁺02, XHL05]. In fact, on all our machines we have noticed that the bottleneck of the system was the CPU. However, for systems with different bottlenecks (e.g., disk I/O or network bandwidth) another definition of load may be more appropriate. In fact, the definition of load should account for the bottleneck resource.

We define the *maximum load* of a server as the maximum number of requests that it can handle meeting the 95% of deadlines. The front-end never directs more than the maximum load to a server. The cluster load is defined as the sum of the current loads of all active servers. Therefore, the maximum load that the cluster can handle is the sum of the maximum loads of all servers. At runtime, the cluster load (i.e., both variables N_{static} and $N_{dynamic}$) is computed every *monitor_period* seconds.

4.2 Server information

In order to reduce the global power consumption at run-time, we furnish the front-end with information about the power consumption of each server for any different value of its load (i.e., it is a function of voltage, frequency, and switching activity). Servers can reduce their own power consumption in a number of different ways, such as using DVS and low-power states for the CPU, self-refresh modes for memory, stopping disk spinning after some time of idleness, etc. Moreover, each server may use a different OS or a different scheduling policy (such as a standard round robin, or a real-time policy to give higher priority to static pages with respect to dynamic ones). No assumption is made at the front-end about local PM schemes.

Once the local PM scheme, the OS, and the scheduling policy have been decided for a server, the power consumption as function of the load and the maximum load can be determined through simple measurements.

In our experiments, after choosing the local PM scheme (see Section 5), we measured the average power consumption for a load in 5% increments. Then, we interpolated the points to have values in 1% increments. We measured the total power consumed by the whole machines, not only by their CPUs. In our case recording the average power consumption for a given load over a period of 10 minutes was sufficient, for we measured AC power directly, with a simple power meter with 2% accuracy [Sea]. Hence, the whole process required at most few hours for each machine. Clearly, identical machines need not to be measured twice. The curve representing the power consumption of each server of our cluster is shown in Figure 2. The last point on each curve represents the maximum load that meets our QoS specification (i.e., 95% of deadlines met), normalized to the fastest machine in the cluster. The parameters for each machine are reported in Table 2 (on Page 11).




Figure 2. Power consumption vs. load for servers

The information about power consumption of servers can then be used by the front-end at runtime. Since the front-end controls the load for each server, and the power consumption of each server for a given load is known,

the front-end has enough information for PM decisions. Notice that the cooling costs of the room have not been taken into account. However, since these costs are expected to be proportional to the AC power drawn, they are automatically reduced by minimizing cluster power consumption.

We now present all the information (and the corresponding notation) about each server needed at the front-end level. $boot_time_i$ and $shutdown_time_i$ represent the time to boot and to shutdown server i , including the time to start and finish the (webserver) process of the server. max_load_i is the maximum load of server i that can satisfy the 95% QoS requirement. off_power_i is the power consumed when the server is off (some components, such as the Wake-On-Lan interface used to power up the machine, may not be completely off). Finally, $power_vs_load_i$ is an array with $\lceil \frac{max_load_i}{load_increment} \rceil$ entries recording the measured power consumption of server i for each value of the load in $load_increment$ percents (we used 1%). The first entry of the array denotes the idle power (i.e., no load).

4.3 On/Off Policy

This section describes the key idea behind our cluster-wide PM scheme. The front-end, besides distributing the load to servers to minimize global power consumption, determines the cluster configuration by turning on/off some servers. Below is the algorithm used by the front-end to decide which servers must be turned on/off.

The algorithm turns machines on and off in a specific order, which is based on the power efficiency of servers (i.e., the integral of power consumption versus load). In our case, according to the values of Figure 2, the ordering for our cluster is: *Transmeta, Blue, Silver, Green*. In some situations we may need to change the order at runtime, as explained later.

The front-end turns on servers as the cluster load increases. However, since the boot time is not negligible, we need to turn machines on *before* that they are actually needed. For this reason, the front-end maintains a variable, called $max_load_increase$, which specifies the maximum load variation that the cluster is prepared to sustain during $monitor_period$. This is essentially the maximum slope of the load characterization for the cluster.

