

Capacity Planning and Power Management to Exploit Sustainable Energy

Daniel Gmach, Jerry Rolia, Cullen Bash, Yuan Chen, Tom Christian, Amip Shah, Ratnesh Sharma, Zhikui Wang

HP Labs

Palo Alto, CA, USA

e-mail: {firstname.lastname}@hp.com

Abstract—This paper describes an approach for designing a power management plan that matches the supply of power with the demand for power in data centers. Power may come from the grid, from local renewable sources, and possibly from energy storage subsystems. The supply of renewable power is often time-varying in a manner that depends on the source that provides the power, the location of power generators, and the weather conditions. The demand for power is mainly determined by the time-varying workloads hosted in the data center and the power management policies implemented by the data center. A case study demonstrates how our approach can be used to design a plan for realistic and complex data center workloads. The study considers a data center’s deployment in two geographic locations with different supplies of power. Our approach offers greater precision than other planning methods that do not take into account time-varying power supply and demand and data center power management policies.

Keywords—component; Capacity Planning, Energy Supply Management, Energy Demand Management, Power Capping

I. INTRODUCTION

Significant research is underway to develop technologies that improve the energy efficiency of data centers and reduce their dependence on the power grid. On the demand side, virtualization technology is being used to consolidate workloads and improve IT utilization [1]; cutting-edge cooling technologies such as dynamic smart cooling [2] and air-side economizers further help improve data center energy efficiency. On the supply side, distributed generation and renewable energy sources are increasingly being deployed. However, the joint behavior of these technologies in an integrated supply–demand context is hard to predict.

The goal of this work is to support the design of a power management plan that matches the supply of power with the demand for power in data centers. We define a *power management plan* as a choice for the peak grid power, a mix of renewable energy sources, energy storage, and data center server power management policies. We assume the data center also has a backup plan that may require the use of more grid power or diesel generators when necessary. A sensitivity analysis using our approach can determine how often such a backup plan is expected to be employed. Although this work focuses on design time, the resulting plan and its policies can be used to optimize the real time management of the data center.

Peak grid power usage is often a concern for data centers as it can affect the power infrastructure of grid power providers. Contracts are often heavily influenced by peak usage [3] and can impose penalties for exceeding an agreed upon peak.

On the supply side, power may come from a primary power source such as the power grid, from local renewable sources, and from energy storage subsystems. The supply of renewable power is often time-varying in a manner that depends on the source that provides the power, the location of power generators, and the local weather conditions.

On the demand side, data center power consumption is mainly determined by the time-varying workloads hosted in the data center and its power management policies. We assume that the data center has pools of servers that execute consolidated workloads, and that these servers support power management policies such as capping the amount of power used by the servers. Power may be capped if the demand for power from the grid exceeds the choice for available peak grid power. We consider two types of power capping policies in this work: server power capping and pool power capping. *Server power capping* exploits dynamic processor frequency scaling to adjust power used by servers. *Pool power capping* dynamically varies the number of running servers based on the availability of power and the demand of the workloads. The capping methods are complementary.

We describe a trace-driven, simulation based capacity planning tool that is able to simulate power management activities in data centers and estimate the impact of power management plans on application performance and power usage. Traces provide the detailed information needed to: effectively consolidate workloads onto servers; report on metrics that express how often demands are satisfied; and, to estimate the impact of consolidation on workload performance and power consumption. A chosen plan can then be used during the operation of the data center to achieve the desired behavior.

A case study involving three months of data for 138 SAP applications in a real data center is used to evaluate the effectiveness of a set of power management plans. Our approach gives greater precision than other methods that do not take into account the time-varying resource usage of the data center, the location specific weather data for renewable power sources, the time-varying supply of power, the power management policies, and the impact of energy storage technologies.

The paper is organized as follows. Section II gives an overview of our overall approach. Section III documents the implementation of server and pool power management and data center energy storage in the simulator. Section IV provides characterizations of workload demand variability and renewable power supply variability. The case study is presented in Section V. Related work is described in Section VI. Section VII offers a summary and the concluding remarks along with a description of future research.

