Energy Efficiency Driven by a Storage Model and Analytics on a Multi-System Semantic Integration

Domitille Couloumb*, Charbel El Kaed†, Ayush Garg*, Chris Healey*, Jonathan Healey*, Stuart Sheehan*
Schneider Electric: *IT Division, †Digital Services Platform
Email: {first.last} @schneider-electric.com
(The order of the authors is arranged alphabetically)

***Abstract*—One of the Internet of Things promises is interoperability where data is shared and understood between things and applications. Such interoperability aims to achieve a common goal with better efficiency, optimization, and a better user experience. However, these things produces data in different formats and semantics making interoperability a real challenge still to be tackled. Linked data is currently positioned as a promising technology capable of addressing the heterogeneity challenge. In this paper, we present an efficient energy storage system which relies on a multi-system semantic representation of several data sources. Our approach analyzes data collected from external services such as weather and billing systems and our internal systems such as the building management system, power monitoring system, and data center system. The aim of the analyzed data coupled with our model enables efficient energy storage usage through forecasting and model predictive control for various purposes such as reducing energy consumed by a utility provider. In this present work, we detail the application model and then demonstrate the capability of this approach through a case study in a large office building.**

Keywords*—IoT, Smart Building, Energy Efficiency, Ontology, Semantic, Energy Storage Model

I. Introduction

The Industrial Internet of Things (IIoT) is expected to transform the industry by massively interconnecting things in order to tackle several challenges such as energy efficiency, cost reduction, and a better user experience.

Those things ranging from sensors, actuators to gateways, and systems tend to continuously collect several types of data such as contextual, environmental, or industrial process related. Such data is then handled either on premise by domain specific systems or sent to a remote platform for several purposes such as monitoring, or control to drive energy efficiency or increase a human comfort in a building area. Things are designed for various purposes in the industry, such diversity is reflected in their computational, storage, and communication capabilities. In addition, due to their various purposes, these things communicate data in different formats, syntax, and semantics according either to home grown data models or industry standards. Such heterogeneity of data representation is one of the most challenging problems in the IoT domain, it prevents interoperability and inter-cooperation between things but also add additional burden on data analytics to understand the semantics of the data [1].

Semantic technology, is one of the most promising fields in the knowledge representation domain, expected to enable interoperability in the IoT. The World Wide Web Consortium (W3C) defines a set of standards , such as RDF, OWL and SPARQL [2], [3], [4], to represent semantics and query linked data, offering an ideal ecosystem and opportunity to tackle the heterogeneity challenge in the IoT.

In our previous work [5], we detailed a remote visualization semantic driven multi-system approach. In this paper, we extend our previous work by adding an additional Data Center Management System and we apply an analytics dimension to perform an energy efficiency strategy on our pilot site.

Effective implementation of energy storage in the power grid is a challenging problem both at the utility and customer level. Uncertainties in pricing, demand, and renewable energy sources can create a large variety of possible scenarios to consider. In this paper, we investigate a micro-grid energy storage and an architecture enabling energy peak-shaving by combining data forecasting and model predictive control to determine energy storage control policies.

The rest of the paper is organized as follows. Section II overviews the industrial context for our approach and then we detail our overall architecture in section III. Section IV details our model formulation for an energy storage system and we present our case study in section V. Section VI draws the related work. Then, we conclude in section VII.

II. Industrial Context: Energy Peak Shaving

This work has been applied to our North American headquarters site whose systems include a Building Management System, a Power Monitoring System, and a Data Center Management System.

The role of the BMS is to monitor and control the mechanical equipment in the building in order to provide a safe and comfortable environment for the occupants. The types of mechanical equipment at the site include chillers, fan coil units, heat pumps, chilled beams, and energy recover ventilators. Conditions such as indoor and outdoor temperature, humidity and air quality are monitored by the BMS, and set points for these conditions are based on schedules.

More than 200 power meters have been installed in the building which are managed and monitored by the PM system. The power meters provide various types of measures, such as active and apparent energy and power, demand and power quality. The meters are placed on feeder circuits throughout the electrical distribution system as well as on critical infrastructure, such as the chillers and other HVAC equipment, data center equipment, generator, lighting and lab equipment.

