

# Demo Abstract: Toward Data-driven Demand-Response Optimization in a Campus Microgrid

Yogesh Simmhan, Viktor Prasanna

Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles CA 90089  
{simmhan, prasanna}@usc.edu

Saima Aman, Sreedhar Natarajan, Wei Yin, Qunzhi Zhou

Computer Science Department, University of Southern California, Los Angeles CA 90089  
{saman, sreedhan, weiyin, qunzhizh}@usc.edu

## Abstract

We describe and demonstrate a prototype software architecture to support data-driven demand response optimization (DR) in the *USC campus microgrid*, as part of the Los Angeles Smart Grid Demonstration Project<sup>1</sup>. The architecture includes a semantic information repository that integrates diverse data sources to support DR, demand forecasting using scalable machine-learned models, and detection of load curtailment opportunities by matching complex event patterns.

## 1 Introduction

The need for energy sustainability and security is leading to a greater focus on efficient management of this scarce resource. In the US, 35 % of total energy expenditure goes toward electricity generation and 21 % of the energy is consumed within residential and commercial buildings. Thus building energy management is essential for an effective energy conservation strategy. In this quest, electrical utilities are upgrading their infrastructure to better monitor and respond to changing consumer demand. Advanced Metering Infrastructures (AMIs), commonly known as *Smart Meters*, are a principal component of this improvement, providing the utility with realtime consumer power usage monitoring and bi-directional communication capability to interact with the consumers (or their software/hardware proxies).

The *Los Angeles Smart Grid Demonstration project* is deploying smart meters within the service area of the largest

<sup>1</sup>This material is based upon work supported by the Department of Energy under Award Number DE-OE0000192. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.


public utility in the US, the LA Department of Water and Power [2]. One key thrust of this project is to investigate techniques for demand-response optimization (DR) – curtailing the electricity usage by consumers during periods of high-demand when power generation capacity is insufficient. The University of Southern California (USC), as part of this project and as the largest private electricity consumer in Los Angeles, is serving as a campus microgrid testbed and exploring *informatics-based solutions for effective DR*[4].

The USC campus is a city within a city, having 100+ buildings with diverse profiles: offices, teaching facilities, libraries, dormitories, food courts, conference centers, and hospitals. These cover both residential and commercial use, making the campus representative of a small township of over 50,000 faculty, staff and students. The building and electricity management system at USC has innate instrumentation to sense realtime electricity consumption in every building on campus, accessible through the USC Facility Management Control Center. This is comparable to the monitoring capability available to electric utilities through smart meters. Sensors also measure airflow of HVAC units, ambient and setpoint temperatures in rooms, CO<sub>2</sub> levels, and occupancy. The Control Center can also centrally regulate power usage in buildings by changing the heating/cooling, air flow rate, lighting, and so on. This provides an information rich environment within which to evaluate DR strategies, which can later be scaled to a city power grid.

Here, we describe our prototype *software architecture* that advances data-driven DR optimization in the USC campus microgrid. Our demonstration will showcase the *semantic information repository* that integrates diverse data sources to support DR, *demand forecasting* using scalable machine-learned models, and determine *load curtailment* opportunities by detecting complex event patterns.

## 2 Software Architecture

Our software architecture (Figure 1) has dual goals that help perform DR: *accurate demand forecasting* and *effective load curtailment*. Rather than use prescriptive building and electrical models based on fine grained and complex system analysis, we take a data-driven approach that leverages data



**Figure 1. Campus Microgrid DR Software Architecture**

mining and machine-learning performed over large scale observational data of the electrical systems (e.g. power usage, equipment status) and indirect impactors of energy use (e.g. weather) from diverse information sources.

