

Towards the Design and Operation of Net-Zero Energy Data Centers

Martin Arlitt, Cullen Bash, Sergey Blagodurov, Yuan Chen, Tom Christian, Daniel Gmach, Chris Hyser, Niru Kumari, Zhenhua Liu, Manish Marwah, Alan McReynolds, Chandrakant Patel, Amip Shah, Zhikui Wang, Rongliang Zhou
Hewlett-Packard Laboratories
Palo Alto, CA
USA
Email: cullen.bash@hp.com

ABSTRACT

Reduction of resource consumption in data centers is becoming a growing concern for data center designers, operators and users. Accordingly, interest in the use of renewable energy to provide some portion of a data center's overall energy usage is also growing. One key concern is that the amount of renewable energy necessary to satisfy a typical data center's power consumption can lead to prohibitively high capital costs for the power generation and delivery infrastructure, particularly if on-site renewables are used. In this paper, we introduce a method to operate a data center with renewable energy that minimizes dependence on grid power while minimizing capital cost. We achieve this by integrating data center demand with the availability of resource supplies during operation. We discuss results from the deployment of our method in a production data center.

KEY WORDS: data center, net-zero energy, renewable energy, sustainability

INTRODUCTION

In recent years, progress has been made in the development of techniques to reduce the environmental footprint of data centers. The first wave of such efforts started over a decade ago and was focused on optimizing the energy efficiency of different data center 'silos' (IT, cooling, power delivery). Examples of techniques that emerged from these efforts include the use of numerical modeling tools (like computational fluid dynamics) to minimize energy use by air-conditioners [1][2][3][4]; arrangement of IT equipment in hot and cold aisles [5][6]; energy-efficient server platform designs and robust power control features [7]; and the idea of direct DC power to data centers [8]. A second generation of environmental solutions for data centers, which focused on efficiency gains via the integration of different silos across the data center, began to emerge about five years ago. Examples of solutions from this wave include metrics for measuring holistic energy efficiency [9][10]; dynamic thermal management of air-conditioners based on the workload at the computer racks [11][12][13][14]; aisle containment [15][16]; thermally-aware as well as energy-aware virtualized workload placement [17][18]; and integration of the data center with local (external) ambient conditions through economizers or on-site renewable sources such as wind and solar photovoltaics [19].

All of the above techniques have enabled great strides in reducing the energy required by the data center. However,

these existing solutions continue to treat supply-side constraints such as energy or cooling availability independently from IT workload constraints (or flexibilities) in load scheduling. As a result, data center utilization continues to be low on average, with many machines working at reduced utilization while still consuming significant power and cooling resources [20]. We believe an opportunity exists for a new generation of data center management solutions that are focused around the integration of energy and cooling sources (supply) with the IT workload (demand) across the data center lifecycle.

In this paper, we introduce a 'net-zero energy' data center, designed and managed in a manner that uses on-site renewables to entirely offset the use of any non-renewable energy from the grid. Specifically, we describe how to combine the use of alternative energy sources with dynamic IT workload scheduling and integrated management techniques to improve overall data center utilization while allowing demand to be "shaped" according to resource availability. We illustrate implementation of such policy-based energy balancing through a working prototype that consumes zero net energy from the public utility grid while meeting all performance criteria and incurring minimal retrofit (capital infrastructure) expense.

The remainder of this paper is organized as follows. First, we discuss our design and methodology for achieving data center scale net-zero energy operation. We then describe the experimental testbed used to validate the methodology. Third, we discuss results obtained from the testbed. We conclude with a discussion of our current results and next steps.

DESIGN FOR NET-ZERO ENERGY

Figure 1 is a diagram of a typical data center architecture. Data centers include three pieces of critical infrastructure: power delivery, cooling resource delivery and IT (includes servers, storage and networking). The power and cooling infrastructures comprise the supply-side portion of the data center ecosystem, with the power infrastructure supplying energy resources in the form of electricity (primarily) and gas, and the cooling infrastructure providing cooling resources (e.g., chilled water, cool air, etc.). As the primary consumer of power and cooling resources in the data center, the IT infrastructure comprises the demand-side portion of the ecosystem. Most data centers obtain all of their primary power from the public utility grid. Facilities built with a higher

degree of reliability will also have on-site power generation for backup use—normally in the form of diesel generator sets.

