Final Technical Report

a. **Federal Agency**	Department of Energy	
b. **Award Number**	DE-EE0007164	
c. **Project Title**	SunDial – An Integrated SHINES System to Enable High-penetration Feeder-level PV	
d. **Principal Investigators(co-PIs)**	Matthew Kromer mkromer@fraunhofer.org Lead, Grid Integration 617.353.0067 Kurt Roth kroth@fraunhofer.org Lead, Energy Systems 617.353.1895	
e. **Business Contact**	Andre Sharon Center Director 617.353.8776 asharon@fraunhofer.org	
f. **Submission Date**	12/15/2019	
g. **DUNS Number**	933928103	
h. **Recipient Organization**	Fraunhofer USA Center for Manufacturing Innovation CMI	
i. **Project Period**	**Start**: 1/1/2016	**End**: 10/31/2019
j. **Submitting Official Signature**	Signature	

Acknowledgement: This material is based upon work supported by the Department of Energy, Office of Energy Efficiency and Renewable Energy, Solar Energy Technologies Office, under Award Number DE-EE0007164.

Fraunhofer USA Center for Manufacturing Innovation CMI

SunDial Final Report

REPORT TO THE
DEPARTMENT OF ENERGY SOLAR ENERGY TECHNOLOGY OFFICE (SETO)
January 2020

Matt Kromer
Kurt Roth, Ph.D.
Michael Zeifman, Ph.D.
Tsz Yip
Jim Boch
Samer Arafa
Aram Shishmanian
Justin Woodard

Disclaimer

Our conclusions are the results of the exercise of our best professional judgment, based in part upon materials and information provided to us by different entities, including but not limited to IPKeys and National Grid. Use of this report by any third party for whatever purposes should not, and does not, absolve such third party from using due diligence in verifying the report's contents.

Any use which a third party makes of this document, or any reliance on it, or decisions to be made based on it, are the responsibility of such third party. Fraunhofer USA accepts no duty of care or liability of any kind whatsoever to any such third party, and no responsibility for damages, if any, suffered by any third party as a result of decisions made, or not made, or actions taken, or not taken, based on this document.

This report may be reproduced only in its entirety.

Acknowledgements

The efforts of several people made this project possible. We thank Kemal Celik, Tassos Golnas, and Guohui Yuan of the U.S. Department of Energy Solar Energy Technology Office (DOE-SETO) and Kavita Ravi (formerly of MassCEC) for their support for and oversight of this project. Christian Hoepfner with Fraunhofer USA provided invaluable technical and programmatic expertise without which this project would not have been possible.

We would also like to gratefully acknowledge the contributions of Justin Woodard, Aram Shishmanian, Samer Arafa, Jason Small, Fouad Dagher, Nick Corsetti, Peter Russo, and Paul Wassink with National Grid for their support integrating with the Patterson Road Solar + Storage Site, recruiting C&I facilities to participate in this project, and technical expertise in designing and testing the system.

In addition, we thank Yuriy Kozhynov, John (Yianni) Kozyrakis, Robert Nawy, Victor Sadikov, and Bill Walker of IPKeys and Omri Schwarz (formerly of Fraunhofer) for their contributions to the FLAME development. Jeff McAulay (formerly of EnerNOC) also participated in many discussions early in the project that helped to shape the FLAME.

Finally, we thank John Blood, Kerry Henry, and, in particular, Jon Snyer for their on-site support of the FLAME implementation at their facilities. Without their assent, the field demonstration would not have occurred.

Funding for this project was provided by the U.S. Department of Energy Solar Energy Technology Office (DOE-SETO) under the SHINES program, award DE-EE0007164, and the Massachusetts Clean Energy Center (MassCEC).

Executive Summary

The Project Team of Fraunhofer USA, National Grid, and IPKeys developed and conducted a pilot deployment of the SunDial system, a virtual power plant platform that enables high-penetrations of solar PV to be integrated into the distribution grid. The pilot was conducted over a 15-month period from August 2018 through October 2019 on a National Grid distribution feeder in Shirley, MA. A vendor-agnostic control platform (the "Global Scheduler") optimally shaped the net load for a virtual portfolio of non-co-located DERs based on user-defined policy objectives. The goal of the SunDial project was to simplify and reduce the risk associated with the deployment of solar in high-penetration environments by:

(1) Developing an open-source, vendor-agnostic dispatch platform that can be readily adapted to optimize control of DERs over a variety of use cases;
(2) Developing auto-calibrating load and solar prediction methodologies that can be readily implemented and scaled to new deployments;
(3) Developing a methodology to use demand-side management with traditional electrochemical energy storage to provide "load shaping" services in high solar penetration environments;
(4) Using grid-scale storage to minimize short-term intermittency association with PV production; and
(5) Deploying on the National Grid distribution system to gain experience on the potential for (and limits of) integrated storage with demand-side management.

During the course of the SunDial pilot, the Global Scheduler dispatched a 1MWh energy storage system (ESS) and several flexible loads based on forward-looking predictions of solar production and customer demand to test multiple use cases, including peak shaving, energy cost optimization, peak-power dispatch, and power firming, while minimizing rapid changes in solar production due to passing cloud cover. The objectives for this field demonstration were to:

(1) Conduct a 12-month pilot demonstration of an integrated portfolio of solar PV, energy storage, and controllable loads to shape the net load on the power distribution system and mitigate mismatch between solar production and demand.
(2) Mitigate PV intermittency by limiting net system ramp rates to <10% of PV system nameplate.
(3) Show how a combination of solar forecasting, load forecasting, and flexible loads can increase the effective energy storage capacity within a DER portfolio by 10 to 20%.
(4) Show that the SunDial system can be used to support a levelized cost of electricity of less than $0.14 per kWh.

Field testing results show that integrating highly granular, bottom-up time-series predictions into the dispatch decision-making process effectively increased storage capacity relative to a non-predictive baseline by upwards of 20%. Furthermore, the results highlight the need to manage prediction uncertainty, for example by maintaining sufficient BESS reserve capacity. Flexible loads were dispatched repeatedly (approximately 70 successful events in total) to supplement electrochemical energy storage, comprising the equivalent of approximately 10% of additional storage capacity. Finally, we identify several challenges that impede scaling SHINES-type systems, including a lack of scalable methods for deploying integrated solar + storage plants; practical difficulties integrating flexible loads into a solar-support use case; and uncertainties in load and solar predictions.

1 Introduction

The Project Team of Fraunhofer USA, National Grid, and IPKeys developed and conducted a pilot deployment of the SunDial system, a virtual power plant platform intended to enable integration high-penetration solar PV within the distribution grid. The pilot was conducted over a 15-month period from August 2018 through October 2019 on a National Grid distribution feeder in Shirley, MA. A vendor-agnostic control platform (the "Global Scheduler") was used to optimally shape the net load for a virtual portfolio of non-co-located DERs based on user-defined policy objectives. The goals of the SunDial project were to reduce the friction and risk associated with deploying solar in high-penetration environments by:

(1) Developing an open-source, vendor-agnostic dispatch platform that can be readily adapted to optimize control of DERs over a variety of use cases;
(2) Developing auto-calibrating load and solar prediction methodologies that can be readily implemented and scaled to new deployments;
(3) Developing a methodology to use demand-side management with traditional electrochemical energy storage to provide "load shaping" services in high solar penetration environments;
(4) Using grid-scale storage to minimize short-term intermittency association with PV production; and
(5) Deploying on the National Grid distribution system to gain experience on the potential for (and limits of) integrated storage + demand-side management.

Prior work conducted during Phases 1 and 2 of this project presented modeled results for the SunDial System that indicated that integrated solar and load forecasting, in combination with flexible loads could support optimal dispatch of DERs and thereby increase the effective storage capacity by 10 to 20 percent, depending on use case and location. A limitation of this prior work is that it assumed "perfect information", i.e., predictions about future PV generation and loads were assumed to be perfectly accurate. These prior results also were based on theoretical/modeled estimates of HVAC load shifting capacity within typical commercial and industrial buildings. Prior work conducted during Phases 1 and 2 also entailed the development and deployment of infrastructure to support integrated testing of the SunDial system, but had not yet entailed testing the SunDial system in a deployed environment.

During the field demonstration phase of the project, the Global Scheduler dispatched a 1MWh energy storage system (ESS) and several flexible loads based on forward-looking predictions of solar production and customer demand to test multiple use cases, including peak shaving, energy cost optimization, peak-power dispatch, and power firming, while minimizing rapid changes in solar production due to passing cloud cover. The objectives for this field demonstration were to:

(1) Conduct a 12-month pilot demonstration of an integrated portfolio of solar PV, energy storage, and controllable loads to help shape the net load on the power distribution system and mitigate mismatch between solar production and demand peaks.
(2) Mitigate PV intermittency by limiting net system ramp rates to less than 10% of PV system nameplate.
(3) Show how a combination of solar forecasting, load forecasting, and flexible loads can increase the effective energy storage capacity within a DER portfolio by 10 to 20%.
(4) Show that the SunDial system can be used to support a levelized cost of electricity of less than $0.14 per kWh.

This report summarizes the results of the SunDial pilot, and is organized as follows:

- Section 2 summarizes prior work to develop the SunDial system, referencing relevant deliverables.
- Section 0 provides an overview of the scope, goals, and methodology for conducting the pilot.
- Section 4 evaluates results from the pilot deployment, including: lessons learned from commissioning and testing an integrated solar + storage plant (Section 4.2); analysis of solar and

load prediction methodologies used to support optimal dispatch of DERs (Section 4.3); full system testing to evaluate optimal dispatch of DERs in support of user-defined strategic (multi-hour) objectives (Section 4.4), and evaluation of the efficacy of multi-hour load shifting to support solar integration (section 4.5).

- Section 5 summarizes overall conclusions and lessons learned from this project.
- Section 6 outlines potential areas for future work.

2 Prior Work

2.1 Prior Phases of the SunDial Project

Prior to the demonstration portion of the project, Fraunhofer worked on and completed several activities essential to performing field testing. A summary of each follows, including a reference to the specific stand-alone reports that provide further details.

(1) *Global Scheduler Development*: Fraunhofer developed the Global Scheduler (GS) platform to optimize SunDial system performance while pursuing optimization objectives. Specifically, the GS develops prediction for future PV generation, evaluates future load predictions and potential managed load profiles from the FLAME, and considers utility/ISO price signals to determine combination of ESS charge/discharge and facility load management actions for a moving 48-hour time horizon that yield the optimal (i.e., lowest cost) system performance for specific use cases. The Global Scheduler Development Report describes the architecture, algorithms and interfaces of the GS and how they were developed.

(2) *Facility Load Aggregation and Management Engine (FLAME) Development:* IPKeys developed the FLAME, which develops baseline load predictions[1] and potential managed load profiles for all participating facilities. The FLAME aggregates those baselines and potential managed load profiles and, in response to price data from the GS, sends the lowest cost potential managed load profiles to the GS. After the GS selects the managed load profile to implement, the FLAME dispatches discrete load-management events at the individual facilities comprising that managed load profile. The Facility Load Aggregation and Management Development Report describes the architecture, algorithms and interfaces of the FLAME and how they were developed.

(3) *Energy Flow Model Development:* Fraunhofer developed a model to simulate the energy flows and economics of different combinations of energy storage (ESS), PV generation, electric loads, and load-management potentials for different control strategies applied to different use cases in different climates. We integrated the algorithms developed and implemented in the Energy Flow Model in the Global Scheduler, and the Global Scheduler Development Report describes them further.

(4) *Solar Plus Storage Integration:* Fraunhofer and National Grid worked together to deploy and commission a 0.5MW/1.0MWh electric energy storage system (ESS) adjacent to two National Grid PV systems totaling 1.5MW. The "Sundial System Design Report" elaborates upon the design and commissioning processes.

2.2 Literature Review

Following is a review of prior research that informed our development of the SunDial system:

System Modeling: We previously presented modeled results using a virtual power plant framework similar to that used in this pilot [28], [29], based on results from Phases 1 and 2 of this project. Results indicated that integrated solar and load forecasting in combination with flexible loads could support optimal dispatch of DERs and thereby increase the effective storage capacity by 10 to 20%, depending on use case and location. However, an important limitation of this prior work is that it assumed "perfect information" –

[1] Fraunhofer developed the facility load prediction algorithms and module for the FLAME.

i.e., predictions about the future state of the system were assumed to be perfectly accurate. Further, these prior results were based on theoretical estimates of HVAC load shifting capacity within typical commercial and industrial buildings. A key focus of the field demonstration was therefore to apply the same methodology within our modeled results to a real-world deployment.

Optimization and Control: The Global Scheduler's optimization algorithm is formulated as an economic dispatch problem in which the system's total demand at each time step over a 24-hour time horizon must be satisfied using a combination of available energy storage, solar generation, shiftable loads, and power imports (i.e., supply=demand). Within this formulation, each resource (and/or combinations of resources) is associated with its own set of cost functions. Resources are further categorized as either *strategic* (complex effort, such as coordinating thermostatically controlled loads or shiftable industrial processes among participating buildings, which can be applied typically once a day) or *tactical* (e.g., a simple shift with no rebound effect, which can be applied more than once a day). We reviewed relevant literature related to this formulation, such as the economic dispatch problem considered in [1] and [2]; both demand and supply (e.g., [3], [4], [5]), ancillary services (e.g., frequency regulation in [6] and [7]) and market regulation with performance-based regulation mechanism ([8], [9]). These approaches often involve nonlinear optimization, due to the nonlinear optimization functions and, at times, nonconvex optimization space [5].

Following most of these references, we adopted a model predictive control (MPC) - based approach that considers a variable time-horizon and time step (typically 24-hour horizon and either one-hour or 15-minute time step) for the design of Global Scheduler. The control possibilities for a day (i.e., 24-h horizon) involve: (1) selection of a strategic load shift, (2) selection of battery charging/discharging power (one per time step), and (3) selection of PV curtailment.

Traditional convex optimization methods, such as the gradient descent or linear programming [10] are not applicable to the GS because of several challenges. The optimization space for strategic load shifting is not only nonconvex but also is not continuous. In principle, since our proposed solution to characterize numerous strategic load shifting possibilities involves just a few options corresponding to the major clusters, we could apply a scheme to optimize "tactical" resources (i.e., the remaining arrays of battery charging/discharging and building load sinking/shedding) for each cluster separately and then select the best solution among the found few "tactical" solutions. However, convex optimization within the battery charging/discharging space is still challenging, because the charging/discharging power available at a given hour depends on the battery charge that, in turn, depends on the charging/discharging history. This potentially makes the optimization space nonconvex. A recursive approach, such as, e.g., dynamic programming [11], can be a good solution tool for the "tactical" optimization problem. However, this method is applicable to only memoryless systems where the current system state depends on the previous state but not on the path leading to the previous state. For example, a potential optimization scenario can preserve a specific geometric shape of the 24-hour net energy profile (e.g., making it as flat as possible). In this case, we cannot assume the system to be memoryless, because the profile's shape ultimately depends on the entire profile.

This diversity of ownership/optimization scenarios in the SunDial system required a versatile optimization approach. As such, we selected the simulated annealing (SA) approach that is known to result in an approximately optimal solution to large combinatorial problems, does not require convexity or continuity of the search space and is applicable to a broad range of optimization problems [12]. Although there are numerous discussions in the literature on a particular design of the SA algorithm, many authors agree that the Corana's based algorithm implemented by Goffe [13] performs well for a wide variety of problems.

Accordingly, we implemented Goffe's design of the SA approach for GS. To deal with the non-continuous control space of strategic DR in FLAME, we run separate optimizations for each of the clusters and then select the one with the minimal value of the optimization function.

Load Predictions: A large body of prior work has developed methodologies for predicting loads and load shapes from individual commercial and industrial buildings, particularly for automated demand response applications. One example of such a methodology, developed by LBNL, is described in [14] and [15].[2] The simplest version of the LBNL model uses only a time series of electric load data. If no other data are provided, then the baseline will be very simple, i.e., the predicted load for a given time of the week simply equals the weighted average load at that time of the week. The model includes the ability to define the number of weighting days that the model will consider. The default value for weighting days is 14, which puts more statistical weight on most recent two weeks of data, which is consistent with [16]. The LBNL model also can use outside air temperature (T_o) as a variable. The temperature-adjusted model assumes that HVAC loads are significant for that facility, and attempts to determine when the building uses electric power to condition the facility. When the facility is heated or cooled, the building considered to be in "occupied mode", and occupied and unoccupied modes are modeled separately. In both modes, the total electric load equals the sum of a time-of-week effect and a temperature-dependent effect. The temperature dependence is piecewise-linear, i.e., the load is assumed to increase or decrease linearly with temperature.[3]

We found that the baseline LBNL approach can be further enhanced by incorporating same-day data into baseline load prediction algorithms as described in [16] and [26]. Specifically, adding morning adjustment factors for weather-sensitive commercial and institutional buildings significantly increases facility load prediction accuracy for almost all models. These factors look at *actual loads that day* and compare them to the predicted loads from a conventional baseline prediction using a prior 14- or 28-day period. The literature shows that both absolute offset (difference between actual and predicted kWh) and multiplicative (ratio of actual to predicted kWh) factors have been evaluated and improve prediction accuracy. Importantly, these factors need to be calculated during periods of time well before any DR (or load management) event occurs to avoid impacts from DR preparations (e.g., precooling, industrial process modifications). For example, for the California DR markets a rolling 10-day window with additive adjustment based on two hours prior to event start was found to provide a solid baseline [16].

Our load prediction methodology adopts a modified version of the approaches outlined above: i.e., prior day / same hour averages are used in combination with same-day adjustments, and facility-specific calendar adjustments. The specific combination of parameters and adjustments was calibrated to the specific facility. We explored the use of hourly temperature adjustments, but found the resulting predictions were less accurate. We speculate that this is explained by the specific customer portfolio included within our study: industrial processes dominate the facility electric loads, so the team has decided not to use the temperature-adjusted model for industrial facilities; and schools tend to have decreased demand during cooling season due to summer schedules.

Solar Predictions: As described within the body of this report (Section 4.3), we found the predictions from satellite weather forecasts tended to have systematic biases that over-estimated production. As such, we

---

[2]A software implementation of this approach is also available on-line at https://bitbucket.org/berkeleylab/eetd-loadshape

[3] Note that [16] cites references that note that the assumed a linear relationship between T_o and electric loads may become more tenuous as T_o increases, so a quadratic function may make more sense.

explored hybrid solar prediction approaches, in which satellite-based predictions were re-calibrated over time based on feedback from actual field data. Prior examples of this approach in practice include [17], which uses ex post analysis of data from California to evaluate multiple approaches for incorporating satellite-based solar production forecasts and production data to improve prediction accuracy; and [18], in which a so-called hybrid black-box model uses weather data and historical production measurements to extract best-fit data about PV sites (orientation, tilt, module conversion efficiency, and shading).

Modeling Load Flexibility: To support implementation of a robust load shifting strategy, we reviewed recent literature related to the modeling and implementation of shedding, sinking, and shifting flexible loads in commercial and industrial (C&I) facilities. A large portion of these references reflect work done by the Demand Response Center at Lawrence Berkeley National Laboratory (LBNL), e.g., see the literature review [19] and the publications website for the Grid Integration Group at LBNL[4], which characterizes demand response potentials for numerous classes of C&I facilities in the United States. In addition, several references ([20], [21], [22], [23], [24], [24]) were reviewed to develop our understanding of how to model demand response potentials in C&I buildings. It should be noted that much of the DR literature focuses on load *shed* events only, whereas we are most interested in loads that can be (1) productively *shifted* to match solar generation (as distinguished from a pure load shed); (2) which can be called repeatedly; and (3) which provide multiple hours of response. For this reason, thermostatic load shifting strategies (e.g., for HVAC) such as that described in [20] were of particular interest.