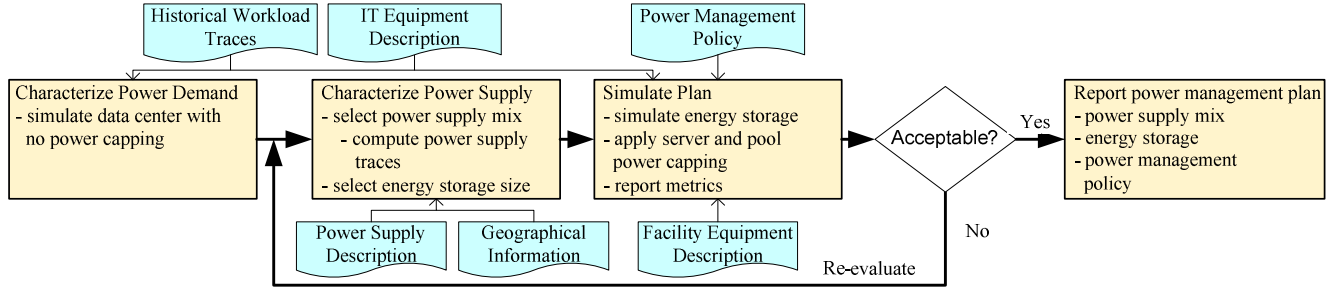


Figure 1: Capacity planning to exploit sustainable energy

II. APPROACH

Figure 1 shows our approach for design and evaluation of capacity planning and power management processes. The first step uses historical workload demand traces and IT equipment descriptions to compute time-varying IT power demands for given data center workload on a particular IT infrastructure. The time-varying IT power demands are used to characterize the total data center power demand assuming no constraints on power supply.

Next, the power supply infrastructure and its location are selected based on the estimated power demand. This allows us to use historical information to produce representative traces of power for each of the power supplies. In general, we want to consider many traces so that the impact of representative set of behaviors can be observed for a sensitivity analysis.

In the third step, we simulate data center operation with both power supply and demand management. Various plans can be evaluated using the tool and compared with performance criteria using standardized metrics. Acceptable plans are finally reported to data center designers and operators to support the design of the data center and the implementation of power management policies.

The following sections describe each step in more detail.

A. Characterize Power Demand

We estimate each server’s power consumption within the trace driven approach as follows. As we simulate the workload placement and server utilization over time, we employ the following simple linear power model:

$$P_{server} = P_{idle} + u * (P_{full} - P_{idle}),$$

where P_{idle} is the idle power of the server and P_{full} is the power consumption of the server when it’s fully utilized. The term u represents the CPU utilization of the server. This simple model has been shown to be a good estimator for power usage [3]. Other formulas could be used as well.

The total IT power demand is the sum of the power per server. Finally, the power demand for the whole data center is calculated through a Power Usage Effectiveness (PUE) value [6], which is the ratio of the total power demand over the IT power demand of the data center. We do not model power for networking and storage in the current version but we plan to incorporate that in future work.

B. Characterize Power Supply

For a selected mix of power sources, we use geographic information such as historical weather data [7] to compute

power supply traces. In our case study, we use traces of measurements made every 5 minutes for sunshine (global irradiance) and wind data (average wind speed) for two different locations in the US: New Mexico and Colorado. The traces cover six weeks, three weeks of June in 1998 and three weeks of June in 1999.

The global irradiance and average wind speed impact power supply for sources, such as PV panels or wind turbines. To determine how much power they provide, we consider the nature of the power source (e.g., efficiency, size, operational profile) and time-varying geographic data.

C. Simulate Power Management Plan and Report

To evaluate which workloads can be consolidated to which servers, we employ a trace-based approach [4][5][10] that assesses permutations and combinations of workloads in order to determine a near-optimal workload placement that provides specific qualities of service for applications. After consolidation, the simulation based capacity planning tool walks forward over additional traces of workload demands to simulate workload placement policy. A migration controller [8] recognizes over and under-utilized resources, causes migrations to improve application QoS, and shuts down as well as starts up servers as needed. We have explored the impact of various data center management policies using the simulation environment [5]. The work presented in this paper enhances the simulation environment to implement server power capping and pool power capping policies and to take into account the time-varying supply of power.

The simulator reports metrics including the *CPU violation penalty* for the workloads and *power deficit*—the difference between required and available power. Though the impact of aggressive consolidation and power deficits upon application performance cannot be modeled precisely using the information available to us, we rely upon the *CPU violation penalty* metric presented in [5] to estimate the sustained impact of resource deficits:

$$pen = I^2 \max_{i=1}^I (w_{pen,i}), \text{ with } w_{pen,i} = 1 - \frac{1 - u_{a,i}^k}{1 - u_{d,i}^k}$$

$u_{a,i}$ and $u_{d,i}$ are the actual and desired CPU utilizations for interval i and k denotes the number of CPUs. The metric *pen* is based on the number of successive intervals I where a workload’s demands are not fully satisfied, i.e., $u_{a,i} > u_{d,i}$, and the expected impact on performance $w_{pen,i}$ as predicted using a queuing theoretic response time estimate for a unit of demand.

