

Figure 5: Impact of parameter T .

ing/discharging is $\eta = 0.8$ [20]. We set the grid energy limit as $P_{grid} = 2MW$ [36].

Algorithms for Comparison: We compare MultiGreen with the offline optimal algorithm (*Optimal*) and an online algorithm *Green* that solely leverages renewable energy production, without exploiting time-varying electricity prices. The Green algorithm also tries to maximize the usage of renewable energy, *i.e.*, leveraging UPS battery to store excess renewable energy production for future need. However, the Green algorithm ignores the two-timescale grid markets, and does not store grid energy in UPS when the electricity prices are low and supply energy when the electricity prices are high.

5.2 Analysis of Sensitivity on Critical Factors

From Theorem 2, we note that the performance of MultiGreen depends on parameters V and T , battery capacity and the energy prices in the two-timescale grid markets. We conduct sensitivity analysis on these critical factors to characterize their impact on the DPSS operational cost.

5.2.1 Impact of Control Parameter V

As shown in Fig. 4, to simulate a 1-day-ahead power market, we fix T to be 24 time slot and each fine-grained time slot is 1 hour, *i.e.*, $N_T = 24$. We conduct experiments with different values of V , which show that as V increases from 0.5 to 100, MultiGreen achieves a time-averaged cost that becomes closer to the optimal solution. This quantitatively confirms Theorem 2 that *MultiGreen can approach the optimal solution within a diminishing gap of $O(1/V)$* . In contrast, the Green algorithm has a constant cost that is irrelevant with V . Interestingly, the crossover between cost curves of MultiGreen and Green clearly captures the trade-off between the average operational cost and constraint satisfaction. When $V < 7.48$, MultiGreen has a higher cost and a higher level of constraint satisfaction than Green. On the other hand, when $V > 7.48$, due to more frequent battery charging/discharging, MultiGreen has a lower cost and a lower level of constraint satisfaction than Green. By choosing an appropriate value of V , *e.g.*, $V = 10$, MultiGreen can achieve a significantly lower cost compared with Green while guaranteeing acceptable satisfaction of constraints on datacenter availability and UPS lifetime.

5.2.2 Impact of Coarse-grained Time Frame T

In Fig. 5, we fix V to be 10 and vary T from 3 time slots (3 hours) to 144 time slots (6 days), which is a sufficient-long range for exploring the impact of different timescales of the grid’s long-term market. We observe that T has rel-




Figure 6: Impact of battery capacity and grid market structure. Two Markets—TM, Real-Time Market—RTM, No Battery—NB.

atively less impact on the cost of operating the DPSS. The fluctuation of the time-averaged cost is more notable when T becomes longer. The rationale is that the term B_ϵ in Theorem 3 is proportional to T , which means that the uncertainties of energy demand and renewable energy increase with the increase of T . Nevertheless, the time-averaged cost only fluctuates within $[-9.7\%, +8.5\%]$. This corroborates Theorem 3 that, even with *infrequent decisions* of the DPSS operations, MultiGreen can still achieve significant cost reduction.

5.2.3 Impact of Battery Capacity and Grid Markets

In Fig. 6, we compare the time-averaged total cost under different battery sizes ($M_{UPS} \in \{0, 0.25, 0.5, 1\}MWh$) over the 31-day period with $V = 10$ and $T = 24$. It shows that the time-averaged total cost decreases with the increase of the UPS battery capacity. The rationale is that an UPS with larger capacity can store more superfluous renewable energy generated, or more energy purchased from the grid when the price is low, to serve the demand, resulting in lower overall costs.