Based upon this review, our modeling approach divides DR techniques into *strategic* (complex effort coordinating thermostatically controlled loads or shiftable industrial processes among participating buildings, can be applied typically once a day) and *tactical* (either simple shift or no rebound effect, can be applied more than once a day). Each technique can be characterized by physical parameters, such as shed or sink time and value, by cost. We considered implementing a “fatigue” parameter that quantitatively expresses degree of reluctance of building occupants/managers to participate in future events incurred by repeated load management events, but did not find that this was relevant for the loads that we considered. Further, we focused primarily on strategic load-shifting, as this was most feasible given the constraints of our customer portfolio.

Use Cases: The use cases modeled during the course of the field pilot are based on numerous current and potential future “Solar+X” tariff structures and value streams. A good summary of potential applications considered by our team is included in the Austin SHINES 2017 program review [30], and in RMI’s review of the Economis of Battery Energy Storage [31]. We focused our efforts on a subset of applications identified in [30], including several “Customer” applications (Demand Charge Reduction, time-of-use pricing); “Energy Market Operations” (Peak Load Reduction, Energy Arbitrage, LMP); “Renewables Integration” (Solar Variance); and Distribution Operations Support (Congestion Management). In addition, the GS platform was developed with the capability to provide several other services that have not yet been implemented: Voltage Support, Power Factor Correction, and Up/Down Regulation.

3 SunDial System Deployment and Field Test Overview

3.1 System Description

The SunDial System was deployed on a National Grid distribution feeder located in Shirley, Massachusetts (Figure 1) and tested for a period of 15 months, beginning in August 2018 and concluding in October 2019.

[4] https://gig.lbl.gov/publications?page=1 .

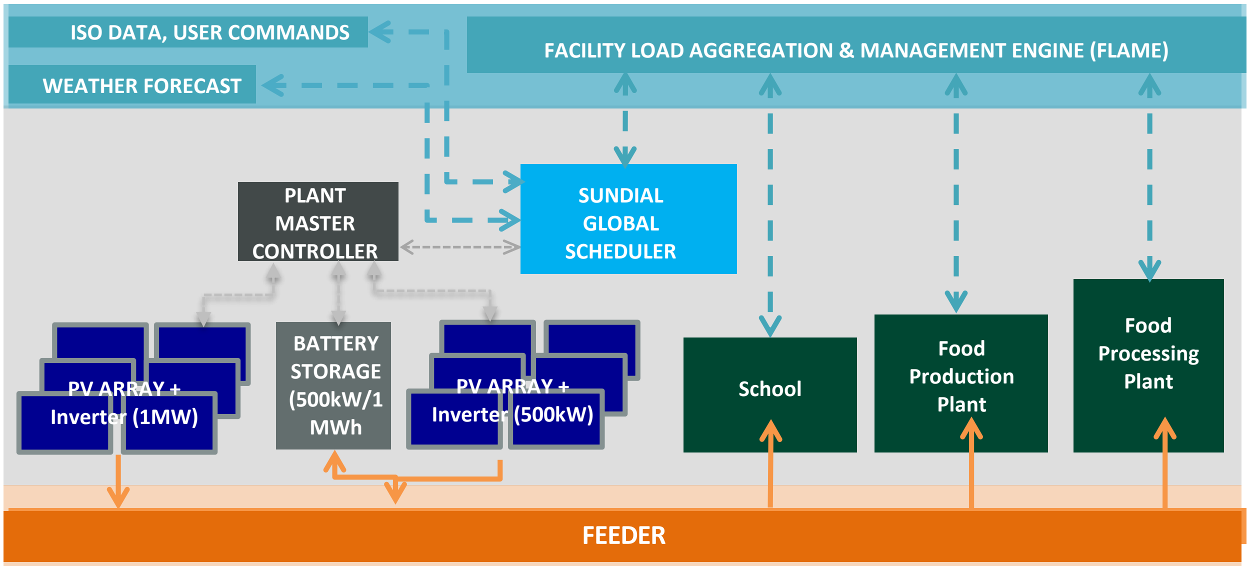

Figure 1: Block Diagram of SunDial System

The system under test consisted of:

(1) A Solar + Storage site (Figure 2), consisting of 500 kW of Solar PV, for which we had the ability to monitor and control power flow (including real power curtailment, reactive power control, and up-ramp limiting); and a 500 kW/1MWh Li-Ion Energy Storage System (Tesla Power Pack), for which we had the ability to monitor and control real and reactive power

(2) 1MW of Solar PV located at an adjacent PV field, for which we had the ability to monitor power, but did not have access to directly control. Due to several issues at this PV site, the 1MW PV plant was unavailable for much of the pilot, so results from integrated system testing are restricted to the 500kW solar plant.

(3) Approximately 3.8MW (peak) of customer loads spread across three commercial and industrial (C&I) facilities: a school, a food-processing facility, and a food-production facility. Due to large differences in the peak loads among the facilities and the total load (3.8MW) relative to the PV 0.5MW) and ESS resources (0.5MW/1.0MWh) available, we applied different scale factors to individual loads to approximately equalize peak demand across the three sites and normalize total demand to approximately equal total solar plus ESS capacity.[5]

(4) Several flexible loads at the participating C&I sites, consisting of: (1) shifting of cooling loads at a school (e.g., by pre-cooling during periods of excess generation and shedding load during other periods) (approximately 5-20% of total site load, depending on time of year); (2) delaying the overnight charging of electrical lift trucks at a food processing facility (approximately 2.5 to 5% of site load); and (3) shifting usage patterns of a drag conveyor at a food production facility (approximately 5% of site load).

[5] A scale factor of 0.1 was applied to the food production facility, 0.2 to the food processing facility, and 1.0 to the school.

Figure 2: SunDial pilot solar + storage site in Shirley, MA

Overall control of the SunDial System was orchestrated by the Global Scheduler (GS), a vendor-agnostic optimal dispatch controller developed by Fraunhofer as part of this project.[6] The Global Scheduler was implemented as an extension of the VOLTTRON platform, an open-source DER operating system developed by Pacific Northwest National Laboratory (PNNL),[7] and was deployed on a Linux server at the Solar + Storage plant. Monitoring and control of endpoint devices at the PV and Storage sites was implemented via a local Modbus connection to a site plant master controller. These devices included the PV inverter; an ESS master controller (which aggregates data and controller for the ESS); meters located at the output of the ESS inverter, PV inverter, and the point of common coupling (PCC); a weather station; and the site's switchgear. Monitoring and control of demand side resources was implemented via connection to a cloud-based server (the "Facility Load Aggregation and Management Engine", or FLAME) developed and operated by IPKeys. Participating C&I facilities were instrumented with telemetry to provide 1-minute power data to the Global Scheduler in near-real time (typical latency of 5 to 15 minutes), and to enable actuation of flexible loads based on commands from the Global Scheduler. Data from individual customer facilities were collected by the FLAME and communicated to the Global Scheduler. The Global Scheduler was also configured to query external data sources, including a satellite-based solar prediction service, ISO New England's web API, a power meter installed on the local distribution system through which we could monitor feeder power flow, and an SSH connection accessible by the project team to enable remote monitoring and control. Data from the system were logged to a local database at a time resolution down to 1 second (e.g., for solar production and ESS data), and was replicated on a remote Fraunhofer server nightly.

3.2 Field Test Approach

The SunDial field test consisted primarily of using the Global Scheduler to monitor and control the system's DER portfolio to simultaneously mitigate short-term intermittency of solar production (e.g., due to passing clouds) and to strategically dispatch DERs to address mismatches between solar production and customer loads. This entailed using the Global Scheduler to monitor device end points, generate predictions about future facility loads and PV generation, and, in response, remotely dispatch the ESS, communicate with the FLAME server to initiate load-management events, and modify PV system set points. Over the course of the field demonstration, the system's strategic objectives were periodically modified to test different use cases and to improve performance of the GS software. The system was designed to enable the rapid addition or removal of resources from the system's portfolio, so devices or classes of devices were included or removed depending on availability and the objectives of the current test regime without affecting our evaluation of other system components.

---

[6] A detailed description of the Global Scheduler is available in the Global Scheduler Development report
[7] VOLTTRON. www.volttron.org

3.3 Field Test Timeline

A timeline summarizing implementation of SunDial system components is shown in Table 1. The pilot deployment included two periods of extended testing when we received telemetry from participating customer facilities and the PV+ESS site and actively controlled the energy storage (ESS) and PV systems. The extended gap in testing was caused by a failure with the PV site's recloser that kept the PV+Storage site offline from Dec 15, 2018 through May, 24, 2019. Several delays were encountered in fully integrating with flexible loads at customer facilities, delaying repeated testing of manual load shifting event calls until May 2019 and implementation of fully automated load shifting until August 2019.[8]

Table 1: Timeline of SunDial System implementation. Lightly shaded boxes indicate that a given feature was active for a portion of the period in question.

GS Capability	2018							2019									
	J	J	A	S	O	N	D	J	F	M	A	M	J	J	A	S	O
Load Forecast & Telemetry																	
Solar Production Forecast & Telemetry																	
Control of DERs																	
Optimization & Strategic Dispatch of DERs																	
Load Shifting																	

4 Results

4.1 Summary of System Operation

The SunDial System operated for 6,700 hours over 275 full days of testing; and operated in "application control" mode, which enables autonomous dispatch of DERs, for approximately 5,900 hours over 241 full days of testing. The system was offline for approximately 4,275 hours over 182 days. Over the course of the pilot, total ESS throughput was approximately 354MWh, or the equivalent of 177 full charge / discharge cycles on the 1 MWh ESS. Approximately 70 total demand response events were called, comprising 32 AC and 35 EV charging load shift events. Figure 3 breaks down system uptime by month. In addition to the recloser failure alluded to above, sources for system downtime included: (1) power outages at the site, which frequently required personnel to manually reset the site's RTAC (real time automation controller) and/or switchgear; (2) planned site maintenance to support commissioning and customer acceptance at the site, that occurred in parallel to the SunDial field demonstration; and (3) loss of communication between the GS and the RTAC. Excluding equipment failures and planned site maintenance, overall uptime for the system was approximately 98%.

[8] "Manual" load shifting refers to instances in which events were remotely called by an individual through the IPKeys cloud platform. "Automated" load shifting refers to events that were initiated through data exchanges between the FLAME and the Global Scheduler based on pre-configured objectives.

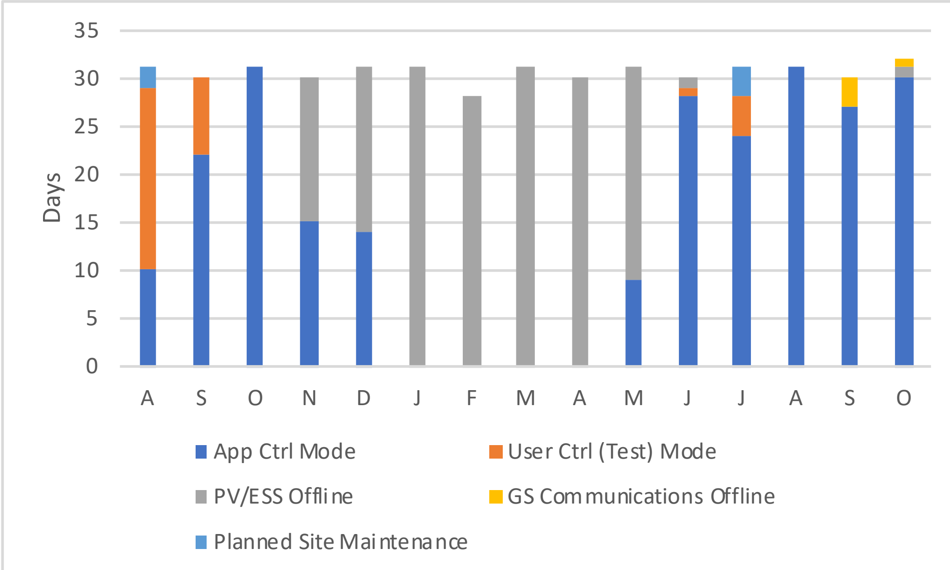

Figure 3: Summary of SunDial system uptime

4.2 Solar + Storage Integration and ESS Performance

As National Grid's first solar + storage project, a key objective of the field demonstration was to gain experience with commissioning and operating grid-scale storage and identify lessons learned for future deployments. These findings are summarized below:

ESS Performance: In total, approximately 354MWh of energy was charged or discharged by the ESS, or the equivalent of full 177 charge / discharge cycles on the 1 MWh ESS. The vast majority (87%) of ESS energy throughput was used for bulk load shifting and the remainder to mitigate rapid changes in irradiance. An estimate of ESS roundtrip efficiency as a function of power output is shown in Table 2. The nominal roundtrip efficiency, defined as energy in to the ESS divided by energy out of the ESS, as measured at the output of the ESS inverter (but before the site's point of common coupling), ranges from 54% at charge levels less than 50kW to 82-85% for higher charge levels. Notably, the bulk of charge/discharge operation at low power is associated with the plant's ramp-rate control function used to smooth short-term intermittency, which contributes to outsized fraction of the ESS energy loss. An optimal charge strategy should accurately model operational efficiency as a function of power level.

Table 2: ESS Roundtrip Efficiency as a function of power

ESS Power (kW)	RT Eff.	% of Energy Throughput
<50	54%	12%
50-100	82%	19%
100-200	85%	37%
200-300	85%	20%
300-400	84%	8%
>400	83%	5%
Total	**80%**	**100%**

Solar + Storage Integration: We observed two different instances when the ESS would trip off on a routine basis and come back online after several minutes. The first instance of this behavior was a false trip caused by the solar inverter's DC breaker closing as the site energized at the start of each day. The second instance was a frequency anomaly that occurred daily at 11AM UTC due to a capacitor bank switching at a nearby industrial facility. Both issues were eventually mitigated – the former by adjusting the ESS trip settings; the latter via a firmware modification by the ESS vendor – but neither issue was readily apparent.

Cold Weather Performance: Due to the site being offline for most of the winter, we had a limited window in which to assess the ESS's performance in cold weather. The ESS incorporates functionality to actively heat the ESS so as to maintain charge and discharge capacity in cold weather conditions, but this capability was not enabled until later in the demonstration phase. During the limited period in which the site was online during cold weather, our ability to charge the ESS was greatly constrained. Figure 4 shows an example of how, after several hours of below-0°C temperature, the ESS effectively cannot accept charge.

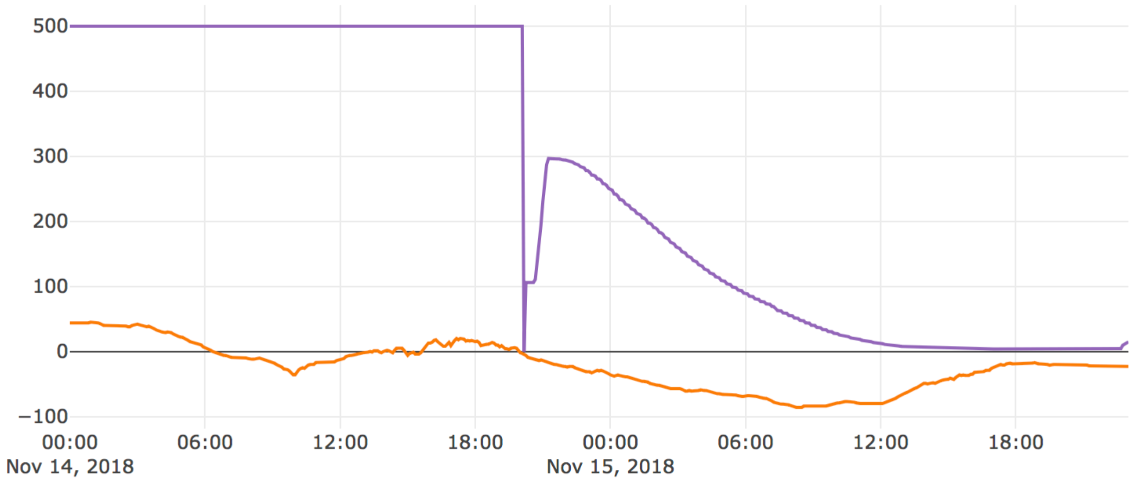

Figure 4: Max ESS Charge power (kW, Purple) and Ambient Temperature x 10 deg C (Orange)

Mitigation of Solar Intermittency:

The SunDial system had a target objective to limit changes in plant power output to less than 10% of nameplate PV capacity per minute. This functionality was initially specified to reside in the site's RTAC, as the RTAC has lowest latency communication with end point devices. However, the ramp rate control function as implemented in the RTAC introduced unstable behavior if enabled at the same time as a non-zero ESS set point. This behavior is illustrated in Figure 5 (properly working ramp rate control with ESS set point equal to zero) and Figure 6 (unstable ESS output + curtailed PV when set point is non-zero). The vendor could not correct this issue, so ramp limiting functionality was instead implemented directly in the Global Scheduler. An example of typical system response to variable irradiance conditions is shown in Figure 7. As shown in Table 3, implementation of ESS feedback eliminated 100% of ramp events greater than 20% of nameplate, 97% of ramp events greater than 10%, and 91% of ramp events greater than 5% of nameplate.

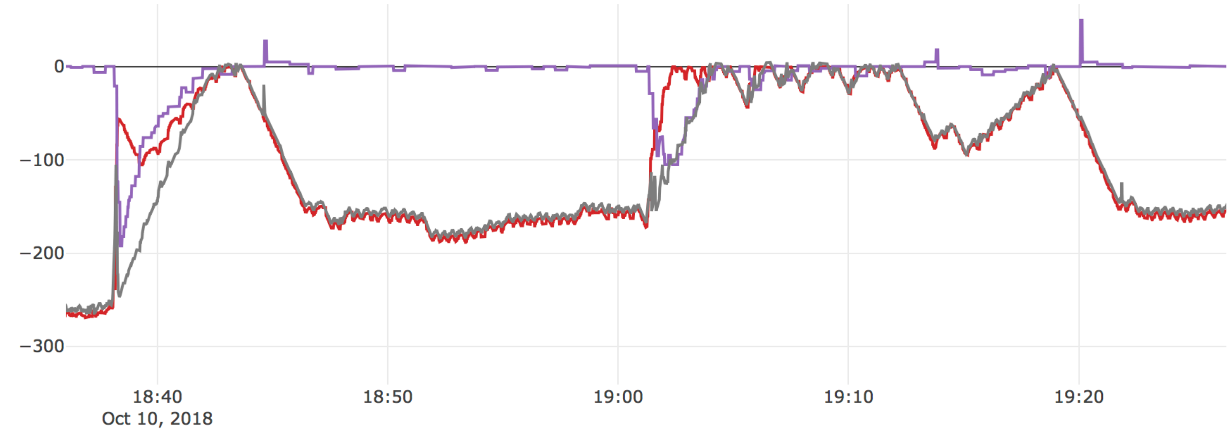

Figure 5: Ramp limit example with no ESS set points commanded. Red – Solar generation; Purple - ESS output; Grey – Net (PV+ESS) output. In this case, ramp limiting function works as designed.

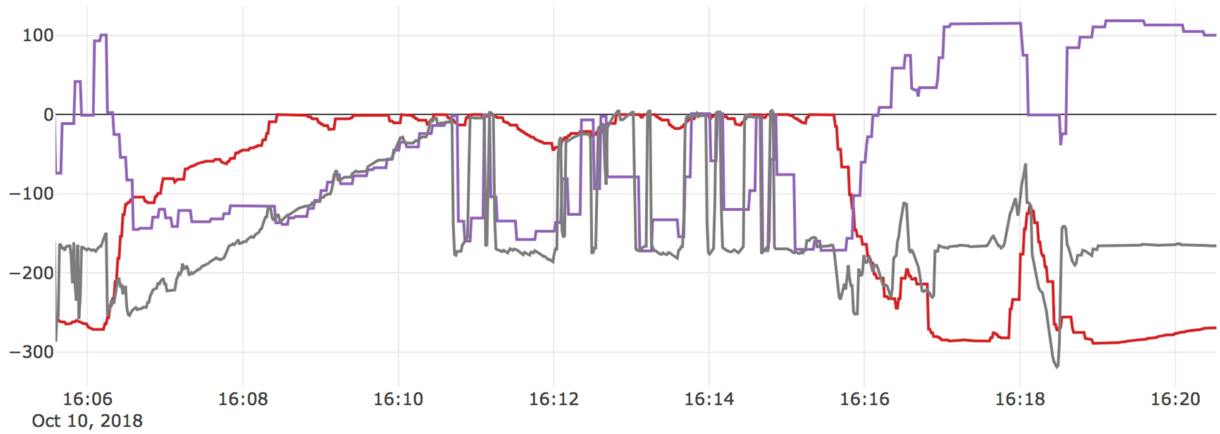

Figure 6: Ramp limit example with simultaneous ESS set points commanded. Red – Solar generation; Purple - ESS output; Grey – Net (PV+ESS) output. In this case, ramp limiting function appears to adversely interact with ESS set points, causing the system to fail to manage system ramp rates to <10% of power/minute.