Many representative power supply traces can be evaluated during a simulation to consider a wide range of power supply conditions and their resulting impact. The process is repeated until a suitable mix of power sources is found. Details of the simulator, including the server and pool power capping policies and simulation of energy storage, are described further in Section III.

III. SIMULATING POWER MANAGEMENT AND STORAGE

This section describes extensions to the simulator that implement server and pool power capping policies and simulate the use of local energy storage.

A. Server Power Capping

Server power capping controls the power demand of each server in a server pool. When a server's power consumption exceeds a given threshold, the power consumption can be reduced through dynamic frequency scaling that scales down the CPU frequency and reduces both power and the capacity available to the applications. We implemented an algorithm that simulates the behavior of server power capping for the server pool. A central server power coordinator considers workload and power demands on all servers. If the aggregate power demand exceeds a threshold, then the utilization is scaled back equally on all servers until the threshold for aggregate power usage is satisfied. This approach is "ideal" in that it assumes perfect coordination among the servers—no server constrains its workloads more than necessary. We chose to implement an ideal approach to better compare the effectiveness of server and pool capping mechanisms.

We have one additional parameter for server capping that limits its effect. For example, we may restrict the impact of the scaling to some percentage of total workload demand on a server, e.g., the total workload demand for power on a server may be reduced up to a limit between 10% and 50%. This protects the workloads from being starved for capacity and generating an ever increasing backlog of unsatisfied demand.

The CPU scheduler employed by the simulator decides what fraction of resource demands can be satisfied for each interval. Any unsatisfied demands are carried forward to the next interval. Our current implementation for scheduling CPU resources simulates the Xen credit scheduler [9] which allocates CPU cycles to workloads based on a weight and capping factor [10].

B. Pool Power Capping

Pool power capping affects the power demand of a server pool by controlling the number of servers used. For this work, we employed a simple policy for pool power capping. The workload migration controller continuously monitors the overall power consumption of the server pool in addition to server utilization, and if the pool power consumption approaches the peak grid limit, then it attempts to further consolidate workloads and shuts down a server. The algorithm identifies the least loaded server and attempts to migrate its workloads to the remaining servers. The migration of a workload to a remaining server is only permitted if there are sufficient memory resources on the

server to support all of its workloads. CPU resources are permitted to be over booked. If all workloads can be migrated off the least loaded server, then the migrations proceed and the server is shut down. This enables the remaining servers to increase their power budgets. Finally, the migration controller has been modified so that it only adds servers when there is sufficient power available.

Pool power capping does not guarantee that power will be brought below a given threshold, but it can be complemented by server power capping to better bound power usage.

C. Energy Storage

Energy storage technologies can be used to store and smooth the supply of power for a data center. A variety of technologies are available [11] including flywheels, batteries, and other systems. Each has its costs, advantages and disadvantages. For example, a certain flywheel may be able to store energy for up to two hours, but is subject to a certain capacity loss over time. The simulator is able to simulate the general behavior of energy storage systems. For each interval in the simulation, when the sum of renewable and available grid power exceeds the demand for power, surplus power is added to the energy storage. The maximum size of energy storage is a factor in the experimental design. At present we place no limit on the rate at which energy can be stored or retrieved. For example, the authors in [11] state that a NaS battery installation of total 7.2 MWh is able to provide 1.2 MW power. The power demands for the simulated data center in our case study are much smaller. So we focus on the energy storage size as a limiting factor rather than the power bandwidth or losses.

IV. CHARACTERIZING POWER SUPPLY AND DEMAND

This section demonstrates the statistical but repetitive nature of enterprise workloads and hence power demands, and that the supply of power from renewable sources is also statistical but repetitive in nature.

A. Server Power Demand

Server power consumption can be estimated using a linear model as shown Section II.A. The P_{idle} term causes power to be used by a server whether or not there is any work to do. For today's servers and blades, the P_{idle} term can represent a significant portion, e.g., 50% or more, of a server's peak rated power usage. As a result, technologies such as frequency scaling have a bounded effect. This suggests that server power capping alone will have a bounded impact and that pool power capping is likely to be advantageous. However, reducing the number of servers will eventually over-subscribe the servers, causing CPU violation penalties, which is reported with our approach and must be taken into account when making planning decisions.