The on/off policy relies on two tables computed offline. The first table, called $mandatory_servers$, keeps the load at which to turn servers on and is used to determine the lowest number of servers needed at a certain load to satisfy the QoS requirement. For example, consider a cluster with three servers having maximum loads $max_load_0 = 0.5$, $max_load_1 = 1.5$ and $max_load_2 = 1.0$, respectively. Suppose that $monitor_period$ is 5 seconds, $max_load_increase$ is equal to 0.05, and the boot time is 10 seconds for every machine. Ideally, we need only one server when the cluster load is less than 0.5, two servers when load is between 0.5 and 2, and all servers when load is higher than 2. However, if we account for the time to boot a new machine and we suppose that the cluster load is checked periodically every $monitor_period$ seconds, the table becomes $mandatory_servers = \{0, 0.35, 1.85\}$. Thus, the first server is always on, whereas second and third servers are turned on when the cluster load reaches 0.35 and 1.85, respectively. In fact, if we consider the boot time of a new server, we have to account for a load reduction equal to $boot_time \frac{max_load_increase}{monitor_period}$. Moreover, if we suppose that the load is checked periodically every $monitor_period$ seconds, we have to introduce an additional interval of time to account for the error when measuring the current load. In general, server i is turned on when the cluster load reaches

$$\sum_{j=0}^{i-1} max_load_j - (boot_time_i + monitor_period) \frac{max_load_increase}{monitor_period}$$

The second table, called $power_servers$, precomputes the number of servers needed to minimize the power consumption for a given load. Unlike the previous table, this table is computed considering the power consumption of servers, and is used to distribute the current load among active servers. The table is built as follows. For a given value of N , we compute the power consumption of the cluster for any possible distribution of the load to N servers

— i.e., for any possible configuration having N machines on. We start considering a load equal to zero, and we increase its value in *load_increment* increments. For any increment of the load, we evaluate which server can handle it in order to minimize the overall energy consumption.

To determine the load at which N servers become more power efficient than $N - 1$, we follow this procedure considering both cases of $N - 1$ and N machines, respectively. The load at which N servers consume less power than $N - 1$ servers is the value after which the N^{th} server is turned on.

The complexity of computing the two tables is $O(N)$ (where N is the number of servers) for *mandatory_servers* and $O(NM)$ for *power_servers*, where $M = \sum_{i=1}^N \lceil \frac{max_load_i}{load_increment} \rceil$. In our cluster, the time to compute these two tables was less than 1msec, which was negligible compared to *monitor_period* (that is in the range of seconds). Thus, this computation can also be performed online. For example, a new ordering of the servers and an online recalculation of the tables become necessary when a server crashes.

A high-level view of the front-end on/off policy is presented in Figure 3. Every *monitor_period* seconds the load is estimated according to Equation 1, then the request counters are reset. The number of mandatory servers $N_{mandatory}$ is determined by a lookup in the *mandatory_servers* table. If $N_{mandatory}$ is higher than the current number of active servers $N_{current}$, all needed servers are immediately turned on.

Each server can be in one of the following states: Off, Boot, On, or Shutdown. After receiving the “boot” command (such as a Wake-On-Lan packet), the server i moves from the Off to the Boot state. It stays in this state for *boot_time_i* seconds (i.e., until it starts the server process), then informs the front-end that it is available for processing, moving to the On state. When server i is shutdown, it stays in the Shutdown state for *shutdown_time_i* seconds, after that the front-end changes its state to Off.

The variable *Cmd* in Figure 3 can have 3 different values: None, Boot or Shutdown. This variable allows to describe the use of thresholds when turning on/off servers. If no server is in transition (i.e., all servers are in the On or Off states) a server may be turned on or off, as decided after a lookup in the *power_servers* table. To be conservative, only one server at a time is turned on or off. Server i is turned off if the system is in state *Cmd* = *shutdown* for at least *time_shutdown_i* + *time_boot_i* consecutive seconds, which is the rent-to-own threshold (see Step 3, Figure 3). Similarly, server i is turned on if *Cmd* = *boot* for *time_boot_i* consecutive seconds (see Step 2, Figure 3). Notice that these thresholds do not apply to the mandatory servers, which are started immediately. The running time of the online part of the algorithm (every *monitor_period* seconds) is negligible because it is in the microsecond range; the complexity is $O(N)$, but can be improved to $O(1)$ by increasing the table size from N to M (that is, storing all entries in an array).