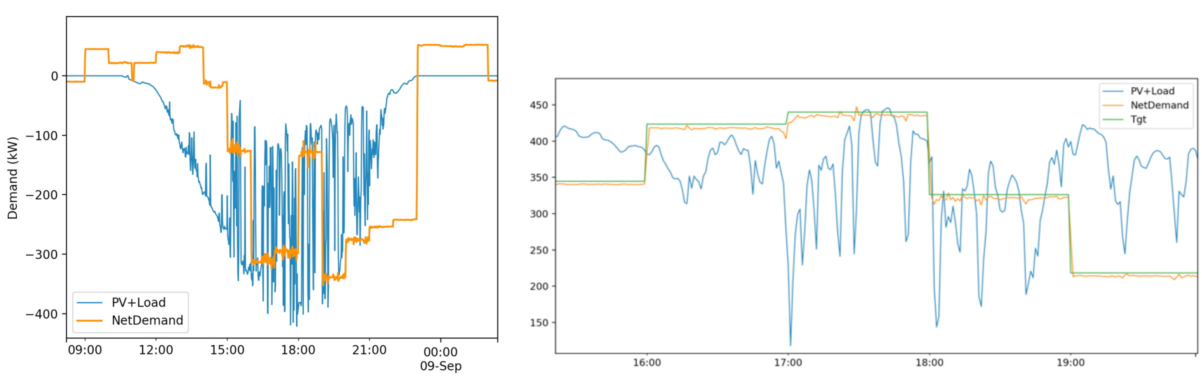

Figure 7: Controlled vs uncontrolled output under variable irradiance conditions.

Table 3: Summary of PV+ESS System Ramp Events, Aug 2019

Rate of Change (% Nameplate per min.)	Control Scenario	
	PV-Only	*PV+ESS*
>40%	1	0
>30%	4	0
>20%	71	0
>10%	436	14
>5%	1,642	180

4.3 Time-Series Solar and Load Predictions

During the initial stages of the SunDial system deployment, we found that our ability to effectively generate strategic dispatch schedules was greatly constrained by shortcomings in both the load and solar time-series prediction methodologies as initially implemented. Said another way, the benefit from strategic DER dispatch was greatly constrained by uncertainty in the predicted state of the system. We therefore iterated multiple times to develop a more accurate method for characterizing load shape and robust mitigation strategies to account for uncertainty in these predictions and to account for real-world issues associated with, e.g., loss of communication.

Multiple methodologies were evaluated for predicting time-series loads and solar production within the system's resource portfolio. In the case of load predictions, analysis showed that a regression-based approach using site-agnostic predictors (in particular, time- and calendar-based schedules, site-agnostic operational classifiers, historical + real-time load data and prediction errors) produced prediction accuracy of 5% root-mean square error (RMSE) for 1-hour ahead predictions to 10% RMSE for 24-hour ahead predictions over the course of the trial deployment. In the case of solar predictions, weather-based production forecasts from satellite data were enhanced by learning from the site's production over time, using classifiers such as time of day, total daily forecast irradiance, and recent production data to refine these results. Implementation of a revised solar model reduced RMSE from 13-20% to 11-14%, depending on time to prediction (see Figure 8).

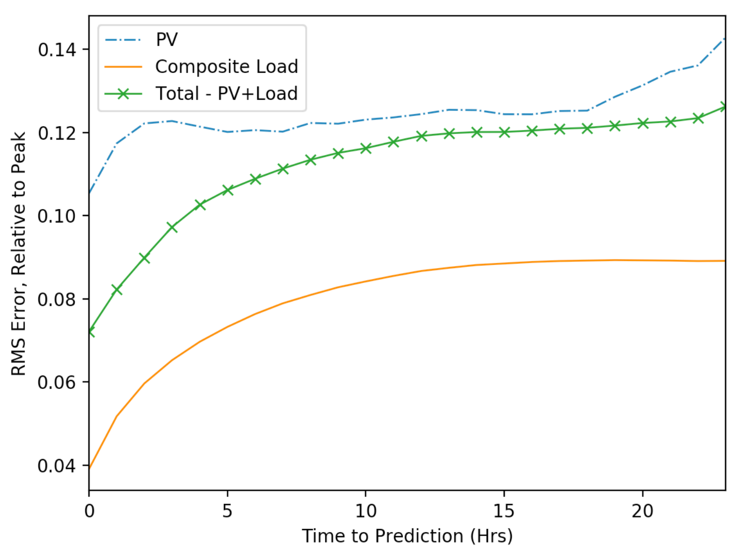

Figure 8: Prediction error as a function of time to prediction

Aggregate load predictions were summed together with solar production forecasts to generate a composite net load prediction for the system (Figure 8), with aggregate net load prediction RSME ranging from 7 to 12%, depending on time to forecast. A snapshot of net load forecast (5-hour ahead prediction) compared to actual net load over a multi-week period is shown in Figure 9. As shown the net load prediction generally captures the composite load shape. Because we minimize RMS error, one limitation of our approach is that it is not tuned to capture peaks, so it will tend to underestimate peak production on a given day. The prediction tool also tends to perform worse on weekends (when the load portfolio has its largest errors) and on days with variable weather (when solar production forecasts are least predictable).

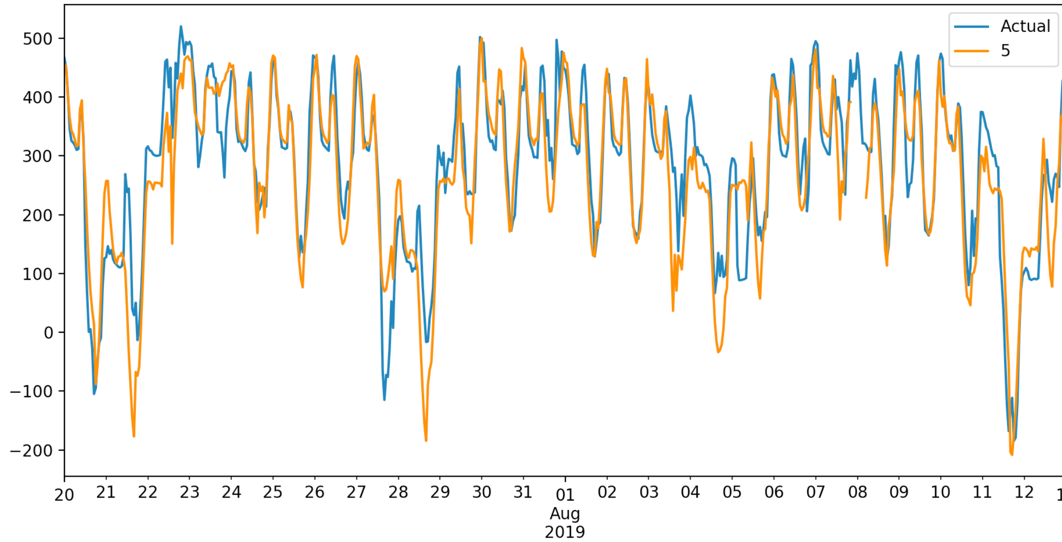

Figure 9: Snapshot of Actual vs 5-Hour ahead predicted net load

4.3.1 Load Prediction Methodology and Results

The goal of the load prediction tool developed for this project is to minimize RMS-error (RMSE) for a continuously updating hourly time-series prediction over a 24-hour time horizon for a virtual portfolio of customer facilities. Predictions for the portfolio of customers were calculated by generating predictions from individual participating facilities using schedule information and historical data, and summing load across facilities to produce an aggregate portfolio forecast.

Three facilities participated in the SunDial pilot: a school, a food-processing facility, and a food-production facility. Due to differences in the magnitude of the participating facilities relative to each other and to the underlying PV and ESS resources available, we applied different scale factors to individual loads to approximately equalize peak demand across the three sites, and to normalize total demand to approximately equal total solar plus ESS capacity.[9] The maximum peak coincident demand for the scaled portfolio was 677 kW.

Initial implementation of a facility load prediction algorithm used facility-specific daily- or weekly customer-provided process schedules in combination with historic load data to generate a rolling 24-hour prediction based on heuristic rules about individual customer sites. We had hypothesized that detailed site-specific customer production schedules could be used to achieve higher prediction accuracy than more generic classifiers. However, our field deployment highlighted several shortcomings with this approach, including: (1) inaccuracies in customer-provided schedules; (2) lack of timely access to schedule information; and (3) high level of effort required to develop detailed process knowledge and to maintain up-to-date schedules. Specifically, we found that customers frequently did not know actual production schedules ahead of time and getting customers to reliably communicated these schedules to the project team was challenging. Hence, it was challenging to ingest data in a timely fashion, and we had no good way to audit the quality of incoming data.

[9] A scale factor of 0.1x was applied to the food production facility; 0.2x to the food processing facility; and 1.0x to the school.

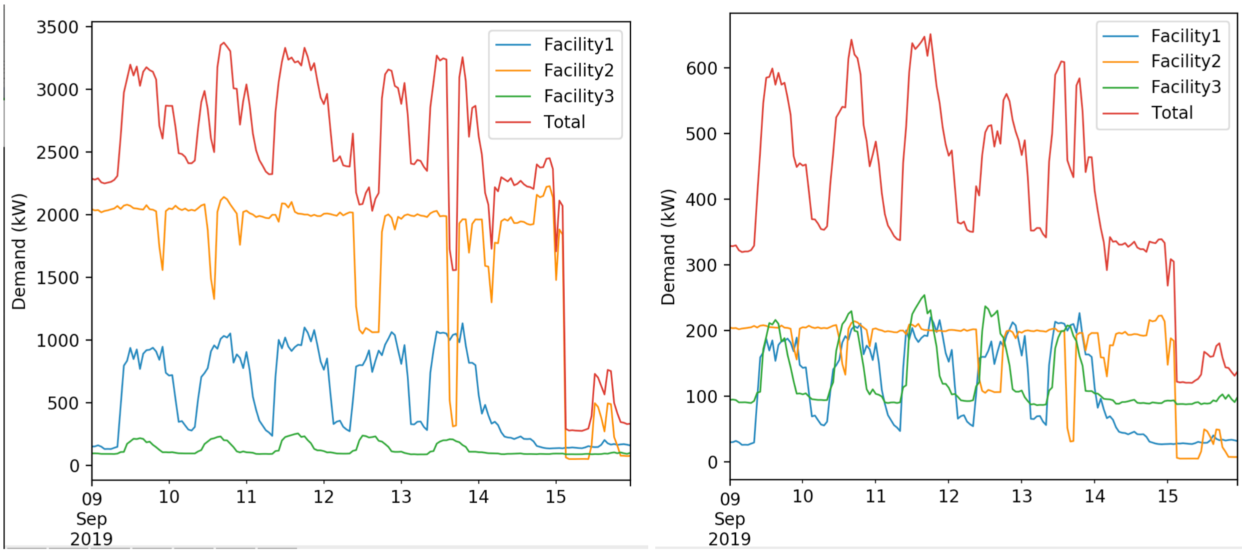

Figure 10: (Left) Unscaled demand for C&I facility portfolio for a typical week; (Right) Demand scaled to approximately equalize individual peak facility loads and normalize aggregate load to PV + ESS capacity.

Based on this initial feedback, subsequent parametric testing was used to evaluate the sensitivity of prediction accuracy to prediction methodology (e.g., site-specific heuristics, linear regression, and multiple machine learning approaches, including dynamic tree classification, k-nearest neighbors, and logistical regression) and predictors employed, with emphasis on predictors that can be auto-calibrated with minimal site-specific knowledge to facilitate scalability for future implementations. A summary of predictors investigated is shown in Table 4.

Table 4: Sample predictors investigated

Category	Predictor
Expected State	On/Off/Part Load Process load type (*food production only*)
Calendar	Hour of day, hour of week, weekday Weekend / Holiday / Summer Schedule?
Detected State	Average Load over 1/2/4/8 hours Predicted load for *same-day, like hours* Predicted load for *previous-day, like hours* Same-day error Shutdown detected?
Weather	Temperature, rH (*not implemented*)

Our analysis showed that a regression-based approach using site-agnostic predictors (in particular, time- and calendar-based schedules, site-agnostic operational classifiers,[10] historical + real-time load data and prediction errors) offered prediction accuracy for the full portfolio of 5% root-mean square error (RMSE) for 1-hour ahead predictions to 10% RMSE for 24-hour ahead predictions over the course of the 16-month trial (see Figure 11). Notably, including site-specific production schedules (approximately 5-8% RMSE) offered only marginal improvements in accuracy relative to the site-agnostic approach, while requiring significantly less effort to implement and maintain.

[10] Refers to generalized classification of facility operational state – for example, “Online”, “Part Load”, “Shutdown”, etc

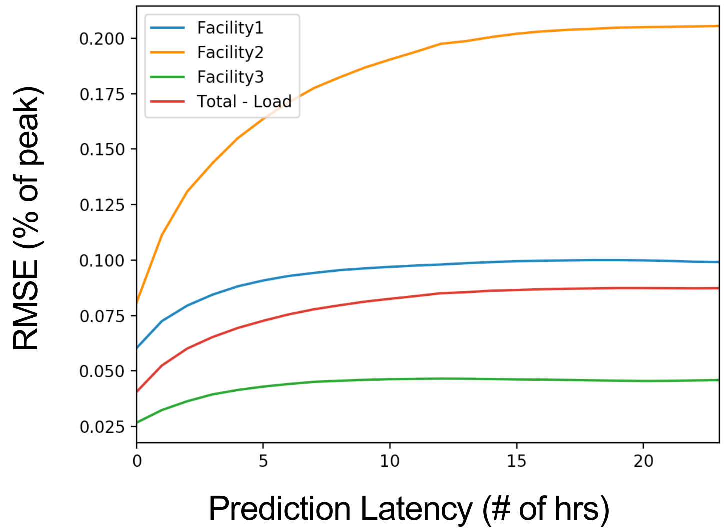

Figure 11: RMSE of load prediction as a function of time-to-prediction.

Figure 12 compares load predictions, as implemented in the GS software, to predictions generated by applying our final algorithmic implementation to the full data capture. Several major revisions were implemented over the course of the pilot, as illustrated by decreasing magnitude of error over time. Specifically, changes between fall 2018 and spring 2019 reflect phasing out reliance on customer-provided process schedules in favor of auto-calibration. Subsequent refinements throughout Summer 2019 modified the specific set of predictors used, e.g., implementing a combination of same-day and previous-day/like-hour corrections. The prediction accuracy "as implemented" in the Global Scheduler for 5-hour ahead predictions over the Fall 2018 period was approximately 16% RMSE after filtering periods when telemetry was offline; in comparison, 5-hour prediction accuracy from Aug-Sep 2019 "as implemented" was 8%; applying final algorithms over the full data capture was 7.5%.

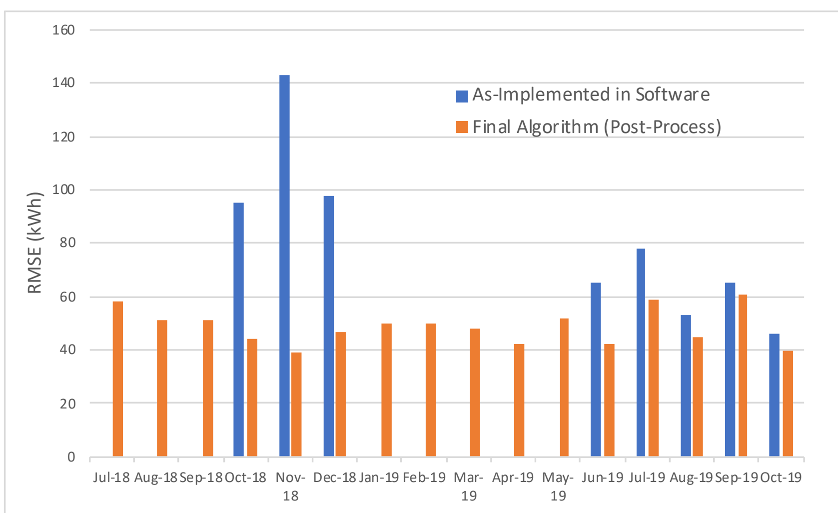

Figure 12: Load prediction accuracy, by month, for five-hour ahead predictions of all buildings. Blue bars show predictions captured by the Global Scheduler in real-time. Orange bars show predictions generated by retroactively applying the final prediction algorithmic implementation across the full data capture.

4.3.2 Solar Prediction Methodology and Results

Like the load prediction tool, the Global Scheduler's solar prediction tool seeks to minimize RMS-error (RMSE) for a continuously updating, hourly time-series forecast of solar production over a 24-hour time horizon. During initial deployment, we used unmodified production forecasts generated by Clean Power Research's SolarAnywhere forecasting service,[11] which applies irradiance estimates derived from weather data to a physics-based model of the solar plant.

Comparison of the first several months of raw production estimates generated by the SolarAnywhere service to actual production showed a systematic bias towards over-estimating production. Over the 5-month observation period from August 2018 through Dec 2018, we found that the solar forecast systematically over-estimates production by approximately 30%, and that days with a high degree of variability in irradiance (e.g., due to intermittent cloud cover) tended to exacerbate this forecast uncertainty. These results are illustrated in Figure 13, which shows total actual vs predicted irradiance (upper and lower left plots), and total actual vs prediction production (upper and lower right plots) on cloudy days (top row) and sunny days (bottom row). As shown, the predicted GHI estimates closely tracks actual GHI (<2% error on sunny days, <10% on cloudy days). Production estimates are within 10% on sunny days, but off by more than 40% on cloudy days. We did not identify a root cause for this discrepancy, but speculate that it is due to variance in specifications used to model the solar site (e.g., unaccounted for shading of modules, variability in solar panel orientation, etc.).

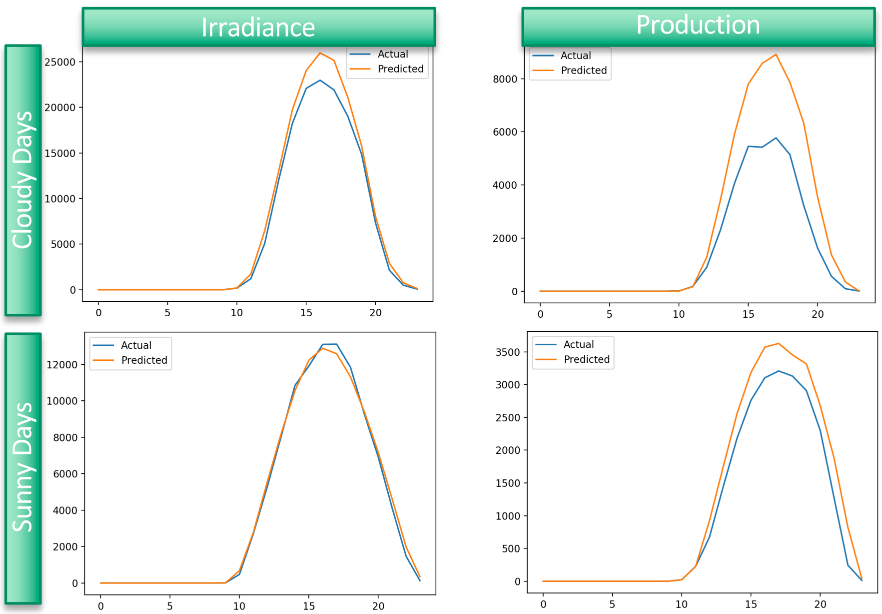

Figure 13: Comparison of total actual vs total 5-hour ahead prediction of GHI and solar production as a function of time of day and cloud cover.