B. Workload Power Demand

Figure 2 illustrates the average and range of power usage for servers that support 138 enterprise applications using between 10 and 20 servers. The workloads are dominated by interactive work but also include batch processing. The x -

axis gives the time of day for seven days. The y -axis gives the power usage in KW. The dark line shows the mean power usage as averaged over six weeks. The shading shows the range of power values around the mean, i.e., the span between the minimum and maximum power demand for each interval. The figure shows that workload demands and hence power demands are indeed statistical but repetitive in nature. This is expected for enterprise workloads since much of an enterprise's work follows the working days of its employees and daily, weekly and monthly cycles of reporting. From more detailed results, variability is due to weekday holidays. Such information is known in advance and can be taken into account in planning exercises.




Figure 2: Variability of data center demand

C. Power Supply

The power supply side for the data center comprises a mix of one more energy sources. For our case study we consider PV panels and wind turbines as time-varying power sources combined with power from the electrical grid.

We calculate the generated power from PV panels based on historical sunshine traces and data on global irradiance for two locations: New Mexico and Colorado. For the panels, we assume an average efficiency of 13%, i.e., including all losses, 13% of the energy from the sun reaching the panel will be transformed into electrical power usable by the data center. Further, we consider PV panels that are able to produce up to 143 Watt per square meter. These numbers are typical for today's PV systems and can be adjusted as technology improves. Finally, the PV installation is sized to produce the desired maximum power. Similarly, for wind power, we consider specifications of an available wind turbine [12] to related wind to power. We used the given profile of the wind turbine to calculate the supply power for historical wind speed traces for a location in Colorado.

Figure 3 shows the power generated from a 100 KW PV installation for a location in New Mexico for six weeks of data. Three weeks of data are from June 1998 and three weeks are from June 1999. The figure illustrates the variability of power supplied by laying the power traces for the six weeks on top of each other. The average curve illustrates the power for each 5 minute interval of a week averaged over the six weeks of data. Further, the figure shows the observed range, i.e., the span between the minimum and maximum power supplied for each interval. From the figure, we see that as expected the mean is very repetitive—clearly the time of day effects show the panels

only perform when the sun shines. We note that 2.5 days out of 7 have wide ranges for generated power over the 6 weeks.

Similar power generation traces and charts are available for other locations and sources of renewable energy.




Figure 3: Variability of 100 KW PV power for New Mexico

V. CASE STUDY

A case study involving ten weeks of real data for 138 SAP applications is used to evaluate the effectiveness of several policies. The traces are obtained from a data center that specializes in hosting enterprise applications such as customer relationship management applications for small and medium sized businesses. The application resource demand traces capture average CPU and memory usage as recorded every 5 minutes. The first 4 weeks of data is used to decide on an initial placement for the workloads upon servers and the following six weeks of data are used to evaluate the effectiveness of power management plans.

The workloads typically required between two and eight CPUs and had memory sizes between 6 GB and 32 GB. The server pools consist of servers having 8 x 2.93-GHz CPUs, 128 GB of memory, and two dual 10 Gb/s Ethernet network interface cards for network and virtualization management traffic. Each of these servers consumes 695 Watts when idle (i.e., P_{idle}) and 1013 Watts when fully utilized (P_{full}).

While the 138 SAP applications present an interesting set of enterprise workloads, modern data centers typically contain many more servers and require much more power than needed for these applications. For this reason we scale the power demand by considering 10 instances of the 138 applications, with each instance in its own server pool, to represent a data center with 1380 SAP workloads. Further, we assume a PUE of 1.2 to calculate the total power demand of the data center. This is a reasonable value for state of the art data centers.

Our experimental design is as follows. Section V.A considers a data center that uses grid power alone to power its servers. The peak grid power is limited, and we demonstrate the impact of server and pool power capping mechanisms on several metrics including CPU violation penalties per hour, peak power used, and power deficit. Section V.B considers several different mixes of solar and wind power for different locations in combination with energy storage.

A. Impact of Power Management Policies

Figure 4 shows the power demand from the workloads when no power capping is in place. The power demands are

computed by the simulator using the baseline simulation behavior as we traverse the historical workload demand traces. During the simulation process, the migration controller increases and decreases the number of servers used based on load. From the figure, the power usage is generally below 200 KW with a few spikes above 220 KW. The average power usage is 166 KW. The figure also illustrates a horizontal black line at 170 KW. The following scenarios evaluate the effectiveness of server and pool power capping at limiting peak power to this value of 170 KW.