For convex and linear power functions, tables *mandatory_servers* and *power_servers* contain the optimal transition points (in the discrete space; for continuous space, see [SB05]). In practice, however, power functions may have concave regions. This means that a server with an abrupt power increase at some load x may not be allocated more than x load, even though the power may become flat above $x + \epsilon$, making it a good target for load allocation. A simple fix to the problem is to consider the average power consumption over a larger interval, rather than the exact value at each load. This effectively results in smoothing the power functions. In our case, although the measured power functions have concave regions, we have found that no smoothing was necessary.

4.4 Request Distribution Policy

The front-end distributes the incoming requests to a subset of the current servers that are in the On state. *load_allocation* is a table containing the estimated load allocated to each server and is computed with the same procedure used to determine the *power_servers* table, in $O(MN)$ time. The load allocation is computed every *monitor_period* seconds, after the on/off decisions.

Another table, called *load_accumulated*, stores the accumulated load of each server, and is reset after comput-

```

1 Every monitor_period seconds
  1.1 Compute the load according to Equation 1
  1.2 Reset the counters:
       $N_{static} = 0$   $N_{dynamic} = 0$ 
  1.3 Compute the minimum number of servers
      that can handle the load:
       $N_{mandatory} = \text{mandatory\_servers}(\text{Load})$ 
  1.4 if ( $N_{mandatory} > N_{current}$ )
      turn on the servers
      set  $N_{current} = N_{mandatory}$ 
      return
  1.5 Compute the number of servers needed to re-
      duce the energy consumption:
       $N_{power} = \text{power\_servers}(\text{Load})$ 
  1.6 if ( $N_{power} > N_{current}$ ) and ( $Cmd \neq \text{Boot}$ )
      Set  $Cmd = \text{Boot}$ 
      Find the next server  $i$  to boot
      Set  $N_{current} = N_{current} + 1$ 
      return
  1.7 if ( $N_{power} < N_{current}$ ) and ( $Cmd \neq \text{Shut-}$ 
       $\text{down}$ )
      Set  $Cmd = \text{Shutdown}$ 
      Find the next server  $i$  to shutdown
      Set  $N_{current} = N_{current} - 1$ 
      return
2 If  $Cmd = \text{Boot}$  for a period of time equal to
   $time\_boot_i$ 
  2.1 Turn on server  $i$ 
  2.2 Set  $Cmd = \text{None}$ 
  2.3 return
3 If  $Cmd = \text{Shutdown}$  for a period of time equal to
   $time\_boot_i + time\_shutdown_i$ 
  3.1 Turn off server  $i$ 
  3.2 Set  $Cmd = \text{None}$ 
  3.3 return

```

Figure 3. On/Off policy

ing *load_allocation*. The server i with the minimum weight

$$w_i = \frac{\text{load_accumulated}_i}{\text{load_allocation}_i} \quad (2)$$

gets the next request. Notice that w_i can be higher than 1 when the load is underestimated. The server that receives the request updates its accumulated load (and thus increases its weight), by adding $A_{static}/\text{monitor_period}$ or $A_{dynamic}/\text{monitor_period}$, depending on the request type. The complexity to find the server with minimum weight is $O(N)$ with a straightforward implementation, but can be improved to $O(\log N)$ using a tree.

4.5 Implementation Issues

We implemented our PM scheme in the *Apache* 1.3.33 Web server [Apa]. We created an Apache module, called *mod_on_off*, which makes on/off decisions. Moreover, we extended an existing module, *mod_backhand* [bac], to

support our distribution policy.

mod_backhand is a module responsible for load distribution in Apache clusters. It allows servers to exchange information about their current usage of resources. It also provides a set of *candidacy* functions to forward requests from an overloaded server to other less utilized servers. Examples of such functions are *byLoad*, which selects as candidate the least loaded server, and *byCost*, which considers a cost for each request.