To address this issue, we implemented a linear regression module that ingests weather-based predictions (such as those generated from the SolarAnywhere model) in combination with other predictors (time of day, total daily forecast irradiance, recent production data, and recent prediction error) to correct these results. The improved performance in prediction accuracy over the course of the pilot is shown in Figure 14. In the aggregate, implementation of revised solar model reduced RMSE from 13-20% for the unmodified SolarAnywhere model to 11-14% for the modified prediction approach, depending on time to prediction.

[11] Solar Anywhere. https://www.solaranywhere.com/

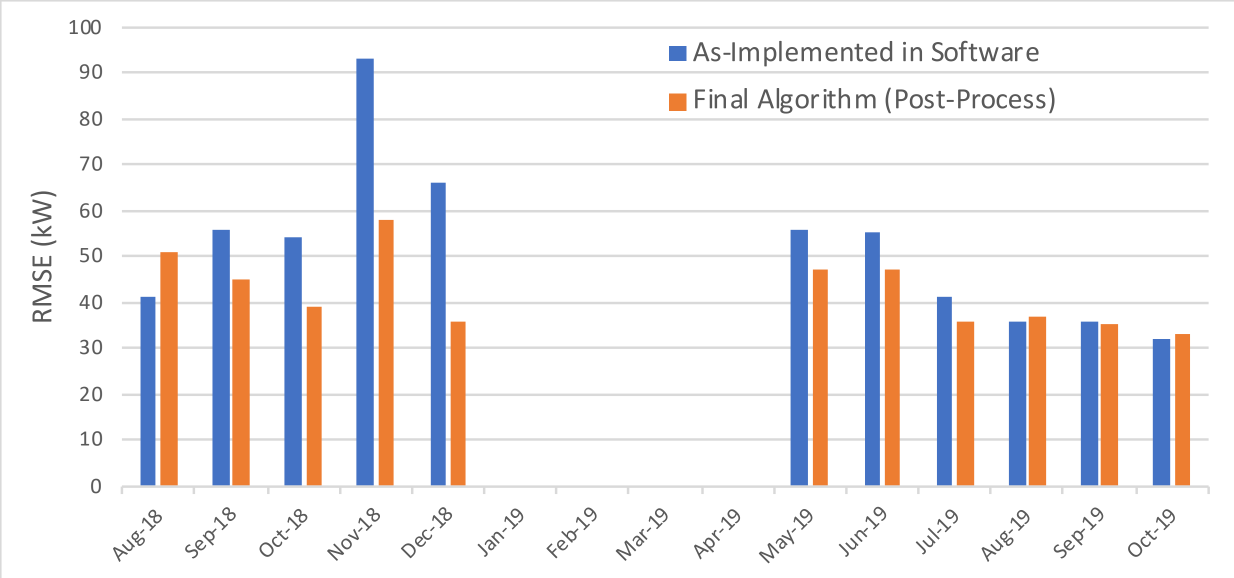

Figure 14: Solar prediction accuracy, by month, for five-hour ahead predictions. Blue bars show predictions captured by the Global Scheduler in real-time. Orange bars show predictions generated by retroactively applying the final prediction algorithmic implementation across the full data capture.

In addition to improving prediction accuracy, as measured by RMSE, we identified conditions that correlate with high degrees of error. Specifically, our analysis shows that the uncertainty in solar production forecasts strongly correlates with the variability of cloud cover (Figure 15). Therefore, strategies such as thresholding days based on the gross sum of predicted irradiance provides a strong indication of forecast uncertainty, which can be incorporated into decisions such as how much battery reserve capacity to maintain.

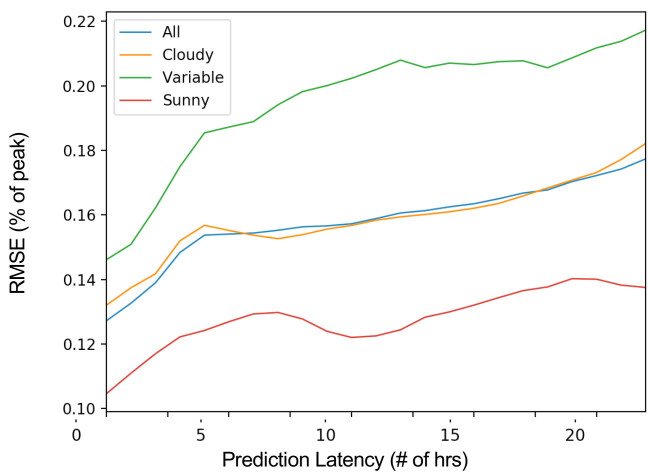

Figure 15: Effect of cloud cover on prediction accuracy

4.4 Evaluation of Strategic Dispatch

4.4.1 Overview

A summary of use cases tested over the course of the SunDial pilot is shown in Table 5.

Table 5: Summary of use cases tested during the SunDial pilot

Use Case Description	Days tested	Description
Peak minimization + Energy Cost Optimization	104	Minimizes peak load and uses excess energy to shift energy from low to high cost periods

Backfeed Minimization + Energy Cost Optimization	15	Minimizes reverse power flow
Virtual Peaker Plant	53	Minimizes and firms net load from 3pm to 7PM
Load Shape	32	System minimizes error relative to a target power output
Power Firming	12	System matches actual net load to predicted net load as predicted at a prior point in time
Cost Optimization	10	Shifts energy from low to high cost periods

Results from this testing are presented below. Data collected from the actual field deployment are supplemented with results generated by re-running the Global Scheduler software given identical exogenous conditions to those observed during our field trial (e.g., solar production, demand, and prediction data based on observations as known at the time of prediction), but with modified control parameters, prediction algorithms, and optimization algorithms. This approach is used to:

(1) Extend results of use-case testing to cover additional time periods, e.g., to enable evaluation of a single use case over longer periods than was feasible during field deployment.

(2) **Compare Global Scheduler performance relative to alternative control strategies:** in particular, we can explore performance relative to a "Perfect Information" case, in which the Global Scheduler is provided with perfectly accurate predictions of PV production and load; and a "Heuristic" case, in which DERs are dispatched using a simple rules-based strategy. The results for the perfect information case reflect the near-optimal dispatch strategy; the results for the heuristic cases illustrate outcomes from a rules-based ESS dispatch strategy.

(3) **Test software updates:** One challenge encountered during the pilot deployment is that we frequently encountered issues, either with software or data collection, that were subsequently corrected. As such, the Global Scheduler software represented a continually improving target. In particular, our load and solar predictions and several aspects of our peak minimization strategy were updated multiple times throughout the program. Hence, later versions of the software give the best end-of-project metric.

4.4.2 Peak Shaving + Energy Cost Optimization

The primary objective in this use case is to minimize peak net demand over a defined time horizon, while the secondary objective is to use excess storage capacity to optimize energy cost given a time-differentiated cost of energy. This use case was tested over several different time periods throughout the pilot demonstration while varying (1) time horizon over which peak demand is minimized (daily vs monthly) and (2) a baseline threshold above which an additional kW of demand incurs a non-zero cost.[12] In addition, the software in use underwent multiple revisions throughout the demonstration period to address issues as they arose. Results from two discrete periods during which this use case was implemented are shown below; these results are broadly illustrative of the overall system behavior.

Results from a two-week period in May-June 2019 are shown in Figure 16. As shown, over this time period peak load is reduced by 16% relative to a non-storage baseline. On days where a peak condition is not anticipated, the system maximizes exports during a late afternoon increase in price.

For this use case, there is significant down-side risk in trying to regulate around an overly aggressive (low) peak, as even a minor miscalculation can cause the system to prematurely discharge the storage system and therefore set a new peak. By way of example, testing the system with similar parameters over a different time frame yielded results as shown in Figure 17. In this case, the prediction algorithm in use at that time underestimated net load, which caused the overall system to seek to regulate to a more aggressive (lower) peak load than was feasible given the energy available for load shifting. Once this initial peak was set, the

[12] An initial baseline threshold is useful to insure that the system does not try to regulate to an unrealistically low target threshold.

ESS was primarily used for the secondary objective (time shifting of energy). The reverse scenario – i.e., a case where the system over-estimates demand, and hence is overly conservative in how it strategically deploys available energy – also presents a challenge, as it increases the likelihood that generation will exceed demand.

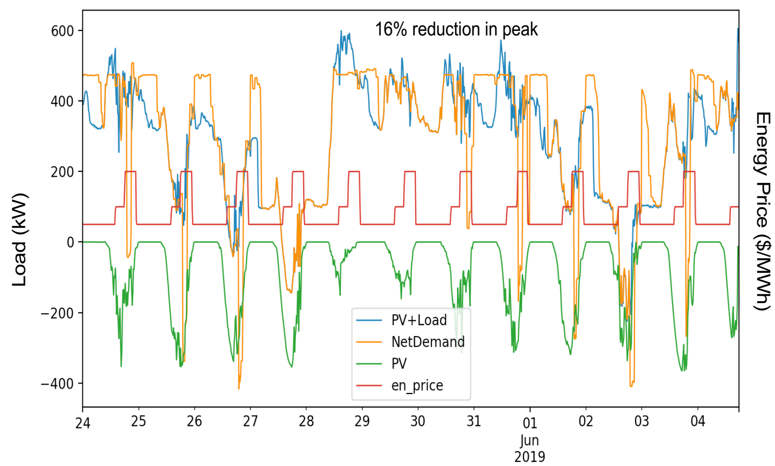

Figure 16: Illustrative example of peak shaving + energy cost optimization, May-June 2019 (13 days)

These issues were mitigated first by revisiting and improving the underlying prediction algorithms to improve accuracy, and second, by implementing additional control objectives to account for forecast uncertainty. Specifically, two additional objectives functions were implemented:

(1) Imposing a non-linear cost on using ESS energy at lower states of energy (SOE), such that as the BESS approaches low SOE conditions, it will only be used for scenarios with a high-marginal value of discharge. For example, at low SOE, if 1 kW of demand reduction requires 5 kWh of storage, the GS would opt not to select

(2) Favoring control strategies that delay *when* costs are incurred for as long as possible within a planning window. For example, for a scenario in which the system expects to set a new peak within its scheduling window, all other factors being equal, this control objective would favor a control strategy that sets the new peak later in the scheduling window over one that sets the new peak earlier in the scheduling window.

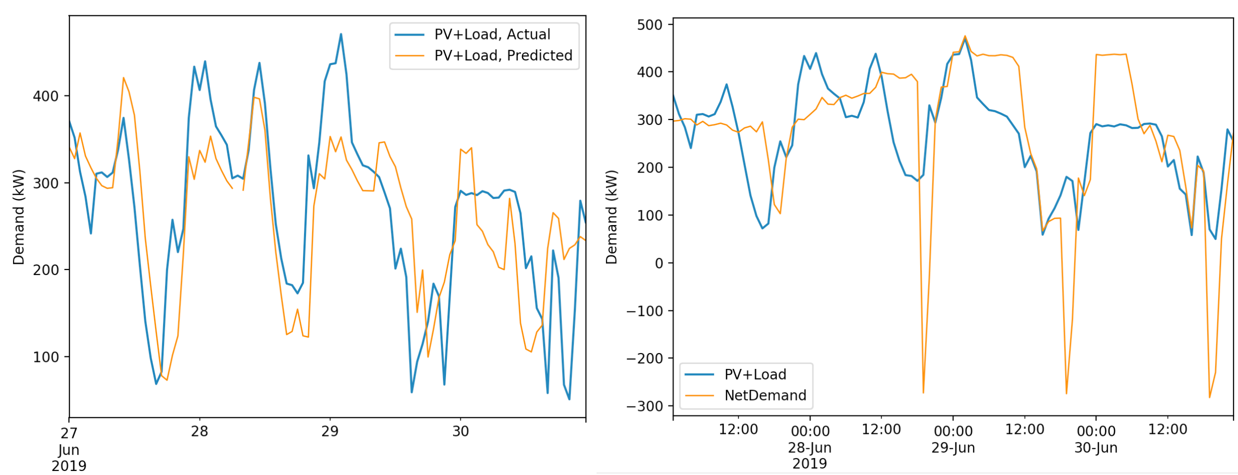

Figure 17: Peak Shaving + Energy Cost Optimization, Jun 27-July 1. (Left) Comparison of predicted to actual net load. Note significant forecast error on 6/28 and 6/29 (Right) Resulting net system load. Note that a new peak is set on 6/29 (note: all times shown are EST + 5).

Re-visiting the same scenario with improved prediction algorithms and having implemented the above hedges against uncertainty yields results as shown in Figure 18. In this case, the prediction closely tracks the actual net load (left side), and the system avoids setting a new peak (right side). In this case, peak demand was 368kW vs 476kW prior to making the above adjustments.

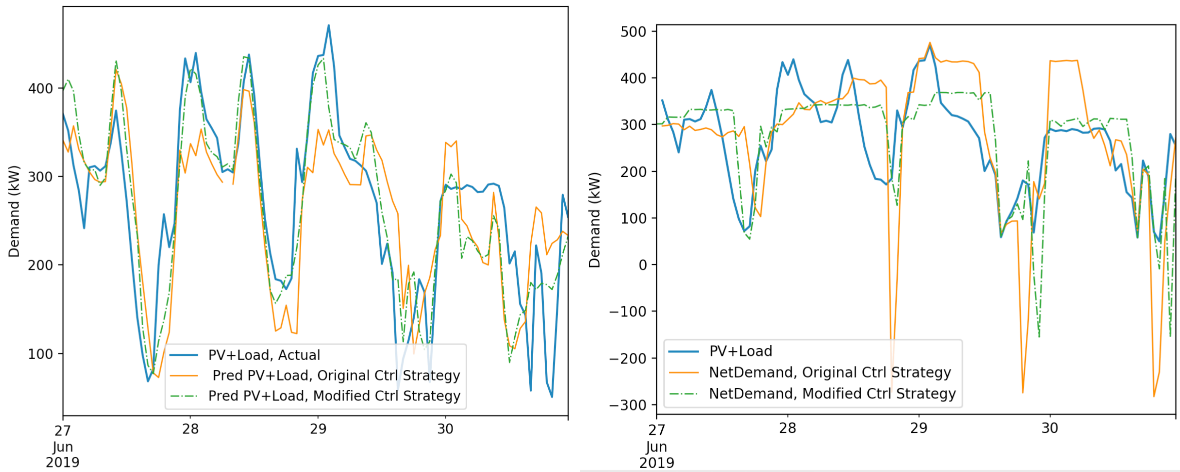

Figure 18: Peak Shaving + Energy Cost Optimization, Jun 27-July 1. (Left) Comparison of predicted to actual net load using our original prediction algorithms (as implemented at that time), and revised prediction algorithm (end-of-project baseline). (Right) Comparison of resulting net system load for the two scenarios.

Results for this use case were evaluated for a period from September 2018 to December 2018, and from June 2019 to October 2019 for several alternate control scenarios:

(1) "Heuristic": Uses a control strategy in which the system regulates net power output below a target threshold, but does not use best-available predictions to assess whether the system is likely to be able to meet this threshold. As such, this strategy runs the risk of trying to satisfy a target objective that is not feasible, and hence prematurely discharges its BESS. Results from this case are illustrative of a rules-based strategy that does not incorporate forecasting.
(2) "Perfect Information": Assumes Global Scheduler had perfect predictions about the future loads and PV generation. Results from this case represent a near-optimal strategy.
(3) "Baseline": Represents results generated by the Global Scheduler using actual prediction algorithms, but prior to implementing prediction uncertainty.[13]
(4) "Uncertainty Mitigation": Represents results generated by the Global Scheduler using actual prediction algorithms and objectives tuned to mitigate prediction uncertainty.

A comparison of the monthly peak demand for each of these control cases relative to a 'No-ESS' case is shown in Figure 19. The resulting monthly reduction in peak demand (Table 6) shows that the modified "uncertainty mitigation" strategy outperforms the "Baseline" approach in every month with an average increase in monthly peak load reduction of 92%, and outperforms the "Heuristic" approach in 7 out of 9 months, with an increase in average peak load reduction of 41% . The average monthly load reduction for the modified control strategy closely approximates that of the "Perfect Information" case with an average peak load reduction of 13% below that of the "Perfect Information" case (113 vs 129 kW).

[13] Results for this case were only generated for the period from June-Oct 2019

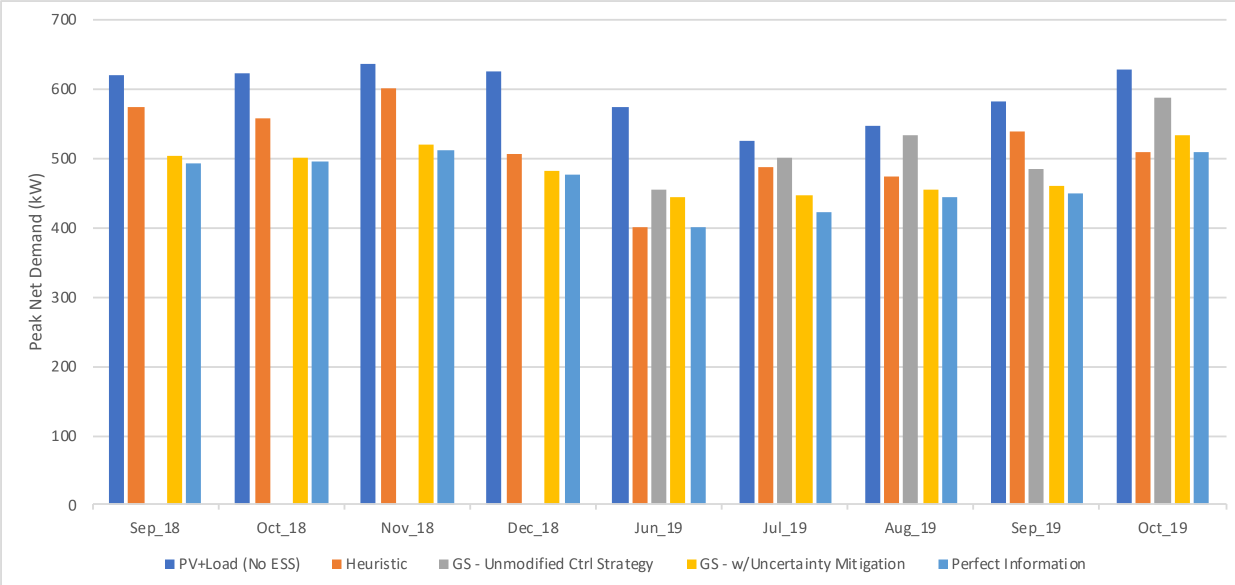

Figure 19: Comparison of monthly peak load as a function of ESS control strategy, relative to no-ESS case ("PV+Load")

Table 6: Monthly peak load reduction relative to no-ESS case

Month	Heuristic	GS – Unmodified Ctrl Strategy	GS - w/Uncertainty Mitigation	Perfect Information
Sep-18	48	n/a	118	128
Oct-18	67	n/a	122	128
Nov-18	35	n/a	116	123
Dec-18	121	n/a	145	149
Jun-19	175	121	130	175
Jul-19	37	23	77	102
Aug-19	73	14	93	104
Sep-19	41	96	121	132
Oct-19	120	41	95	119
Average	**80**	**59**	**113**	**129**

4.4.3 Backfeed Minimization + Peak Load Reduction

This backfeed minimization use case is a variant of the Peak Load Reduction use case presented in Section 4.4.2: in this case, the system is configured with objectives that simultaneously minimize reverse power flow on potential solar surplus days while also seeking to minimize peak system load on high demand days, an example of which is shown in Figure 20. In total, we observed 24 days in which generation exceeded demand for a portion of the day, with a total of 9,450kWh of energy backfed to the distribution system; under a perfect information scenario, the system would have been forced to shed 299 kWh of energy, and would be able to productively shift the balance of 9,155kWh that would otherwise be curtailed. The Global Scheduler was modeled to capture approximately 70% this total potential load shed.