Figure 4: Power Demand Trace with no Power Capping




Figure 5: Power Demand with Server Power Capping (50%)

Figure 5 shows power demand over time when server power capping is introduced. During the simulation process, servers are still removed from the pool when demands drop; however, they are only added back to the pool when demands increase and there is sufficient power available. When power demands exceed the peak of 170 KW, then we simulate a perfect frequency scaling policy that reduces the power for all workloads equally, up to 50%, in an effort to keep power demands below 170 KW. Workload demands that are not satisfied are carried forward to be processed as more power becomes available. The figure shows that power is kept below 170 KW.

Figure 6 presents results for the pool power capping policy. During the simulation process, servers are not only removed when demands decrease but also when demand for power approaches the peak power limit of 170 KW. The idea here is that by consolidating workloads, fewer servers are contributing their P_{idle} components of power demand to power usage. The figure shows that power demand is generally below 180 KW, with a few spikes above 180 KW.

Figure 7 presents the results for combined pool and server power capping with a threshold of 50%. For the simulated data set, the integrated scenario again keeps the demands below the 170 KW peak.




Figure 6: Power Demand with Pool Power Capping




Figure 7: Power Demand with Integrated Server (50%) and Pool Capping

Power management policies also impact the application performance behavior. Table 1 shows the impact of the no capping and power capping scenarios on application performance behavior and power deficit metrics. The table shows the CPU violation penalties per hour, the peak power used, and the percentage of intervals where a power deficit occurred. The table shows that for no capping there is an average of 5 violation penalty units per hour. The violations inevitably arise when workload demands exceed the capacity of a server. The value of 5 acts as a baseline for the desirable application quality of service.

Table 1: CPU Violation Penalties and Power Metrics for Power Capping Policies

Scenario	CPU Violation Penalties per Hour	Peak Power KW	Power Deficit Intervals
No Power Capping	5.10	223.43	39.63%
Pool Power Capping	263.44	191.17	6.24%
Server P. Cap (50%)	4566.66	170.09	0.02%
Pool and Server P. Cap (50%)	531.01	170.09	0.06%

From Table 1, we see that out of the different capping policies, pool power capping has the lowest CPU violation penalties per hour of 263. However, its peak power usage is 191 KW, which is well above our threshold of 170 KW. Approximately 6% of the 5 minute intervals had power demands that above the limit. The server power capping scenario reduced peak power usage to 170 KW by limiting available compute resources as much as 50%. It incurred fewer intervals where workload demands could not be satisfied, but had much more sustained and severe violations

hence giving a much higher CPU violation penalty per hour. We deem the combined pool and server power capping scenario as most effective. It reduced the peak power to 170 KW and had much lower CPU violation penalties than server power capping alone. The combined methods are effective because server capping reacts quickly to changes in power needs, while pool power capping controls the impact of P_{idle} .

For the remainder of the case study, we use pool and server power capping together. We chose a threshold of 50% for aggressive server power capping.

B. Impact of Power Supply Mix and Energy Storage

This section reports the CPU violation penalty metric for scenarios that augment peak power from the grid with renewable power and stored energy. In our experiments, we evaluate different power management plans with unlimited grid power, and with grid power limits ranging from 140 KW to 200 KW. The figures only show results for cases where the maximum power deficit for any interval is less than 5 KW.

Figure 8 shows the CPU violation penalties per hour for a data center running on a combination of grid and stored energy only—no renewable power sources are used. For peaks of 200 KW and 190 KW, the ability to store 10 KWh of energy is sufficient to bring the CPU violation penalties to the ideal level of 5. For a peak limit of 180 KW doing so requires the ability to store 50 KWh. Adding 100 KWh of stored energy reduced the CPU violation penalties per hour for the 170 KW limit by a factor of 10, but was not enough to reduce the violations to ideal levels. We note that the figure shows several examples where adding energy storage capacity increases the CPU violation penalties per hour instead of decreasing them. This is because the way in which the system reacts to fluctuating power supply and demand differs as more storage is added, possibly delaying some decisions to consolidate workloads and shutdown servers. In general, when available power is too close to the demand for power, we see unstable behavior as illustrated by CPU violation penalties per hour that oscillate with changes in storage capacity.