The extensive metering provides valuable insight into energy consumption patterns. The PM system also monitors digital protective circuit breakers which themselves include power monitoring and communication capabilities.

Our site, also hosts a Data Center solution along with a 300 kWh battery which is used as a backup solution for our data centers to prevent any power outage.

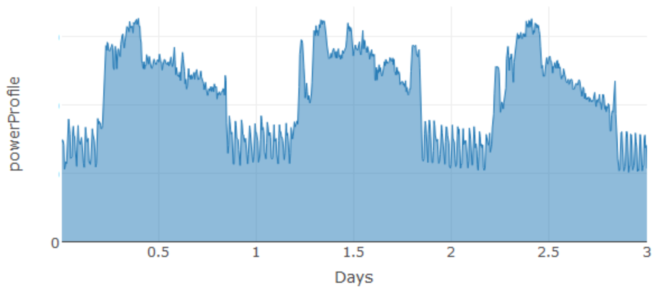

Fig. 1. Example Facility Load

In our continuous effort regarding energy management saving, we opted to focus on energy peaks occurring on our site, as shown in Figure 1. For industrial and commercial customers, electricity prices are based on two elements: actual usage (kWh), and a penalty based on the highest point of demand (or peak) within the billing period (weekly, monthly, or annually). Demand is often calculated using demand intervals, which are typically a short timeframe (e.g., 15 minutes) during which overall usage is aggregated and tracked as a total. The average calculated is the kW demand for this period. Charges based on demand can be controlled or reduced by the customer by modifying energy usage. For example, peak shaving is the ability to control usage from an energy supplier during intervals of high demand for purposes of limiting or reducing demand penalties for the billing period.

Determining efficient control policies for (battery) energy storage systems have been the subject of many papers, such as Zhu and Hug-Glanzmann [6], Baker et al. [7], [8], Salas and Powell [9] and Jiang et al. [10]. In these energy storage control approaches, constraints on the size, charge rate, and discharge rate of the energy storage employed are combined with forecasting of variable loads and pricing and optimization methods to achieve energy and cost efficiency.

In the next section, we demonstrate our approach of applying analytics based on a semantic representation approach of the external data sources and the three on premise systems in the building in order to detect and predict those energy peaks. Then, we show how by relying on our energy storage model coupled with a battery backup system, we can effectively shave these peaks in energy consumption with building-specific forecasts.

III. Overall Architecture

Our approach is based on a semantic representation of the data retrieved from the existing systems on site, as shown in Figure 2. Each of these systems captures two types of data : topological data and timeseries data. Topological data expresses various information regarding the system, such as the devices' connectivity, geo-location. The timeseries data is a key, value pair of the collected readings along with a timestamp indicating the time of collect.