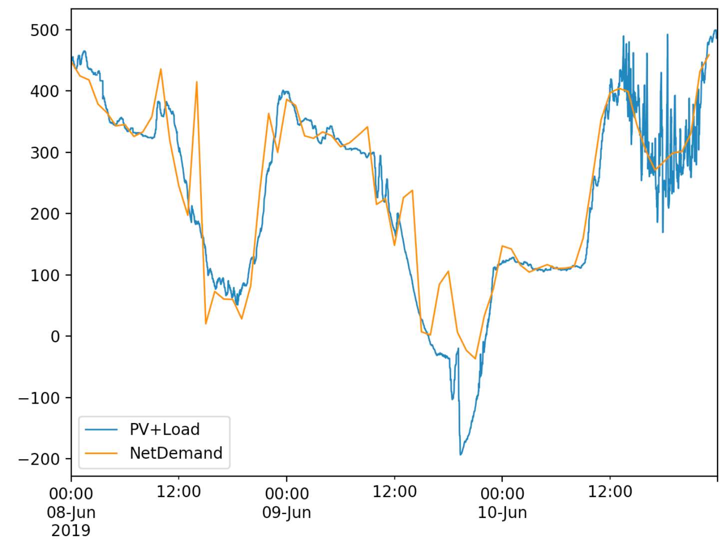

Figure 20: Backfeed mitigation example

4.4.4 Virtual Peaker Plant

This use case implements a virtual power plant that provides firm capacity during periods of peak stress on the power distribution system. The system's primary objective is to present a predictable, constant, and minimized demand profile during the four-hour period from 3PM to 7PM eastern daylight time (6PM to 10PM UTC). This use case is representative of tariffs established by the Massachusetts Clean Peak Standard, which offers incentives for firm production from renewable generation during designated peak hours. The system had two secondary objectives: (1) to minimize overall peak for the virtual power plant, and (2) within these constraints, to minimize energy imports (i.e., charge using local solar).

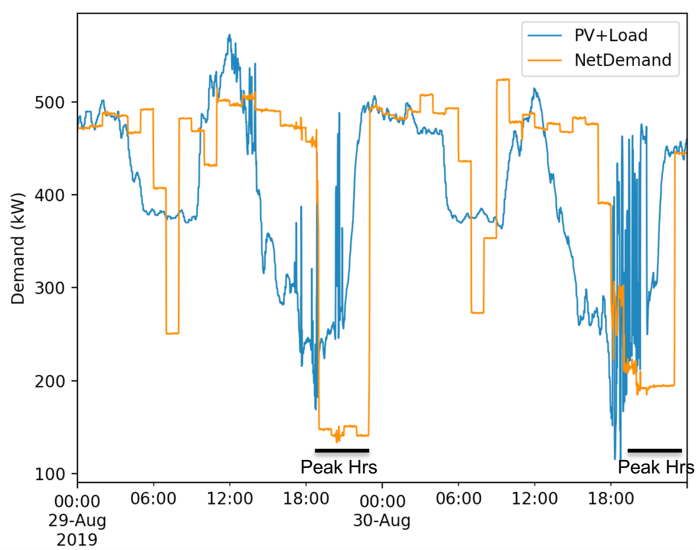

Figure 21: Virtual peaker plant example, Aug 2019. Black lines indicate targeted periods for peak load reduction

An example of net system load after implementing the virtual peaker case is shown in Figure 21. Over a 10-day trial implementation (8/22 to 8/31), the virtual peaker plant implementation realized an average peak reduction of 245 kW over the targeted four-hour window relative to a "no storage" case. Relative to an identically dimensioned ESS using a constant discharge strategy to minimize afternoon peak, the virtual peaker plant implementation reduces afternoon peak by an average of 35kW. Hence, replicating the virtual peaker plant's afternoon peak-load reduction using a constant discharge strategy would require approximately 140kWh additional energy storage (i.e., 35kW x 4 hours), or 20% of the ESS size. In addition, the secondary peak-load reduction objective reduces the peak over the full day by approximately 23 kW.

Subsequent to the field test, modifications were implemented to improve performance relative to a "perfect information" case: (1) implementation of a non-linear cost to battery discharge at low SOE to mitigate forecast uncertainty (see Section 4.4.2 for additional discussion), and (2) revised prediction algorithms. Implementing these changes further improved performance, providing an average of 48kW reduction in the afternoon peak, equivalent to 196kWh (i.e., 4 hrs x 48kW) of additional energy storage, relative to a constant discharge case; and reducing overall peak by approximately 28kW. Results for four different control scenarios are summarized relative to a "No ESS" case in Table 7.

Table 7: Comparison of Virtual Peaker Plant afternoon peak load reduction scenarios

Scenario	Avg Afternoon Peak (kW)	Average Peak Load Reduction (kW)	
		Relative to No ESS	Relative to Constant Discharge
PV+Load Only (No ESS)	421	0	-207
Constant Discharge Strategy	214	207	0
GS, As-Implemented for field test	178	243	36
GS, Modified Control Objectives	166	255	48
Perfect Information	140	281	74

4.4.5 Output Firming

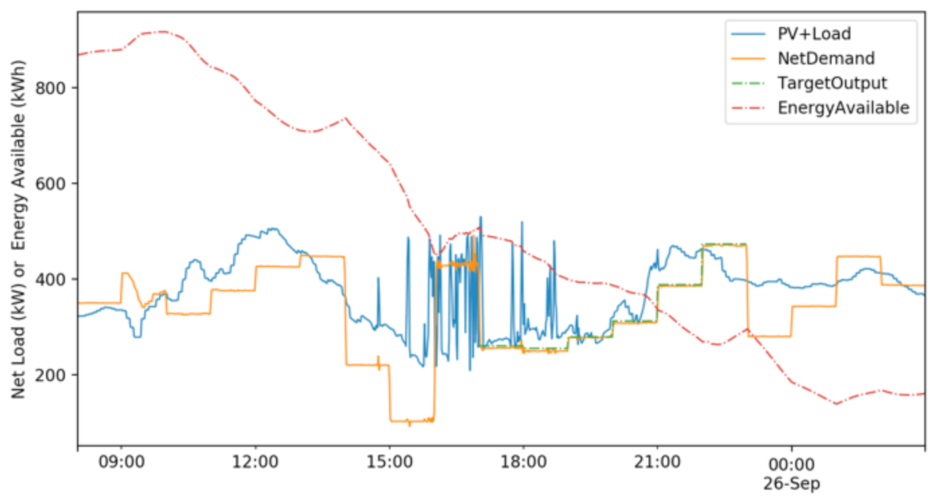

Figure 22: Firmed system output over a six hour period, based on net load prediction generated seven hours prior.

An alternative way to use the system's predictive capability is to provide firm, predictable demand for a defined time window. To support this use case, the Global Scheduler was configured to match the system's actual load shape to the load shape as predicted at 6AM (EDT) / 10AM (UTC) during the hours from 1PM to 7PM (EDT) / 5PM to 11PM (UTC), an example of which is shown in Figure 22. In this case, the ESS is set to a midpoint state of charge (500 kWh) at the start of the firming period, and then regulates the net system output over the next six hours to match the target output defined by the predicted net load. In this case, the prediction underestimated net load over the target time period, so the ESS state of charge decreases over the firming window, but the cumulative error is less than the available ESS reserve.

Results from this use case were modeled using data gathered over the full pilot period to estimate the energy storage required to reliably follow predicted system output under various scenarios. The energy storage required was estimated by calculating the cumulative error in net load (predicted – actual) for each day over the target time window, and then adjusting for the ESS round-trip efficiency. Results were generated for scenarios that varied the time of prediction relative to the start of the firming period (1 hour / 4 hours / 6 hours prior) and the duration of the firming period (4 hours / 6 hours), and we compared these cases against a base case scenario in which the system regulates its output to equal the average historical net load over the target time period.[14]

Table 8 and Table 9 show the percentage of hours that the specified BESS would firm for each of the four prediction scenarios. As shown, for the 6-hour firming case (Table 8) a 1MWh BESS firms nearly 99% of hours for the 1-hour prior prediction; the 6-hour-prior case requires in excess of a 1.5MWh BESS to achieve the same performance; and the base case requires >2.5MWh. The majority of the uncertainty occurs on weekends and holidays (when C&I loads within the project's virtual portfolio are typically low, but occasionally and unpredictably remain in a production state). If we limit analysis to weekdays-only, the 1-hour ahead case becomes 99.7% reliable at 1MWh, and the 6-hour ahead case becomes 98% reliable.

Table 8: Percent of all hours firmed by target battery size for a 6-hour firming window

BESS Size (kWh)	Prediction Time, Relative to Start of Firming Period			
	11AM EDT	8AM EDT	6AM EDT	Base Case
1,000	98.9%	96.8%	95.3%	81.7%
1,500	99.7%	99.3%	98.3%	93.6%
2,000	99.9%	99.8%	99.5%	97.0%
2,500	100.0%	100.0%	100.0%	98.7%

Table 9: Percent of all hours firmed by target battery size for a 4-hour firming window

BESS Size (kWh)	Prediction Time, Relative to Start of Firming Period			
	11AM EDT	8AM EDT	6AM EDT	Base Case
1,000	99.6%	98.9%	96.3%	95.1%
1,500	99.9%	99.7%	99.1%	98.9%
2,000	99.9%	99.9%	99.8%	99.7%
2,500	100%	100%	100%	99.9%

One implication of this analysis is that approximating the performance of one of the predictive cases with the base case requires significantly more storage capacity – for example, the 4 hour prior / 4-hour firming window case can achieve ~99% reliability with a 1MWh ESS, while the base case would require a 1,500MWh ESS to achieve the same level of performance. In a similar vein, we find that there is a marked difference in efficacy as the prediction time and firming windows increase.

4.5 FLAME Field Performance

4.6 Overview

The SunDial System's baseline energy storage capacity was supplemented by several flexible customer loads that were identified as part of initial site screening during Phases 1 and 2 of this project. These flexible load resources consisted of:

HVAC cooling loads: We had the ability to remotely modify zone temperature cooling setpoints for packaged rooftop units (RTUs) via the school's building automation system (BAS) by +/- 3ºF, shifting

[14] To account for variations in load, separate averages were used for weekend and weekdays.

HVAC consumption to pre-cool the building during periods of excess solar generation (increasing load) and/or ride through periods of high demand (thereby decreasing load). Cooling loads were found to comprise approximately 20% of total site load during summer months and ranged from 0 to 15 percent of total site load during the spring and fall (depending on T_{out}).. In real terms, this equates to up to approximately 50kWh of pre-cool and shed during summer and fall events (again, depending on event duration and T_{out}), typically over a 1- to 3-hour period.

Charging of electrical lift trucks: Charging loads at the food processing facility could be remotely delayed by actuating contactors installed on the circuits servicing the plant's EV charging stations. Without intervention, charging typically occured at night (most typically from midnight to 4AM), but could be delayed to begin as late as 10AM and complete by 2PM. For this usage profile, the primary utility of this load in a high-penetration solar context is to shift load into daytime hours to absorb excess generation on sunny days. In total, approximately 54kWh of load could be shifted over a 4-hour period, or 11kWh after applying the scale factor for the scaled portfolio used in end-to-end testing. The majority of EV power consumption occurs in the first two hours (25 and 19 kWh/h, respectively), comprising approximately 2.5% of peak load, or 5% of average daytime load.

Drag conveyor: Operation of a drag conveyor at the production plant could be shifted by up to five hours a day. Typical weekday usage for the conveyor drag entailed operating for approximately 5 hours from 9AM to 1PM. This could be shifted to start any time within a window from 7AM to 12PM. The conveyor drag provided approximately 90kW (approximately 5% of total site load) of load shift for up to five hours (450kWh total load shift), or 9kW/45kWh for the scaled portfolio used in end-to-end system testing. The FLAME server did not have the ability to remotely actuate the conveyor drag; rather, it was able to provide indicator lights to plant operators indicated "favorable" or "unfavorable" times to operate.

Table 10: Summary of load shift resources and underlying load shift potential for the SunDial pilot

Flexible Load Resource	Total Average Load increase per event (shoulder seasons / backfeed conditions) (kWh)		Total Average Load decrease per event (summer / peak load conditions) (kWh)		Notes
	Raw	Scaled	Raw	Scaled	
EV Charging	54	10.8	0	0	Typical charging pattern is a four hour cycle which peaks in hours 1-2. Baseline operation recharges at night, so load decrease unavailable to support solar.
Conveyor Drag	450	45	-450	-45	Baseline Load is 90kW x 5 hr. Available to be shfited through most of the day. Non-automated asset, relies on customer response.
HVAC	40	40	-40	-40	Highly temperature dependent. Typical summer time (peak events) typically allowed for two hours of load shed at 20kW; typical fall (backfeed-limiting) events typically allowed for two hours of pre-cool at 20kW.
Total (theoretical)	544	95.8	-490	-85	"Realized" load shift incorporates EV and HVAC loads only, as the conveyor was typically non-responsive; "theoretical" load shift includes all three loads
Total (realized)	94	50.8	-40	-40	

A summary of the overall load shift portfolio is shown in Table 10.

Load response events were requested by the Global Scheduler and actuated by the FLAME: To initiate an event, the Global Scheduler communicated user-defined objectives in the form of price signals to the FLAME, which responded with a series of multiple alternate load profiles over a 48-hour window that support the desired objective. The Global Scheduler could, in turn, choose to accept any potential load profiles; a selection by the Global Scheduler would cause the FLAME to schedule events at the corresponding participating facilities. Example price signals sent by the Global Scheduler, and the corresponding predicted load shift events for two days in September 2019 are shown in Figure 23: in both cases, a low price signal early in the day (indicated, for example, a solar surplus) is followed by a higher price later in the day. This price signal triggers an EV charging event (increasing load during the low price period early in the day); and an AC load shift event (increasing load early, and shedding load later).

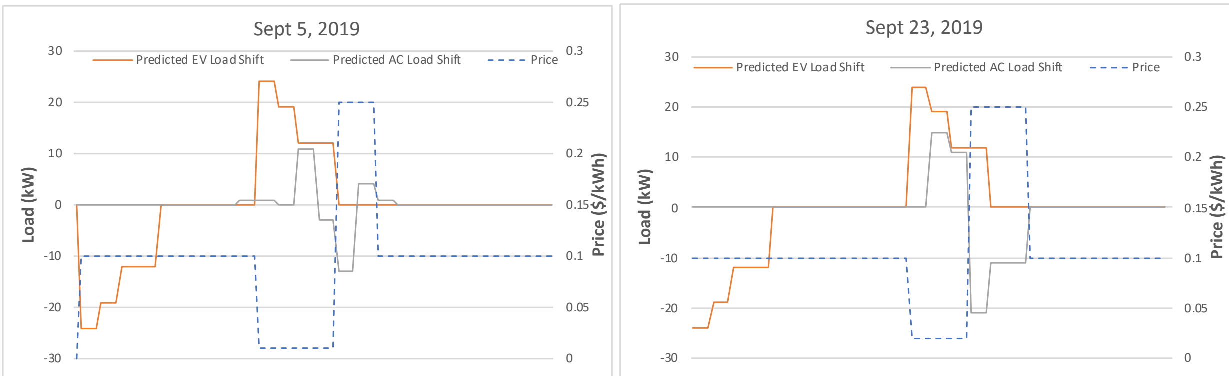

Figure 23: Price signal sent by Global Scheduler and resulting predicted load shift events for two example days

A number of issues delayed full implementation of the SunDial System's load shift capability, including issues with control equipment at customer sites and issues with the FLAME server's load actuation software. However, we were able to call repeated events – typically daily - from August 6th through October 9th, 2019. In total, we called 35 successful EV charging events and approximately 32 AC shifting events. Due to non-responsiveness by the customer and lack of visibility into specific loads, we were not able to verify whether any conveyor drag events were successfully executed during this pilot.

4.7 EV Charging Performance Assessment`

The IPKeys FLAME server called a total of 35 successful[15] load-shifting events between August 6 and October 9, 2019 based on commands from the Global Scheduler. On average, metered data indicate[16] that the average event shifted 54kWh, relative to the 67kWh[17] predicted (see Figure 24). Load shift predictions had a mean average error (MAE) of 25kWh and a root-mean squared error (RMSE) of 30kWh.

[15] Events called on Sundays or holidays when charging would not typically occur were not considered successful events. A separate Fraunhofer analysis of EV charging data (n=90 charging cycles) yielded an average of 53kWh and 58kWh over four- and six-hour time windows, respectively.
[16] We multiplied the submetered power for two EV chargers by six to estimate kWh for all twelve pallet riders.
[17] The EV charging model predicted hourly kWh values of 24, 19, 12, 12 for a four-hour event.

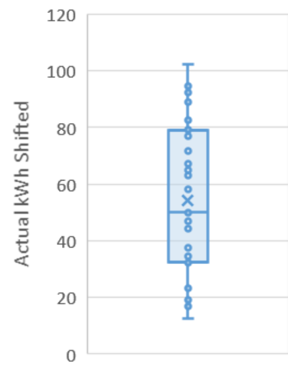

Figure 24: Whisker plot of EV charging energy shifted (n=35); the center box delineates the middle quartiles, the middle line the median, and the "X" the average.

Clearly, the high variability in energy shifted (max = 102 kWh, min = 13kWh) reflects variability in EV usage patterns, including that of the two EVs on the submetered charging circuit. Figure 25 plots the number of hours each of the 12 pallet-rider EVs spend *away* from the charging station, as determined from reviewing a video feed of the EV chargers over an 11-day period in 2019. Although only a snapshot of time, the daily usage averaged over all riders varies from 3 to 14 hours per day. Assuming that time out correlates at least somewhat with EV usage and, hence, energy required to recharge the EV, this represents approximately a 4.5-fold range in charging energy observed over an 11-day stretch. Over this period, the average away hours per day for the two rider EV charging circuits we submetered is very similar to that of the entire rider fleet, i.e., 8.2 vs. 8.1, with a MAE of 1.3 h/day (1.4 excluding Saturday).

About half of the events had appreciable (i.e., 10%+ of the average load shift) additional energy shifted outside the four-hour event window. As shown in Figure 26, this typically occurred after the four-hour window. This increased per-event charging energy by 15kWh for an average load shifted over *all* events of 61kWh (MAE of 26kWh and a RMSE of 30kWh), about 10 percent less than the predicted 67 kWh.

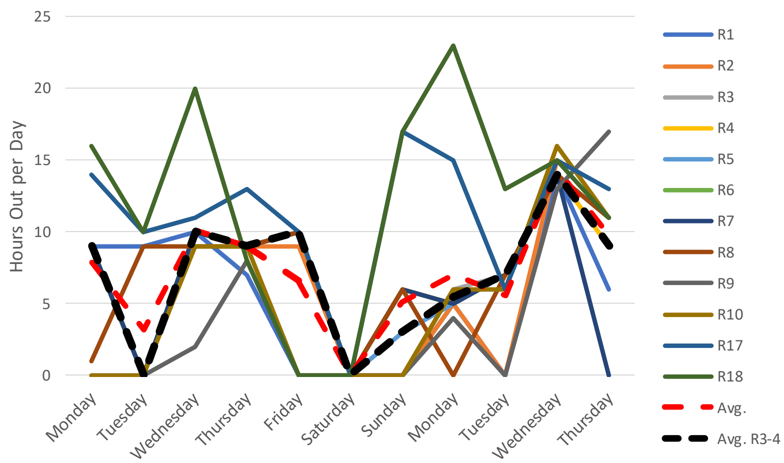

Figure 25: Hours per day away from the EV charging station observed from video analysis over an 11-day period in 2019.

