

# Real-world, Full-scale Validation of Power Balancing Services from Packetized Virtual Batteries

Kate Desrochers, Vanessa Hines,  
Forrest Wallace, John Slinkman, Andrew Giroux  
Packetized Energy  
Burlington, VT 05401  
Email: kate@packetizedenergy.com

Adil Khurram, Mahraz Amini,  
Mads Almassalkhi, Paul D.H. Hines  
University of Vermont  
Burlington, VT 05405  
Email: paul.hines@uvm.edu

**Abstract**—There is increasing consensus that flexible demand is critical to solve challenges associated with the rapid growth of variable renewable generation and aging transmission, distribution and generation infrastructure. Conventional direct load control programs are largely insufficient to address these issues. This paper presents results from validation tests of a new approach to demand side management, in which an aggregated fleet of devices is managed as a virtual battery, using principles that are found in communication networks: packetization and randomization. Validation results from a cyber-physical testbed with 5000 devices and a field-trial with 82 customer-owned water heaters show that the packetized virtual battery system can effectively solve a number of different problems. Customer satisfaction survey results illustrate that the system is able to maintain a high level of service quality.

## I. INTRODUCTION

Both academic research and industry experience have shown that flexible, responsive demand is critical to managing infrastructure costs and balancing the variability of renewable supply. The core concepts and basic enabling technology behind flexible demand go back to work by Morgan [1], Scheppe [2] and others in the late 1970s. But existing full-scale industry implementations have largely focused on direct load control (remotely disconnecting loads during peak hours) or behavioral demand response (asking customers to curtail demand), both of which require active engagement from customers and grid operators and can negatively impact quality of service.

Motivated by rapid growth in variable wind and solar supply and the availability of Internet of Things (IoT) technology for active sensing and control, new approaches are emerging that aim to bring the vision of flexible demand to reality. Transactive (auction-based) methods have been proposed [3] and demonstrated [4], but making auction-based systems easy to use for residential and small-commercial customers is difficult, since most do not have a sufficiently detailed understanding of electricity to effectively form bids. Others have proposed schemes that aim to simultaneously provide customer quality of service and grid services using hierarchical

decision models [5], [6] and dynamic ordering of resources based on fitness based assignments [7].

A number of researchers have recently shown that randomization can be a powerful tool for making fleets of loads continuously dispatchable [8]–[10], while avoiding potential synchronization effects that can negatively impact grid stability. While these approaches do employ sensing and decision-making at the device-level, the devices all respond to largely uni-directional broadcast control signals from the coordinator/aggregator that crucially depend on assumptions about devices being relatively homogeneous, which may lead to barriers for full-scale field implementation.

Coordinating large fleets of diverse devices to actively balance supply and demand in real time, while still being easy to use for both grid operators and end users is difficult. The approach employed in this paper, Packetized Energy Management (PEM), [11]–[13] addresses this difficulty by leveraging concepts that are key to internet communications. Specifically, PEM makes use of two powerful concepts: packetization to divide the delivery of energy into manageable chunks and randomization to spread packets of energy demand over time to align with a desired, dynamic schedule. Here we refer to a dispatchable fleet of devices operating under PEM (packetized devices) as a Packetized Virtual Battery (PVB).

Our prior work has shown, via software simulations, that a diverse fleet of packetized devices (e.g., water heaters, electric vehicles and residential-scale batteries [14]) can match demand to variable supply signals. But there are many examples in the literature of promising academic ideas that face practical challenges when transitioning to full-scale, hardware-in-the-field implementation. Similarly, the existing literature on flexibility does not provide evidence of how flexible demand algorithms can provide practical, usable solutions to real-world electricity industry challenges. The key contribution of this paper is the description and presentation of a full-scale cyber-physical validation and real-world commercial implementation of the Packetized Virtual Battery system.

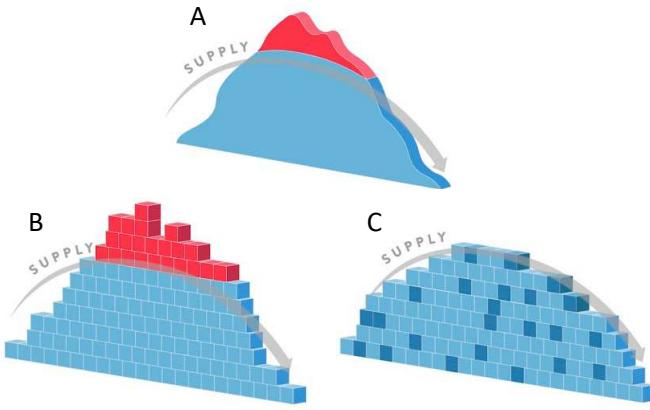


Fig. 1. Illustration of packetization and randomization applied to reducing excess peak load. **A:** Initial load profile. **B:** Packetization. **C:** Randomization process moves the energy packets in time.