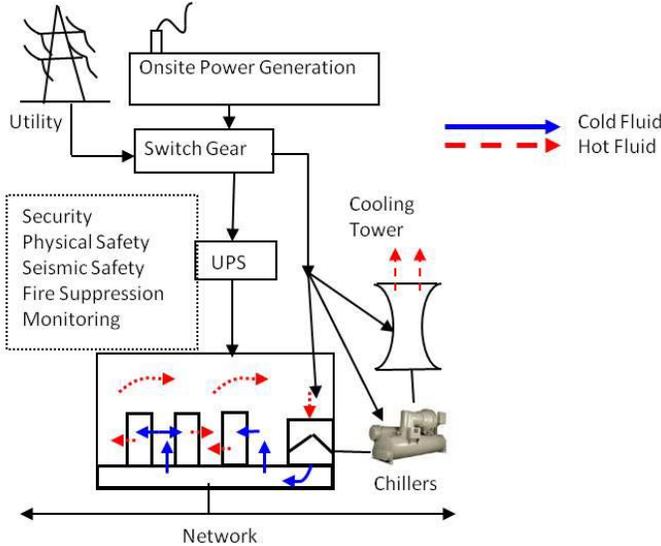


Figure 1: Typical Data Center Infrastructure

Cooling resources are typically provided via mechanical refrigeration and, although we place them within the supply-side portion of the ecosystem, can clearly have significant impact on the overall power demand of the data center. In order to reduce energy use and dependence on the grid simultaneously, a growing number of data centers today rely on micro-grids to both power and cool the data center [21]. A micro-grid is an interconnected grid that provides multiple means for generating and distributing a resource. It generally consists of one or more on-site generation means (e.g., photovoltaic panels, air-side economization) that can work in tandem with a more traditional source (e.g., public utility, mechanical refrigeration). Figure 2 is a block diagram that describes a supply-side micro-grid architecture for a data center. Energy storage can be added to further reduce grid-dependence and may be necessary for “islanded” operation (i.e., off-grid operation).

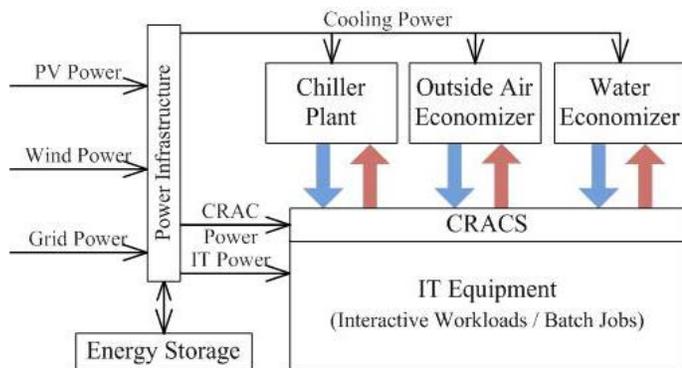


Figure 2: Micro-grid Based Infrastructure

In the remainder of this paper, we define a net-zero energy data center as one that consumes no net energy from a public

utility grid over the lifecycle of the data center. Equation 1 describes the net energy consumed by a data center.

$$E_{Net} = E_{Embedded} + E_{Op} - E_{Op, Renewable} \quad (1)$$

In Eq. 1, $E_{Embedded}$ describes the energy used to manufacture the data center. It includes the embedded energy in the materials that comprise the data center facility and the infrastructure contained therein. E_{Op} refers to the total energy consumed during operation of the data center and includes consumption from the power, cooling and IT infrastructures. $E_{Op, Renewable}$ refers to the renewable energy generated on-site that is used to offset or augment the energy used by the grid. Our goal is to drive E_{Net} to zero over the complete lifecycle of the data center.

METHODOLOGY FOR NET-ZERO ENERGY

There are two key considerations for achieving net-zero energy over a data center lifecycle according to Eq. 1. Lifecycle embedded energy must be low to reduce renewable offset requirements (and thereby reduce capital costs), and a management architecture must exist to balance data center runtime energy demand with supply-side constraints. This section will describe each of these areas.

Embedded Energy and Lifecycle Design

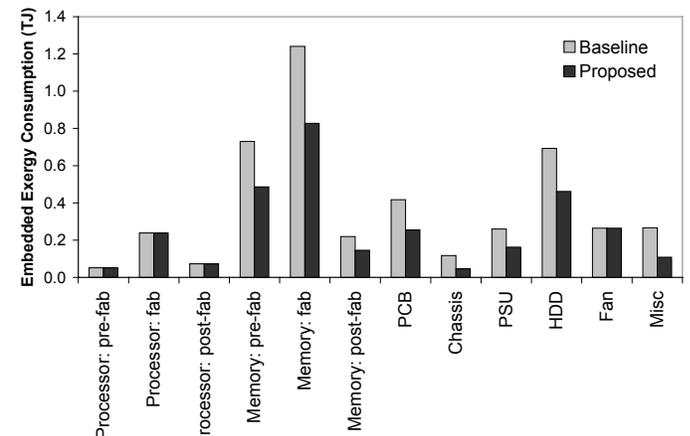


Figure 3: Comparison of embedded exergy consumption for a baseline bricks-and-mortar data center relative to a facility designed from a lifecycle perspective [30].

Life-cycle assessment (LCA) has been in practice for several decades [22][23]. It involves taking an end-to-end approach to assess the environmental impact from cradle-to-cradle, including the extraction of raw materials, manufacturing, transportation, operation, and disposal. Prior work has successfully developed lifecycle exergy consumption models for select IT systems [24] and has shown that optimizations based on lifecycle exergy consumption can map fairly well to optimizations based on other types of environmental criteria [25][26]. More recently, Shah et al. [27] presented an input-output model to obtain a rapid but approximate estimate of the end-to-end environmental footprint of data centers.

