Profiling Sustainability of Data Centers

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

***Abstract*—Today's data centers consume vast amounts of energy, leading to high operational costs, excessive water consumption, and significant greenhouse gas emissions. With the approach of micro grids, an opportunity exists to reduce the environmental impact and cost of power in data centers. To realize this, demand side power consumption needs to be understood and co-managed from the perspectives of both supply and demand. We present an approach to achieve this via data center power profiling and demonstrate its applicability for an enterprise data center.**

***Index Terms*—Data Center, Energy Profiling, Resource Management**

I. INTRODUCTION

IT is estimated that the information and communication technology sector (ICT) is responsible for 2% of the global carbon emissions [1]. Much of this is due to the energy consumption of data centers. Significant research is underway to develop technologies that reduce energy demands and the environmental impact of data centers. Virtualization technology is being used to consolidate workloads and improve IT utilization, and leading-edge cooling technologies such as air-side economizers further help improve data center efficiency. On the supply side, distributed generation and renewable energy sources are increasingly being deployed. However, the impact of these technologies on sustainability, particularly in an integrated supply-demand context, is not well understood. For example, the interaction of energy supply with a data center's unique workloads and hence energy demands is hard to predict. Micro grids can help to enable co-management of energy supply and demand. Our data center profiling method exploits micro-grids to enable fine-grained co-management of supply and demand.

II. OUR SOLUTION

We have developed an analysis tool that quantifies the overall impact of different energy management approaches during design and operation of data centers, and helps to determine the most cost effective and sustainable approach. During design, it quantifies the impact of alternatives that can be used to power a data center. During operation, it considers power demand management and evaluates the impact of alternative IT and facilities management policies on time-varying workloads and makes recommendations regarding how much power is required, what mix of power sources is desirable, and how power sources should be allocated across the data center infrastructure. To achieve this, our approach takes into account detailed knowledge of time-varying power demands and location-specific, time-varying power supplies.

A. *Overview of Our Approach*

Given a certain workload demand for a data center, we generate a detailed energy profile for the data center, including energy supply and demand traces as well as time-varying estimates of the net energy used from the grid. Our solution considers different demand management policies as well as equipment and location choices. The energy supply side comprises one or more energy sources, such as photovoltaic panels, municipal solid waste facilities, wind turbines and the grid. An outline of the profiling process and data flow is shown in Fig. 1.

As shown in Fig. 1, our approach undertakes the task in a systematic way following a series of steps. It begins with IT workload traces. A workload simulator incorporates specified workload management policies and IT equipment information to generate the IT resource utilization for given workload traces. Once we have obtained the IT resource utilization trace, we then apply a power model to determine the time-varying power demand for the entire data center. Next, we use a power supply simulator to assess different energy supply solutions and generate the estimated power supply trace. Given certain goals and the results of the previous steps (i.e., the power supply and demand traces), a candidate solution is generated and then analyzed to determine its operating characteristics, including costs, sustainability metrics and quality of service (QoS) metrics. Multiple solutions can be compared and a solution that meets the specified goals can be selected.

B. *Workload Simulator*

The first step is to estimate the energy demand of the data center. Given resource demand traces for workloads under study, based on historical data or future demand predictions, we employ a simulator that incorporates specified workload management policies and IT equipment information to generate the resource utilization for the data center [2]. The simulator allows the evaluation of various configuration and management policies and provides time-varying information on IT requirements, utilization values and QoS metrics. Energy demands are inferred from the utilization values. This is described further later in the paper.

The architecture of the workload simulator is illustrated in Fig. 2. The simulator takes as input historical workload

---

All authors are with Hewlett Packard Laboratories, Palo Alto, CA 94306, USA (email: {daniel.gmach, yuan.chen, amip.shah, jerry.rolia, cullen.bash, tom.christian, ratnesh.sharma}@hp.com).