We added a new candidacy function, called *byEnergy*, to implement our request distribution policy. Notice that only front-end machines use this function. In addition, servers provide some feedback about their current real-time utilization (value of the variable U as explained in Section 5) to front-ends. We used this feedback to prevent the overloading of the servers. In particular, the server with the minimum w_i is selected, providing that it is not overloaded.

The *mod_on_off* module communicates with *mod_backhand* through shared memory. On initialization, *mod_on_off* acquires server information and computes both *mandatory_servers* and *power_servers* tables. *mod_on_off* executes periodically every *monitor_period* seconds. On each invocation it performs the following tasks: (a) computes the current load based on the counters N_{static} and $N_{dynamic}$ (that are incremented in the “Apache *post read* request” phase), (b) looks up in the table to determine the number of servers needed for the next period, (c) computes the *load_allocation* table for the active servers (not shown in Figure 3), (d) turns on (by sending Wake-On-Lan packets) and off (by invoking special CGI scripts) servers, and finally (e) resets the counters N_{static} , $N_{dynamic}$ and *load_accumulated*. In addition, it displays at runtime the estimated power and energy consumption of each server, based on the *power_vs_load* and *load_accumulated* tables.

5 Server Power Management

In addition to front-end directed cluster reconfigurations (i.e., turning on/off machines), the servers perform their own local PM to reduce power consumption of unutilized or underutilized resources. We present an example of a QoS-aware DVS scheme and we discuss an implementation using the Apache Webserver [Apa].

5.1 Local DVS Policy

We rely on a local real-time scheme, where each request is an aperiodic task and is assigned a deadline. Each request *type* has a deadline [RXMM04, XZR⁺05, HDC⁺05] to allow for more accurate load estimation.

We consider a *soft* real-time system, in which the schedule is not generated by a real-time scheduler and C_i is the average execution time (i.e., A_{static} or $A_{dynamic}$), not the worst-case. Let D_i be the time remaining to the deadline, then the real-time utilization of a server is defined as $U = \sum_i \frac{C_i}{D_i}$.

If the CPU is the bottleneck of the system (as in our case), the CPU frequency to handle this rate of requests is $U f_{max}$, where f_{max} is the highest possible frequency of the CPU. Each server periodically computes its utilization U and sets the CPU frequency to the closest value higher than $U f_{max}$.

We note two peculiarities of this system. First, since we are considering soft real-time systems, we could set the frequency to the closest value *lower* than $U f_{max}$ at the detriment of response time. Second, DVS architectures may have inefficient operating frequencies [SR03], which are when there are higher frequencies that consume less energy. A simple online tool for inefficient frequency elimination has been provided in [PAR]. Removal of inefficient operating frequencies is the first step in any DVS scheme. This was not necessary in our servers, because surprisingly all frequencies were efficient, although we had a different experience with other systems we tested [XZR⁺05].

5.2 Implementation Issues

We implemented an Apache module, called *mod_cpufreq*, responsible for CPU speed settings at the user level. On Athlon machines, the CPU speed was changed by writing to the `/sys/` file system, using the *CPUfreq* inter-

face [cpu]. On the Transmeta machine the speed was changed by writing to a model-specific register (MSR). Since the register cannot be written from user-level we added two system calls for setting and reading its value [Gle]. After detecting the available frequencies, our module creates an Apache process that periodically sets the CPU frequency according to the current value of U . We chose as period $10ms$ to match *any* default Linux kernel; the measured overhead for changing voltage/frequency in the Athlon64 machines is approximately $50\mu s$.

To compute U , the module needs to know the type (i.e., static or dynamic) and the arrival time of each request. At every request arrival (called “Apache *post-read* request” phase), the arrival time and the deadline are recorded with μs accuracy and stored in a hash table in shared memory. The request type is determined from the name of the requested file. Thus, a single queue traversal is necessary to compute U . In fact, the current value of U depends on all queued requests, therefore the complexity is $O(R)$ where R is the number of requests queued; the overhead is negligible. Requests are removed from the queue after being served (called “Apache *logging* request” phase).