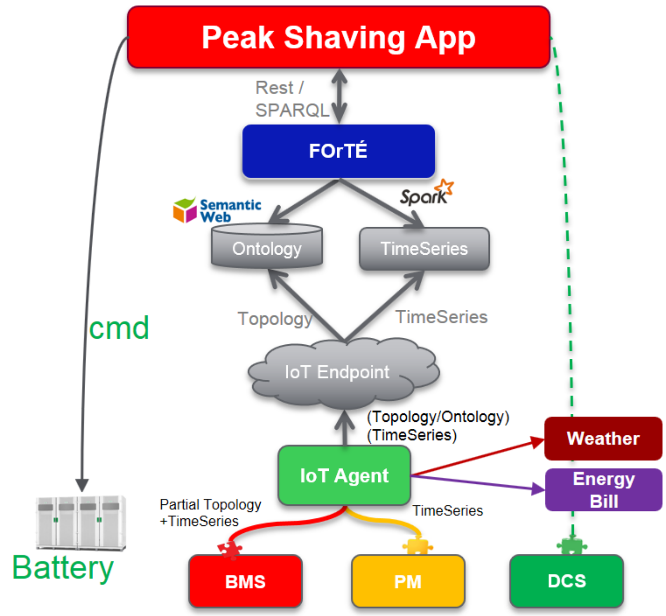

Fig. 2. Overall Architecture

The IoT agent retrieves data from our weather service such as outside temperature and humidity. We also collected the energy bills from the utility specific to our building and integrated them to our IoT agent to be sent to our cloud repository. In addition, our IoT agent as detailed in [5] collects the topological information and timeseries from the three on premise systems. The topological information is merged into one ontology representation as shown in Figure 3. The semantic representation of the three systems enables the interconnectivity of elements into one single view. For example, in Figure 3, `BMS1` is a Buildings Management System, it is connected to an Automation Server `AS-11` which controls a Compressor `Comp1`. This energy consumption of the compressor is measured by the Power Monitor system `PM1`. In addition to the interconnectivity, the semantic representation provides a formal representation of the systems with reasoning capabilities. For example, the relation `IsLocatedIn` shown in Figure 3 declared as transitive will be relied upon by a inference engine part of the ontology store to answer queries such as "Find all temperature sensors in Building A". The inference engine will rely on such properties and their nature to answer such query, for example, in Figure 3, `sensor1` is located in Room `202A` which is located in the `West2` area, on the floor `L2` of the building, since `IsLocatedIn` is declared as transitive, the inference engine will deduce that the sensor1 is actually located in the building. Transitivity is given as an example, however, other relations are deduced based on symmetricity and inversability.

Once extracted, the topological and timeseries data are pushed to an IoT cloud endpoint which routes the data according to the type. The topological data is inserted in an ontology store and the timeseries is persisted in a tabular store. The data in the two stores are linked based on the timeseries identifier. Our choices for separating the data into two stores is detailed in our previous work [11], [5].

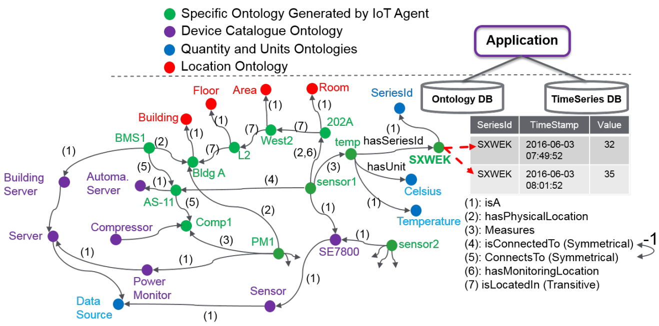

Fig. 3. Contextual and Timeseries Data Extracted from the Systems

FOrTÉ [11], our Federated Ontology and Timeseries query Engine bridges the gap between the two stores and provide one interface to query the data based on SPARQL [4] the standard query language promoted by the W3C. FOrTÉ takes as an input a SPARQL query and a desired W3C standard returned format such as JSON [12] or CSV.

We strongly argue that a semantic representation of the underlying systems provide a more coherent data of the overall system and provide a multi dimension views to be exposed to various contributors and actors in the overall system. For example, with a semantic representation, a maintenance operator can now track the connectivity issues between the various components of the system in addition to their physical location for replacement. Data scientists can now focus and extract the data they only need without the need to understand how a building management system, power monitoring system, and a data center system operate and is interconnected. They would only need to know how to extract data points based on abstraction relations (monitors, controls) provided by the ontology representation.

Our Peak Shaving Application performs semantic queries through FOrTÉ in order to retrieve the data points of interest such as external temperature, loads on the HVAC system, and total load on the building. Once the data points of interests are retrieved along with their timeseries information, our application runs to predict the energy consumption for the next 12 hours and utilizes model predictive control to send a command to the on site battery to be discharged in order to inject power on the site to compensate the peaks or to charge the battery for future use. As discharge energy is retrieved from our local battery instead of the utility, we therefore avoid penalties related to the energy surges.

We detail next our model for an energy storage system.

IV. Forecasting and Control for an Energy Storage System

We split the formulation as shown in Figure 4 into the following parts: energy storage model, control policies, data acquisition, and forecasting.