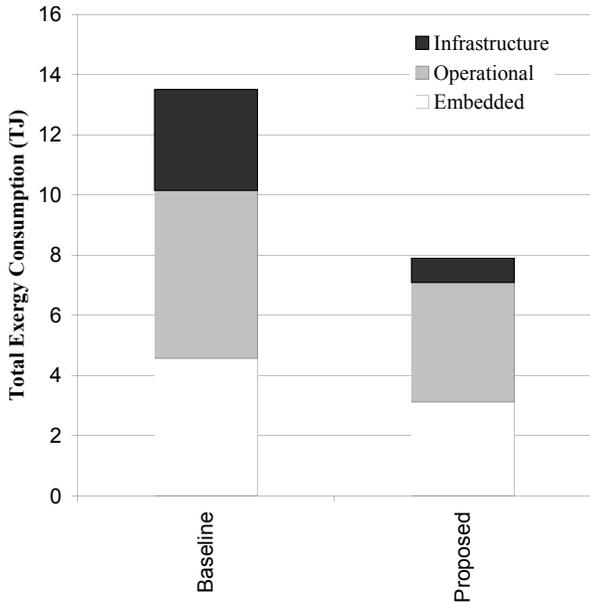


Figure 4: Overall exergy consumption across the lifecycle for a baseline and a lifecycle-based design.

In this paper, we propose using an LCA-based approach to dematerialize the data center as much as possible, by eliminating redundant materials and maximizing the reuse of materials. Doing so reduces the available energy required to build and maintain the data center thereby reducing E_{embedded} . As an example of the impacts from such a design philosophy, Figure 3 summarizes the results of a lifecycle-based data center design developed in prior work [28], which shows approximately 31% lower exergy consumption relative to typical bricks-and-mortar data centers. Disaggregating traditional IT systems so that common architectural functions such as memory and disk can be pooled together and shared allows for a reduced number of components, which in turn enables smaller printed circuit board sizes and corresponding reductions in the materials and energy used to manufacture the boards. Such disaggregated systems also allow for improved upcycling (i.e., reuse of components) at end-of-life. Similarly, eliminating sheet metal from the system packaging helps to reduce materials in the chassis. Interestingly, even though we did not explicitly focus on reducing operational exergy consumption, such a design also yielded benefits in runtime energy consumption (Figure 4). In particular, the IO hub and NIC components saw lower energy needs due to reduction in the number of components via dematerializing of systems. While the results summarized here are specific to the case study described in Meza et al. [28], such lifecycle-based optimization during the design stage is critical to meeting net-zero operational targets at minimal costs.

Operational Architecture

Figure 5 describes our architecture for the operation of a net-zero energy data center. There are four primary modules consisting of a) Prediction, b) Planning, c) Execution, and d) Verification and Reporting.

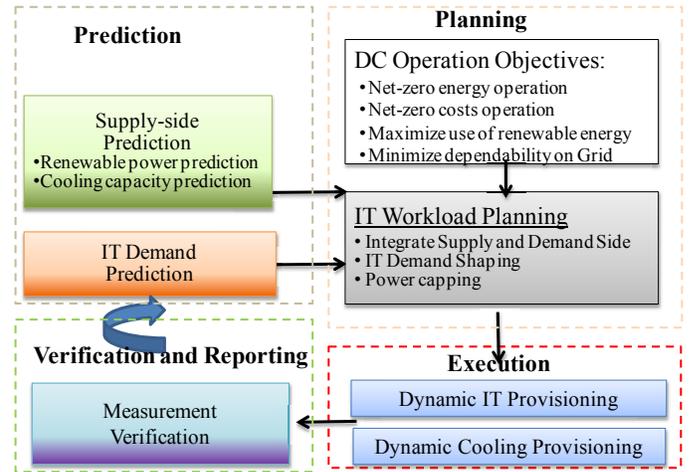


Figure 5: Net-Zero Energy Data Center Architecture

The Prediction module is comprised of a supply-side resource forecaster and a demand-side forecaster. The supply-side forecaster is used to predict the availability of resources within the power and cooling micro-grids. As an example, historical power traces, weather information and configuration information are fed into the Prediction module that then predicts future electricity generation capacity up to a day in advance. Similarly, external weather conditions (primarily temperature and humidity) and the cooling infrastructure configuration are used to predict the capacity of various cooling micro-grid components like air or water side economization [29]. Along with the supply-side predictor, a demand-side predictor uses historical workload traces to predict IT workload demand up to a day in advance. We categorize demand according to critical workloads that need to be executed upon arrival, and non-critical workloads that are delay tolerant (e.g., batch jobs, workload subject to spot pricing, etc.).

The Planning module combines output from the Prediction module with a) data center infrastructure considerations like IT capacity and cooling infrastructure capacity, b) high level operational goals like achieving net-zero energy operation, and c) performance goals defined through Service Level Agreements (SLAs). Optimization algorithms are then used to generate a workload schedule and resource provisioning plan that meets the operational goals subject to resource availability and performance constraints.

The Execution module is responsible for deploying the schedule developed in the Planning module. IT workload is managed via a Dynamic IT Provisioning component that manages workloads in real-time according to performance requirements, operational objectives and operational efficiency (including machine computational efficiency and cooling efficiency) [18]. It is tightly coupled to a Dynamic Cooling Provisioning component that manages the cooling micro-grid to optimize thermal performance and the operational efficiency of the data center cooling infrastructure [13].

A Verification and Reporting module is also included and is responsible for reporting results and insuring that the actual

execution is aligned with the plan. Any misalignment is addressed by updating the plan.

