# Sustainable Computing Lab
# Research Overview

David Irwin
University of Massachusetts Amherst
College of Engineering
Department of Electrical and Computer Engineering
July, 2022
https://www.sustainablecomputinglab.io

Broadly, the lab's research focuses on designing experimental computer systems with an emphasis on improving efficiency and sustainability. Specifically, the lab's research addresses problems with improving energy-efficiency, carbon-efficiency, and cost-efficiency in two major domains: cloud computing and energy systems. Modern society's dependence on massive amounts of, largely "dirty," energy is not sustainable. Thus, meeting society's future energy needs, while mitigating its economic and environmental impact, is a critical challenge. Computer systems represent both a problem and potential solution to developing a sustainable society. For example, the cloud data centers and platforms that power our information-based economy are one of the fastest growing segments of industrial energy usage, and now account for an estimated 1-2% of U.S. electricity consumption. Yet, computer systems are also key to developing a "smart" grid that optimizes energy-efficiency, minimizes energy costs, and accommodates a high penetration of renewable energy.

The lab's research methodology is experimental in nature and focuses on designing, modeling, deploying, and analyzing system prototypes, as well as collecting and analyzing real-world data, to both identify and solve problems in the design and operation of existing systems. The importance and prominence of the lab's research has steadily increased over time, especially recently, as the impacts of climate change have become more visible and acute. There is now a broad societal consensus that, moving forward, designing efficient and sustainable cloud and energy systems will be critical to satisfying the strict emissions targets necessary to mitigate the worst outcomes of a warmer climate, i.e., by preventing earth's average temperature from rising more than 1.5-2°C. Computing research in sustainability is particularly important, since computer systems increasingly drive the automated operation and management of large-scale societal infrastructure, e.g., data centers, buildings, electric grids, etc.

The lab's research approach has been unique in that I have structured it "horizontally" across a wide range of applications in efficiency and sustainability that span many technical "verticals," including distributed systems and networking, operating systems and virtualization, security and privacy, applied data science and machine learning, sensor networks and the Internet-of-Things, analytical performance modeling, and clean energy systems. Our research also often combines domain-specific knowledge from other disciplines, such as power systems and building science, with the technical computing areas above to improve upon domain-agnostic solutions. We have pursued the research agenda above through a large number of funded projects of varying size and scope, ranging from small focused projects to large multi-institutional inter-disciplinary projects. Our results have yielded experimental prototype systems, techniques that have been deployed at large scales to improve the efficiency of production cloud data centers, widely-used open-source software and datasets, a large volume of publications in high-quality peer-reviewed conferences and journals, and significant technical impact on the community by any quantitative measure, including awards and nominations, funding, citations, and successful Ph.D. students. The following provides a brief overview of the lab's recent research contributions and future plans. This overview focuses on our research over roughly the past 5 years.

## 1 Research Overview

A common challenge underlies much of the lab's work: designing systems that are capable of gracefully handling uncontrollable resource "dynamics," which require them to continuously adapt to changes in their resources' availability and characteristics. These dynamics arise from uncontrollable aspects of our physical or economic environment, such as changes in weather, carbon emissions, user behavior, or prices. As one example, the mix of generators the electric grid uses to satisfy its variable demand both changes over time and differs by location. Since generators have different carbon-efficiencies, i.e., the amount of carbon they emit for each unit of energy they produce, the carbon emissions from generating grid power also varies over time and by location. Thus, systems that manage infrastructure, e.g., for data centers or buildings, can improve their carbon-efficiency by responding to these dynamics to shift their energy usage to when and where low-carbon energy is available, e.g., by storing and releasing energy from batteries or migrating load. Similar types of uncontrollable resource dynamics manifest themselves in numerous other ways, including variable spot prices for energy and computation, as well as variable renewable energy generation. Computation is particularly well-positioned to respond to these resource dynamics, since it often has significant spatial, temporal, and performance flexibility, which enables shifting the location, time, and intensity of its execution to better align with low carbon emissions, low prices, or high renewable generation. In addition, computation can also leverage numerous software-based fault-tolerance techniques, including checkpointing, replication, and recomputation, to continue execution despite an unexpected lack of resource availability, e.g., due to high carbon emissions, high prices, or low renewable generation, which may require deactivating servers.