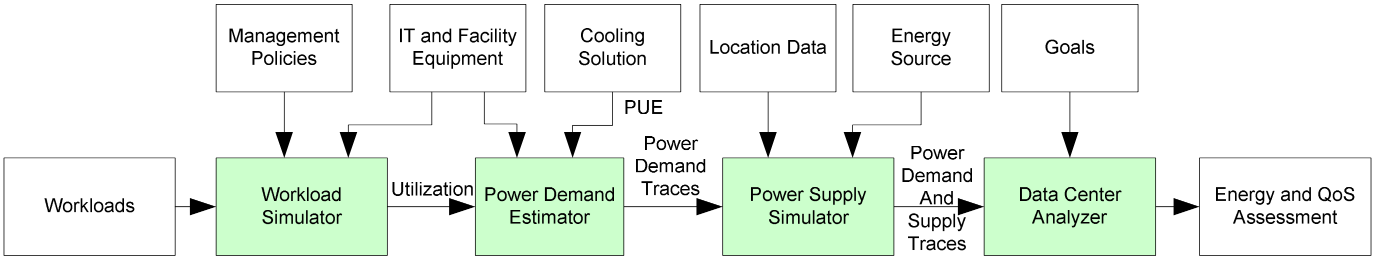

Fig. 1: Process to assess power supply and demand for a data center

demand traces, an initial workload placement, server resource capacity descriptions, and a management policy. The server descriptions include numbers of processors, processor speeds, real memory size, and network bandwidth. A routing table directs each workload's historical time varying resource requirement data to the appropriate simulated server. Each simulated server uses a fair-share scheduling strategy to determine how much of the workload demand is and is not satisfied. The central pool sensor makes time varying information about satisfied demands available to management controllers via an open interface. The interface also is used to integrate different controllers with the simulator without recompiling its code.

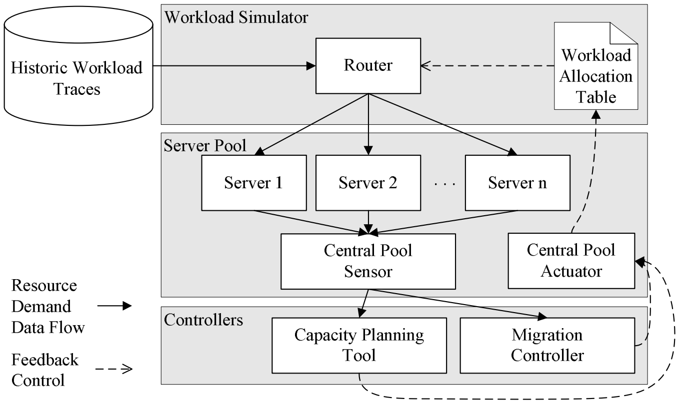

Fig. 2: Architecture of the Workload Simulator

The controllers periodically gather accumulated metrics and make decisions about whether to cause workloads to migrate from one server to another. Migration is initiated by a call from a controller to the central pool actuator. In our workload simulation environment this causes a change to the routing table that reflects the impact of the migration in the next simulated time interval. During the simulation process various metrics are gathered, e.g., the time-varying utilization of servers, the number of running servers, quality of service for each workload, number of workload migrations, etc. Different controller policies cause different behaviors that we observe through these metrics. Similarly, the different policies have an impact on energy demands.

C. *Power Demand Estimator*

The power demand estimator estimates the aggregate power demand from server, storage, and networking equipment and then includes the facility power demand. We estimate the total IT power usage in a data center based on the average CPU utilization across all servers, and the aggregate power consumed by all servers at idle and in full utilized states as follows:

$$P = P_{idle} + \left(P_{busy} - P_{idle}\right) * u + P_{others},$$

where u is the average CPU utilization across all servers, P_{idle} and P_{busy} are the aggregate power consumed by all servers at idle and fully utilized states respectively, and P_{others} is a fixed offset which captures the power consumption of storage and networking switches. Further information on the power model can be found in [3]. The dynamic power consumption by storage and networking equipment will be considered in our future work. The model is calibrated from historical data or a set of experiments. This simple model has proven to be surprisingly accurate in modeling power consumption of a large number of servers since other components' activities are either static or correlate well with CPU activity. Given the resource utilization traces generated by the workload simulator, we use the model to track the aggregate power demand of the data center.