The microgrid is a dynamic and always-on system. Hence, information relating to it needs to be monitored continuously. *Information integration* is central our architecture and helps build a *semantic database* to support DR data analysis [3]. Our information integration pipeline abstracts common operations to provide a modular and extensible framework that can easily incorporate new information sources. The first stage of the pipeline is a *transport agent* responsible for acquiring raw data from remote sources using different access protocols. For example, weather data is retrieved using *HTTP REST* web service and equipment status from sensors using the *building control system's protocol*. Thus, it suffices to specify the location of a new information source rather than write custom code for data retrieval. Next, a *parser agent* extracts specific data attributes from the raw data. This agent supports parsing of diverse formats such as *HTML*, *Comma Separate Value* and *Excel* files to construct normalized data tuples. Lastly, a *semantic agent* annotates and maps these tuples into meaningful concept instances that are described by a *semantic ontology*. The ontology relates concepts from domains such as electric grids, equipment, building plans, room schedules, etc. The resulting semantic instances are stored in a Jena semantic database.

Demand forecasting predicts future electricity consumption in the microgrid for advance warning of a potential supply shortage. We use *machine-learning techniques* to build load forecasting models for the campus [1]. Our *regression tree* modeling technique uses historical electricity consumption information for campus buildings, recorded from sensors at 15 min intervals, as the feature to be predicted using features such as maximum outside temperature, building age and end-use, and the academic calendar. Prediction models are trained at the campus- and building- scales, and for daily and 15-min intervals of power demands. Once trained, the models use the current or future values of the predictor features to forecast the current or future electricity demand. The training and prediction data for the models come from the semantic database. The learning algorithms are data and compute intensive. We implement the learning models using *Map-Reduce* and run it across nodes of a *Hadoop* cluster.

Direct load curtailment often uses static, multi-level curtailment policies based on the load reduction required, for e.g., by reducing the cooling for all rooms in a building. However, such globally applied static policies can be punitive to user comfort. We use *complex event processing (CEP)* to be more discerning and detect patterns of events that in-

dicade curtailment opportunities at a lower cost. For e.g., a pattern to detect a HVAC unit that is cooling an office-room with no effect in reducing the room temperature over time may indicate an open doorway/windows leaking energy, thus allowing us to switch off the cooling unit or notify the occupant. Our semantic CEP system allows facility managers to specify such *load curtailment patterns* using higher level semantic concepts based on the ontology used by the semantic database. Patterns are defined using filters, aggregation and time windows defined over multiple event sources. The CEP system sources streams of events in realtime from the information integration pipeline directly and locates these patterns. A pattern to detect the earlier example would use events from active HVAC units and ambient temperature sensors, join them based on their co-location in the same office-space, and filter them based on a lack of temperature change within a time period.

The last component of the software architecture is a *policy engine*, currently under investigation. The engine will utilize the machine-learned demand forecast models for the microgrid to determine the need for load reduction, select the most suitable curtailment strategy based on static policies and dynamic opportunities detected by the CEP system, effect these curtailments, and monitor their impact using the semantic database.

### 3 Demonstration

Our demonstration will showcase three major components of our architecture: semantic information integration, machine-learned demand forecasting models, and detecting load curtailment patterns using CEP, all operating in realtime on the USC campus microgrid. These components will run on servers present in the Electrical Engineering department (EEB) on the USC campus and interactively run and display results on a portal interface accessed from a web browser running on a local laptop at the conference venue. The information integration pipeline will ingest data from multiple HVAC, ambient temperature, and power usage sensors in three buildings on campus along with the current weather from the NOAA's National Weather Service into the semantic database, which can be queried from the portal. The portal will be used to launch a Hadoop job on-demand to train a machine-learned model for campus demand forecasting. The resulting model will be used to plot and display the power usage prediction for the buildings and the campus inside the browser. Lastly, the portal will display realtime load curtailment patterns that are detected by the CEP system for these buildings.

### 4 References

- [1] S. Aman, Y. Simmhan, and V. K. Prasanna. Improving Energy Use Forecast for Campus Micro-grids using Indirect Indicators. In *Intl Workshop on Domain Driven Data Mining*, 2011.
- [2] www.smartgrid.gov. Los Angeles Department of Water and Power: Smart Grid Regional Demonstration, December 2010.
- [3] Y. Simmhan, Q. Zhou, and V. K. Prasanna. *Green IT: Technologies and Applications*, chapter Semantic Information Integration for Smart Grid Applications. Springer, 2011.
- [4] Y. Simmhan, et. al. An Informatics Approach to Demand Response Optimization in Smart Grids. Technical report, University of Southern California, 2011.