In Fig. 6, we also compare the case with energy purchase in two-timescale markets with the case where only the real-time market exists, both with $V = 10, T = 24, M_{UPS} = 0.5MWh$. We can observe that the existence of the grid’s long-term market can bring in additional cost reduction. The reason is that DPSS can purchase certain amount of energy beforehand in the grid’s long-term market with relative lower prices

In addition, we can observe that even without the UPS battery, the MultiGreen algorithm with the two-timescale markets can reduce the cost by 10.06%, compared to the Green algorithm. With two markets, when we increase the battery size from 0 to $1MWh$, the average operational cost reduction ranges from 10.06% to 34.21%. The benefit brought by energy storage is higher than that of exploiting the two markets. When the battery size is large enough, MultiGreen can approach the optimal offline algorithm.

5.3 Characterizing Algorithm Robustness

As mentioned in Sec. 3, our MultiGreen algorithm approximates the future queue statistics as the current values. Now we explore the influence of approximation errors on the performance of MultiGreen. We add a random approximation error to the datacenter energy demand, solar energy generation and energy prices, *e.g.*, uniformly distributed $\pm 50\%$

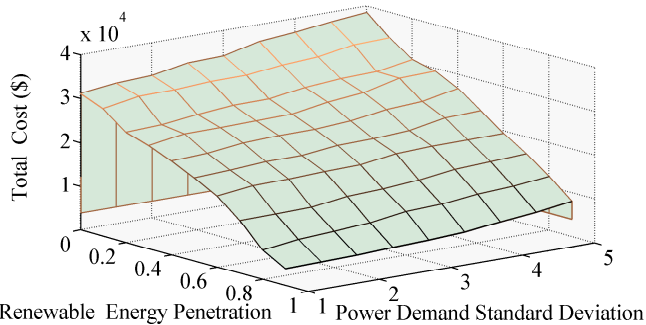


Figure 7: DPSS operational cost at various levels of renewable energy penetration and energy demand variation.

errors [39]. We let MultiGreen make all the control decisions based on the data set with such random errors under different values of V . In Fig. 8, we show the differences in percentage between the DPSS operational costs achieved with approximated values and the results we obtained using the original traces. We observe that the difference fluctuates within $[-1.3\%, 2.1\%]$ for all values of V . Thus, MultiGreen is robust to inaccurate future information.

Further, we study the impacts of renewable energy penetration (the percentage of renewable energy in the total datacenter energy supply) and the variation of datacenter energy demand on the total cost. In Fig. 7, x-axis represents *renewable energy penetration* in the range of $[0, 100\%]$. Y axis represents the standard deviation of the demand, *i.e.*, $\sqrt{\sum_{t=0}^{KT-1} [d(t) - \mathbb{E}(\vec{d})]^2 \times p_{d(t)}}$, where $\mathbb{E}(\vec{d})$ is the expectation of the series of demand $d(t)$ over time length $t \in [0, KT]$, and $p_{d(t)}$ is the distribution probability of $d(t)$. We assume that the random variable of the datacenter energy demand is uniformly distributed ($p_{d(t)} = 1/KT$). As expected, Fig. 7 shows that with the increase of penetration of renewable energy, the DPSS operational cost decreases significantly. The rationale is that renewable energy is harvested cost-free (we do not consider the construction cost). In contrast, as the variation of demand increases, the operational cost increases slightly. The rationale is that intensive variation incurs large approximation errors.

6. RELATED WORK

In this section, we discuss the research most pertinent to this work as follows. The first category of works is exploiting renewable energy in datacenters. Many large IT companies recently consider greening their datacenters with renewable energy [7, 8, 10, 22–24, 35, 41]. However, the intermittent nature of renewable energy poses significant challenges to make use of them. Some works make the traffic “follow the renewables” to execute workload when/where renewable energy is available [10, 22–24, 41] or carbon footprint is low [7]. However, these approaches require prediction of renewable energy production when scheduling workload, or sacrifice performance to avoid wasting renewable energy. Other works supply renewable energy to deferrable loads to align demand with intermittent available renewable energy [27, 29, 30]. But they are from the prospective of renewable energy providers and do not consider energy storage and multiple markets in the smart grid.