A problem we encountered during the implementation was that our scheme worked very well except for fast machines serving a large amount of small static pages. In this case, those machines were not increasing their speed, resulting in a large number of dropped requests. A further investigation revealed that the value of U was close to zero. We did not see this phenomenon on slower machines (such as Transmeta) nor using bigger pages. The problem was that the requests were served too fast (in approximately $150\mu s$). Such short requests were queued, served, and removed from the queue before other requests were added to the queue and served. Thus, at any time only a few requests (usually just one) was in the queue, and when *mod_cpufreq* recomputed the utilization, it resulted in an underestimation of U . In other words, even though the requests were received, queued, and serviced at the OS-level, Apache was not able to see them because it is a user-level server and it has not information about requests stored at the OS level. We called this problem the “*short request overload problem*” phenomenon.

A simple fix was to compute the utilization also over a recent interval of time *interval_size* (we used $200ms$):

$$U_{recent} = \frac{(N_{static}A_{static} + N_{dynamic}A_{dynamic})}{interval_size}$$

Our module sets the CPU speed to $\max(U, U_{recent})f_{max}$ except for when $U_{recent} > U$ and $U_{recent} > 80\%$; in that case, the module selects the next speed above the one that would be selected according to the value of U_{recent} . In this situation, Sharma et al.’s work with a kernel webserver (kHTTPd [STA⁺03]) aware of small requests at the OS-level has a nice synergy with our approach and could be used in lieu of our scheme. Exploring the composition of our cluster configuration and Sharma’s (or other similar DVS) work is left for future work. The problem with including such work in our scheme is exactly the reason why the authors discontinued the development of kHTTPd: it difficulty of maintaining, developing and debugging a kernel-level server.

6 Evaluation

To evaluate our QoS-aware PM scheme we used a small cluster composed by one front-end and 4 different servers. Every machine ran *Gentoo Linux 2.6* as operating system and *Apache 1.3.33* servers. The parameters of the machines are shown in Tables 1 and 2.

The cluster has been tested using 2 clients connected to the cluster with a GbE interface and Gbps switch; the clients generate up to 3,186 requests per second, which corresponds to a total maximum cluster load equal to 2.95 (all loads were normalized to that of *Silver* machine). A total cluster load of 0.05 (or 5%) corresponds on average to 54 requests/second. Considering request types, however, greatly improves the prediction, as 54 requests/second may correspond to a load ranging from 0.02 (if $N_{dynamic} = 0$) to 1.32 (if $N_{static} = 0$). We assigned deadlines of $50ms$ and $200ms$ for requests of static and dynamic pages.

We set *max_load_increase* = 0.005, therefore we had *mandatory_servers* = {0.000, 0.062, 1.012, 2.012} and *power_servers* = {0.000, 0.100, 1.050, 2.040}.

Transmeta

Frequency (MHz)	333	400	533	667	733
Idle (W)	8	8.5	8.5	9	9
Busy (W)	9	9.5	10.5	12	12.5

Blue

Frequency (MHz)	800	1800	2000	2200
Idle (W)	68	73	76	80.5
Busy (W)	74.5	93.5	105.5	120.5

Silver

Frequency (MHz)	1000	1800	2000	2200	2400
Idle (W)	70	74.5	78.5	83.5	89.5
Busy (W)	80.5	92.5	103.5	119.5	140.5

Green

Frequency (MHz)	1000	1800	2000
Idle (W)	68	79	87
Busy (W)	77	108	131

Table 1. Idle/busy power consumption (in Watts) for each server at each frequency

Machine name	Processor model	RAM memory size	Cache size	Wake-On-Lan support	Boot time (sec)	Shutdown time (sec)	Off power (W)	Max load
Transmeta	Transmeta Crusoe TM5800	256 MB	512 KB		100	60	1	0.10
Blue	AMD Athlon 64 Mobile 3400+	1GB	1 MB	✓	33	11	8	0.95
Silver	AMD Athlon 64 3400+	1GB	512 KB	✓	33	12	8	1.00
Green	AMD Athlon 64 3000+	1GB	512 KB	✓	33	11	8	0.90
Front-end	AMD Athlon 64 Mobile 3400+	1GB	1 MB	Not applicable				

Table 2. Parameters of the machines of the cluster

6.1 DVS policy

As first experiment, we evaluated the effectiveness of our local DVS scheme. We compared our *mod_cpufreq* module with the default PM in Linux (i.e., HALT instruction when idle) and with Sharma’s DVS scheme for QoS-aware web servers proposed in [STA⁺03]. This scheme adjusts the speed of the processor to the minimum speed that maintains a quantity called *synthetic utilization* below the theoretical utilization bound ($U_{bound} = 58.6\%$) that ensures that all deadlines are met [AL01].