Once we obtain the IT power demand, the total power demand is estimated using the PUE (power usage effectiveness) metric [5], which is the ratio of the total power used by the data center facility to the power used by the IT equipment itself.

$$PUE = \frac{P_{IT} + P_{Facility}}{P_{IT}}$$

PUE can be estimated from simulation or historically averaged data. If the PUE can be measured empirically, then the facility power can be determined as the difference between the total and the IT power (top-down approach). Conversely, if the infrastructure description is known based on a model or experiment, then the facility power can be determined from cooling and power delivery models by simplifying the facility power into cooling and power delivery components(bottom-up approach). In doing so, we obtain

$$PUE = 1 + \frac{1}{COP_G} + \%Power\ Delivery\ Losses,$$

where COP_G is the coefficient of performance of the ensemble [10], a measure of the cooling power required to remove heat from the IT equipment.

D. *Power Supply Simulator*

The power supply simulator assesses different energy supply solutions. This step considers location data, climate information like sunshine hours or wind speeds, and various supply solutions like photovoltaic panels or wind turbines. We determine time-varying traces for the estimated power supply for combinations of such power sources. Further, it determines the mean and variability of energy supply, i.e., the power changes between consecutive measurement intervals.

E. Data Center Analyzer

The data center analyzer evaluates the data center with respect to acquisition costs, operational costs, sustainability metrics and QoS metrics. Sustainability metrics include embedded footprint, CO_2 emissions, and water consumption. QoS metrics measure whether or not QoS objectives, e.g., application response time, have been met, and if not by how much they fall short. We are especially interested in what impact power capping policies have on the QoS.

III. CASE STUDY

We have implemented a first version of the data center simulator and determined power demand profiles for a production data center that comprises 138 enterprise resource planning workloads. The workload demand 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. Each workload was hosted on its own server, so we use resource demand measurements for a server to characterize the workload's demand trace. The measurements were originally recorded using vmstat [4]. Traces capture average CPU and memory usage as recorded every 5 minutes. A characterization of the workloads in [2] shows their bursty nature, e.g., for more than 40% of the workloads, the top 3% of demand values are between 10 and 2 times higher than their remaining demands.

For the data center we consider a resource pool of servers with 8 x 2.93-GHz processor cores, 128 GB of memory, and two dual 10 Gb/s Ethernet network interface cards for network traffic and virtualization management traffic, respectively. Each of these servers consumes 695 Watts when idle and 1013 Watts when it is fully utilized.