Designing systems capable of handling uncontrollable resource dynamics raises many research questions that have animated, and continue to animate, my research. For example: How much information about these resource dynamics should systems expose to applications? How should applications respond to rapid, frequent, and uncontrollable variations in their resources' characteristics and availability? What abstractions should systems support to enable applications to better respond to resource dynamics? How well can we model and predict resource dynamics, even though we cannot control them? How can we use such models and predictions to improve systems' performance, and do they have any privacy implications? In addressing these questions, the lab's work has made substantial contributions to optimizing the efficiency and sustainability of cloud computing and energy systems. Below, we summarize the lab's recent contributions in both areas.

### 1.1 Optimizing Cloud Computing

As noted above, the lab's research has addressed improving cloud data center efficiency and sustainability.

**Efficiency.** The installed capacity of cloud data centers is continuing to grow rapidly. Thus, improving data center efficiency is important in maximizing the benefit of the existing capacity, as well as delaying the need to build new capacity. The lab has made significant contributions to improving the cost- and energy-efficiency of cloud data centers and applications. Note that cost- and energy-efficiency are related metrics, since energy is a major cost in operating a data center. Much of the lab's work has focused on enabling cloud users to optimize their costs by leveraging the different resource purchasing options offered by cloud platforms, which differ in their cost, performance, and availability. In particular, we have designed many systems and applications that exploit "transient" virtual machines (VMs), which have a low price but are highly unreliable [25, 24, 27, 45, 33]. For example, we designed TR-Kubernetes, a modified version of Kubernetes, a popular software framework for managing data centers, that optimizes the cost of executing mixed interactive/batch workloads on transient VMs, while also enforcing arbitrary availability requirements specified by interactive services despite transient VM unavailability. This work showed that TR-Kubernetes required minimal extensions to Kubernetes, and was capable of lowering the cost (by 53%) and improving the availability (99.999%) of representative workloads on a cloud platform when using transient compared to on-demand VMs [25].

While transient VMs are highly unreliable, cloud platforms also enable users to reserve VMs for long periods of time, i.e., multiple years, which offers higher reliability. Cloud platforms sell these "reserved" VMs at discount, e.g., 40-60%, relative to "on-demand" VMs, which users can rent on-demand on a per-hour basis. As a result, optimizing cloud costs also requires users to determine how many fixed reserved resources to buy versus rent based on their workload. In recent work, we introduced the concept of a waiting policy for cloud-enabled schedulers, which is the dual of a scheduling policy, and showed that optimizing cloud cost depends on it. While a scheduling policy determines which jobs run when fixed resources are available, a waiting policy determines which jobs wait for fixed resources when they are not available (rather than run immediately by renting on-demand resources). Waiting policies are important for cloud-enabled schedulers because they dictate the tradeoff between job performance and cost. In initial work [28, 29], we defined multiple waiting policies and developed analytical models to reveal their tradeoff between fixed resource provisioning, cost, and job waiting time. We evaluated the impact of these waiting policies on a year-long production batch workload consisting of 14M jobs run on a 14.3k-core cluster, and showed that a compound waiting policy decreases the cost (by 5%) and mean job waiting time (by 7×) compared to the current cluster. In subsequent work, we showed that our waiting policies could be applied in practice without *a priori* knowledge of job running and waiting times [30]. As part of this work, we also collected, analyzed, and publicly released nearly two years of data on Amazon's Reserved Instance Marketplace [26].

Finally, we have also designed techniques for improving data center efficiency from a cloud operator's perspective. In particular, as part of a collaboration with Google, we designed multiple policies for "over-committing" data center resources, i.e., by allocating resources that exceed the physical capacity, and showed that these policies increase machines' usable CPU capacity by 10-16% [17]. Variants of these policies have since been deployed in production in Google's data centers. We have also been active in many projects developing systems for managing academic cloud and network research testbeds, including GENI [4], CloudLab, and, most recently, the Open Cloud Testbed [46]. These testbeds are effectively smaller-scale versions of commercial clouds that enable a higher level of resource visibility and control that is necessary for research.

**Sustainability.** While improving data center efficiency is important, even highly efficient data centers may not be sustainable if they consume a significant fraction of their energy from carbon-intensive sources. As the negative impacts of climate change have become more visible, improving cloud data centers' sustainability by reducing their carbon emissions has become an increasingly important focus of research in academia and industry, especially recently. The lab has been working on improving data center sustainability for over a decade, long before its recent increase in prominence, and have made substantial contributions both over that time and recently. For example, we published some of the first papers in major conferences, including ASPLOS and NSDI, on managing server clusters and applications on intermittent power from renewable energy sources and demand response programs [11, 12, 32, 34, 37, 44].