The measured power consumption of each scheme on the *Blue* machine is shown as function of the load in Figure 4. The graph shows that our scheme outperforms the other schemes, especially for the mid-range load values. Higher savings are obtained on machines with a more convex power function (the power function of the *Blue* machine is rather linear — see Figure 2). In fact, for a rate of 300 requests/sec (approximately 28% load) the average processor frequency is 1.25GHz using our scheme and 1.5GHz using Sharma’s scheme, but the amount of energy saved is only 3%. Importantly, we observed that both schemes maintained the QoS level above 99% even at the highest load.

6.2 Overall scheme

To evaluate the overall scheme, we performed many experiments with and without the cluster-wide PM scheme (On/Off scheme), and with and without the local PM scheme (DVS scheme). For each load value, we measured the power consumption of the entire machine (not only CPU) for each scheme independently (see Figure 5).

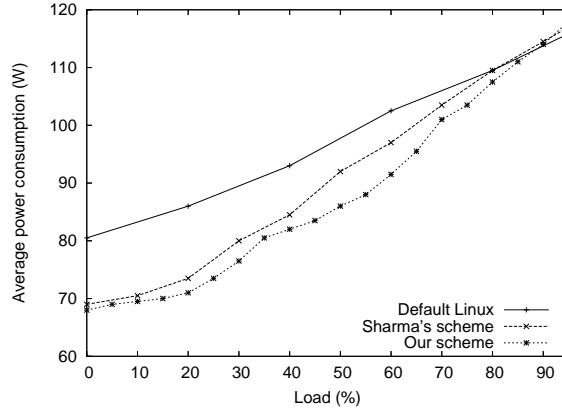


Figure 4. Comparison of DVS policies

The On/Off policy allows a striking reduction of the energy consumption for low values of the load, because (obviously!) it allows to turn off unutilized servers. In Figure 5 we can see that when load = 0, the cluster consumption is around 40W because each Athlon server consumes 8W when in the Off state, and the Transmeta also consumes 8W when in the On state. The DVS technique, instead, has its biggest impact whenever a new server is turned on, since not all active servers are fully utilized. However, its importance decreases as the utilization of the active servers increases. For high values of the load (in our case, at 70% or higher) all servers are on, therefore the On/Off technique does not allow to reduce energy consumption. In those situations, however, there is still room for the DVS technique, that becomes more important than the On/Off technique.

The energy consumption of all servers without any power management scheme was $1.32KWh$. On average, we measured a reduction of energy consumption equal to 17% using DVS, 39% using On/Off, and 45% using both schemes.

It is worth noting that the front-end estimation of the total energy consumed when using DVS was extremely accurate: the difference from the actual values was less than 1%. For example, when using the on-off scheme, the measured value was $.72KWh$, while the front-end estimated value was $.725KWh$ (the resolution of our power/energy meter [Sea] is $.01KWh$).

To measure the impact of cluster-wide and local PM schemes in the loss of QoS, we ran many four-hour experiments with workloads derived from actual webserver traces, and generated with the same shape of statistics taken from our *cs.pitt.edu* domain (see Table 3). The average delay without any PM scheme was $8.29ms$; a small response time is due to all machines being on at all times, and running at maximum frequency. Adding DVS (local PM) had a very small impact on the delay, with the average delay measured at $8.77ms$. However, with On/Off scheme, we measured an average delay equal to $12.29ms$ without DVS and $12.83ms$ with DVS. In both cases, the average delay was not higher than 50% of the no-PM delay and was quite small with respect to deadlines.

7 Conclusions and Future Work

We have presented a new QoS-aware power management scheme that combines cluster-wide, On/Off scheme and local power management (DVS) techniques in the context of heterogeneous clusters. We have also described and evaluated an implementation of the proposed scheme using the *Apache* Webserver in a small realistic cluster.