PROTOTYPE IMPLEMENTATION

The Net-Zero Energy architecture of Figure 5 is deployed in a production data center located in Palo Alto, CA. The data center is comprised of 85 racks of IT equipment and eight Computer Room Air Conditioning units (CRACs) of varying manufacture, as shown in Figure 6. The CRACs are provided chilled water via mechanical chillers. The supply temperature and blower speed of each CRAC are controlled via data obtained from a temperature sensor network distributed throughout the data center, as described in [13]. Each rack is outfitted with ten temperature sensors (five inlet, five exhaust) that report data every 15 seconds.

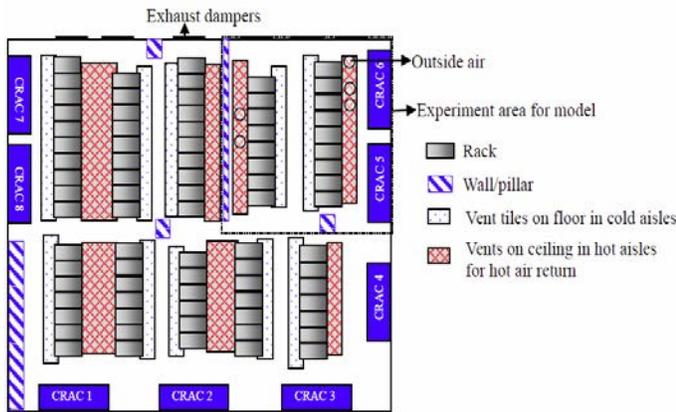


Figure 6: Data Center Layout

The data center is also equipped with an air-side economizer. The economizer delivers outside air to the CRAC returns preferentially based on CRAC return air temperature. Control of the economizer is integrated into the control of the CRAC units [13]. Along with a cooling micro-grid, the data center contains a power micro-grid consisting of a 134 kW peak photovoltaic array that is grid-tied.

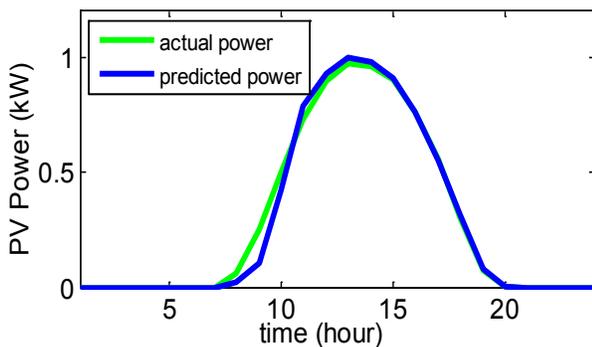


Figure 7: PV Generation Prediction

Our prototype deployment includes the Prediction, Planning and Execution modules as described above. Our initial experiments were conducted with our supply-side resources scaled down to the size of our IT test bed. Our IT test bed consisted of four BL465c G7 servers, each with two 12-core 1.8 Ghz processors and 64 GB of memory and a total of 48 KVM virtual machines. Our IT demand consisted of critical

and non-critical workloads. The critical demand was comprised of eight 3-tier Web applications (RUBiS—an e-Bay-like online auction), and the non-critical demand was comprised of 24 batch jobs that included scientific computing, animation and image processing, and financial analysis applications. The PV, cooling data, and interactive workload traces are scaled to our IT testbed capacity. We next illustrate the net-zero energy workload flow described in Figure 5 by showing results obtained from each module of our prototype implementation.

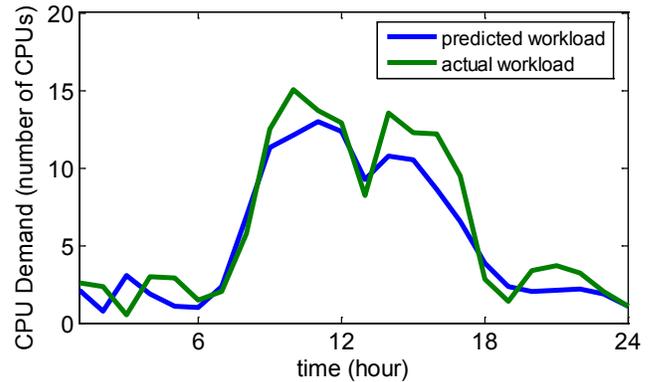


Figure 8: Workload Prediction

Figure 7 shows the predicted and actual values for the PV supply for September 10, 2011 scaled to our test bed. A k-nearest neighbor algorithm was used along with historical data and weather forecasts to develop the prediction [31]. The average PV prediction errors (i.e., the average difference between actual and predicted values) typically range from 5% to 20% and are dependent upon the occurrence of similar weather conditions in the past and accuracy of the weather forecast. A similar curve, not shown, is developed for the prediction of cooling capacity for the air-side economizer. In addition to supply-side resource prediction, we need information about the expected demand in order to complete a planning schedule. Our demand prediction algorithms can accommodate interactive and batch workloads and are based partially on historical data [29]. Figure 8 shows the predicted and actual workload based on CPU demand. Average prediction error is around 20%. Results show that our planning module can accommodate these levels of supply and demand prediction accuracy [29].