More recently, with the increasing focus on this area, we have co-led two large related successful projects on i) developing new software systems and applications that manage carbon emissions as a first class metric (in collaboration with VMware), and ii) leveraging these systems to design a testbed for enabling research on carbon-efficient applications. The foundation of both projects relies on new system software for virtualizing energy resources, similar to how current data centers virtualize computing resources. Our work on both developing this energy virtualization layer, and demonstrating how applications can leverage it was published at ASPLOS 2023 [36]. In addition, we have also published papers at prominent venues on our vision for sustainable computing research moving forward based on current trends in computing and our experience from working in this area [19, 23]. Indeed, the lab's research focus is centered on sustainable computing because we believe that academic research in this area is particularly important, as industry lacks strong financial incentives to address problems in this area.

### 1.2 Optimizing Energy Systems

As summarized below, the lab's research has also made significant contributions to sensing, modeling, and optimizing the operation of energy systems to improve their efficiency and sustainability.

**Sensing.** Society's energy system, primarily the electric grid, is essentially a massive distributed system spread across each continent that delivers power from tens of thousands of energy producers to millions of energy consumers. As a result, accurate and fine-grained monitoring of environmental and operational data at large-scales and low-cost is challenging. Thus, we have designed a number of approaches for indirect, low-cost, at-scale sensing that are useful for improving infrastructure efficiency and sustainability. In general, these approaches have focused on leveraging data analytics to infer valuable, but hard-to-get, sensor data from other more readily available data sources. For example, since the dominant energy consumer in buildings is Heating, Ventilation, and Air Conditioning (HVAC), providing access to fine-grained spatial information about a building's interior temperature can enable HVAC optimizations that ensure comfort while using minimal energy. To provide such data, we developed approaches for inferring ambient temperature using a smartphone's CPU temperature sensor [39, 38, 7], which effectively transforms every smartphone in a building into a sensor that provides location-specific temperature data. Similarly, we also developed an approach to inferring the real-time power usage of individual devices from a building-wide energy meter by leveraging device power signatures. We used this approach to define "virtual" power meters for devices without requiring the installation of many physical power meters [31]. Finally, while the approaches above sense raw data, we have also developed "virtual" sensors for detecting higher-level criteria, such as system faults. For example, we have leveraged techniques developed from the modeling work, discussed below, to develop data analytics approaches for detecting and classifying faults in solar farms [9, 10].

**Modeling.** We have also done substantial work in data-driven modeling of energy systems, particularly for solar arrays and buildings. Specifically, we developed Solar-TK, an open-source publicly-available toolkit for data-driven solar performance modeling of solar sites [18]. Solar-TK produces accurate solar performance and forecast models from a small amount of generation data, and captures the impact of solar geometry, location, and weather on solar output. Solar-TK incorporates much of my prior work on solar modeling, such as modeling the impact of snow on solar generation [21]. In addition, we have also developed Peak-TK, a similar type of open-source toolkit for predicting peak usage in energy systems [6], since knowing peak usage is important for energy system operations [35]. The primary source of error in Solar-TK's performance models is quantifying the impact of cloud cover on solar output. To address this issue, we have recently focused on leveraging high-resolution, fine-grained multispectral data from the latest generation of GOES-R satellites, which is made publicly available in near real-time [1, 3]. This work has shown that satellite data can substantially improve the accuracy of solar performance models and near-term solar forecasts. Finally, we have also made contributions in modeling the energy-efficiency of buildings. For example, WattScale analyzes energy usage data from a large population of buildings to identify both the least energy-efficient buildings and the underlying cause of the inefficiency [14]. In recent work, we have also analyzed the socioeconomic profiles of energy-efficient and energy-inefficient buildings to better understand how efficiency relates to income, which can enable more equitable programs for subsidizing energy-efficiency improvements [40].