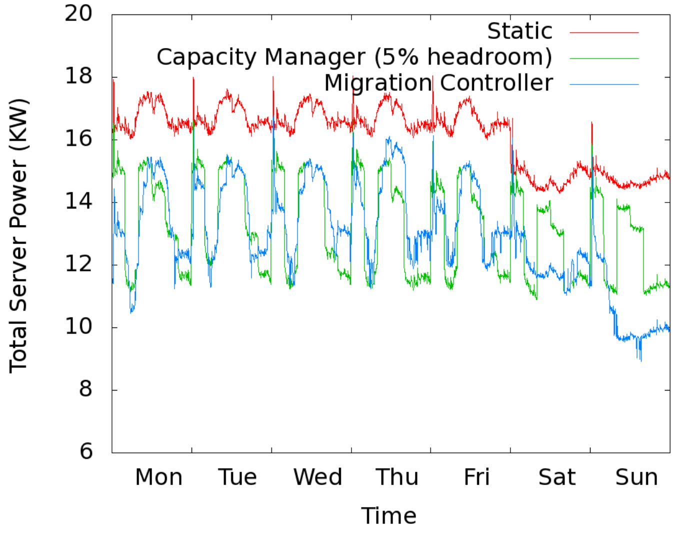

Fig. 3: IT power demand

A. Data Center Simulator

Fig. 3 shows power traces from one week for the IT infrastructure for three management policies. In the *Static* scenario, a capacity planning tool [6] determined workload placement for the week based on the demands in the previous three weeks. This scenario represents a traditional virtualized data center where workloads are static during runtime. In the *Capacity Manager* scenario, the capacity planning tool is invoked every four hours. It uses workload trace information from the past three weeks to predict workload demands into the next four hours. The predicted workload demands are then used to aggressively consolidate workloads onto physical servers. This behavior might cause a high number of workload migrations every four hours. We note that there exist approaches limiting the maximum number of migrations [7] which can be considered to reduce the migration overhead. Nonetheless in this paper, we consider the extreme case allowing all workloads to migrate. Furthermore, it allocates 5% headroom onto each physical server to diminish overload situations. The third scenario, *Migration Controller* scenario, employs a dynamic migration controller that continuously monitors resource consumption, migrates workloads off overloaded servers, and consolidates workloads from under-loaded servers. Furthermore, it turns servers on/off as needed. As shown in Fig. 3, compared to the S*tatic* approach, both *Capacity Manager* and *Migration Controller* reduced power consumption substantially. Further, in the *Migration Controller* scenario, less power is consumed on weekend days than in the *Capacity Manager* scenario. The power demand curve of the *Capacity Manager* scenario further exhibits high changes in power demand whenever between two consecutive 4-hour intervals the number of required servers changes. In contrast to that, the curve of the *Migration Controller* scenario is smoother.

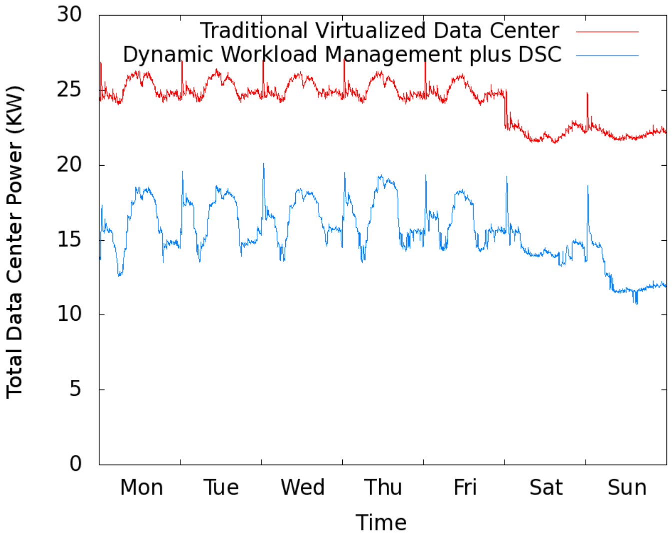

Fig. 4: Total power demand

B. Power Model

We further evaluated the total power from both IT and facility for two different data center solutions: a traditional virtualized data center with a typical PUE of 1.5, and a state of the art data center with dynamic workload migration and Dynamic Smart Cooling (DSC) [9] corresponding to a PUE of 1.2. DSC adjusts the blower speed and the supply air temperature (SAT) set points of the CRACs based on data from temperature sensors that measure the inlet temperature of IT equipment. This allows DSC to supply the requisite cooling to the IT equipment without over-cooling the room. The results in Fig. 4 show that data center energy efficiency as well as workload management policies have a huge impact on the total power demand trace.