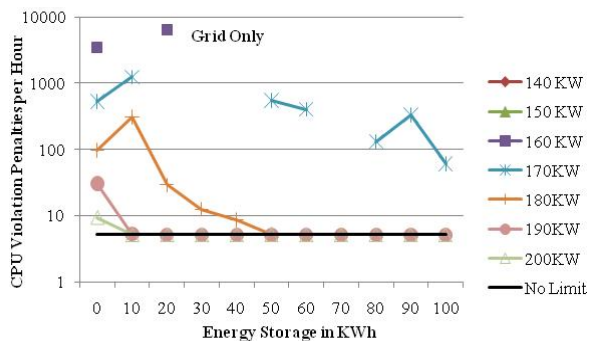


Figure 8: QoS for a Data Center Running on Grid Power and Stored Energy Only

Figure 9 augments the grid and energy storage of Figure 8 with 100 KW of PV for a data center in New Mexico. Limiting the grid power impacts the total power supply limit for the data center, which is the sum of the grid power limit

and the available on-site power. As the available on-site power changes over the day, the power supply for the data center is also time dependent. Pool and server power capping is then used to limit the data center power demand to the time-varying power supply. By simulating the data center, we estimate the amount of power consumed from the grid over time. With the additional power source and 40 KWh of stored energy, peak grid power can be reduced to 160 KW to achieve ideal CPU violation penalties of 5 per hour. Furthermore, peak power can be reduced to 150 KW with 80 KWh of stored energy. (Recall that Figure 8 had showed that without this additional supply of on-site power, energy storage of even 100 KWh was not sufficient with a 170 KW peak grid power to attain a CPU violation penalty of less than 100 per hour.




Figure 9: QoS for a Data Center in New Mexico with 100 KW PV

Figure 10 shows the impact of increasing the supply of solar power from 100 KW to 250 KW. With energy storage of 30 and 90 KWh, the data center could limit peak grid power to 160 KW and 140 KW, respectively. This shows that a larger fraction of the data center power demand can be efficiently supplied by solar energy for a location in New Mexico, though more energy storage is required.




Figure 10: QoS for a Data Center in New Mexico with 250 KW PV

Figure 11 and 12 give planning results for a data center in Colorado. Figure 11 augments the data center grid power with a 225 KW wind turbine. Though the turbine has a highly rated peak output, the actual supply of power from wind at the Colorado location is much lower and is highly variable. Figure 11 shows that achieving ideal CPU violation penalties per hour still requires a peak grid power of 180 KW

with 20 KWh of energy storage. Figure 12 shows that if an additional 100 KW of local PV power is added to the facility, then an ideal CPU violation penalty per hour level can be achieved with a peak power of 150 KW and 100 KWh of energy storage.




Figure 11: QoS for a Data Center in Colorado with a 225 KW Wind Turbine

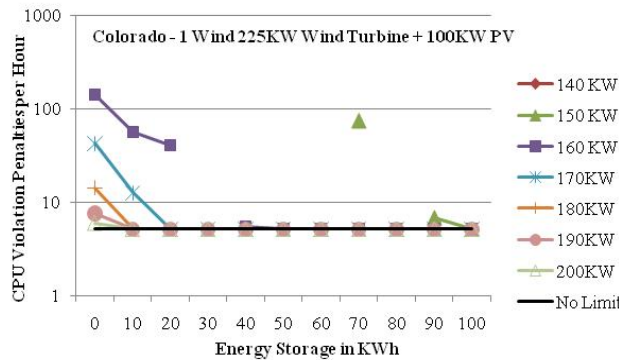


Figure 12: QoS for a Data Center in Colorado with a 225KW Wind Turbine and 100 KW PV

To summarize the results of the planning scenarios, the New Mexico data center can limit peak grid power to 150 KW with 100 KW PV and 80 KWh of energy storage. The Colorado data center can limit peak power to 150 KW by using one 225 KW wind turbine, 100 KW of PV, and 100 KWh of energy storage.

VI. RELATED WORK

Significant efforts have been made to improve data center energy efficiency. These efforts predominantly fall into two categories: efforts around assessment of energy use within data centers; and efforts to reduce the energy used by data centers.

Accurately portraying the energy used by data centers has been a topic of widespread discussion within the literature. For example, nearly 8 years ago, Mitchell-Jackson et al. [13] surveyed data centers within California and found a widespread discrepancy between reported estimates of power use and actual measurements of power use. This discrepancy was attributed largely to the lack of standard metrics for communicating energy use within data centers. Based on this finding, Mitchell-Jackson et al. [14] estimated the regional and nationwide energy use within

data centers. More recent studies are also available [15] [16]. The Green Grid [6], an industry consortium, has proposed standardized metrics to enable consistent and comprehensive measurement and tracking of energy efficiency in data centers across the industry.

In terms of efforts to reduce the energy used by the data center, an area of growing focus has been virtualization of IT resources [17][18]. In addition, numerous efforts have been related to developing energy-efficient IT hardware [19][20]. Fan et al. proposed a power model to estimate the power used by IT equipment in data centers [3]. But the IT equipment often only uses around 50% of the energy in the data center. Therefore, much research has gone into management of cooling resources within the facility, including the use of outside air [21] or real-time dynamic thermal management based on sensors and control of data center air-conditioners [2]. Early work that integrates the IT and facility has also begun [22], for example, by considering the energy efficiency associated with cooling in choosing where to place workloads within the data center.