Once we have supply and demand side predictions, the Planning module develops a schedule for workload execution using workload demand shaping. Figure 9(a) shows the predicted power consumption of a workload demand prior to demand shaping over a 24 hour period. The figure shows power estimates from critical (interactive) and non-critical (batch) workloads, and for the cooling micro-grid. Note that the batch jobs are weighted towards the night-time when servers tend to be less utilized but were otherwise executed upon arrival. Figure 9(b) shows the expected power consumption according to a net-zero energy execution plan. The workload power consumption is displayed along with our predicted PV output. The total amount of workload is equivalent between Figure 9(a) and 9(b). Only the non-critical

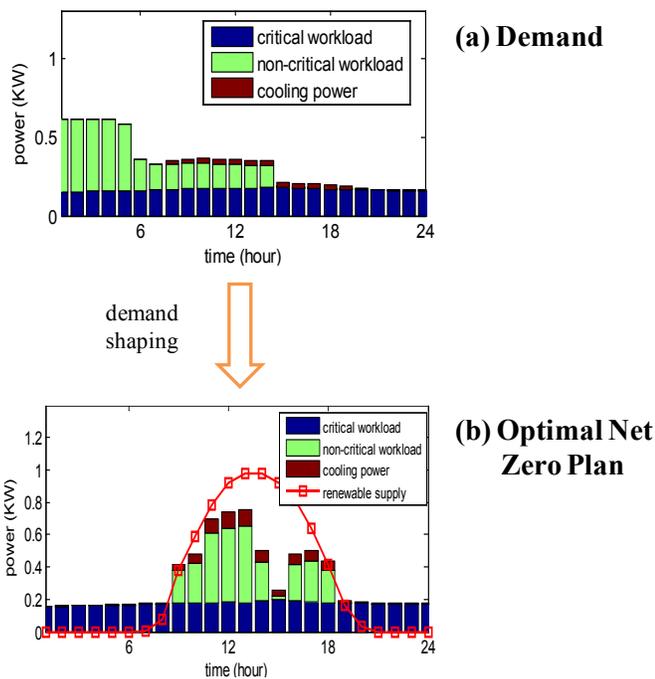


Figure 9: Workload Demand Shaping

workload is subject to demand shaping in Figure 9(b). The critical workload is executed upon arrival and in the absence of renewable power must utilize non-renewable resources (e.g., between 7pm and 7am). A renewable resource surplus (seen as the white space under the “renewable supply” curve in Figure 9(b)) in the plan accounts for the non-renewable use and enables achievement of a net-zero result. The reduction of workload execution at 3 PM seen in Figure 9(b) is a result of elevated external ambient temperatures that make cooling resources more expensive. Also, although embedded energy must be considered to achieve net-zero energy consumption across the data center lifecycle, we do not include it in this plan. As such, Figure 9(b) describes an *operational* net-zero energy plan. We note that a different ratio of non-critical to critical workload will result in a different plan and perhaps a different optimal supply-side mix. For example, if there are significantly more critical workloads than non-critical, then a mix of constant and intermittent renewable energy might be chosen to power the facility.

The schedule is then submitted to the data center run time workload manager for execution. Figure 10 shows the power trace of the plan vs. the actual power consumption of the IT and cooling infrastructure when applying the plan on our test bed. We note that the power consumption of the experiment lags the planned power consumption by a few minutes due to the time it takes to distribute or consolidate workloads when the number of physical servers change. We intend to incorporate this in the planning process as well as to speed up the boot process of physical machines and virtual machine migrations in our testbed in the future.

RESULTS AND DISCUSSION

To investigate the impact of supply-aware demand shaping on

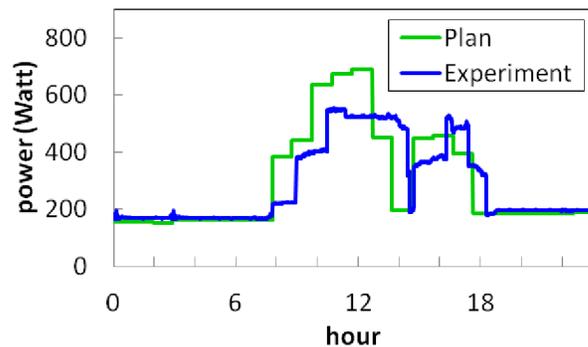


Figure 10: Net-zero Execution

overall resource consumption and to highlight the benefits of our solution, we examine different workload schedules through the Prediction and Planning stages of our solution (i.e., unlike the previous section, where the schedule was implemented and performance actually measured in our testbed, the below discussion is based on simulation results). The PV, cooling and workload traces are from our data center as described in the Prototype Implementation Section. Interactive workloads are deemed critical and their resource demands must be met. Non-critical workloads (batch jobs) can be rescheduled as long as they finish before their deadlines. The plan period is 24-hours and the capacity planner creates a plan for the next 24-hours, including the hourly capacity allocation for each workload.

Net-Zero Energy Case Study

The goal is to achieve net-zero operation (i.e., total power consumption \leq total renewable supply) while minimizing the energy exchange with the grid or storage energy. In other words, keep the dependence on grid low to reduce the recurring power cost or keep the dependency for energy storage low to reduce capital expense. In each case, the overall 24 hour workload was identical and the non-critical workload was adjusted via demand-shaping according to the plan. Figure 11 shows three different net-zero plans. *Optimal*—a net-zero energy plan that reshapes non-critical workloads to take full advantage of available renewable energy supply (Figure 11(a)). This schedule is generated by our capacity planner by adding a net-zero constraint to our optimization problem. This schedule does more work during the day when renewable energy is available. Additionally, some renewable energy is reserved to offset non-renewable energy used at night for critical workloads. This excess renewable power is reserved at 3pm when the outside air temperature peaks, thus minimizing the energy required for cooling. The other plans are *Night*—a “traditional” plan where non-critical workload is executed at night to ensure no interference with critical workloads occurs and to take advantage of idle machines (Figure 11(b)); *Flat*—a plan where the non-critical workload is averaged across the entire day (Figure 11(c)). Again, only non-critical demand is subject to demand shaping. In the *Optimal* plan, non-critical workloads are scheduled during the day time when PV power is available. Therefore, the IT and facility power demand matches the PV power generation, allowing the direct use of renewable energy. The *Night* plan schedules the non-critical