**Operations.** Finally, we have made substantial contributions to optimizing the operation of energy systems, in many cases, by leveraging the underlying sensing and modeling work above. For example, we leveraged Solar-TK's probabilistic forecasts to determine the amount of solar energy to commit in day-ahead electricity markets that maximizes revenue [20, 22]. We also leveraged Solar-TK in evaluating the potential benefits of interconnecting solar home systems in developing countries with an unreliable electric grid, which showed that interconnecting existing homes could increase electrification rates by more than 25% and reduce average costs by up to 30% per household [8]. Similarly, we leveraged Peak-TK to design VPeak, an approach that uses residential loads volunteered by their owners for coordinated control by a utility for grid optimizations [5]. Since the use of volunteer resources often comes with hard limits on how frequently they can be used by a remote utility, we leveraged Peak-TK's techniques for carefully selecting which days to operate these loads based on expected peak demand. Our results show that VPeak was able to reduce up to 26% of the total demand when selectively shaving peaks at local hotspots and up to 46.7% of the demand for grid-wide peak shaving. In addition, we have also made contributions to scheduling grid energy storage to reduce carbon emissions [15], managing shared community solar and energy storage [16], analyzing the incentives for grid defection [2], inferring optimal thermostat schedules from energy data [13], quantifying the potential for solar-powered electric bikes to reduce carbon emissions [41, 43], and analyzing the potential for reducing carbon emissions by replacing gas-powered heating with heat pumps [42].

## References

[1] **Akansha Singh Bansal**, T. Bansal, and D. Irwin. A Moment in the Sun: Solar Nowcasting from Multispectral Satellite Data using Self-Supervised Learning. In *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems (e-Energy)*, pages 1–12, June 2022. Acceptance Rate: 27/92 = 29%.

[2] **Akansha Singh Bansal** and D. Irwin. On the Feasibility, Cost, and Carbon Emissions of Grid Defection. In *Proceedings of the IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 1–7, Beijing, China, November 2019.

[3] **Akansha Singh Bansal** and D. Irwin. See the Light: Modeling Solar Performance using Multispectral Satellite Data. In *Proceedings of the Seventh ACM International Conference on Systems For Energy-Efficient Buildings, Cities, and Transportation (BuildSys)*, pages 1–10, August 2020. Acceptance Ratio: 38/108 = 35%.

[4] I. Baldine, Y. Xin, A. Mandal, C. Heermann, J. Chase, V. Marupadi, A. Yumerefendi, and D. Irwin. Networked Cloud Orchestration: A GENI Perspective. In *Proceedings of the Second IEEE Workshop on Management of Emerging Networks and Services (MENS)*, pages 573–578, Miami, Florida, December 2010.

[5] P. Bovornkeeratiroj, J. Wamburu, D. Irwin, and P. Shenoy. VPeak: Exploiting Volunteer Energy Resources for Flexible Peak Shaving. In *Proceedings of the Eighth ACM International Conference on Systems For Energy-Efficient Buildings, Cities, and Transportation (BuildSys)*, pages 1–10, November 2021. Acceptance Ratio: 28/107 = 26%. **Best Paper Award Runner-up.**

[6] P. Bovornkeeratiroj, J. Wamburu, D. Irwin, and P. Shenoy. PeakTK: An Open Source Toolkit for Peak Forecasting in Energy Systems. In *Proceedings of the ACM Conference on Computing Systems and Sustainable Societies (COMPASS)*, pages 1–11, Seattle, Washington, June 2022. Acceptance Rate: 37/100 = 37%.

[7] J. Breda, A. Trivedi, C. Wijesundara, P. Bovorkeeratiroj, D. Irwin, P. Shenoy, and J. Taneja. Hot or Not: Leveraging Mobile Devices for Ubiquitous Temperature Sensing. In *Proceedings of the Sixth ACM International Conference on Systems For Energy-Efficient Buildings, Cities, and Transportation (BuildSys)*, pages 1–10, New York City, New York, November 2019. Acceptance Ratio: 38/131 = 29%.

[8] S. Correa, **Noman Bashir**, A. Tan, D. Irwin, and J. Taneja. Extend: A Framework for Increasing Energy Access by Interconnecting Solar Home Systems. In *Proceedings of the Third ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS)*, pages 1–9, Guayaquil, Ecuador, June 2020. Acceptance Ratio: 26/84 = 31%.