In DSC, IT and cooling management are integrated through a data model to link information in the thermal model such as sensor readings with information in the IT model such as server location. We further automatically collect IT and facility data, e.g., local workload placement index (LWPI), and deliver it to controllers at run time. LWPI is a measure of how efficiently a location in the data center can be cooled [8]. LWPI is the sum of three components, has units of temperature, and is described as:

$$\begin{aligned} LWPI_{server} = {} & (Thermal\ Management\ Margin)_{server} \\ & + (AC\ Margin)_{server} \\ & - (Hot\ Air\ Recirculation)_{server} \end{aligned}$$

The first term in the equation is the *Thermal Management Margin* at a server. This is the difference between a server's inlet air temperature and its specified set point, which is usually given by the manufacturer as a maximum recommended temperature, less some margin. The second term is *Air Conditioning (AC) Margin*. AC Margin is the sum over all CRACs of the difference between the CRACs current supply air temperature and their minimum possible set point, weighted by the degree to which the particular CRAC can actually supply cold air to the server and influence its inlet air temperature [9]. The final component is *Hot Air Recirculation* which quantifies the amount of hot air recirculation at the server and is the difference between the server's inlet air temperature and the temperature of the air supplied through vent tiles in close proximity to the server. Higher LWPI values indicate higher thermal efficiency. LWPI is only well-defined at thermal sensor locations near the inlets of compute servers. Temperature measurements can be obtained through an external sensor network or from inlet air temperature sensors incorporated into recent server hardware. LWPI is computed directly for servers with sensors. Linear interpolation is used for other servers. The workload management controllers use LWPI to choose the best location to place the workloads and DSC is regularly updated with server state information (e.g., server power status), which is received from the controllers. DSC then adjusts the CRACs accordingly to react to the changes in IT workload demand. The implementation of dynamic smart cooling, its integration with workload management, and usage of modern cooling equipment results in a low PUE of 1.2.

C. Power Supply Simulator

The power supply profile highly depends on the technology and location. Fig. 5 shows two power generation traces for a week in May of a 4KW PV array for two locations: Houston and Seattle. While the PV power generation in Houston is generally predictable, the power trace from Seattle is lower, more variable, and less predictable. We note that some power sources, such as municipal solid waste (MSW) facilities, can offer base load power throughout the day. Even though differences in energy supply and demand can be compensated by the grid to some extent, workload management policies can partially address this gap by considering power availability and implementing policies to cap the energy demand during forecasted periods of low supply-side availability.

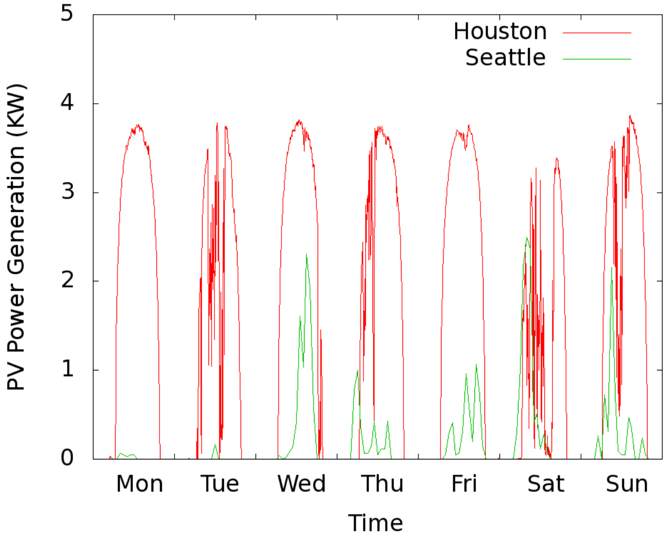

Fig. 5: Power supply traces for photovoltaic

Fig. 6 combines supply-side and demand-side management for a location in Houston. Power is provided by a mix of 4KW photovoltaic and 15KW base load from a municipal solid waste facility. Additional power demand is served through the grid, and we assume that surplus supply-side power can be sold back to the grid.