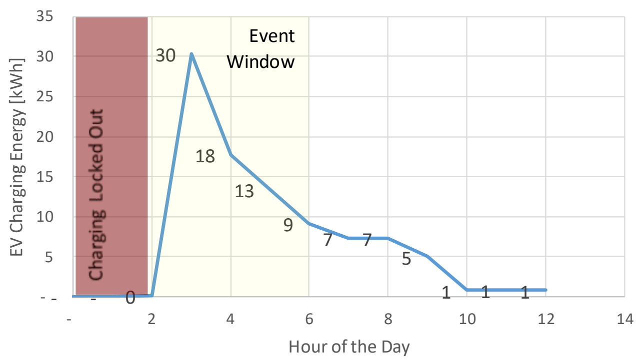

Figure 26: Example of EV charging with appreciable post-event charging. The 1kWh values represent standby mode.

4.7.1 Customer Acceptance

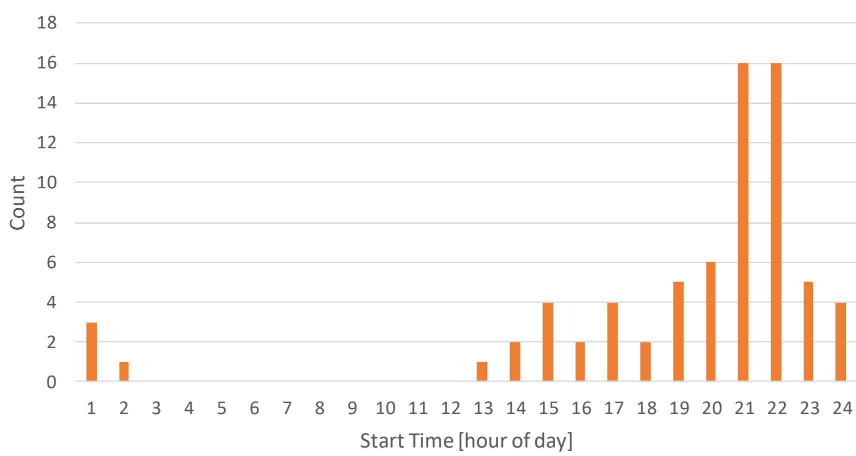

Figure 27: Pallet rider EV charging start time distribution (n=71)

The food production facility manager indicated that the EVs typically require 6-8h to charge, followed by a similar cool-down period prior to use to maximize battery lifetime. Since EV charging typically began in the evening (see Figure 27), the EVs batteries would have plenty of time to complete charging prior to the resumption of pallet rider use the following afternoon. In contrast, we called many events over a two-month period where we locked out charging until 10 or 11AM, i.e., they would finish charging right around the time when they would next be used. Nonetheless, we did not receive any complaints from the facility manager about problems with the EVs.

4.8 AC Load Shifting Assessment

Evaluating the load-shifting performance of the school AC systems proved to be quite challenging for two main reasons. First, the magnitude of AC power draw on many days was a similar fraction (10%) of the whole-building power draw as the uncertainty/variability in whole-building power draw, even after taking into account how power draw varies as a function of time of day. Second, although we learned in May 2019 that the school has SiteSage power submetering for numerous circuits, only four of the seven RTUs were submetered. In late May, data were no longer recorded for those circuits because the school converted their wireless networks from 2 to 5GHz. Although we quickly procured the necessary hardware to submeter all seven RTUs (i.e., CTs for the additional three RTUs and new wireless bridges), the school took several months to bring the systems online, i.e., the three RTUs came online in mid-September and the other four RTUs in October. Consequently, we acquired a very limited set of cooling-season data. Nonetheless, we were able to demonstrate that the FLAME called and actuated events over a significant portion of the summer and early fall of 2019.

4.8.1 Basic Approach – Whole-building Power Data

We evaluated the AC load shifting performance using whole-building power data by first establishing a baseline power draw profile based on days without events and then comparing power draw during events to the baseline. The baseline approach modeled hourly building energy consumption as a function of time of day and cooling-degree hours (CDH), where CDH equals the difference between the outdoor temperature, T_{out}, and the balance temperature, T_{bal}[18]:

$$CDH = T_{out} - T_{bal}.$$

Using whole-building power, P_{shool}, data and SiteSage T_{out} data[19], we calculated the coefficients for a linear CDH regression for each hour, *t*, of the school day, based on all normal school days without an event, yielding the coefficients shown in Table 1.

$$P_{school} = C_1(t) + C_2(t)*CDH.$$

Table 11: CDH Regression Coefficients for the School

	Fall 2018 – T_{bal}=60°F Weather Underground T_{out}			Spring 2019 – T_{bal}=52°F SiteSage T_{out}			Summer 2019 – T_{bak} = 60°F SiteSage T_{out}		
Time	$C_2(t)$	$C_1(t)$	R^2	$C_2(t)$	$C_1(t)$	R^2	$C_2(t)$	$C_1(t)$	R^2
8-9	3.2	177	0.80	0.82	191	0.25	0.96	77.4	0.16
9-10	2.8	181	0.85	0.74	195	0.27	1.18	76.8	0.27
10-11	2.3	188	0.83	0.87	192	0.35	0.68	85.2	0.12
11-12	2.4	185	0.90	1.00	192	0.44	0.91	81.1	0.23
12-13	2.4	182	0.77	1.06	192	0.49	0.85	83.5	0.35
13-14	2.3	161	0.85	1.04	172	0.58	0.9	84.1	0.18
14-15	2.2	155	0.92	0.96	165	0.55	0.97	81.8	0.32
15-16							0.93	77.6	0.35
16-17							0.96	67.4	0.49

[18] The balance temperature is T_{out} when the building requires neither space heating or cooling. In practice, this can vary by zone in a building, so effectively we define T_{bal} as the highest T_{out} when school AC consumption equals zero.

[19] WeatherUndergroud.com data for Leominster/Fitchburg airport, KFIT, yielded similar coefficients and R^2 values.

It is not clear why the regressions for fall 2018 had a much higher goodness of fit than those in spring 2019. The rather poor accuracy for the summer 2019 regressions likely reflecting the smaller sample size due to several event days and a period of time when the AC did not work.

We observed that some days had appreciably higher or lower power draw than expected. To address this, we also using a baseline incorporating a same-day adjustment based on the average power in the hour prior to an event, P(t-1):

If $T_{out}(t\text{-}1) > T_{bal}$: $P_{adj} = P(t\text{-}1) - (C_1(t\text{-}1) + C_2(t) * (T_{bal} - T_{out}(t\text{-}1))$

$$P(t) = C_1(t) + C_2(t) * (T_{bal} - T_{out}(t\text{-}1)) + P_{adj}.$$

If $T_{out}(t\text{-}1) \leq T_{bal}$: $P_{adj} = P(t\text{-}1) - C_1(t\text{-}1)$,

$$P(t) = C_1(t) + P_{adj}.$$

Consequently, we used these equations to calculate the baseline power values for each hour and evaluated changes in whole-building power due to events relative to the baseline values.

Unfortunately, the uncertainty in the hourly load predictions (~±5 to 10%, depending on time to prediction) is similar to the magnitude of the *total* load shift predicted by IPKeys from AC events (an average of 23kWh for 14 valid events once school started on August 28, 2019). As a result, the estimated *actual* load shift from the whole-building data (10kWh) has a high (15kWh) mean absolute error (MAE) and root-mean square error (RMSE, 16kWh). This precludes meaningful comparisons of the actual and predicted load shift[20] and the net energy impact (i.e., net kWh relative to baseline power) for individual events.

During the summertime, whole-building power levels are appreciably lower.[21] Although the hourly baseline power regressions have low goodness of fit values, the events yielded reasonably consistent results among eight precool events and six shed events that followed a precool. Figure 28 shows incremental energy consumption during the first hour of precool events, while Figure 29 shows the incremental energy consumption for the *entire* event as a function of total CDH during events. The incremental first hour kWh does not vary greatly among events; this is expected since the *incremental* AC load should be similar among events.[22] Moreover, the total precool energy usually increases approximately linearly with total CDH, as longer events and higher T_{out} increase cooling loads.[23]

This suggests that, with enough baseline performance data and data from functional tests (i.e., events), meaningful assessments of the load management potentials of commercial building AC system can be developed.

[20] For a precool followed by a shed, the load shift equals the predicted (or actual) load increase during precool plus the load decrease during the shed, divided by two. For example, a precool that increased consumption by 15kWh followed by a shed that decreased load by 21kWh has net impact of (15 + 21)/2 = 18kWh.

[21] This is due to the fact that the AC loads in our portfolio from a school, which is out of session during summer months

[22] To a first order, the incremental AC loads are due to the (up to 3ºF) greater temperature difference between T_{out} and T_{in} (increases ventilation air, infiltration, and shell loads) and heat flowing out of the effective thermal mass of the school as T_{in} decreases by 3ºF.In addition, the AC coefficient of performance (COP) decreases as T_{in}-T_{out} increases, increasing AC consumption during events as well.

[23] This will hold until baseline AC runtime approaches 100%, in which case the AC cannot run much longer and, eventually, $T_{set,zone}$ cannot be maintained. The school did not operate in that regime during the summertime.

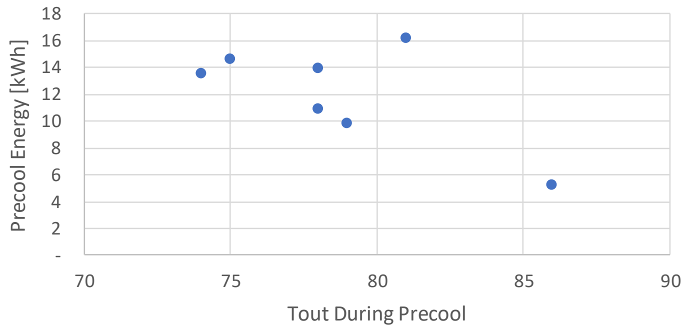

Figure 28: Incremental energy impact of summer precool events during the first hour.

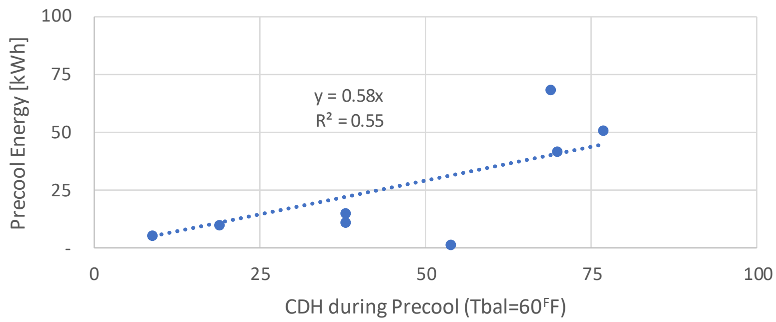

Figure 29: Total incremental energy impact of summer precool events as a function of total event CDH.

The load sheds exhibited similar trends (see Figure 30 and Figure 31).

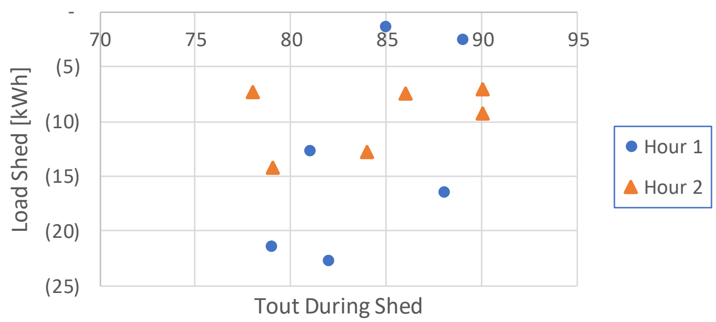

Figure 30: Incremental energy (negative = decrease) during summer shed events for the first and second event hours.

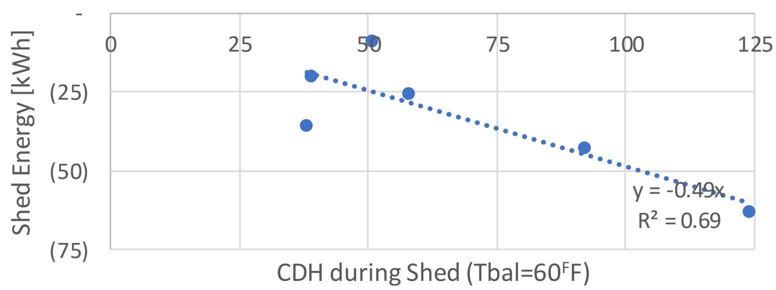

Figure 31: Total incremental energy impact (negative = less decrease) of summer shed events as a function of total event CDH.

Unfortunately, the whole-building power data are more challenging to interpret after school resumed in late August. Crucially, the T_{in} data logged indicate that many of the RTUs did respond as expected to events called around and after the start of the school year (for example, see Figure 32, in contrast to Figure 33). In addition, this is likely exacerbated by the appreciable uncertainty in the baseline power estimated described earlier combined with higher facility-level power draw (school in session) and more moderate T_{out} values. Most shed events appear to be directionally correct (Figure 34), but the analysis finds that a majority of the precooling events have a negative power impact, akin to a load shed.

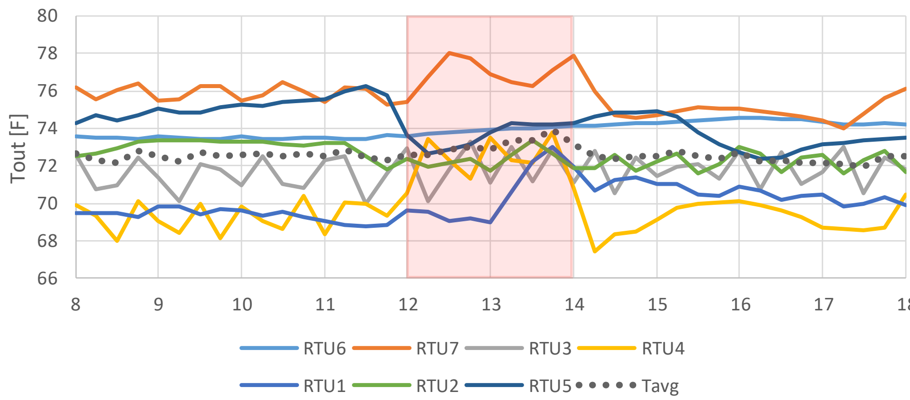

Figure 32: T_{in} data logged in the zones served by RTUs during a 3°F shed event from 12-14h; T_{out} ~78-80°F during event.

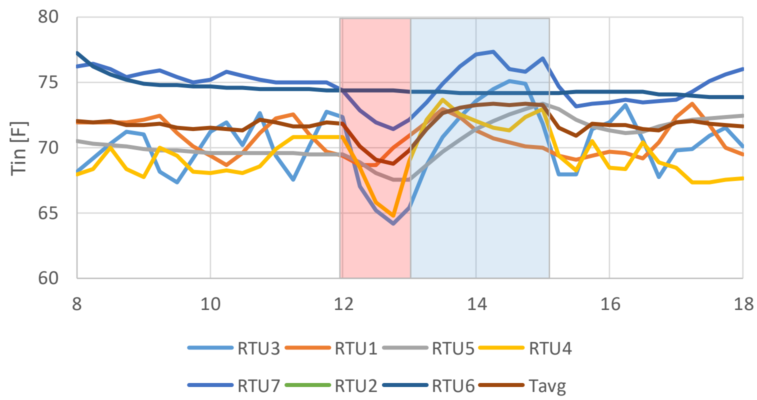

Figure 33: Zone T_{in} profiles for summer 2019 event with 3°F precool from 12-13h and 3°F shed from 13-15h.

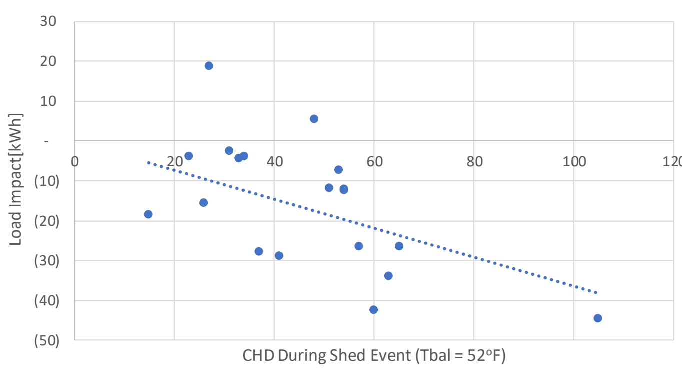

Figure 34: Fall shed event net load impacts as a function of event CHD.

4.8.2 Fall RTU Power Data

We did have packaged rooftop unit (RTU) power data coincident with some events that allowed us to analyze the net energy impact of the controlled units. Figure 35 shows the performance of four RTUs during an event called in the summer of 2018. Three of the four RTUs consume appreciably more energy during the 3°F decrease in T_{set} during the event. However, the total energy consumption includes both ventilation fan and AC-related (compressor, condenser fan) energy, so we established a baseline power draw level for the RTUs when the AC was not running[24] and also evaluated the changes in AC-related RTU energy consumption as well. As a result, the 11.5kWh increase in HVAC energy consumption during the event represents a 2.5-fold increase in AC-related energy consumption during the event.

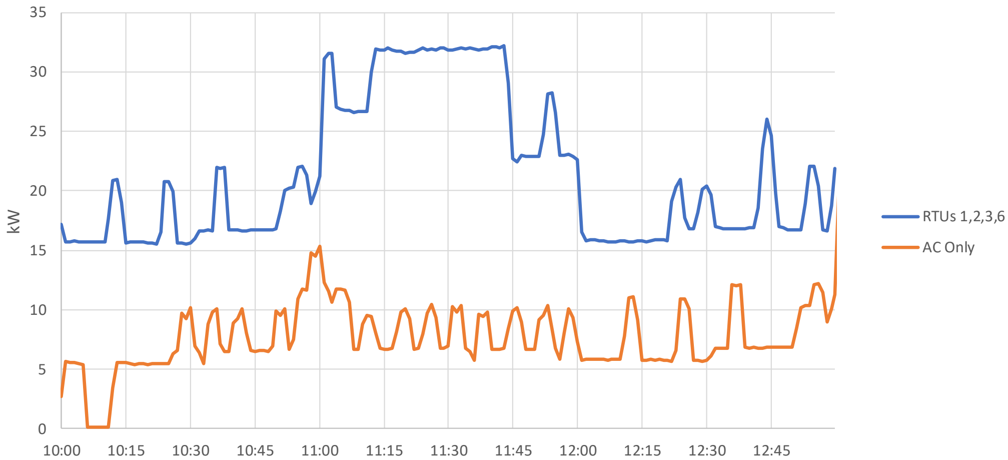

Figure 35: Power for RTUs 1,2,36 on a day with a 3ºF T_{set} decrease from 11-12h, T_{out} = 80ºF.

Similarly, for the 3ºF shed shown in Figure 36, the ~50% decrease in total RTU consumption represents a ~87% decrease in AC-only power draw during the hour-long event.

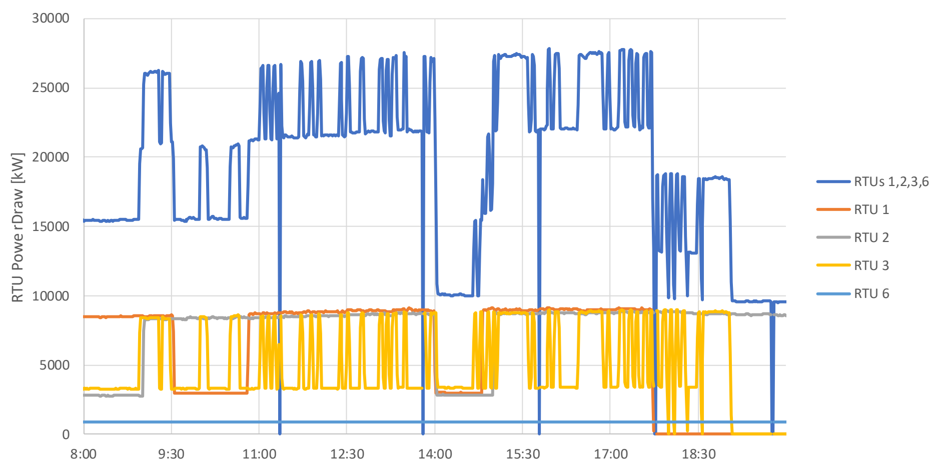

Figure 36: Power for RTUs 1,2,3,6 on a day with a 3ºF increase (shed) from 14-15h, T_{out} = 83ºF.