First, the data is retrieved by relying on our federated query engine FOrTÉ and feed to our two modules : the Forecast and the Model Storage. The Forecast module aims to predict the indoor temperature, in relation to the outdoor temperature and the expected energy consumption used by the HVAC system to heat or cool based on the desired indoor temperature. Then,

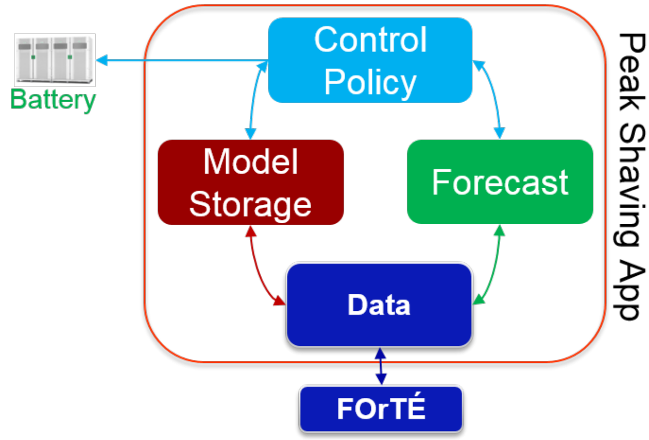

Fig. 4. Energy Storage Application Model

based on the forecast and the retrieved data, the model storage can now populate all the required elements such as the amount of energy at a given time of a device and the energy demand. The control policy module relies on the model storage and forecast data to take a decision regarding activating the battery on site or buying energy from the utility or other energy sources such as solar or wind energy providers in a smart grid environment.

A. Energy Storage Model

To describe our energy storage system, we rely on the general model and notation of Salas and Powell [9] and Jiang et al. [10].

1) Assumptions: We assume that the performance (efficiency, charge rate limits, capacity) of the energy storage is constant regardless of usage, i.e. the efficiency doesn't change with charging/discharging and historical usage doesn't wear down the storage. Models of these characteristics would be different for individual energy storage options, like lead-acid or Li-ion batteries, flywheels, or water tanks. Similarly, we assume that there no cost to holding the energy in storage.

Additionally, we assume that we are a relatively small consumer of grid power. In that case, we can demand as little or as much from the grid as desired and the price of grid power does not change significantly based on our policy to charge or discharge our own energy storage resources.

2) Storage Device: Our energy storage system can be expressed as follows:

- R^c: Energy capacity (MWh)
- η^c, η^d: Charging and discharging efficiency of device, respectively
- γ^c, γ^d: Maximum charge and discharge rate of device, respectively
- β^c, β^d: Minimum charge and discharge rate, respectively

3) State of the System: The state of the system at any given time t can be expressed as follows:

- R_t: Amount of energy in device at time t (MWh)

- E_t: Net amount of renewable energy available at time t (MWh)
- D_t: Aggregate energy demand at time t
- P^t: Price of grid electricity at time t
- $S_t = (R_t, E_t, D_t, P_t)$: The system state at time t
- x_t^{IJ}: Amount of energy transferred from source I to sink J at time t. Sources and sinks include:
    - W = renewable (for example, wind and solar)
    - D = demand
    - R = energy storage
    - G = grid
- x = $(x_t^{WD}, x_t^{GD}, x_t^{RD}, x_t^{WR}, x_t^{GR}, x_t^{RG})$
- ϕ = $(0, 0, -1, \eta^c, \eta^c, -1)$
- $R_{t+\delta t} = R_t + \phi^T x_t$

4) Constraints: These constraints hold for all times t in the set of timesteps T.

- Energy Charge Limit: $x_t^{WR} + x_t^{GR} \leq R^c - R_t$
- Demand: $x_t^{WD} + \eta^d x_t^{RD} + x_t^{GD} = D_t$
- Energy Discharge Limit: $x_t^{RD} + x_t^{GD} \leq R_t$
- Minimum and Maximum Charging Rates: $\beta^c \leq x_t^{WR} + x_t^{GR} \leq \gamma^c$
- Minimum and Maximum Discharging Rates: $\beta^d \leq x_t^{RD} + x_t^{RG} \leq \gamma^d$
- Flow Conservation: $x_t^{WR} + x_t^{WD} \leq E_t$

5) Objective Function: To simply minimize the peak, we use an objective function that minimizes the load flowing from the grid for all timesteps, T, considered.

$$\min_{t \in T} x_t^{GD} + x_t^{GR} \tag{1}$$

We note that more complex objective functions could be used to minimize energy cost consumption under varying energy prices in addition to minimizing the peak grid load. To combine peak-shaving with time of use load shifting, we could also use an objective that balances peak costs and time-of-day costs:

$$\min \sum_{t \in T} P_t \left(x_t^{GR} - \eta^d x_t^{RG} + x_t^{GD} \right) + P_L d \tag{2}$$

where P_L is the peak load cost with an additional constraint for the peak load:

- Peak Load: $x_t^{GD} + x_t^{GR} \leq d$ for all $t \in T$.