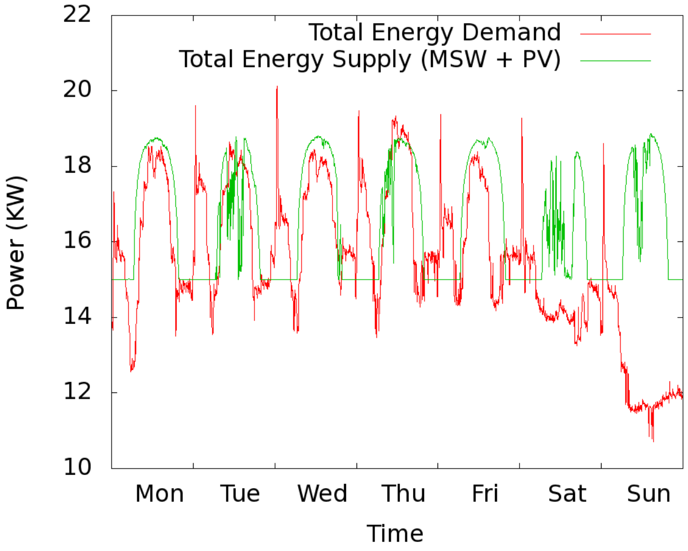

Fig. 6: Energy supply and demand

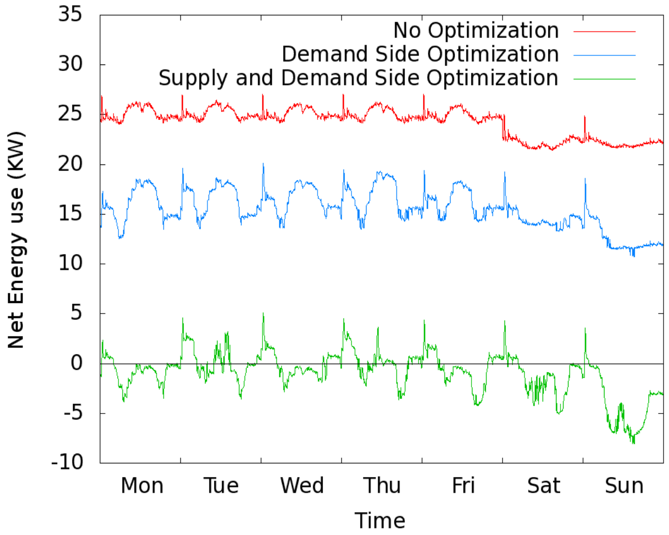

Fig. 7: Net energy

D. *Data Center Analysis*

Fig. 7 shows the net energy balance for three different data center solutions with varied energy management policies: (1) No Optimization: static workload + grid power, (2) Demand Side Optimization, and (3) Supply and Demand Side Optimization. It clearly shows that supply and demand side optimization manages to nearly achieve a data center with net-zero grid energy use. Fig. 8 shows the energy, CO_2, and water footprint for the three scenarios. The CO_2 and water footprint numbers are calculated assuming average numbers for power generation in the U.S. For the water footprint, we consider direct and indirect water consumption. Indirect consumption is the water used during power production and direct water consumption is the water used through evaporative cooling towers at the data center side.

In the *No Optimization* scenario, the total energy demand for the week is 4072 KWh resulting in a footprint of nearly 2.4 tons of a CO_2 and 3697 gallons of water. Optimizing the power demand side helps to reduce the total power demand to 2608 KWh. This reduces the CO_2 footprint to 1.6 tons and the total water consumption to 2633 gallons. Integrating supply and demand side management and using renewable on-site energy sources helps to achieve an overall negative grid power usage by selling back more power to the grid than consuming. The power drawn from the grid over the week accumulates to 61 KWh corresponding to a footprint of 37 kg of CO_2 and 31 gallons of water. In addition, 1329 gallons of water are consumed directly through evaporative cooling of the data center resulting in a total water consumption of 1360 gallons.