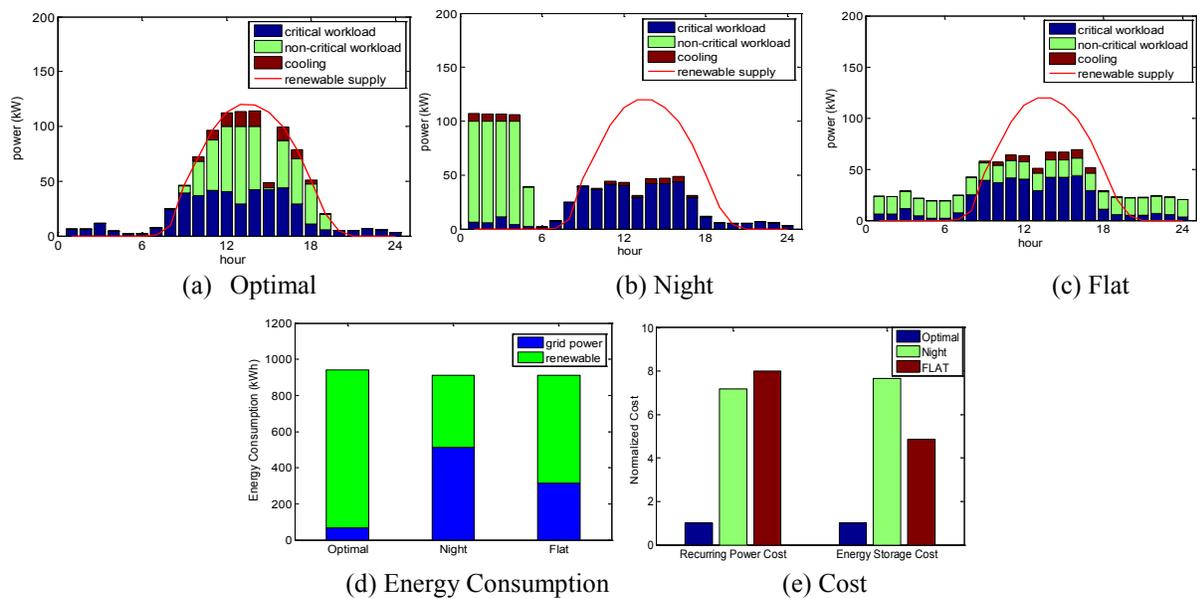


Figure 11: Comparison of Different Net-Zero Plans

workloads at midnight and uses all available IT infrastructures to finish them as early as possible. In contrast to that the *Flat* plan schedules the non-critical workloads across the full day.

Figure 11(d) compares the energy consumption of the three plans over the 24 hour period. All three plans achieve net-zero energy use. The combined renewable and non-renewable energy consumption is almost the same¹, but by increasing the utilization of the renewable supply, the *Optimal* plan uses 83% less grid energy than the *Night* plan and 65% less grid energy than the *Flat* plan. As the *Optimal* plan reduces the dependence of grid energy significantly, the recurring cost of grid power decreases significantly for a grid-tied data center. If we assume that unused renewable energy needs to be stored (e.g., for a non-grid-tied data center), then the demand on energy storage is much lower for the *Optimal* plan than for the *Night* and *Flat* plans. As in the grid-tied plan, a larger portion of the renewable energy is used directly; although being net-zero, *Night* and *Flat* require the storage of a large fraction of the generated PV power. By requiring less energy to be stored, the *Optimal* plan also helps to reduce the capital costs for a net-zero energy data center. Figure 11(e) compares the recurring grid power costs for a grid-tied data center and the energy storage costs for a non-grid-tied data center of each plan (normalized to the costs of the *Optimal* plan). Compared with the *Night* plan, the *Optimal* plan reduces the recurring grid power cost by 86% for the grid-tied data center and the energy storage cost by 87% for the non grid-tied data center. The *Optimal* plan reduces the recurring grid power and storage cost of the *Flat* plan by 88% and 79%, respectively. For non-grid-tied data centers that need energy storage, assuming a cost of 400 \$/kWh for energy storage, the savings on the energy storage expenditure for a data center with 1MW PV will be \$1.4 million over the *Night* plan and \$0.8 million over the *Flat* plan.

¹ There is a very small difference because of different cooling power consumption.

These results illustrate that *sizing* the supply side to achieve net-zero energy is not sufficient. As discussed in our use case, given the same PV installation, *operating* the data center differently can have a significant impact on the operational cost and/or the capacity of supply (e.g., energy storage size) that is required to achieve net-zero. Hence, optimizing workload management according to available supply is a critical step to achieve net-zero energy use while reducing overall environmental impact and cost of data centers .