We obtained data for RTUs 4,5 and 7 coincident with ten AC load-shift events in September and October 2019. RTU 5 did not exhibit any variation in power draw during school hours, even during events, indicating that it was configured to only provide ventilation air. In addition, RTU 4 did not provide space cooling after its AC was disabled by the facility manager in late September due to control problems. Consequently, we evaluated the performance of RTU 7.

[24] Reasoning that all RTUs did not run continuously (they wouldn't be able to keep up at higher T_{out} if that were the case), we found the minimum power draw for a 1-minute period between 8AM and 5PM for each RTU and summed them to find the ventilation-only power draw.

Table 12 summarizes the characteristics and conditions of the AC load-management events analyzed. All consisted of a one- to three-hour 3ºF precool followed by a three-hour 3ºF shed.[25] The similar event designs make it possible to evaluate the performance across events.

Table 12: Characteristics of AC load management events analyzed.

			T_{out} Before and During Event			
Date	*Time*	*Type*	*Before*	*Hour 1*	*Hour 2*	*Hour 3*
9/20/2019	13-14:00	Precool 3F	74	77		
9/20/2019	14-17:00	Shed 3F	77	78	79	78
9/23/2019	12-14:00	Precool 3F	81	82	84	
9/23/2019	14-17:00	Shed 3F	84	87	88	86
9/24/2019	11-14:00	Precool 3F	69	72	72	73
9/24/2019	14-17:00	Shed 3F	73	73	74	72
9/25/2019	12-14:00	Precool 3F	65	67	69	
9/25/2019	14-17:00	Shed 3F	69	70	72	71
9/27/2019	12-14:00	Precool 3F	64	67	70	
9/27/2019	14-17:00	Shed 3F	70	72	72	72
9/30/2019	13-14:00	Precool 3F	59	61		
9/30/2019	14-17:00	Shed 3F	61	63	64	63
10/1/2019	12-14:00	Precool 3F	63	64	67	
10/1/2019	**15**-17:00	Shed 3F	69	70	70	71
10/2/2019	13-14:00	Precool 3F	72	73		
10/2/2019	14-17:00	Shed 3F	73	70	64	63
10/4/2019	11-14:00	Precool 3F	55	57	56	58
10/4/2019	14-17:00	Shed 3F	58	58	57	56
10/7/2019	12-14:00	Precool 3F	72	73	73	
10/7/2019	14-17:00	Shed 3F	73	75	74	72

In all cases, we regressed RTU 7 power data, $P_{RTU\text{-}7}$ as a function of CDH during school operating hours outside of events to develop an equation for calculating baseline power draw, $P_{RTU\text{-}T,b}$ (see Figure 37). We then applied a fixed same-day adjustment, $P_{RTU\text{-}7,adj}$, for each day by subtracting the actual average power in the hour prior to each event from the baseline ($P_{RTU\text{-}7,\,base}$) and adding that difference to all subsequent hours to create an adjusted baseline, $P_{RTU\text{-}7,base,adj}$ for all hours during each event and the portion of post-event hour when RTU 7 continued to run.[26]

$$P_{RTU\text{-}7,base} = 2.6 + 0.137 * (T_{out}\,(^{o}F) - 52^{o}F)$$

$$P_{RTU\text{-}7,adj} = P_{RTU\text{-}7,base}(t_e - 1) - P_{RTU\text{-}7}(t_e - 1)$$

[25] The event with a two-hour precool from 12-14h followed by a load shed from 15-17h is effectively a three-hour load shed, as the compressor did not turn on during the "gap" hour when the zone temperature set point, $T_{zone,set}$, went from 69ºF to 72ºF. A longer gap between the precool and shed would generate similar.

[26] RTU 7 power time-series data did not vary as a function of time of day on days without space cooling and served an administrative space with relatively constant occupancy, so we did not apply a time-of-day adjustment. A larger sample of afternoon hours without events would be necessary to meaningfully evaluate if a time-of-day adjustment is warranted.

The time when RTU7 stopped running after 17h varied from event to event. Although we assumed the model for baseline power draw applied to the 17-18h window, it is quite possible that it is less accurate for that time window.

$$P_{RTU\text{-}7,base,adj}(t) = 2.6 + 0.137 * CDH(t) + (2.6 + 0.137 * CDH(t_e\text{-}1) - P_{RTU\text{-}7}(t_e\text{-}1)).$$

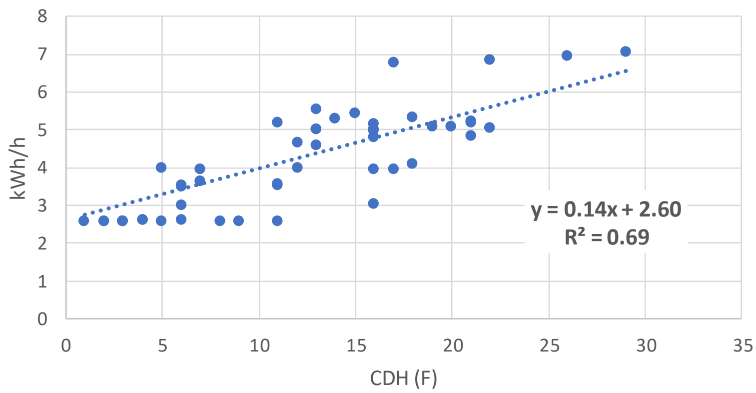

Figure 37: Regression for RTU 7 power draw between 8AM and 5PM as a function of cooling-degree hours (CDH), T_{bal} = 52°F. A regression based on 9AM-3PM yielded a similar slope and had a slightly higher (0.74) goodness of fit.

Based on that, we obtain the results shown in Figure 38 through Figure 43. As expected, the first-hour load shed shows the highest correlation with T_{out}/CHD (Figure 38).

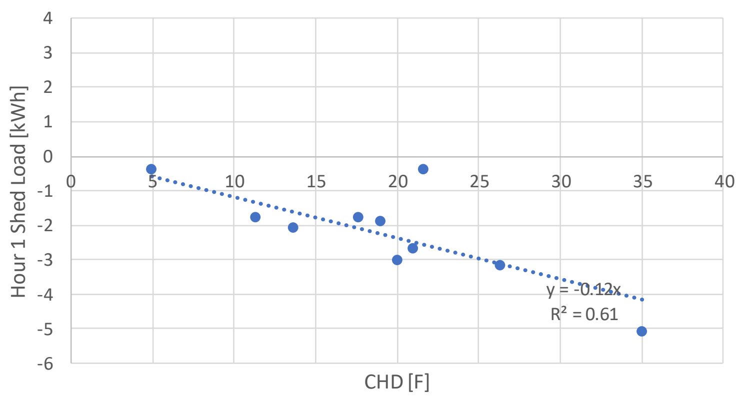

Figure 38: RTU7 first-hour load shed as a function of CDH (T_{bal} = 52°F).

Predicting the total load shed over multiple hours is more complex. All buildings or zones that increase T_{set} will experience a period of time when cooling is not required until the zone temperature, T_{zone}, rises to the new T_{set}. The *time* that it takes for a given zone to reach the new T_{set} can vary appreciably among buildings and depends on several factors, including building thermal mass, outdoor air ventilation rates, and internal and solar heat gains. After T_{zone} reaches the new T_{set}, the AC system will have a somewhat lower power level due to lower cooling loads. Nonetheless, the *total* load shed over a multi-hour shed event shows a similar goodness of fit (0.65) once an outlier data point is removed (Figure 39)

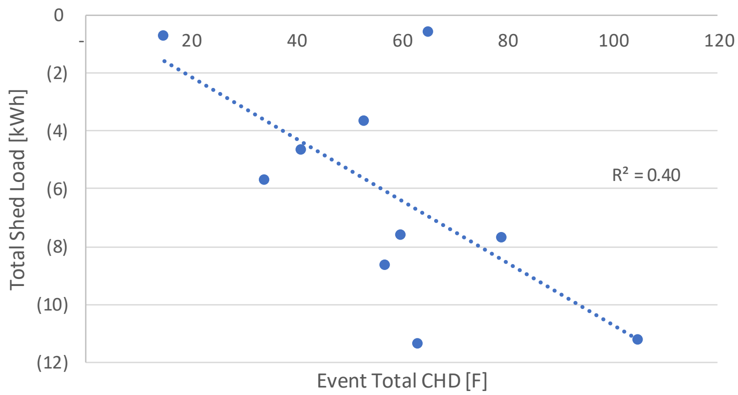

Figure 39: RTU7 total energy impact of shed events versus total event CDH (T_{bal}=52°F). Omitting the outlier increases R^2 to 0.65.

The precool energy is less predictable (Figure 40). Examining events with near-zero or negative precool energy revealed that, in those cases, the AC system was already running at full power for the cooling stage that was engaged during an event (Figure 41), the event caused the AC to run (Figure 42) or that variability in precisely when the AC system ran during the hour prior to the event.

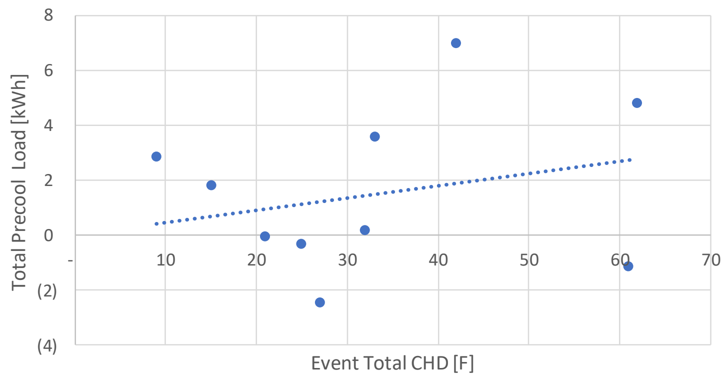

Figure 40: RTU7 total precool load as a function of total CDH during each event (T_{bal} = 52°F).

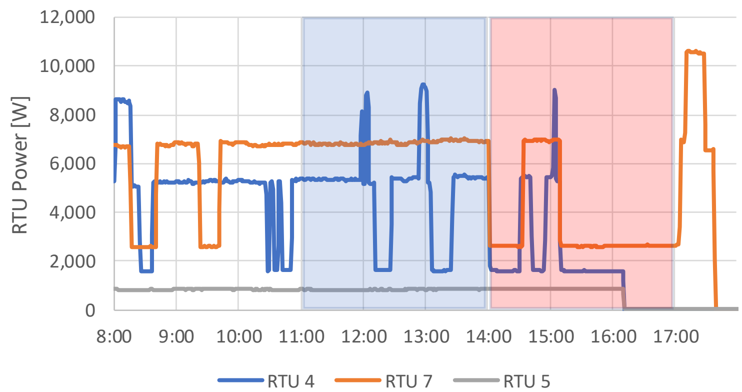

Figure 41: RTU power during a 3°F precool event from 11-14h, followed by a 3°F shed from 14-17h (T_{out}=72-74°F).

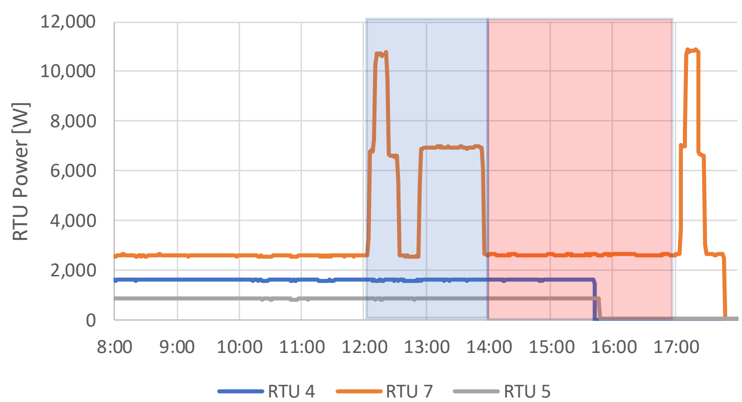

Figure 42: Event on fall day after several cooler days where AC did not run prior to event (T_{out} = 72-75°F).

Unsurprisingly, the calculated net energy impact does not exhibit a clear trend relative to CDH, i.e., it depends on the duration and conditions of each event (Figure 43).

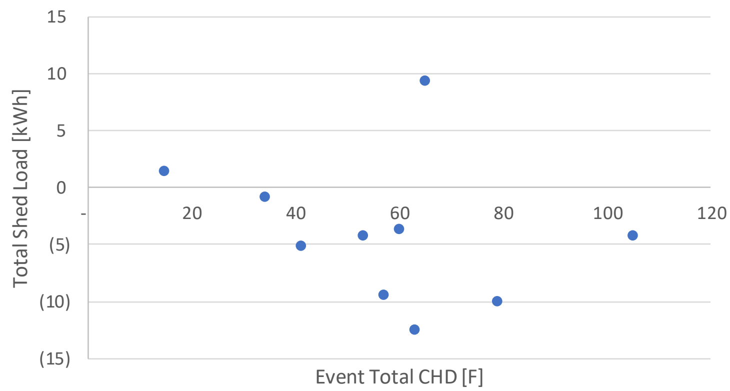

Figure 43: RTU7 net event load impact as a function of CDH (T_{bal} = 52°F).

4.9 Conveyor Drag Performance Assessment

Initial characterization of the conveyor drag indicated the potential to shift approximately 90kW of customer load by up to 5 hours (Figure 44). This resource was not automated: we were able to provide the customer with feedback to indicate "preferred" or "un-preferred" windows of operation, but required the customer to actuate the loads. Further, this circuit was not directly submetered, so we did not have access to direct feedback as to whether or not the load was in use.

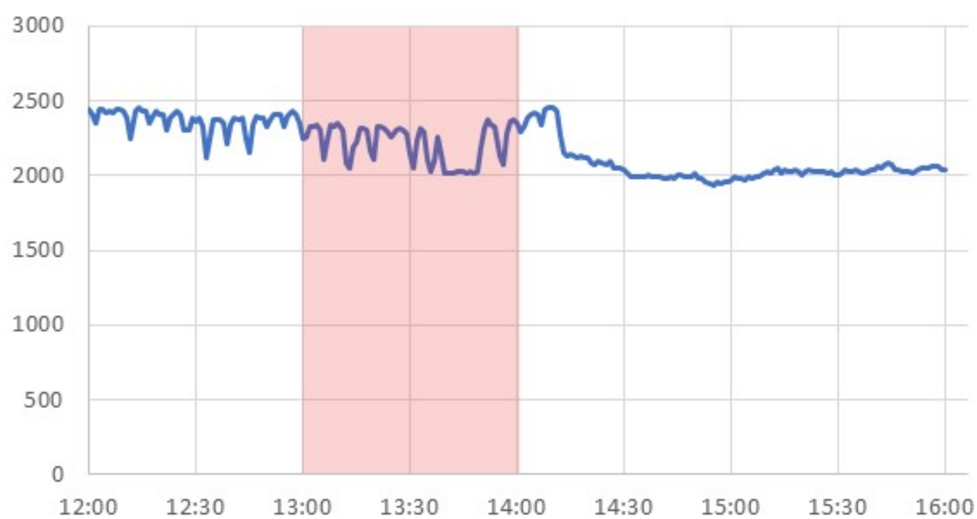

Figure 44: Approximately 130kW of observed load reduction due to pausing conveyor drag during an example event

We attempted to call multiple events throughout July-September 2019, but the customer was non-responsive, so we have no evidence that this resource performed as planned. This experience helps to highlight the importance of automating response for demand-side assets (rather than requiring direct customer intervention); further, the lack of feedback as to the process status highlights the value of submetering key loads to accurately characterize performance of responsive loads.

5 Summary of Results

We tested the ability of a virtual power plant, incorporating solar PV, grid-scale energy storage, and flexible C&I loads, to mitigate short-term solar intermittency while satisfying multiple strategic objective functions in combination, including peak shaving, energy cost optimization, peak-power dispatch, and power firming over a 15-month pilot demonstration. Key results and lessons learned from this deployment are summarized below:

Mitigation of PV Intermittency: The SunDial System was used to minimize short-term changes in solar production plant output while meeting multiple strategic (multi-hour) control objectives. Data collected during the course of the field demonstration indicated that we were able to eliminate 100% of ramp events greater than 20% of system nameplate capacity per minute, 97% of ramp events greater than 10%, and 91% of ramp events greater than 5%, relative to a start-of-project performance target of 10% of nameplate capacity per minute.

Net Load Predictions: A regression-based approach was used to generate time series load and solar predictions with root-mean square error (RMSE) ranging from 8 to 12% of peak load, depending on the time-to-prediction (1 to 24 hours ahead). Most load predictions have focused on accurately predicting building peak loads under peak conditions. However, the use cases envisioned by SunDial require accurate baseline models for most hours of the year. We found that time series load estimates obtained using site-agnostic predictors (in particular, time- and calendar-based schedules, site-agnostic operational classifiers, historical + real-time load data and prediction errors) offered performance that was comparable to results obtained from more detailed site-specific process models, but required significantly less knowledge and effort to implement and maintain. In the case of solar predictions, weather-based production forecasts from satellite data were enhanced by learning from the site's production over time, using classifiers such as time of day, total daily forecast irradiance, and recent production data to correct these results. Implementation of a revised solar model reduced RMSE from 13-20% to 11-14%, depending on time to prediction. Identifying conditions associated with high prediction uncertainty – which tended to be weekends in the case of loads, and variable production days in the case of solar – can be used to mitigate uncertainty by, for example, building in additional reserve margin to dispatch schedules.

Strategic Dispatch and Integration of Solar PV with Flexible Loads:: The SunDial System was used to test several different use cases for the purpose of optimally shaping net load on a distribution system, including peak shaving, backfeed mitigation, power firming, and energy cost optimization using energy storage and flexible C&I loads. In total, the Global Scheduler executed approximately 177 ESS charge/discharge cycles and initiated 35 successful EV charging events and approximately 32 AC shifting events. A newly developed communications protocol was developed and tested to facilitate coordination between the demand management service provider (IPKeys) and the Global Scheduler. Initial results highlighted challenges related to Solar + Storage integration and prediction uncertainty, which caused the system to significantly underperform relative to a "perfect information" case, and in some cases relative to a heuristic (rules-based) strategy. Subsequent revisions to the system's control software, including (1) improvements to net load prediction methodologies and (2) incorporation of objectives to hedge against system uncertainty by leaving additional storage reserve, helped to greatly mitigate these issues.

One way to characterize overall system performance is by comparing the amount of energy storage capacity required by an alternative DER control framework to achieve performance equivalent to that presented here. During the course of this project, we tested two different complementary methods that can help reduce the need for energy storage in high-solar penetration distribution systems: (1) integrating granular, bottom-up time-series predictions and uncertainty mitigation into the dispatch decision-making process; and (2) productively shifting flexible loads.