IV. 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.

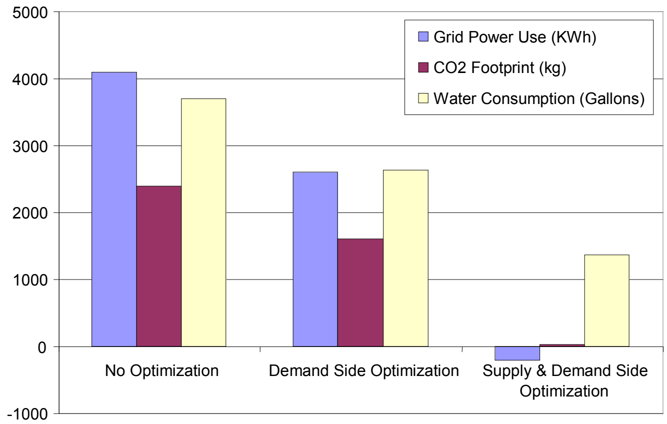

Fig. 8: Energy, CO_2, and water footprint

Accurately portraying the energy used by data centers has been a topic of widespread discussion within the literature. For example, nearly a 8 years ago, Mitchell-Jackson et al. [12] 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. [13] estimated the regional and nationwide energy use within data centers. Aebischer et al. [14] highlighted and proposed metrics to measure the energy use in data centers across the IT equipment while taking into account the energy used by auxiliary resources. Subsequently, the Uptime Institute [15] began surveying IT equipment manufacturers and specific data center operators to report on industry-wide trends in data center energy efficiency. Greenberg et al. [16] also benchmarked about 22 production data centers in the US, and recommended best practices for energy efficiency which have also been reported elsewhere [17][18]. Perhaps the most comprehensive and recent studies around energy use within data centers can be found in a report by the US Environmental Protection Agency [19] and the work of Koomey [20], who respectively estimated the aggregate electricity used by data centers across the US and the world. These widely cited estimates suggest that data centers consume approximately 1.5% of electricity use within the US, and that the energy used by data centers has been doubling about every five years. The Green Grid [5], an industry consortium, has also 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 [21][22]. In addition, numerous efforts have been related to developing energy-efficient IT hardware [23][24]. Fan et al. proposed a power model to estimate the energy 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 [25] or real-time dynamic thermal management based on sensors and control of data center air-conditioners [9]. Early work that integrates the IT and facility has also begun [8], for example, by considering the energy efficiency associated with cooling in choosing where to place workloads within the data center.

Our work is unique from the above in the sense that most of the above efforts are exclusively related to measuring or managing demand-side energy use. Most of the work related to management of IT equipment does not consider the facility; while work on improving facility energy efficiency considers cooling and power delivery only. In addition, prior solutions do not consider the time varying nature of workload demands and power sources when evaluating alternative approaches for powering data centers. Recognizing the growing importance of the energy source for next-generation data centers, our work includes consideration and quantification of the availability of energy on the supply side while including time-varying considerations. Torcellini et al. studied a net-zero energy building [26] 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.

We believe this is the first work that integrates not only IT and facility energy use within the data center, but also the supply and demand energy profiles of the data center.

V. CONCLUSION AND FUTURE WORK

This paper presented an approach to co-manage data center's energy supply and demand side. Our approach achieves this via integrated data center power demand profiling and available power supply assessment. We further demonstrated its applicability for an enterprise data center.

We have implemented the workload simulator and a power model that provide information on the energy demand. We plan to employ the simulator and this process to develop design blueprints for net-zero energy data centers under different circumstances. A key challenge is to determine policies for sizing and delaying workloads with respect to power availability and without significantly affecting their resource access quality. Factoring in the effects of dynamic electricity pricing into our cost model is another area of future work.

Power storage systems can help to overcome sudden power demand peaks or drops in power supply. We plan to evaluate the impact of such power storage systems on the hosted applications' quality of service, availability, the choice of management policies and the data center power consumption.

REFERENCES

[1] "Gartner estimates ICT industry accounts for 2 percent of global CO2 emissions", April 2007.

[2] D. Gmach, J. Rolia, L. Cherkasova, A. Kemper, "Resource Pool Management: Reactive versus Proactive of Let's be Friends", Computer Networks, Special Issue on Virtualized Data Centers, 2009

[3] X. Fan, et al., "Power Provisioning for a Warehouse-sized Computer", ACM Intl. Sym. on Computer Architecture, 2007.