Impact of Workload Mixes and Prediction Errors

Our solution improves energy efficiency and reduces grid power use by shaping non-critical workloads (e.g., batch jobs) at each timeslot. The more non-critical workloads a data center has, the more savings can be achieved. We study the savings of our solution over a traditional *Night* schedule that runs non-critical workloads at night. The impact of workload mixes on the reduction of grid power use is shown in Figure 12. We can see from the figure that the savings improve as the percentage of non-critical workload increases. However, our solution can reduce grid power use significantly even when the non-critical workload is only 20% of the total workload.

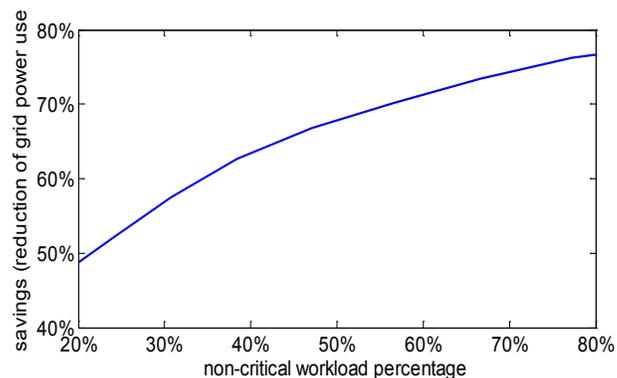
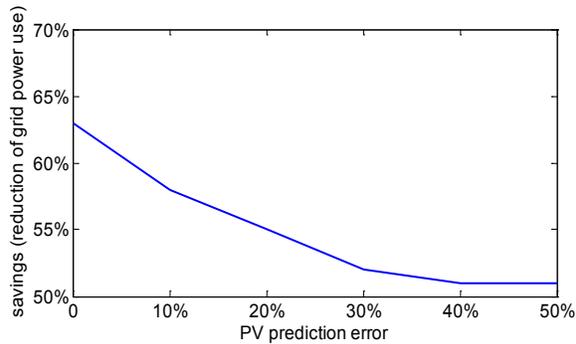
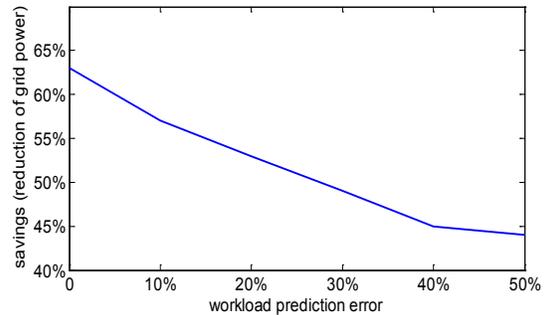


Figure 12: Impact of Workload Mixes



(a) Impact of PV Prediction Errors



(b) Impact of Workload Prediction Errors

Figure 13: Impact of Prediction Errors

We also evaluate the impact of prediction errors on the benefits of our solution. Figure 13(a) shows the savings (i.e., percentage of grid power reduction) with different PV prediction errors. The figure shows that the savings slightly drop as the prediction errors increase, but our solution can still reduce grid power usage more than 50% with a prediction error of 50%. Finally, we evaluate the impact of workload demand prediction errors. The results shown in Figure 13(b) are similar to the PV prediction errors. If predictions don't match the actual demand well, then our runtime controllers in the execution step will mitigate the errors in the plan via a feedback mechanism. Re-planning and adjustment will be done if the verification step indicates that an execution doesn't meet its goal. Further, if more grid power is used because the actual PV generation is less than predicted, the workload planning will reserve renewable energy from the next day to offset the additional non-renewable used.

CONCLUSION AND NEXT STEPS

In this paper we introduced an architecture for operating a data center in a manner that consumes no net energy from the public utility grid. We demonstrated how the architecture was used by applying it to an operational data center test bed. We then compared a net-zero energy solution with other common workload scheduling algorithms and quantified energy and cost savings.

Although our results are very encouraging, additional research is required to meet the overall goal of achieving net-zero energy over a data center's lifecycle. Future work will involve adding embedded energy cost to our architecture and scaling out our solution. We are also actively investigating how to co-locate critical and non-critical workloads simultaneously on a minimal set of IT equipment. This is an important next step in achieving a cost effective net-zero energy data center.