## II. THE PACKETIZED VIRTUAL BATTERY SYSTEM

Packetized Energy Management [11]–[13] enables a fleet of distributed energy resources (DERs) to match demand with variable supply signals (such as renewable resource availability or market signals) through packetization and randomization (see Fig. 1). More specifically, a DER controller (e.g., for a water heater, HVAC system, EV charger, distributed battery) periodically wakes up, measures the local need for stored energy, and then probabilistically requests a packet of energy from a server. The probability of requesting increases as the device’s need for energy increases.

Packetized devices coordinate their actions by interacting with a PVB server, which accepts or rejects packet requests. When the PVB server gets a request from a device, it compares the total power consumption of the fleet with a pre-defined set point and accepts or rejects the request based on this comparison. The server is essentially device agnostic, in that it does not keep track of the state or location of individual devices. Instead the server focuses on managing the fleet as a single virtual battery.

The results in this paper are for the application of the Packetized Virtual Battery system to electric resistive hot water heaters. For this case, devices enrolled in a PVB are enabled to monitor the average tank temperature, using a weighted combination of upper and lower thermostat measurements. When the temperature is above or below the upper or lower temperature limits (e.g.,  $T_H = 140^\circ\text{F}$  and  $T_L = 120^\circ\text{F}$ ) the device opts out of the process and simply turns off or on as needed. When the temperature is between  $T_L$  and  $T_H$  the device will send a request for a packet of energy (e.g., 5 minutes at 4.5kW) with a probability that increases with the estimated tank temperature. As a result, water heaters with tank temperatures closer to the lower bound request packets more frequently than those with tank temperatures closer to the upper bound, thus ensuring that devices that need energy have a higher chance of consuming a packet. The stochastic request mechanism is described in detail in [13]. The opt-out



Fig. 2. The water heater management device used in this work, known as the Mello™ smart thermostat for water heaters. Two temperature sensors allow the device to measure the water/tank temperature near the tank’s upper and lower thermostats. Users can adjust the temperature set point (the middle of the flexible range) using the buttons on the device or through a phone app.

mechanism provides a guaranteed Quality of Service level: if the device’s state is below the lower bound (i.e., the water tank is too cold) it will opt out of the PVB and consume power until the temperature returns to within the acceptable range.

To implement the PVB system within a utility context, we designed a smart thermostat-style controller for electric hot water heaters (see Fig. 2), which creates a secure connection to the cloud-based PVB server using TLS (Transport Layer Security) 1.2 with on-device hardware encryption. The device gives the participating utility customer the ability to control their temperature set point, which subsequently increases or decreases  $T_L$  and  $T_H$  (within safe upper and lower limits). During device installation, the mechanical thermostats on the water heater are adjusted to ensure that the device has primary control over the water heater during normal operations. In order to enable the PVB server to accurately estimate the fleet-wide power consumption, each device is equipped with the ability to measure current, voltage and power factor.

Once formed, a PVB can provide electric utilities with a wide variety of valuable grid services, such as peak load reduction, energy price arbitrage, ancillary services and grid constraint management to defer transmission and distribution investments. The following subsections describe the application of PVBs to these various services.

### A. Peak load reduction

Peak load reduction to mitigate the need for new generation or bulk transmission system upgrades has been the principle use case for demand response programs from the beginning. Most of these programs use one of two formats: voluntary reductions in which utilities send customers messages asking them to manually reduce load, potentially with some financial benefit, and direct load control programs in which devices are remotely disconnected through some sort of communication system. The performance from voluntary or incentive-based programs can be hard to predict and may decay over time. Direct load control programs typically face two key problems (a) running out of stored energy and (b) rapid cold or hot “load