Figure 8: The impact of approximation errors in operational cost reduction.

The second category of works is leveraging energy storage in datacenters. Recently, UPS shows its benefits to reduce electricity costs in datacenters [11–13, 36, 38]. Datacenters can store energy in the UPS when energy prices are low and discharge UPS when prices are high, to reduce the power drawn from the grid [13, 36]. Moreover, UPS can shave peaks [11, 12]. During periods of low demand, UPS batteries store energy, while stored energy can be used to temporarily augment the grid supply during hours of peak load. However, these works focus on studying the benefits of UPS battery for power cost reduction, and no renewable energy and grid markets are considered. On the contrary, we leverage UPS to study how to manage multiple power supplies of a datacenter in an integrated way.

Third stream of works is on multiple timescale dispatch, pricing and scheduling in smart grid. Nair *et al.* [1] studied the optimal energy procurement from long-term, intermediate, and real-time markets under intermittent renewable energy supplies. Jiang *et al.* [19] proposed an optimal multi-period power procurement and demand response algorithm without energy storage. “Risk-limiting-dispatching” is proposed in [37] to manage integrated renewable energy. However, the above three approaches assume that the demand can be known ahead. Jiang *et al.* [21] solved the optimal day-ahead procurement and real-time demand response problem using dynamic programming, while He *et al.* [15] formulated the multi-timescale power dispatch and scheduling problem as a Markov decision problem. Both these approaches need substantial system statistics and are computationally expensive. We mitigate these disadvantages by applying two-stage Lyapunov optimization that makes online control decisions without *a priori* knowledge or any stationary distribution of energy prices, demand and supply. Recently the authors in [13, 32, 33, 39, 41] distributed requests across multiple data centers to reduce electricity costs by leveraging both time diversity and location diversity of electricity prices in the smart grid. In contrast, we study how to reduce the operational cost in a datacenter powered by multiple power sources rather than how to distribute requests across datacenters.

In addition, interest has been growing in power management in smart grids and datacenters using Lyapunov optimization [5, 6, 9, 26, 42]. On smart grids, several works have proposed optimal power management based on single-stage Lyapunov optimization. However, they either focused on managing individual household demand [14] or did not consider the interaction between renewable energy and energy storage [14, 19, 20, 27]. In contrast, we manage the uncertain datacenter demand and multi-source energy supply in

a systematic fashion using two-stage Lyapunov optimization. Although [34, 39] have used two-stage Lyapunov to design a two-timescale algorithm and a T -Step Lookahead algorithm, both of them study how to schedule jobs or distribute requests in solely grid-powered geographical datacenters rather than how to supply multi-source energy in a datacenter with uncertain demand.

7. CONCLUSION

In this paper, we study an important problem of how to minimize the operational cost of datacenters by using multiple energy resources. We propose MultiGreen, an online control algorithm applying the two-stage Lyapunov optimization technique, which optimally schedules multiple energy supply sources to power a datacenter, in a cost minimizing fashion. Without requiring *a priori* knowledge of system statistics, MultiGreen can deliver reliable energy to datacenters while minimizing the operational cost in the datacenter's long-run operation. Both mathematical analyses and trace-driven evaluations demonstrate the optimality and robustness of MultiGreen. Especially, it can approach the offline optimal cost within a diminishing gap of $O(1/V)$, which is mainly decided by the UPS battery capacity, grid market structure and DPSS operation frequency for energy purchasing and UPS charging/discharging.

8. ACKNOWLEDGEMENTS

The Corresponding Author is Fangming Liu. The research was supported by a grant from The National Natural Science Foundation of China (NSFC) under grant No.61133006.