B. Data Acquisition

We rely on FOrTÉ to retrieve data from the various underlying systems in a given facility by formulating semantic queries. The process of retrieving the data from various systems to construct our knowledge graph and timeseries data are detailed in our previous work [5], [11]. Once the data is accessible to FOrTÉ, we retrieve the following data information for our Energy Storage application:

- Power consumption: from the Power Monitoring system which monitors several devices and equipment such as the HVAC, lighting, kitchen equipment, elevator loads. Thus, we retrieve the total power consumption of the facility and a granular consumption by type of equipment.
- Energy consumption: for the total building and by equipment type.
- Indoor Temperature & Humidity: are collected by room and floor from the various temperature sensors deployed in the building and connected to our building management system.
- Outside Temperature & Humidity: retrieved from our weather services in two parts: actual and forecast.
- Energy Prices: sent by the utility company and are integrated into our data repository.

C. Forecasting

To complete the model, our application will require forecasts for all uncertain or variable quantities. In many applications, this may be as simple as forecasting the next day's demand. However, forecasting may be required to capture hour-ahead or day-ahead pricing schedules or to estimate possible renewable energy sources included in the system, such as solar or wind generation. It might also be useful to forecast the energy usage of key components of the overall energy cost, such as chiller or HVAC loads.

D. Control Policy

Determining efficient control policies for (battery) energy storage systems have been the subject of many papers, such as Zhu and Hug-Glanzmann [6], Baker et al. [7] and [8]. We use model predictive control to solve the peak-shaving problem with forecasting. Model predictive control determines an optimal control policy based on the expected value of the forecasted values over a fixed time horizon (in our case 24 hours into the future). This produces a linear program that can be solved using any available linear programming tools, similar to the work of Camacho and Bordons [13]. More computationally intensive and accurate solution options, such as Markov decision processes or approximate dynamic programming, are discussed in Salas and Powell [9] and Jiang et al. [10].

V. Case Study: Energy Storage for a Large Office Building

Our case study consists of three parts: model setup, implementation, and simulated results.

A. Setup

This case study was performed at our North American headquarters facility. The load data was collected at 5 minute intervals for 274 days and used to simulate the deployment of an energy storage system behind the meter at the facility. A Lithium Ion battery of the following characteristics was chosen for this simulation:

- Size (R_c): 300 kWh
- Maximum charge/discharge rate (γ^c, γ^d): 4C - full discharge in 15 minutes
- Efficiency (η^c, η^d): 95% each way

Facility power demand experienced an expected mid day peak on workdays, generally peaking with around twice the demand of the facility after hours base load. On top of the daily peak, the facility experiences more frequent, smaller demand spikes from the starting and stopping of HVAC chillers. The weekends resulted in a flat demand equal to the work day after hours baseload. These days are less interesting and were omitted from the data set.

Additionally, the utility rate structure was assumed to be 3 tiered for energy charges, consisting of peak, off-peak, and part-peak charges. Demand charges were applied in peak and part peak times, in addition to an overall demand charge. Demand charges were assumed to be billed based on the peak usage in the relevant time period over a month. This rate structure is more progressive than some regions, however consistent with a trend towards Time Of Use (TOU) billing.

- Peak Hours:1 PM to 5 PM
- Peak Demand Charge: 18 \$/kW
- Peak Energy Charge: 0.14709 \$/kWh
- Part-peak Hours: 11 AM to 1PM, 5 PM to 7 PM
- Part-peak Demand Charge: 6 \$/kW
- Part-peak Energy Charge: 0.10275 \$/kWh
- Overall Demand Charge: 11 \$/kW
- Off-Peak Energy Charge: 0.07311 \$/kWh

B. Implementation

Our energy storage application is implemented entirely in Python 2.7 with a dependency on the following libraries: `Sklearn`, `Numpy`, `Pandas`, and `PuLP`.