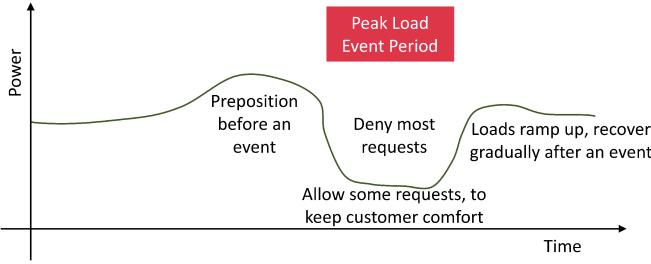


Fig. 3. Illustration of the PeakCrusher service

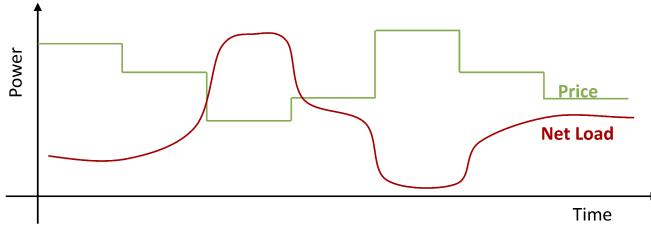


Fig. 4. Illustration of the LoadShaper service

pickup” at the end of demand response events. If the demand response event is long (a few hours), participating devices may run out of stored energy (get too hot or too cold) near the end of the event, which can cause customer discomfort. And then at the end of the event, all of the devices turn on nearly simultaneously, potentially resulting in a new peak load after the event period concludes, reducing the benefit and potentially introducing grid stability challenges if the resource is sufficiently large.

The PVB system deals with these challenges through the use of randomization, packetization and strategic set point adjustments. Specifically, the peak load reduction service (known as PeakCrusher) pre-positions loads before a demand response event, approves some packet requests during the event (enabling longer duration events) and then avoids load spikes after event by gradually restoring the load. Device-level randomization helps with this process by ensuring that devices are staggered in time, which prevents sudden changes in load. The PVB mitigates the potential for synchronized oscillations in load after the event by automatically ramping up the PVB server’s set point after the demand response event. By automatically computing these set point adjustments at the server-level, the only information required from the PVB operator is the timing and duration of the peak load event.

#### B. Energy arbitrage

Because there is very little energy storage in the grid, wholesale energy market prices fluctuate rapidly over a wide range within short periods of time. During periods of high wind and solar production, negative prices are increasingly common. During peak periods, wholesale prices can be 100x higher than average. PVBs can be used to arbitrage these price differences by adjusting net load to match the inverse of electricity prices. Our energy arbitrage service (LoadShaper) uses predictive algorithms to shape net load either to wholesale

market prices or to a pre-defined schedule. Because the server does not need to actively model device-level behavior, the problem of dispatching the PVB for load shaping is substantially simplified, and established algorithms for dispatching batteries for energy arbitrage can be used (e.g., [15]).

#### C. Additional grid services

As shown in Sec. III, the services in Secs. II-A and II-B have been field deployed and tested. In addition, the following two additional services are in development.

Our ancillary service tool (FastTracker) enables groups of grid edge distributed energy resources to provide frequency regulation and spinning reserve services into wholesale electricity markets. In this service, algorithms enable real-time prediction of available capacity, leading to more accurate bids and more accurate tracking, leading to increased ancillary service revenue potential.

Finally, our distribution network tool (GridOptimizer) enables reliable and resilient grid operations by managing grid constraints in real-time, thus mitigating the need for expensive transmission and distribution investments. By fusing data sources such as AMI, micro-PMUs, network models, and SCADA/EMS data, GridOptimizer optimally dispatches PVBs to mitigate overload, under-voltage, and over-current conditions.

#### D. Benefits and costs of virtual battery systems

Based on industry reports (e.g., [16]) and our own analysis, the value of grid services from fleets of DERs, such as electric hot water heaters, can range from \$100-\$400 per device per year, depending on location and device characteristics. A water heater with a 20°F flexible temperature range and an average load of 0.5kW (about a 10% duty cycle) provides services that are similar to a 0.5kW, 2 kWh battery. The device, installation and program marketing costs are typically less than \$300 per device, which leads to an upfront cost of \$150/kWh, which is less than half the cost of grid-scale battery systems (see [17] for a similar analysis). As manufacturers increasingly include “smart device” connectivity in their new appliances, the need for retrofit hardware solutions will decline, which has the potential to reduce these upfront costs substantially.

### III. LARGE-SCALE CYBER-PHYSICAL VALIDATION

The PEM enabled water heaters requesting energy packets from the PVB server constitute a cyber-physical system. This joint communication and computation system presents unique challenges such as communication delays and the need to ensure interoperability between diverse systems [18]. To better understand these challenges, we developed a cyber-physical testing platform that simulates the real-world implementation of the PVB system with a high level of accuracy.