Our experimental results show that: (a) our load estimation is very accurate; (b) the On/Off policy allows a striking reduction of the power consumption; (c) the DVS technique is very important whenever a new server is turned on or, as shown before, when all servers are on; (d) as expected, For high values of the load the On/Off technique does not help to reduce energy consumption but there is still room for the DVS technique.

Using both techniques we saved up to 45% of the total energy consumed by the servers with a limited loss in

Request type	%	Request type	%
4 ms (CGI)	0.10	6-7 KB (html)	2.84
7 ms (CGI)	0.71	7-8 KB (html)	1.58
23 ms (CGI)	0.98	8-9 KB (html)	1.80
40 ms (CGI)	0.23	9-10 KB (html)	1.87
200 ms (CGI)	0.06	10-20 KB (html)	10.74
0-1 KB (html)	37.78	20-30 KB (html)	3.62
1-2 KB (html)	8.86	30-40 KB (html)	1.17
2-3 KB (html)	6.56	40-50 KB (html)	0.67
3-4 KB (html)	4.58	50-60 KB (html)	0.80
4-5 KB (html)	4.94	60-70 KB (html)	1.46
5-6 KB (html)	3.38	above 70 KB (html)	5.27

Table 3. Web server statistics: percentage of accesses, approximate size (for static pages), and running time (for dynamic pages)




Figure 5. Evaluation of cluster-wide and local techniques

terms of QoS. In the worst case, the average delay was increased by at most 50%, and was still very small when compared to the deadlines.

As immediate future work we plan to investigate the use of both suspend-to-disk and suspend-to-RAM techniques to reduce the time to boot and shutdown a server in the On/Off technique. We also plan an integration of our cluster PM schemes with other grid-like or cluster (e.g., Condor) load balancing schemes.

References

- [AL01] Tarek F. Abdelzaher and Chenyang Lu. Schedulability Analysis and Utilization Bounds for Highly Scalable Real-Time Services. In *Proceedings of the 7th IEEE Real-Time Technology and Applications Symposium (RTAS'01)*, Taiwan, June 2001.
- [Apa] Apache. HTTP Server Project. <http://httpd.apache.org/>.
- [bac] The Backhand Project. <http://www.backhand.org/>.
- [BDH03] Luiz André Barroso, Jeffrey Dean, and Urs Holzle. Web Search for a Planet: The Google Cluster Architecture. *IEEE Micro*, 23(2):22–28, 2003.
- [BEK⁺02] Pat Bohrer, Elmootazbellah N. Elnozahy, Tom Keller, Michael Kistler, Charles Lefurgy, Chandler McDowell, and Ram Rajamony. *The case for power management in web servers*. Kluwer Academic Publishers, 2002.