[9] M. Feng, **Noman Bashir**, P. Shenoy, D. Irwin, and D. Kosanovic. SunDown: Model-driven Per-Panel Solar Anomaly Detection for Residential Arrays. In *Proceedings of the Third ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS)*, pages 1–5, Guayaquil, Ecuador, June 2020. **Note Paper**. Acceptance Ratio: 34/84 = 40%.

[10] M. Feng, **Noman Bashir**, P. Shenoy, D. Irwin, and D. Kosanovic. Model-driven Per-Panel Solar Anomaly Detection for Residential Arrays. *ACM Transactions on Cyber-Physical Systems (TCPS)*, 37(4):1–21, August 2021.

[11] P. P. II, B. Varghese, T. Guo, D. Irwin, P. Shenoy, A. Mahanti, J. Culbert, J. Goodhue, and C. Hill. Analyzing the Efficiency of a Green University Data Center. In *Proceedings of the Seventh ACM/SPEC International Conference on Performance Engineering (ICPE)*, pages 1–11, Delft, Netherlands, March 2016. Acceptance Ratio: 19/57 = 33%.

[12] D. Irwin, N. Sharma, and P. Shenoy. Towards Continuous Policy-driven Demand Response in Data Centers. In *Proceedings of the SIGCOMM Workshop on Energy and IT: from Green Networking to Smarter Systems (GreenNets)*, pages 19–24, Toronto, Canada, August 2011. Acceptance Ratio: 8/19 = 42%. **Best Paper Award**.

[13] S. Iyengar, S. Kalra, A. Ghosh, D. Irwin, P. Shenoy, and B. Marlin. Inferring Smart Schedules for Dumb Thermostats. *ACM Transactions on Cyber-Physical Systems (CPS)*, 3(2):17:1–17:29, March 2019.

[14] S. Iyengar, S. Lee, D. Irwin, P. Shenoy, and B. Weil. WattScale: A Data-driven Approach for Energy Efficiency Analytics of Buildings at Scale. *ACM Transactions on Data Science (TDS)*, 2(1):1–24, January 2021.

[15] R. Jha, S. Lee, S. Iyengar, M. Hajiesmaili, P. Shenoy, and D. Irwin. Emission-aware Energy Storage Scheduling for a Greener Grid. In *Proceedings of the Ninth ACM International Conference on Future Energy Systems (e-Energy)*, pages 1–11, Melbourne, Australia, June 2020. Acceptance Ratio: 25/88 = 28%. **Best Paper Award Runner-up**.

[16] S. Lee, P. Shenoy, K. Ramamritham, and D. Irwin. AutoShare: Virtual Community Solar and Storage for Energy Sharing. *Springer Energy Informatics*, 4(10), July 2021.

[17] **Noman Bashir**, N. Deng, K. Rzadca, D. Irwin, S. Kodak, and R. Jnagal. Take it to the Limit: Peak Prediction-driven Resource Overcommitment in Datacenters. In *Proceedings of the Sixteenth ACM European Conference on Computer Systems (EuroSys)*, pages 1–18, April 2021. Acceptance Ratio: 38/191 = 19.9%. **Awarded Artifacts Available Badge, Artifacts Functional Badge, and Results Reproduced Badge. One of 14 accepted papers awarded all three badges.**

[18] **Noman Bashir**, **Dong Chen**, D. Irwin, and P. Shenoy. Solar-TK: A Data-driven Toolkit for Solar PV Performance Modeling and Forecasting. In *Proceedings of the Sixteenth IEEE International Conference on Mobile Ad-Hoc and Smart Systems (MASS)*, pages 1–11, Monterey, California, November 2019. **Invited Paper**.

[19] **Noman Bashir**, T. Guo, M. Hajiesmaili, D. Irwin, P. Shenoy, R. Sitaraman, A. Souza, and A. Wierman. Enabling Sustainable Clouds: The Case for Virtualizing the Energy System. In *Proceedings of the Twelfth ACM Symposium on Cloud Computing (SoCC)*, pages 1–8, Seattle, Washington, November 2021. Acceptance Ratio: 46/145 = 32%.

[20] **Noman Bashir**, D. Irwin, and P. Shenoy. A Probabilistic Approach to Committing Solar Energy in Day-ahead Electricity Markets. In *Proceedings of the Eleventh IEEE International Green and Sustainable Computing Conference (IGSC)*, pages 1–15, October 2020. Acceptance Ratio: 33%.