To retrieve data from FOrTÉ, we relied on the W3C standard SPARQL [4] Query Language over REST. Listing 1. 1 depicts a simple SPARQL query to retrieve the power consumption of the chillers in our facility. The query requests the power meters installed in our facility which monitors the power consumption of the all the chillers in our facility. Additional queries aiming to retrieve both topology and timeseries data are detailed in [11] in addition to a performance evaluation for a scalable data retrieval reaching up to 1 billion timeseries entries in less then 30 minutes. Such scalability is provided through Spark[1] the large scale data processing engine.

[1] http://spark.apache.org

```
Select DISTINCT ?timeseriesId ?chillerName WHERE {
 ?powerMeter a pdevices:PowerMeter.
 ?powerMeter qt:monitors ?chiller.
 ?chiller a bldgs:Chiller.
 ?chiller qt:hasName ?chillerName.
 ?powerMeter bldgs:hasPoint ?point.
 ?point qt:hasMeasureType qt:RealPower .
 ?point qt:hasTimeSeries ?timeseriesId.
}
```

Listing 1: SPARQL Query: Power Consumption of all the Chillers

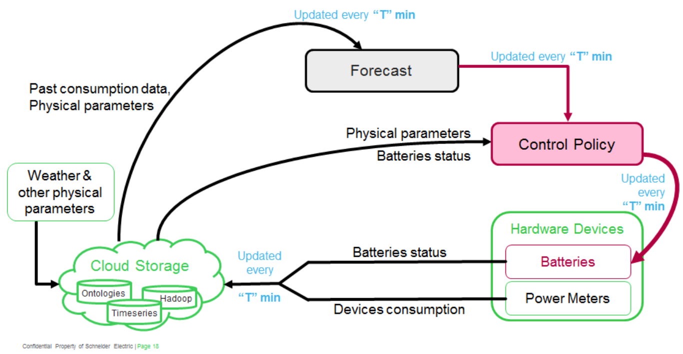

Fig. 5. Model Predictive Control Architecture

Since we are working with data from an office infrastructure, we first did some feature selection based on what factors influence the power consumption on a typical weekday or weekend. We then used K-Nearest-Neighbor (KNN) technique to forecast as we were able to leverage the fact that office space load consumption patterns do not show a lot of variation. This coupled with semantic queries via FOrTÉ enabled us to apply our forecasting algorithms effectively on a real time basis. The idea can be extended to most commercial and industrial sites where the load profile have certain recurring pattern.

C. Results

With the energy system and cost structure defined, a model predictive control (MPC) scheme was set up, as shown in Figure 5, leveraging the semantic data representation of the building data. The load forecast is updated with new data pulled through semantic queries via FOrTÉ. These updated forecasts update the linear program solution and in turn the battery control points.

Over ten test data segments of ten days, the MPC technique up was able to save money through both arbitrage and demand charge reduction, as shown in Table I. The presented results are displayed as savings extrapolated out to a month (30 days), as that is the frequency that demand charges are assumed to be assessed. Overall, the success of the control scheme is highly dependent on the accuracy of the forecast. A perfect forecast allows optimal control of the battery fairly easily. However, in the absence of an optimal forecast, the semantic queries via FOrTÉ allowed for quick updates of the forecast and control parameters which greatly increases the success of the energy storage control. Generally, as the frequency of updating increases the energy storage controls improve, highlighting the importance of the data storage and query system.

Overall, the success of the energy storage is still dependent the accuracy of individual forecasts, however the increased

Polling interval (hrs)	Avg. Utility bill ($/month)	Avg. Savings ($/month)
No Battery	63,867.18	0.00
24	63,253.333	613.85
12	62,802.25	1,064.94
8	63,564.30	302.88
6	63,054.64	812.54
4	62,997.68	869.5
3	63,002.98	864.20
2	62,828.00	1,039.18
1	62,338.78	1,528.40
0.5	62,162.14	1,705.04
0.25	62,168.03	1,699.15