9. REFERENCES

- [1] S. Adlakha and A. Wierman. Energy Procurement Strategies in the Presence of Intermittent Sources. *Technical report*, 2012.
- [2] D. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific Belmont, MA, 1995.
- [3] G. Dantzig. *Linear Programming and Extensions*. Princeton university press, 1998.
- [4] N. Deng, C. Stewart, and J. Li. Concentrating Renewable Energy in Grid-tied Datacenters. In *Proc. of IEEE International Symposium on Sustainable Systems and Technology (ISSST)*, May 2011.
- [5] W. Deng, F. Liu, H. Jin, and X. Liao. Online Control of Datacenter Power Supply Under Uncertain Demand and Renewable Energy. In *Proc. of IEEE ICC*, June 2013.
- [6] W. Deng, F. Liu, H. Jin, and C. Wu. SmartDPSS: Cost-Minimizing Multi-source Power Supply for Datacenters with Arbitrary Demand. In *Proc. of ICDCS*, July 2013.
- [7] P. Gao, A. Curtis, B. Wong, and S. Keshav. It's Not Easy Being Green. In *Proc. of ACM SIGCOMM*, Aug. 2012.
- [8] Y. Gao, Z. Zeng, X. Liu, and P. R. Kumar. The Answer Is Blowing in the Wind: Analysis of Powering Internet Data Centers with Wind Energy. In *Proc. of INFOCOM (Mini Conference)*, Apr. 2013.
- [9] L. Georgiadis, M. Neely, M. Neely, and L. Tassiulas. *Resource Allocation and Cross-layer Control in Wireless Networks*. Now Pub, 2006.
- [10] Í. Goiri, K. Le, T. Nguyen, J. Guitart, J. Torres, and R. Bianchini. GreenHadoop: Leveraging Green Energy in Data-Processing Frameworks. In *Proc. of EuroSys*, Apr. 2012.
- [11] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar. Benefits and Limitations of Tapping into Stored Energy for Datacenters. In *Proc. of ACM ISCA*, June 2011.
- [12] S. Govindan and B. Urgaonkar. Leveraging Stored Energy for Handling Power Emergencies in Aggressively Provisioned Datacenters. In *Proc. of ACM ASPLOS*, Mar. 2012.
- [13] Y. Guo, Z. Ding, Y. Fang, and D. Wu. Cutting Down Electricity Cost in Internet Data Centers by Using Energy Storage. In *Proc. of GLOBECOM*, Dec. 2011.
- [14] Y. Guo, M. Pan, and Y. Fang. Optimal Power Management of Residential Customers in the Smart Grid. *IEEE Transactions on Parallel and Distributed Systems*, 23(9):1593–1606, 2012.
- [15] M. He, S. Murugesan, and J. Zhang. Multiple Timescale Dispatch and Scheduling for Stochastic Reliability in Smart Grids with Wind Generation Integration. In *Proc. of INFOCOM*, Apr. 2011.
- [16] http://en.wikipedia.org/wiki/Google_Energy.
- [17] <http://www.datacenterknowledge.com/archives/2012/10/22/amazon-cloud-outage-affecting-many-sites/>.
- [18] <http://www.greenpeace.org>. How Clean is Your Cloud. 2012.
- [19] L. Huang, J. Walrand, and K. Ramchandran. Optimal Power Procurement and Demand Response with Quality-of-usage Guarantees. *Arxiv preprint arXiv:1112.0623*, 2011.
- [20] L. Huang, J. Walrand, and K. Ramchandran. Optimal Demand Response with Energy Storage Management. *Arxiv preprint arXiv:1205.4297*, 2012.
- [21] L. Jiang and S. Low. Multi-period Optimal Procurement and Demand Responses in the Presence of Uncertain Supply. In *Proc. of IEEE Conference on Decision and Control (CDC)*, 2011.
- [22] C. Li, A. Qouneh, , and T. Li. iSwitch: Coordinating and Optimizing Renewable Energy Powered Server Clusters. In *Proc. of ISCA*. ACM, June 2012.
- [23] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser. Renewable and Cooling Aware Workload Management for Sustainable Data Centers. In *Proc. of SIGMETRICS*, June 2012.
- [24] Z. Liu, M. Lin, A. Wierman, S. Low, and L. Andrew. Geographical Load Balancing with Renewables. In *Proc. of ACM GreenMetrics*, June 2011.
- [25] MIDC. <http://www.nrel.gov/midc/>.
- [26] M. Neely. Stochastic Network Optimization with Application to Communication and Queueing Systems. *Synthesis Lectures on Communication Networks*, 3(1):1–211, 2010.
- [27] M. Neely, A. Tehrani, and A. Dimakis. Efficient Algorithms for Renewable Energy Allocation to Delay Tolerant Consumers. In *Proc. of SmartGridComm*, Oct. 2010.
- [28] NYISO. <http://www.nyiso.com>.
- [29] A. Papavasiliou. *Coupling Renewable Energy Supply*