Server power capping is addressed at the server level in [25] using a feedback controller through dynamic scaling of the P-states of the processors. VirtualPower [26] integrates power management mechanisms of physical servers with virtualization technologies in order to right-provision the amount of CPU cycles and reduce the server's power consumption. In [27], the authors cap server power through alternating between extreme power states that can improve the available capacity even if the power is capped below the budget. In [28], the authors exploit the power capping for individual servers through both DVFS and clock-throttling driven by an adaptive controller, thus achieving integrated control for both power minimization and power capping. The power capping problem for a group of servers is also studied in [28] through proportional sharing of the total power budget among the individual servers; the budget per server is then enforced through local power capping. In [29], the power capping problem is addressed at the server level, the group level and the data center level through clock-throttling, DVFS, workload migration and consolidation enabled through virtual machine migrations. This can happen on a variety of time scales and space granularities.

Our work differs from the above in the sense that we are actively managing the data center power demand according to available power through power management policies. We consider the time-varying nature of workload demands and power sources when evaluating alternative approaches for powering data centers and take better account of the impact of consolidation on application performance via the CPU violation metric. Torcellini et al. studied a net-zero energy building [23] and their approach to integrating supply and demand is similar to ours, but a data center presents unique challenges. Primarily, the energy densities are significantly larger than typical buildings, and demand is more dynamic and complex. Previous work [24] focused on the estimation of supply mixes for given workload demands. This paper goes beyond that and simulates the operation of the data center that incorporates energy supply information into workload management.

VII. CONCLUSIONS AND FUTURE WORK

This paper describes an approach for designing an energy management plan that matches the supply of power with the demand for power in data centers. The power management plan defines a choice for a mix of baseline power, e.g., from the power grid, power from renewable energy sources, and stored energy. The plan matches the energy supply to the power needs of a data center's expected workloads while taking into account the estimated impact on workload performance behavior.

A case study demonstrates the use of our approach to design a plan for a realistic and complex data center workload for two geographic locations. For each plan the simulator reports on the expected impact of power constraints on workload performance behavior. A plan can be designed that minimizes such impact. Our approach gives greater precision than other methods that do not take into account time-varying resource usage, location specific weather data for renewable power sources, the time-varying supply of power, energy management policies, and the impact of energy storage technologies. A sensitivity analysis that applies our approach using years of historical traces could help to determine how often a backup plan, e.g., the use of additional power, is expected to be employed.

Our next steps for the simulator include improving power demand calculations to take into account p-state changes in addition to frequency scaling. Also, the simulation model for energy storage alternatives should take into account power bandwidth and power losses over time. Additional power usage for compute storage and networking should be taken into account in the same trace driven manner. Part of our future work will be a more detailed and quantitative comparison with other power capping and workload managing methods. We further plan to integrate additional performance metrics such as application level performance and will enhance the simulator to support a cost benefit analysis regarding different plans.