- [BR04] Riccardo Bianchini and Ram Rajamony. Power and Energy Management for Server Systems. *Computer*, 37(11):68–74, November 2004.
- [CCSY99] Valeria Cardellini, Michele Colajanni, and Philip S. Yu. Redirection Algorithms for Load Sharing in Distributed Web-server Systems. In *19th IEEE International Conference on Distributed Computing Systems (ICDCS'99)*, pages 528–535, Austin, TX, June 1999.
- [cpu] Linux kernel CPUfreq subsystem. <http://www.kernel.org/pub/linux/utils/kernel/cpufreq/cpufreq.html>.
- [EKR02] E.N. Elnozahy, Michael Kistler, and Ramakrishnan Rajamony. Energy-Efficient Server Clusters. In *Workshop on Power-Aware Computer Systems (PACS'02)*, 2002.
- [FM02] K. Flautner and T. Mudge. Vertigo: Automatic Performance-Setting for Linux. In *Proceeding of the 5th Symposium on Operating Systems Design and Implementation (OSDI'02)*, December 2002.
- [FRKR05] Wes Felter, Karthick Rajamani, Tom Keller, and Cosmin Rusu. A performance-conserving approach for reducing peak power consumption in server systems. In *International Conference on Supercomputing (ICS'05)*, pages 293–302, Cambridge, Massachusetts, June 2005.
- [Gle] Sunny Gleason. Power Aware Operating Systems: Task Specific CPU Throttling. <http://www.cs.pitt.edu/PARTS/implementation/>.
- [HDC⁺03] Taliver Heath, Bruno Diniz, Enrique V. Carrera, Wagner Meira Jr., and Ricardo Bianchini. Self-Configuring Heterogeneous Server Clusters. In *Workshop on Compilers and Operating Systems for Low Power (COLP'03)*, September 2003.
- [HDC⁺05] Taliver Heath, Bruno Diniz, Enrique V. Carrera, Wagner Meira Jr., and Ricardo Bianchini. Energy Conservation in Heterogeneous Server Clusters. In *Proceedings of the 10th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'05)*, June 2005.
- [HQPS98] I. Hong, G. Qu, M. Potkonjak, and M. Srivastava. Synthesis Techniques for Low-Power Hard Real-Time Systems on Variable Voltage Processors. In *Proceedings of the 19th IEEE Real-Time systems Symposium (RTSS'98)*, Madrid, Spain, December 1998.
- [LRR⁺03] Charles Lefurgy, Karthick Rajamani, Freeman Rawson, Wes Felter, Michael Kistler, and Tom W. Keller. Energy management for commercial servers. *IEEE Computer*, 36(12):39–48, December 2003.
- [LS01] Jacob Lorch and Alan Smith. Improving Dynamic Voltage Scaling Algorithms with PACE. In *ACM SIGMETRICS*, June 2001.
- [MSS⁺04] Justin Moore, Ratnesh Sharma, Rocky Shih, Jeff Chase, Chandrakant Patel, and Parthasarathy Ranganathan. Going Beyond CPUs: The Potential of Temperature-Aware Data Center Architectures. In *First Workshop on Temperature-Aware Computer Systems*, Munich, Germany, June 2004.
- [PAR] PARTS. Power Efficiency Test. <http://www.cs.pitt.edu/PARTS/demos/efficient>.
- [PBCH01] Eduardo Pinheiro, Ricardo Bianchini, Enrique V. Carrera, and Taliver Heath. Load Balancing and Unbalancing for Power and Performance in Cluster-Based Systems. In *Workshop on Compilers and Operating Systems for Low Power (COLP'01)*, Sept 2001.

- [RXMM04] Cosmin Rusu, Ruibin Xu, Rami Melhem, and Daniel Mossé. Energy-Efficient Policies for Request-Driven Soft Real-Time Systems. In *Euromicro Conference on Real-Time Systems (ECRTS'04)*, Catania, Italy, July 2004.
- [SB05] Claudio Scordino and Enrico Bini. Optimal Speed Assignment for Probabilistic Execution Times. In *Second Workshop on Power-Aware Real-Time Computing (PARC'05)*, Jersey City, NJ, September 2005.
- [Sea] Seasonic. Power Angel. <http://www.seasonicusa.com/products.php?lineId=8>.
- [SL04] Claudio Scordino and Giuseppe Lipari. Using resource reservation techniques for power-aware scheduling. In *Proceedings of the 4th ACM International Conference on Embedded Software*, pages 16–25, Pisa, Italy, September 2004.
- [SR03] S. Saewong and R. Rajkumar. Practical Voltage-Scaling for Fixed-Priority RT-Systems. In *Proceedings of the 9th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS'03)*, May 2003.
- [STA⁺03] Vivek Sharma, Arun Thomas, Tarek Abdelzaher, Kevin Skadron, and Zhijian Liu. Power-aware QoS Management in Web Servers. In *Proceedings of the 24th IEEE Real-Time Systems Symposium (RTSS'03)*, Cancun, Mexico, December 2003.
- [XHL05] Ming Xiong, Song Han, and Kam-Yiu Lam. A Deferrable Scheduling Algorithm for Real-Time Transactions Maintaining Data Freshness. In *IEEE Real-Time System Symposium*, Miami, Florida, December 2005.
- [XZR⁺05] Ruibin Xu, Dakai Zhu, Cosmin Rusu, Rami Melhem, and Daniel Mossé. Energy-Efficient Policies for Embedded Clusters. In *ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'05)*, June 2005.
- [YDS95] F. Yao, A. Demers, and S. Shankar. A Scheduling Model for Reduced CPU Energy. In *IEEE Annual Foundations of Computer Science*, pages 374–382, 1995.