[21] **Noman Bashir**, D. Irwin, and P. Shenoy. DeepSnow: Modeling the Impact of Snow on Solar Generation. In *Proceedings of the Seventh ACM International Conference on Systems For Energy-Efficient Buildings, Cities, and Transportation (BuildSys)*, pages 1–10, November 2020. Acceptance Ratio: 38/108 = 35%.

[22] **Noman Bashir**, D. Irwin, and P. Shenoy. A Probabilistic Approach to Committing Solar Energy in Day-ahead Electricity Markets. *Elsevier Sustainable Computing: Informatics and Systems (SUSCOM)*, 29:1–10, March 2021.

[23] **Noman Bashir**, D. Irwin, P. Shenoy, and A. Souza. Sustainable Computing - Without the Hot Air. In *Proceedings of the First Workshop on Sustainable Computer Systems Design and Implementation (HotCarbon)*, pages 1–7, San Diego, California, July 2022.

[24] **Pradeep Ambati** and D. Irwin. Optimizing the Cost of Executing Mixed Interactive and Batch Workloads on Transient VMs. *Proceedings of the ACM on Measurement and Analysis of Computing Systems (POMACS)*, 3(2):28:1–28:24, 2019.

[25] **Pradeep Ambati** and D. Irwin. Optimizing the Cost of Executing Mixed Interactive and Batch Workloads on Transient VMs. In *ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pages 1–14, Phoenix, Arizona, June 2019. Acceptance Ratio (Winter Deadline): 23/127 = 18%.

[26] **Pradeep Ambati**, D. Irwin, and P. Shenoy. No Reservations: A First Look at Amazon's Reserved Instance Marketplace. In *Proceedings of the Twelfth USENIX Workshop on Hot Topics in Cloud Computing (HotCloud)*, pages 1–6, Boston, Massachusetts, July 2020. Acceptance Ratio: 22/95 = 23%.

[27] **Pradeep Ambati**, D. Irwin, P. Shenoy, L. Gao, A. Ali-Eldin, and J. Albrecht. Understanding Synchronization Costs on Transient Cloud Resources. In *IEEE International Conference on Cloud Engineering (IC2E)*, pages 1–11, Prague, Czech Republic, June 2019. **Invited Paper**.

[28] **Pradeep Ambati**, **Noman Bashir**, D. Irwin, and P. Shenoy. Waiting Game: Optimally Provisioning Fixed Resources for Cloud-enabled Schedulers. In *Proceedings of the Thirty-third ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 1–12, November 2020. Acceptance Ratio: 68/380 = 18%. **Best Paper Award Finalist** and **Best Student Paper Award Finalist**.

[29] **Pradeep Ambati**, **Noman Bashir**, D. Irwin, and P. Shenoy. Modeling and Analyzing Waiting Policies for Cloud-Enabled Schedulers. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 32(12):3081–3100, December 2021.

[30] **Pradeep Ambati**, **Noman Bashir**, D. Irwin, and P. Shenoy. Good Things Come to Those Who Wait: Optimizing Job Waiting in the Cloud. In *Proceedings of the Twelfth ACM Symposium on Cloud Computing (SoCC)*, pages 1–13, Seattle, Washington, November 2021. Acceptance Ratio: 46/145 = 32%.

[31] **Sean Kenneth Barker**, S. Kalra, D. Irwin, and P. Shenoy. Building Virtual Power Meters for Online Load Tracking. *ACM Transactions on Cyber-Physical Systems (CPS)*, 3(2):23:1–23:24, March 2019.

[32] N. Sharma, **Sean Barker**, D. Irwin, and P. Shenoy. Blink: Managing Server Clusters on Intermittent Power. In *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 185–198, Newport Beach, California, March 2011. Acceptance Ratio: 32/152 = 21%.

[33] P. Sharma, S. Lee, T. Guo, D. Irwin, and P. Shenoy. Managing Risk in a Derivative IaaS Cloud. *IEEE Transactions on Parallel and Distributed Computing (TPDS)*, 29(8):1750–1765, August 2018.

[34] R. Singh, D. Irwin, P. Shenoy, and K. Ramakrishnan. Yank: Enabling Green Data Centers to Pull the Plug. In *Proceedings of the Tenth USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 143–155, Lombard, Illinois, April 2013. Acceptance Raio: 38/171 = 22%.