TABLE I. RESULTS: TEST DATA SEGMENT 1

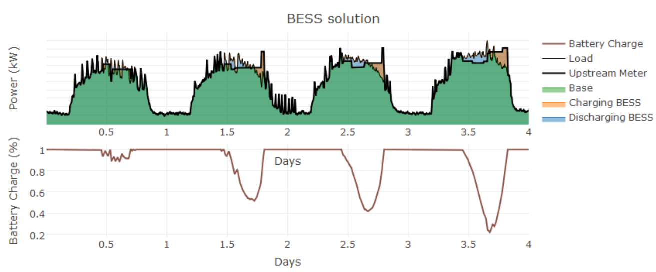

Fig. 6. Four Day Energy Storage Deployment.

frequency of updating allows the controls to be made in the shorter and generally more accurate region of the forecasts. Figure 6 shows a four day period as an example. Batteries are held in reserve early in the day, discharged in the middle of the day, and recharged later (when the prime time for peak load is over and energy prices are cheaper). This helps avoid under-prescribing and over-prescribing of the energy storage capacity, which lead to reduced economic savings. Over-prescribing is especially damaging, as the energy storage capacity usually runs out during peak demand of the day with the greatest demand, resulting in a utility meter peak at the worst time.

In the next section, we overview previous works related to our approach.

VI. Related Work

Efforts regarding energy management systems targeting optimization and efficiency in facilities have been around for many years now. In this section, we outline related work similar to our approach integrating an energy model storage and a semantic data representation.

Reegle[2] is a popular information portal in the renewable energy domains, it exposes clean energy data through a semantic representation in RDF [2]. OpenEI positions itself as a collaborative knowledge sharing platform with open access to energy information. Datahub[3] provides a variety of energy related data in linked and open data formats which can be integrated in our approach such as the energy prices by country and region. In addition, some energy demand historical data can be integrated for model training and accuracy improvement can be used.

Weise et al in [14] summarizes several initiatives around providing an ontology language to capture energy related data in buildings through a set of ontology alignment techniques to enable interoperability in smart city context. These models can also be integrated in our semantic representation to enable interoperability with other systems on a smart grid network.

In [15] an ontological model is proposed and validated on airport facilities to retrieve high level information regarding the performance of energy consumers. The Bat-MP [16] platform is an ontology-based energy management platform which relies on a semantic representation allowing several building automation protocols. The aim of the Bat-MP platform is to enable data sharing in a given building for research purposes related to the energy efficiency challenges.

Zhu and Hug-Glanzmann [6], Baker et al. [7], [8], Salas and Powell [9] and Jiang et al. [10] provide an energy storage control approaches, constraints on the size, charge rate, and discharge rate of the energy storage employed are combined with forecasting of variable loads and pricing and optimization methods to achieve energy and cost efficiency, our work differentiate mainly in the semantic data representation layer and our storage model.

Barton et al. [17] propose a similar approach to ours where they rely on a probabilistic method to predict the energy penetration on a weak electrical grids including renewable generation such as wind generators. However, no semantic representation is provided by the underlying systems.

Derguech et al. [18] propose an approach very similar to ours where predictive analytics is applied on semantic data representation from sensor data represented by the Semantic Sensor Network Ontology [19] combined with weather observation and prediction. The aim of their approach is to provide consumer with energy consumption and production predictions through a user interface. In comparison to our approach, we integrate a battery storage system to allow a system reconfiguration for energy saving.

VII. Conclusion & Future Work

In this paper, we investigated a micro-grid energy storage and an architecture enabling energy peak-shaving by combining demand forecasts and model predictive control to determine energy storage control policies. Our energy storage model utilizes inputs of outside/inside temperature and humidity in addition to the power consumption of the HVAC chiller equipment to predict the overall consumption of the facility.

The proposed approach was applied on our North American Headquarters facility instrumented with three systems: building management, power monitoring, and data center systems along with their associated equipment and sensors. The overall system is represented through semantic technology which provides a coherent view of the overall system explicitly detailing the interconnectivity between the various elements of the system. Such representation allows the easy access of information for various actors and contributors to the system. Thus, data scientists can now focus on the extraction of information needed for their application without the need to understanding in detail how the underlying system operates. We found that employing our method to a large commercial building facility can achieve substantial savings in utility costs.

In our future work, we intend to connect our model to the solar panels recently installed in our facility and integrate it in

[2]http://www.reegle.info
[3]https://datahub.io

the overall system to charge our on-site battery, consume the produced energy locally or inject it in a smart grid ecosystem. The solar load will require an additional forecasting model. Additionally, we will look to refine our load forecasting model and integrate other systems such as the meeting rooms reservation system and the occupancy rate of the building to derive better insights and better control decisions of our building management system.