REFERENCES

1. M. Toulouse, D. Lettieri, V. Carey, C. Bash, A. Shah, "Computational and Experimental Validation of a Vortex-Superposition-Based Buoyancy Approximation for the COMPACT code in Data Centers", Proceedings of the ASME Intl. Mechanical Engineering Congress & Exposition (IMECE), 2011.
2. C. Patel, C. Bash, C. Belady, "Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications", Proceedings of InterPACK, Kauai, HI, 2001.
3. J. VanGilder, "Real-Time Data Center Cooling Analysis", Electronics Cooling, September 2011.
4. R. Schmidt, "Effect of Data Center Characteristics on Data Processing Equipment Inlet Temperatures", Proceedings of InterPACK, Kauai, HI, 2001.
5. S. Shrivastava, S. Sammakia, R. Schmidt, M. Iyengar, "Comparative Analysis of Different Data Center Airflow Management Configurations", Proceedings of InterPACK, San Francisco, CA, 2005.
6. C. Patel, R. Sharma, C. Bash, A. Beitelmal, "Thermal Considerations in Cooling Large Scale High Compute Density Data Centers", Proceedings of ITherm, 2002.
7. 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, 2008.
8. M. Ton, B. Fortenbery, W. Tschudi, "DC Power for Improved Data Center Efficiency", Lawrence Berkeley National Laboratory, <http://hightech.lbl.gov/dc-powering/>, 2008.
9. A. Shah, C. Patel, C. Bash, R. Sharma, R. Shih, "Impact of rack-level compaction on the data center cooling ensemble", Proceedings of ITherm, 2008.
10. C. Patel, A. Shah, "Cost Model for Planning, Development and Operation of a Data Center", Hewlett-Packard Laboratories Technical Report, HPL-2005-107, June 2005.
11. C. Bash, C. Patel, R. Sharma, "Dynamic Thermal Management of Air Cooled Data Centers", Proceedings of ITherm, San Diego, CA, 2006.
12. M. Patterson, R. Weidmann, M. Leberecht, M. Mair, R. Libby, "An Investigation Into Cooling System Control Strategies for Data Center Airflow Containment Architectures", Proceedings of InterPACK, Portland, OR, 2011.
13. R. Zhou, Z. Wang, T. Christian, A. McReynolds, C. Bash, "Optimization and Control of Cooling Microgrids for Data Centers", Proceedings of ITherm, San Diego, CA, 2012.
14. V. Sundaralingam, P. Kumar, Y. Joshi, "Server Heat Load Based CRAC Fan Controller Paired With Rear Door Heat

- Exchanger”, Proceedings of InterPACK, Portland, OR, 2011.
15. T.J., Breen, E. Walsh, J. Punch, A. Shah, C. Bash, B. Rubenstein, S. Heath, N. Kumari, “From Chip to Cooling Tower Data Center Modeling: Influence of Air-Stream Containment on Operating Efficiency”, Proceedings of ASME/JSME AJTEC, 2011.
 16. P. Kumar, Y. Joshi, M. Patterson, R. Steinbrecher, M. Mena, “Cold Aisle Air Distribution in a Raised Floor Data Center”, Proceedings of InterPACK, Portland, OR, 2011.
 17. H. Chen, M. Kesavan, K. Schwan, A. Gavrilovska, P. Kumar, Y. Joshi, “Spatially-Aware Optimization of Energy Consumption in Consolidated Data Center Systems”, Proceedings of InterPACK, Portland, OR, 2011.
 18. Y. Chen, D. Gmach, C. Hyser, Z. Wang, C. Bash, C. Hoover, S. Singhal “Integrated management of application performance, power and cooling in data centers”, Proceedings of NOMS, 2010.
 19. B. Watson, A. Shah, M. Marwah, C. Bash, R. Sharma, C. Hoover, T. Christian, C. Patel, “Integrated Design and Management of a Sustainable Data Center”, Proceedings of InterPack, San Francisco, CA, 2009.
 20. L. Barroso, U. Holzle, “The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines”, *Synthesis Lectures on Computer Architecture*, Morgan and Claypool, 2009.
 21. R. Sharma, C. Bash, M. Marwah, C. Patel, T. Christian, “Microgrids: A New Approach to Supply-Side Design for Data Centers”, Proceedings of the ASME Intl. Mechanical Engineering Congress & Exposition (IMECE), 2009.
 22. ISO 14040: “Environmental management—Life Cycle Assessment—Principles and framework”, Intl. Organization for Standardization, 2006.
 23. H. Bauman, A.-M. Tillman. „The Hitch Hiker’s Guide to LCA”, Studentlitteratur AB, 2004.
 24. C. R. Hannemann, V. P. Carey, A. J. Shah, C. D. Patel. “Lifetime exergy consumption of enterprise servers”, Intl. Journal of Exergy, Vol. 7, No. 4, pp. 439-453, 2010.
 25. A. J. Shah, C. D. Patel, V. P. Carey. “Exergy-based metrics for sustainable design”, Fourth Intl. Exergy, Energy and Environment Symposium (IEEEES-4), Sharjah, UAE, 2009.
 26. I. Dincer, M. Rosen. “Exergy: Energy, Environment and Sustainable Development”, Elsevier, 2007.
 27. A. Shah, C. Bash, R. Sharma, T. Christian, B.J. Watson, C. Patel, “Evaluating Life-Cycle Environmental Impact of Data Centers”, Journal of Electronic Packaging, Vol. 133, Article 031005, 2011.
 28. J. Meza, R. Shih, A. Shah, P. Ranganathan, J. Chang, C. Bash, “Lifecycle-Based Data Center Design”, Proceedings of the ASME Intl. Mechanical Engineering Congress & Exposition (IMECE), 2010.
 29. Z. Liu, Y. Chen, C. Bash, D. Gmach, A. Wierman, M. Marwah, Z. Wang, “Renewable and Cooling Aware Workload Management for Sustainable Data Centers”, Proceedings of ACM Sigmetrics 2012, June, 2012.
 30. J. Meza, R. Shih, A. Shah, P. Ranganathan, J. Chang, C. Bash, “Lifecycle-Based Data Center Design”, Proceedings of the ASME Intl. Mechanical Engineering Congress & Exposition (IMECE), 2010.
 31. Dudani, S. A. 1976. The Distance-Weighted k-Nearest-Neighbor Rule. IEEE Transactions on Systems, Man and Cybernetics SMC-6(4):325–327, 1976.