[35] A. Soman, A. Trivedi, D. Irwin, B. Kosanovic, and P. Shenoy. Peak Forecasting for Battery-based Energy Optimizations in Campus Microgrids. In *Proceedings of the Ninth ACM International Conference on Future Energy Systems (e-Energy)*, pages 1–5, Melbourne, Australia, June 2020. **Note Paper.** Acceptance Ratio: 10/37 = 27%.

[36] A. Souza, **Noman Bashir**, J. Murillo, W. Hanafy, Q. Liang, D. Irwin, and P. Shenoy. Ecovisor: A Virtual Energy System for Carbon-Efficient Applications. In *Proceedings of the Twenty-Eighth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 1–15, Vancouver, Canada, March 2023.

[37] **Supreeth Subramanya**, **Zain Mustafa**, D. Irwin, and P. Shenoy. Beyond Energy-Efficiency: Evaluating Green Datacenter Applications for Energy-agility. In *Proceedings of the Seventh ACM/SPEC International Conference on Performance Engineering (ICPE)*, pages 1–12, Delft, Netherlands, March 2016. Acceptance Ratio: 19/57 = 33%.

[38] A. Trivedi, P. Bovorkeeratiroj, J. Breda, P. Shenoy, J. Taneja, and D. Irwin. Phone-based Ambient Temperature Sensing Using Opportunistic Crowdsensing and Machine Learning. In *Proceedings of the Eleventh IEEE International Green and Sustainable Computing Conference (IGSC)*, pages 1–15, June 2020. Acceptance Ratio: 33%. **Best Paper Award**.

[39] A. Trivedi, P. Bovorkeeratiroj, J. Breda, P. Shenoy, J. Taneja, and D. Irwin. Phone-based Ambient Temperature Sensing Using Opportunistic Crowdsensing and Machine Learning. *Elsevier Sustainable Computing: Informatics and Systems (SUSCOM)*, 29:1–10, March 2021.

[40] J. Wamburu, E. Grazier, D. Irwin, C. Crago, and P. Shenoy. Towards Equity in Energy Efficiency Analyses. In *Proceedings of the First ACM SIGEnergy Workshop on Fair, Accountable, Transparent, and Ethical AI for Smart Environments and Energy Systems (FATESys)*, pages 1–5, Coimbra, Portugal, November 2021.

[41] J. Wamburu, S. Lee, M. Hajiesmaili, D. Irwin, and P. Shenoy. Ride Substitution Using Electric Bikes: Feasibility, Cost, and Carbon analysis. In *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)*, pages 1–27, March 2021.

[42] J. Wamburu, **Noman Bashir**, D. Irwin, and P. Shenoy. Data-driven Decarbonization of Residential Heating Systems. In *Proceedings of the Ninth ACM International Conference on Systems For Energy-Efficient Buildings, Cities, and Transportation (BuildSys)*, pages 1–11, November 2022. Acceptance Rate: 34/105 = 32%.

[43] J. Wamburu, C. Raff, D. Irwin, and P. Shenoy. Greening Electric Bike Sharing Using Solar Charging Stations. In *Proceedings of the Seventh ACM International Conference on Systems For Energy-Efficient*

*Buildings, Cities, and Transportation (BuildSys)*, pages 1–10, August 2020. Acceptance Ratio: 38/108 = 35%.

[44] C. Wang, M. Zink, and D. Irwin. Optimizing Parallel HPC Applications for Green Energy Sources. In *Proceedings of the Sixth IEEE International Green and Sustainable Computing Conference (IGSC)*, pages 1–8, Las Vegas, Nevada, December 2015. Acceptance Ratio: 24/67 = 36%.

[45] G. Zhao, L. Gao, and D. Irwin. Sync-on-the-fly: A Parallel Framework for Gradient Descend Algorithms on Transient Resources. In *Proceedings of teh IEEE International Conference on Big Data (BigData)*, pages 1–10, December 2018. Acceptance Ratio: 103/518 = 20%.

[46] M. Zink, D. Irwin, E. Cecchet, H. Saplakoglu, O. Krieger, M. Herbordt, M. Daitzman, P. Desnoyers, M. Leeser, and S. Handagala. The Open Cloud Testbed (OCT): A Platform for Research into New Cloud Technologies. In *Proceedings of the IEEE International Conference on Cloud Networking (CloudNet)*, pages 1–8, November 2021. **Invited Paper**.