REFERENCES

[1] D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, and D. Dennison, "Hidden technical debt in machine learning systems," in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, ser. NIPS'15. Cambridge, MA, USA: MIT Press, 2015, pp. 2503–2511.

[2] W3C, "Resource description framework," http://www.w3.org/RDF, 1999.

[3] B. Sean, v. H. Frank, H. Jim, H. Ian, M. Deborah, P.-S. Peter, and A. Stein, "Web ontology language," http://www.w3.org/TR/owl-features/, 2004.

[4] P. Eric and S. Andy, "Sparql query language for rdf," www.w3.org/TR/rdfsparql-query, 2004, sPARQL Query Language for RDF, W3C.

[5] C. E. Kaed, B. Leida, and T. Gray, "Building management insights driven by a multi-system semantic representation approach," in *IEEE WFIoT*, 2016.

[6] D. Zhu and G. Hug-Glanzmann, "Real-time control of energy storage devices in future electric power systems," in *PowerTech, 2011 IEEE Trondheim*, June 2011, pp. 1–7.

[7] K. Baker, G. Hug, and X. Li, "Optimal integration of intermittent energy sources using distributed multi-step optimization," in *Power and Energy Society General Meeting, 2012 IEEE*. IEEE, 2012, pp. 1–8.

[8] ——, "Optimal storage sizing using two-stage stochastic optimization for intra-hourly dispatch," in *North American Power Symposium (NAPS), 2014*, Sept 2014, pp. 1–6.

[9] D. F. Salas and W. B. Powell, "Benchmarking a scalable approximate dynamic programming algorithm for stochastic control of multidimensional energy storage problems," Department of Operations Research and Financial Engineering, Princeton University, Tech. Rep., 2013.

[10] D. R. Jiang, T. V. Pham, W. B. Powell, D. F. Salas, and W. R. Scott, "A comparison of approximate dynamic programming techniques on benchmark energy storage problems: Does anything work?" in *Adaptive Dynamic Programming and Reinforcement Learning (ADPRL), 2014 IEEE Symposium on*. IEEE, 2014, pp. 1–8.

[11] C. El Kaed and M. Boujonnier, "Forte: A federated ontology and timeseries query engine," in *The 3rd IEEE International Conference on Smart Data*. IEEE, 2017.

[12] W3C, "Serializing sparql query results in json)," https://www.w3.org/TR/rdf-sparql-json-res/, 2013.

[13] E. F. Camacho and C. Bordons, *Model predictive control*. Springer, 1999.

[14] W. Mathias and et al, "Ontologies and datasets for energy management system interoperability," 2014.

[15] N. M. Tomašević, M. v. Batić, L. M. Blanes, M. M. Keane, and S. Vraneš, "Ontology-based facility data model for energy management," *Adv. Eng. Inform.*, vol. 29, no. 4, pp. 971–984, Oct. 2015. [Online]. Available: https://doi.org/10.1016/j.aei.2015.09.003

[16] J. Caffarel, S. Jie, J. Olloqui, and R. Martínez, *Bat-MP: An Ontology-Based Energy Management Platform*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 9–16.

[17] J. Barton and D. Infield, "Energy storage and its use with intermittent renewable energy," *IEEE Transactions on Energy Conversion*, vol. 19, no. 2, pp. 441–448, 6 2004.

[18] W. Derguech, E. Bruke, and E. Curry, "An autonomic approach to real-time predictive analytics using open data and internet of things," *2014 IEEE 11th Intl Conf on Ubiquitous Intelligence Computing and 2014 IEEE 11th Intl Conf on Autonomic Trusted Computing and 2014 IEEE 14th Intl Conf on Scalable Computing and Communications and Its Associated Workshops (UIC-ATC-ScalCom)*, vol. 00, pp. 204–211, 2014.

[19] A. Sheth, C. Henson, and S. S. Sahoo, "Semantic sensor web," *IEEE Internet Computing*, vol. 12, no. 4, pp. 78–83, Jul. 2008. [Online]. Available: http://dx.doi.org/10.1109/MIC.2008.87