[4] vmstat. http://www.linuxcommand.org/man_pages/vmstat8.html

[5] 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.

[6] J. Rolia, L. Cherkasova, M. Arlitt, A. Andrzejak, "A Capacity Management Service for Resource Pools", HP Labs Technical Report, HPL-2005-1, 2005.

[7] D. Gmach, J. Rolia, L. Cherkasova, G. Belrose, T. Turicchi, and A. Kemper (2008). "An Integrated Approach to Resource Pool Management: Policies, Efficiency and Quality Metrics", in Proceedings of the 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). Anchorage, Alaska, USA.

[8] C. Bash and G. Forman, "Data center workload placement for energy efficiency," in Proceedings of ASME InterPACK Conference, Vancouver, Canada, 2007.

[9] C. Bash, C. Patel, and R. Sharma, "Dynamic thermal management of aircooled data centers", Proc. of the Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM), San Diego, CA, 2006.

[10] C. Patel, et al., "Energy Flow in the Information Technology Stack: Coefficient of Performance of the Ensemble and its Impact on the Total Cost of Ownership," Proc. ASME International Mechanical Engineering Congress and Exhibition (IMECE), Chicago, IL, 2006.

[11] R. Sharma, A. Shah, C. Bash, T. Christian, C. Patel, "Water Efficiency Management in Datacenters: Metrics and Methodology", Proc. of the IEEE International Symposium on Sustainable Systems and Technology (ISSST), 2009.

[12] J. Mitchell-Jackson, et al., "Data Center Power Requirements: Measurements from Silicon Valley," Energy, Vol 28, No 8, pp. 837-850, 2003.

[13] J. Mitchell-Jackson, et al., "National and Regional Implications of Internet Data Center Growth in the US," Resources, Conservation and Recycling, Vol 36, No 3, pp 175-185, 2002.

[14] B. Aebischer, et al., "Energy Efficiency Indicators for Data Centers", International Conference on Improving Energy Efficiency in Commercial Buildings (IEECB), Frankfurt, Germany, 2004.

[15] The Uptime Institute, 2005-2010 Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment.

[16] S. Greenberg, et al., "Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers," ACEEE Summer Study, Pacifica Grove, CA, USA, 2006

[17] PG&E, "High Performance Data Centers: A Design Guidelines Sourcebook," Pacific Gas & Electric Company, 2006.

[18] W. Tschudi, et al., "Measuring and Managing Data Center Energy Use," HPAC Engineering, pp. 45-51, March 2006.

[19] EPA, "Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431," US Environmental Protection Agency, EnergyStar Program, August 2007.

[20] J. Koomey, "Estimating Total Power Consumption by Servers in the U.S. and the World," available online http://enterprise.amd.com/Downloads/svrpwrusecompletefinal.pdf, February 2007.

[21] P. Barham et al., "Xen and the Art of Virtualization," Proc. ACM Symposium on Operating Systems Principles (SOSP), Bolton Landing, New York, USA, 2003.

[22] D. Marshall, et al., "Advanced Server Virtualization," Taylor & Francis, New York, NY, 2006.

[23] K. Lim, et al., "Server Designs for Warehouse-Computing Environments," IEEE Micro, pp. 41-49, January/February 2009.

[24] T. Burd, et al., "A Dynamic Voltage Scaled Microprocessor System," IEEE Journal of Solid State Circuits, Vol 35, No 11, pp. 1571-1580, November 2000.

[25] A. Kumar, Y. Joshi, "Use of Airside Economizer for Data Center Thermal Management," Proc. Thermal Issues in Emerging Technologies (ThETA 2), Cairo, Egypt, 2008.

[26] P. Torcellini, S. Pless, M. Deru, and D. Crawley, "Zero Energy Buildings: A Critical Look at the Definitions", ACEEE Summer Study, Pacifica Grove, CA, USA, 2006