- with Deferrable Demand*. PhD thesis, University of California, Berkeley, 2012.
- [30] A. Papavasiliou and S. Oren. Supplying Renewable Energy to Deferrable Loads: Algorithms and Economic Analysis. In *Proc. of Energy Society General Meeting*, July 2010.
- [31] A. Qureshi. *Power-Demand Routing in Massive Geo-Distributed Systems*. PhD thesis, Massachusetts Institute of Technology, 2010.
- [32] L. Rao, X. Liu, L. Xie, and W. Liu. Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi-electricity-market Environment. In *Proc. of INFOCOM*, Mar. 2010.
- [33] L. Rao, X. Liu, L. Xie, and Z. Pang. Hedging Against Uncertainty: A Tale of Internet Data Center Operations Under Smart Grid Environment. *IEEE Transactions on Smart Grid*, 2(3):555–563, 2011.
- [34] S. Ren, Y. He, and F. Xu. Provably-Efficient Job Scheduling for Energy and Fairness in Geographically Distributed Data Centers. In *Proc. of ICDCS*, June 2012.
- [35] C. Stewart and K. Shen. Some Joules Are More Precious Than Others: Managing Renewable Energy in the Datacenter. In *Proc. of the Workshop on Power Aware Computing and Systems*, 2009.
- [36] R. Urgaonkar, B. Urgaonkar, M. Neely, and A. Sivasubramanian. Optimal Power Cost Management Using Stored Energy in Data Centers. In *Proc. of ACM SIGMETRICS*, June 2011.
- [37] P. Varaiya, F. Wu, and J. Bialek. Smart Operation of Smart Grid: Risk-limiting Dispatch. in *Proc. of the IEEE*, 99:(1):40–57, 2011.
- [38] D. Wang, C. Ren, A. Sivasubramanian, B. Urgaonkar, and H. Fathy. Energy Storage in Datacenters: What, Where, and How much? In *Proc. of SIGMETRICS*, June 2012.
- [39] Y. Yao, L. Huang, A. Sharma, L. Golubchik, M. Neely, et al. Data Centers Power Reduction: A Two Time Scale Approach for Delay Tolerant Workloads. In *Proc. of IEEE INFOCOM*, Mar. 2012.
- [40] R. Zavadil, J. King, N. Samaan, J. Lamoree, M. Ahlstrom, B. Lee, D. Moon, C. Finley, D. Savage, R. Koehnen, et al. Final Report–2006 Minnesota Wind Integration Study. *The Minnesota Public Utilities Commission*, 2006.
- [41] Y. Zhang, Y. Wang, and X. Wang. GreenWare: Greening Cloud-Scale Data Centers to Maximize the Use of Renewable Energy. In *Proc. of ACM Middleware*, Dec. 2011.
- [42] Z. Zhou, F. Liu, H. Jin, B. Li, B. Li, and H. Jiang. On Arbitrating the Power-Performance Tradeoff in SaaS Clouds. In *Proc. of IEEE INFOCOM*, Apr. 2013.