Incorporation of solar- and load- prediction methodologies into the dispatch decision-making process offered marked improvements relative to a heuristic (rules-based) approach, which was used to define a baseline performance level. For example, several use cases under test evaluated the system's ability to manage peak load over different time horizons – e.g., afternoon peaks, overall daily peaks, and monthly peaks – in combination with one or more additional objectives. The average peak load reduction for each of these sets of objectives for three different scenarios ("Heuristic", "Global Scheduler", and "Perfect Information") is shown in Table 13. The peak reduction achieved by the Global Scheduler case ranges from 83 to 90% of that achieved by the Perfect Information case, and is 1.25 to 2.8 times higher than the Heuristic case. Table 13 also estimates the amount of additional energy required for the "Heuristic" case to match the target performance of the "Global Scheduler" case.[27] Using this methodology, it is calculated that the storage capacity required to replicate the performance of the Global Scheduler using a heuristic approach ranges from 145 to 400 kWh. These peak load reduction use cases were also used tested in combination with each other (e.g., to simultaneously manage afternoon peaks that stress the power system, and overall peaks that drive costs for many C&I customers), and in combination with other objectives. For example, the monthly peak load reduction use case was shown in combination with a backfeed minimization case to eliminate approximately 70% of total reverse power flow.

Table 13: Estimated reduction in peak demand relative to a no-ESS case, and additional ESS capacity required for Heuristic case required to match demand reduction of the Global Scheduler case.

Peak Reduction Period	Avg Reduction in Peak Demand (kW)			ESS Energy-equivalent (kWh)
	Heuristic	Global Scheduler	Perfect Information	
Monthly	80	113	129	400
Daily	30	84	101	145
Afternoon ('Virtual Peaker')	207	255	281	192

An alternative way to use the system's predictive capability is to provide firm, predictable demand for a defined time window. Testing this use case showed that a 500kW/1000MWh storage system was able to provide committed capacity for 99% of hours over a 4-hour window for a predicted commitment generated 4-hours prior to the start of the commitment window, or for a 6-hour window for a 1-hour ahead commitment. Achieving this same level of performance was estimated to require an additional 500kWh of storage in the former case, and 1,000 kWh of additional storage in the latter case.

LCOE Characterization: During prior phases of this project, we developed and applied toolsets to calculated an "LCOE-equivalent" metric for evaluating the levelized cost of ownership for the system. The metric is defined as "Levelized Cost to Provide Electricity divided by the Energy Consumed by SunDial Loads". Our prior analysis compared costs between a "solar-only" scenario and the SunDial system (assuming perfect information). We estimated that a deployment in the northeast was approximately $0.08 per kWh for both the solar-only and SunDial case for estimated 2020 prices; and was $0.09 per kWh for the solar only case and $0.10 per kWh for the SunDial case using estimated 2017 prices. Using a similar methodology, we recalculated LCOE using the above metric based on actual generation, load, and peak load reduction data gathered during the course of the pilot, and applying the same 2020 cost assumptions as used previously. The resulting calculations broadly replicate the prior results: the solar-only case was estimated to cost $0.10 per kWh, and the SunDial case was estimated to cost $0.104 per kWh (a 4% increase relative to the solar-only case). The heuristic case was estimated to have a levelized cost of $0.107 per kWh (7% higher than the solar-only case). The primary driver for the differences from the Phase 1 characterization are due to a reduction in annual solar production compared to our initial (Phase 1) estimates.

Performance of Flexible Loads: We found that the portfolio of loads under investigation were able to repeatably realize approximately 51kWh of energy from shiftable loads that were capable of absorbing solar generation (i.e., increasing load) on a typical day during shoulder seasons; and approximately 40kWh of load reduction energy from shiftable loads on a typical summer day that were capable of reducing peak load. These flexible load resources did impose significant constraints as to the time of day and duration of available load shifts. An additional large and highly flexible load (conveyor drag) would have increased

[27] These estimates were generated by iteratively increasing the capacity of storage in the Heuristic case until the average peak load reduction approximated that of the Global Scheduler case.

total shiftable load to 96kWh shoulder season and 85kWh during the summer season; however, we were not able to effectively utilize this resource due to lack of customer responsiveness.

6 Lessons Learned and Future Work

Effective use of energy storage for strategic dispatch is highly dependent on prediction accuracy and uncertainty mitigation: The biggest improvements in Global Scheduler performance over the course of the project were from refining prediction methodologies and incorporating methods to account for prediction uncertainty by building in additional storage reserve capacity. For example, for the peak load management use cases, incorporating a predictive, adaptive controls approach had benefits equivalent to 20% to 40% of the total storage system capacity relative to a rules-based approach. We found that leaving approximately 30% of storage capacity in reserve offered an effective hedge given the uncertainty, but this approach could be refined by adjusting the reserve margin based on the current state of the system and the specific control objectives.

Recommendation: Future work should focus on developing toolsets that are capable of more robustly and dynamically accounting for PV generation and load forecast uncertainties based on the state of the system.

Commercial and industrial demand side resources can effectively supplement energy storage for managing high penetration of solar PV: For the cases studied, flexible loads comprising approximately 10% of the energy storage capacity were used to support strategic (multi-hour) load shaping objectives. It is notable that we were able to call these resources on a repeatable basis – i.e., nearly every weekday during August-October 2019 – using automated toolsets without notably impacting customer operations. However, constraints around the availability and predictability of these load resources made it challenging to treat them as firm capacity. These issues can be partially mitigated by incorporating load submetering, increasing the size of (i.e., diversifying) the load management portfolio, and improving models of dynamic HVAC loads.

Energy storage can be used to simultaneously manage short-term PV intermittency and support strategic load shifting objectives, but duty cycles needs to account for real-world ESS performance: We successfully used the ESS to simultaneously mitigate short-term intermittency changes and achieve daily performance objectives. However, a sizeable portion of ESS charge / discharge throughput that was used to smooth changes in plant production occurred at low load (<10% of nameplate capacity), with measured real-world roundtrip efficiency on the order of 40 to 50%. In comparison, we measured real-world roundtrip efficiencies between 82-85% at higher charge / discharge levels, which are roughly in line with manufacturer specifications. In similar vein, although we were only to collect a limited amount of data in cold weather conditions, we found that roundtrip efficiency significantly increased/decreased as outdoor temperature decreased. These results clearly indicate that optimal ESS dispatch strategies need to incorporate accurate models for operational efficiency as a function of power level and outdoor temperature to .

Accurate, reliable, and timely telemetry is critical to effectively leveraging distributed resources: A recurrent theme throughout this project was the value of reliable access to real-time data for loads and generation. Ideally (as discussed below), sub-metered load data could provide still further value. This type of telemetry is critical for developing accurate predictions, particularly for rapidly detecting major deviations from expected generation or loads and modifying operating decisions, and for detecting and addressing problems as they arise (ongoing commissioning).

There is a need for streamlined and scalable processes and technologies for integrating solar PV with energy storage: We encountered numerous delays and challenges related to defining and testing plant controls, adverse interactions between vendor devices, and detecting and responding to plant outages. Issues were attributable to multiple factors, including difficulty integrating and coordinating multiple vendor devices; coordination between multiple stakeholders in the plant procurement and deployment process; lack of clarity in specification for plant-level controls; and extensive customization in the development of site-level controls.

Recommendation: More mature, configurable control platforms designed specifically for integration of solar and storage sites would offer a lot of value relative to project-specific controls development. Such a product should incorporate straightforward methods for defining and configuring solar + storage set points and low-level controls; event detection notifications; and common solar + storage use cases. In a similar vein, development of model specifications for utility solar + storage procurements and acceptance testing would greatly facilitate deployment and operation of integrated PV + Storage sites.

Incorporating auto-calibrating, self-learning techniques into solar and load predictions is a critical enabler for scalability of integrated approaches: The accuracy of solar generation and load predictions were significantly improved by feeding back near-realtime and historical data into the prediction framework. In addition, we found that time-series load estimates obtained using site-agnostic predictors offered performance comparable to results obtained from more detailed site-specific process models, but required far less knowledge and effort to implement and maintain. This latter point is critical, as it can enable rapid deployment to fleets of buildings. Flexible load predictions could also theoretically leverage this type of self-learning, auto-calibrating capability.

Recommendation: Develop machine learning algorithms (MLAs) that continuously refine baseline solar and load predictions and load management potentials using meter and environmental data. This differs from most statistical methods currently used for DR programs, which typically consider the last month of data. This process must be automatic and generate predictions without knowledge of the inner workings of the building, i.e., it needs to determine dependence on external variables such as T_{out}, day of the week, season and other variables that affect energy consumption.

Cost-effectively integrating flexible loads to support high-penetration solar will requires a highly streamlined process: Even allowing for the potential additional value streams that SunDial could access, the pilot took weeks of analysis to develop accurate prediction values. In addition, curtailment service providers (CSPs) have a limited amount of customer attention for DR/Load Management Implementation, so excessive demands on customer time, e.g., more than a few hours total, will likely lead to nonresponsive customers or customer attrition. A commercially viable implementation would require onboarding sites in less than half a day, and site implementation should take less than one day. Specific challenges and recommendations for achieving this goal are outlined below:

- ***Automated Event Implementation is Essential for Use Cases with Numerous Events:*** The FLAME called events on a majority of weekdays during the summer and fall. The events were usually detectable at the two facilities where controls actuation was automated, but not at the one facility that forewent automating controllable loads. It is not realistic for most facilities to manually implement events on many/most workdays of the year.
- ***Predicting AC Load Management Potentials for a Wide Range of Conditions and Durations is Challenging*:** AC loads and the magnitude of AC power that can be shifted over different time horizons depend on factors that are difficult or time-intensive to model, including building thermal mass, AC capacity, internal heat gain profiles, and AC efficiency as a function of T_{out}. Developing accurate models is also currently not a core competency of most CSP, who focus on predicting load sheds for occasional events at peak conditions. Consequently, it is likely beyond the capability of most CSPs.

  Recommendation: Develop standard – and, ideally, automated – functional test and modeling and analysis procedures to acquire and analyze building data to generate reasonably accurate models to predict AC load flexibility with minimal time investment. AC submetering (see below) also can provide valuable inputs. The facility-specific model should be automatically updated as additional events occur; this will require integration of the model platform with both the BAS (T_{zone}, cooling status data) and electric meter and any cooling-related submeters. These could be integrated with model-predictive control (MPC) approaches, e.g., as commercialized (Q-Coefficient[28]).

- ***OpenADR can be Expanded to Communicate Information Needed for a Wider Range of Use Cases:*** IPKeys developed and implemented a new communications protocol as an expansion of OpenADR, that specifies communications between the FLAME and Global Scheduler to support the delivery load-shaping services as a means to facilitate high-penetration of solar.
- ***Submeter Data for Major Loads:*** Submetering major loads greatly facilitates evaluation of load-management potentials and baseline load predictions, particularly for loads that are of a similar magnitude as the uncertainty of baseline load predictions.

  Recommendation: Major C&I loads could be made "load-management ready" by installing them with on-board or submetered power measurement capabilities.
- ***Controls Integration is Often Challenging and Inconsistently Reliable:*** We encountered multiple problems with the BAS programming, including incorrect measurement point references and disabling of data acquisition functionality by the controls contractor despite using the controls contractor requested by the school.
- ***High-Fidelity AMI Data are Essential to Model and Quantify Load Impacts*:** High-resolution (<15 minutes) can help to quickly quantify the load impact of process loads. IPKeys routinely installs infrastructure to obtain and communicate in near-real time 1-minute power data, at an approximate cost of two hours and $2,000 per meter. In IPKeys' experience, if the utility needs to install a KYZ board, that process can take as long as six months.

  Recommendation: Install metering infrastructure than makes high resolution data readily available on site to customers.

7 Publications and Presentations

Table 14 summarizes SunDial-related publications and presentations to date.

Table 14: SunDial-related Publications and Presentations to Date

Event	Title	When & Where	Presenter
ACEEE Summer Study on Energy Efficiency in Buildings *(submitted)*	Evaluation of Time-Series Load-Prediction Methodologies to for Optimal Energy Storage Dispatch in Solar-Dominant Distribution Grids	August, 2020 Pacific Grove, CA	M. Kromer, K. Roth
2020 IEEE Power & Energy Society General Meeting (PESGM) *(submitted)*	Optimizing DER Dispatch in a Renewables Dominant Distribution Network Using a Virtual Power Plant	August, 2020 Montreal, Canada	M. Kromer
DOE-SETO Seminar	SunDial: Enabling High Penetration Solar with Integrated Energy Storage and Demand Management	2/26/2019	M. Kromer, K. Roth
IEEE WCPEC-7	"The SunDial Framework: Enabling High Penetration Solar through the Integration of Energy Storage, Demand Management, and Forecasting"	6/10/2018 Waikoloa, HI	Tsz Yip
Solar Power International 2018	"Sundial: Enabling High Penetration Solar through Integrated Energy Storage and	9/25/2018	M. Kromer

[28] See: http://qcoefficient.com/ .

	Demand Management. Session: Smart Energy Management Systems.”		
InterSolar 2018	“Enabling High Penetration Solar with Integrated Energy Storage, Demand Management, and Forecasting”	7/14/2018	Matt Kromer
SETO Portfolio Review	SunDial	2/12/2018	Kurt Roth
SHINES LCOE Review	Fraunhofer CSE – SHINES LCOE Methodology Overview	2/5/2018	Matt Kromer
Eighth Conference on Innovative Smart Grid Technologies	“Integrated System to Enable High-Penetration Feeder-Level PV: Preliminary Design and Simulation Results”	4/25/2017	M. Zeifman, M. Kromer, K. Roth
Consortium for Energy Efficiency Winter Meeting	“SunDial: An Integrated PV + Energy Storage + Load Management System Enabling High-Penetration Feeder-Level PV”	1/19/2017, San Francisco, CA	K. Roth
ACEEE Intelligent Energy conference	“SunDial: An Integrated SHINES System Enabling High Penetration Feeder-Level PV”	12/6/2016, Austin, TX	K. Roth
Intersolar USA	“The SunDial Framework: Enabling High Penetration of PV through Integrated, Feeder-Scale Control of DERs”	7/13/2016, San Francisco, CA	M. Kromer
SHINES Kickoff Meeting	“SunDial –An Integrated SHINES System to Enable High-penetration Feeder-level PV”	5/18/2016, Washington, DC	M. Kromer, K. Roth

8 References

[1] C.A. Correa, G. Marulanda, A. Garces, "Optimal microgrid management in the colombian energy market with demand response and energy storage," Power and Energy Society General Meeting (2016).
[2] A.I. Selvakumar, K. Thanushkodi, "A new particle swarm optimization solution to nonconvex economic dispatch problems," IEEE Trans. Power Systems, vol. 22, pp. 42-51 (2007).
[3] Z. Yang, R. Wu, J. Yang, K. Long, P. You, "Economical operation of microgrid with various devices via distributed optimization," IEEE Trans. Smart Grid, vol. 7, pp. 857-867 (2016).
[4] Q.D. La, Y.W.E. Chan, B-H. Soong, "Power management of intelligent buildings facilitated by smart grid: a market approach," IEEE Trans. Smart Grid, vol. 7, pp. 1389-1400 (2016).
[5] H. Abdi, E. Dehnavi, F. Mohammadi, "Dynamic economic dispatch problem integrated with demand response (DEDDR) considering non-linear responsive load models," IEEE Trans. Smart Grid, vol. 7, pp. 2586-2595 (2016).
[6] E. Vrettos, et al, "Experimental demonstration of frequency regulation by commercial buildings – part I: modeling and hierarchical control design," IEEE Trans. Smart Grid, in press.
[7] E. Vrettos, et al, "Experimental demonstration of frequency regulation by commercial buildings – part II: results and performance evaluation," IEEE Trans. Smart Grid, in press.
[8] G. He, Q. Chen, C. Kang, P. Pinson, Q. Xia, "Optimal bidding strategy of battery storage in power markets considering performance-based regulation and battery cycle life," IEEE Trans. Smart Grid, vol. 7, pp. 2359-2367 (2016).
[9] X. Zhang, G. Hug, J.Z. Kolter, I. Harjunkoski, "Model predictive control of industrial loads and energy storage for demand response," Power and Energy Society General Meeting (2016).
[10] Boyd, S. and Vandenberghe, L., "Convex Optimization," Cambridge University Press (2004)
[11] Bertsekas, D.P., "Dynamic Programming and Optimal Control," Athena Scientific (2005)
[12] A. Corana, et al, "Minimizing Multimodal Functions of Continuous Variables with the Simulated Annealing Algorithm," ACM Transactions on Mathematical Software 13, 262-280 (1987).
[13] W.L.Goffe,, G.D. Ferrier, and J. Rogers, "Global optimization of statistical functions with simulated annealing," Journal of econometrics 60, 65-99 (1994).
[14] Mathieu et al. "Quantifying Changes in Building Electric Load, With Application to Demand Response." IEEE Transactions on Smart Grid 2:507-518, 2011.
[15] Price P. "Methods for Analyzing Electric Load Shape and its Variability." Lawrence Berkeley National Laboratory Report LBNL-3713E, May 2010.
[16] Goldberg, M.L. and G.K. Agnew. 2013. "Measurement and Verification for Demand Response". Report Prepared by DNV KEMA for the National Forum on the National Action Plan on Demand Response: Measurement and Verification Working Group. February.
[17] Monforte, Frank, Fordham, C, Blanco, J, Hannam, D, Peri, S, Barsun, S, Kankiewicz, A, and Norris, B. "Improving Solar & Load Forecasts by Reducing Operational Uncertainty." California Energy Commission. Publication number: CEC-500-2019-023. 2019.
[18] Chen, D and Irwin, D. "Blackbox Solar Performance Modeling: Comparing Physical, Machine Learning, and Hybrid Approaches". University of Massachusetts Amherst. Greenmetrics '17, Urbana-Champaign, IL. June 5, 2017.
[19] Kiliccote, S., D. Olsen, M.D. Sohn, and M.A. Piette. 2016. "Characterization of demand response in the commercial, industrial, and residential sectors in the United States". *Energy and Environment*. Vol. 5 no. 3, May/June.
[20] R. Yin, et al, "Study on Auto-DR and Pre-cooling of Commercial Buildings with Thermal Mass in California," Energy and Buildings, v. 42, pp. 967-975 (2010).
[21] Hirsch, A., et al, "Pilot Testing of Commercial Refrigeration-Based Demand Response," NREL report (2015).
[22] G. Chatikar, et al, "Distributed Energy Systems Integration and Demand Optimization for Autonomous Operations and Electric Grid Transactions," LBNL Report (2015).
[23] D.S. Watson, et al, "Fast Automated Demand Response to Enable the Integration of Renewable Resources," California Energy Commission (2012).
[24] D.J. Olsen, et al, "Grid Integration of Aggregated Demand Response, Part I: Load Availability Profiles and Constraints for the Western Interconnection," LBNL Report (2013).
[25] M. Starke, Alkadi, N., Ma, O., "Assessment of Industrial Load for Demand Response across U.S. Regions of the Western Interconnect," ORNL Report (2013).
[26] D. Scott, et al, "The Impact of Control Technology on the Demand Response Potential of California Industrial Refrigerated Facilities," LBNL Report (2012).
[27] Coughlin, K., M.A. Piette, C. Goldman, and S. Kiliccote. 2008. "Estimating Demand Response Load Impacts: Evaluation of Baseline Load Models for Non-Residential Buildings in California". LBNL Report, LBNL-63728. January.
[28] Kromer, M, Zeifman, M, Roth, K, and Yip, T. "The SunDial Framework: Enabling High Penetration Solar through the Integration of Energy Storage, Demand Management, and Forecasting". 2018 IEEE 7th World Conference on Photovoltaic Energy Conversion (WCPEC-7). June 10-15, 2018. Waikoloa, HI.
[29] Zeifman, M, Kromer, M, Roth, K. "Integrated system to enable high-penetration feeder-level PV: Preliminary design and simulation results," IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT). April 2017.
[30] Martin, L, Williams, T. "Austin SHINES Project". SHINES Program Review 2017. Jan 2017. https://www.energy.gov/sites/prod/files/2017/05/f34/SHINES%20Program%20Review_Austin%20Energy_0.pdf

[31] Fitzgerald, G, Mandel, J, Morris, J, Touati, H. "The Economics of Battery Energy Storage: How multi-use, customer-sited batteries deliver the most services and value to customers and the grid." Rocky Mountain Institute, September 2015.