REFERENCES

- [1] Y. Chen, D. Gmach, C. Hyser, Z. Wang, C. Bash, C. Hoover, and S. Singhal, "Integrated Management of Application Performance, Power and Cooling in Data Centers", in Proc. of the 12th IEEE/IFIP Network Operations and Management Symposium (NOMS), 2010.
- [2] C. Bash, C. Patel, R. Sharma, "Dynamic thermal management of aircooled data centers", Proc. of the Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM), 2006.
- [3] X. Fan, W. Weber, L. Barroso G. Bramlett, "Power Provisioning for a Warehouse-sized Computer," in Proc. of the ACM Int. Symposium on Computer Architecture, 2007.
- [4] D. Gmach, J. Rolia, L. Cherkasova, A. Kemper, "Capacity Management and Demand Prediction for Next Generation Data Centers", in Proc. of the IEEE Int. Conf. on Web Services, 2007.
- [5] D. Gmach, J. Rolia, L. Cherkasova, A. Kemper, "Resource Pool Management: Reactive versus Proactive or Let's be Friends", Computer Networks, Special Issue on Virtualized Data Centers, Vol 53, No 17, 2009.
- [6] C. Belady, et al., "Green Grid Data Center Power Efficiency Metrics: PUE and DCiE", Available at http://www.thegreengrid.org/gg_content/TGG_Data_Center_Power_Efficiency_Metrics_PUE_and_DCiE.pdf, 2007.
- [7] National Renewable Energy Laboratory, "CONFIRM Solar Energy Resource Data", http://rredc.nrel.gov/solar/new_data/confirm/
- [8] D. Gmach, S. Krompass, A. Scholz, M. Wimmer, A. Kemper, "Adaptive Quality of Service Management for Enterprise Services," in ACM Transactions on the Web (TWEB), Vol 2, No 1, 2008.
- [9] Xen Wiki, "Credit Scheduler: Credit-Based CPU Scheduler," <http://wiki.xensource.com/xenwiki/CreditScheduler>, 2007.
- [10] D. Gmach, "Managing Shared Resource Pools for Enterprise Applications", Ph.D. thesis, Technische Universität München, 2009.
- [11] R. Walawalkar, J. Apt, "Market Analysis of Emerging Electric Energy Storage", DOE/NETL report 2008/1330.
- [12] J.P. Saylor & Associates, Technical Data Series Vestas V27-225 kW, http://www.windturbinewarehouse.com/pdfs/vestas/Vestas_V_27_SA_C_DSM_7_8_05.pdf
- [13] J. Mitchell-Jackson, J. Koomey, B. Nordman, M. Blazek, "Data Center Power Requirements: Measurements from Silicon Valley," Energy, Vol 28, No 8, pp. 837–850, 2003.
- [14] J. Mitchell-Jackson, J. Koomey, M. Blazek, B. Nordman, "National and Regional Implications of Internet Data Center Growth in the US," Resources, Conservation and Recycling, Vol 36, No 3, pp 175–185, 2002.
- [15] EPA, "Report to Congress on Server and Data Center Energy Efficiency: Public Law 109–431," US Environmental Protection Agency, EnergyStar Program, 2007.
- [16] J. Koomey, "Estimating Total Power Consumption by Servers in the U.S. and the World," available online <http://enterprise.amd.com/Downloads/svrpwusecompletefinal.pdf>, 2007.
- [17] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, A. Warfield, "Xen and the Art of Virtualization," in Proc. of the 19th ACM Symposium on Operating Systems Principles (SOSP), 2003.
- [18] D. Marshall, W. Reynolds, D. McCrory, "Advanced Server Virtualization," Taylor & Francis, 2006.
- [19] K. Lim, P. Ranganathan, J. Chang, C. Patel, T. Mudge, S. Reinhardt, "Server Designs for Warehouse-Computing Environments," IEEE Micro, vol. 29, no. 1, pp. 41–49, 2009.
- [20] T. Burd, T. Pering, A. Stratakos, R. Brodersen, "A Dynamic Voltage Scaled Microprocessor System," IEEE Journal of Solid State Circuits, Vol 35, No 11, pp. 1571-1580, November 2000.
- [21] A. Kumar, Y. Joshi, "Use of Airside Economizer for Data Center Thermal Management," Proc. Thermal Issues in Emerging Technologies (ThETA 2), Cairo, Egypt, 2008.
- [22] C. Bash, G. Forman, "Data center workload placement for energy efficiency," in Proceedings of ASME InterPACK Conference, Vancouver, Canada, 2007.
- [23] P. Torcellini, S. Pless, M. Deru, D. Crawley, "Zero Energy Buildings: A Critical Look at the Definitions", ACEEE Summer Study, Pacifica Grove, CA, USA, 2006.
- [24] D. Gmach, Y. Chen, A. Shah, J. Rolia, C. Bash, T. Christian, R. Sharma, "Profiling Sustainability of Data Centers", in Proc. of the Int. Symposium on Sustainable Systems and Technology (ISSST), 2010.
- [25] C. Lefurgy, X. Wang, M. Ware. "Power capping: A prelude to power shifting", Cluster Computing, 11(2):183–195, 2008
- [26] R. Nathuji, K. Schwan, "VirtualPower: coordinated power management in virtualized enterprise systems", in Proc. of the 21st ACM Symposium on Operating Systems (SOSP), pp. 265–278, 2007.
- [27] R. Das, A. Gandhi, J. O. Kephart, M. Harchol-balter, C. Lefurgy, "Power Capping Via Forced Idleness", Workshop on Energy Efficient Design, 2009.
- [28] Z. Wang, X. Zhu, C. McCarthy, P. Ranganathan, V. Talwar., "Feedback control algorithms for power management of servers", in 3rd International Workshop on Feedback Control Implementation and Design in Computing Systems and Networks (FeBID), 2008.
- [29] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, X. Zhu. "No power struggles: Coordinated multi-level power management for the data center", in Proceedings of ASPLOS, pages 48–59, 2008.