Fig. 5 shows the resulting cyber-physical test bed. In this system, the PVB server is modeled with a Python-based web server that receives energy requests, as HTTP POST messages. The server replies to requests with yes or no depending on the difference between the actual power being consumed

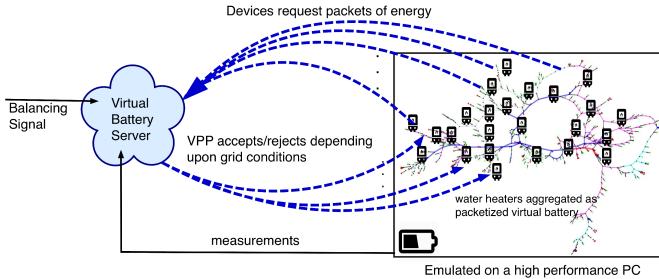


Fig. 5. Illustration of the cyber-physical testbed used for large-scale validation.

by the fleet of devices and real-time balancing signal (the set point). Each request message contains information about the requested power and the time length of the packet. If conditions are favorable, the server accepts incoming requests. Devices (water heaters) are simulated in separate processes (implemented in C++ code for computational efficiency), each of which includes a differential equation model of a water heater and the ability to send HTTP post messages to the server. These simulated devices are designed to operate in real time (one second of simulation time = one second of wall-clock time) in order to maintain realistic conditions. Both the server and the simulated devices run on a high performance multi-core PC.

To test the system, we simulated a fleet of 5000 packetized electric resistive hot water heaters, with an ON-state power consumption of 4.5 kW and a mean tank size of 200 litres (see Fig. 6). The packet size and mean time to request were set to be 3 minutes each. The device settings were initialized to ensure that tank temperatures remain within  $T_L = 49^\circ\text{C}$  and  $T_H = 61^\circ\text{C}$  range, with the set point at 55°C. Fig. 6 shows the PVB fleet tracking a balancing signal, approximately representing variable power availability from wind or solar resources. In this trajectory, the balancing signal for the first 50 minutes is relatively low, which decreases tank temperatures and the fleet-wide stored energy; essentially the virtual battery is discharging. After  $t = 50\text{min}$ , the balancing signal shifts to a higher mean value, representing an increase in wind or solar generation. The PVB fleet then charges for 80 minutes, increasing the average temperature of the fleet (see Fig. 6b). Through the opt-out mechanism, the mean temperature of the entire population is maintained well within the pre-established acceptable bounds, thus ensuring quality of service.

The results in Fig. 6 show that the PVB can charge or discharge to track a balancing signal, while simultaneously maintaining quality of service for end users. The results also suggest that the 5000-device PVB can provide nearly  $\pm 2.5\text{MW}$  of flexibility, i.e.  $\pm 0.5\text{kW}$  of flexibility per device.

#### IV. VALIDATION WITH A FLEET OF CUSTOMER-OWNED WATER HEATERS

An initial deployment of 300 Packetized water heater management devices (Fig. 2) began in March of 2018 in

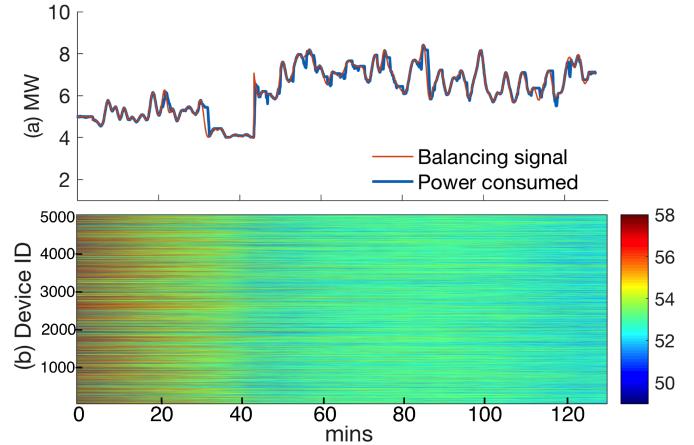


Fig. 6. (a) Packetized virtual battery, simulated in the cyber-physical testbed, tracking a balancing signal representing variable wind or solar availability, (b) Temperature profile of water heaters over time.

partnership with a rural Vermont utility. As of September 1, 2018, 82 devices were installed in the utility service area.

This utility operates the PVB for energy arbitrage and peak reduction purposes using the LoadShaper and PeakCrusher modules. When operating in the LoadShaper module, as described in Sec. II-B, the predictive algorithms incorporate five-minute and day-ahead pricing signals from the regional wholesale grid operator, ISO New England, to shape the net load. Fig. 7 provides results from the LoadShaper module deployed from midnight to 3:30AM on August 11th. Locational Marginal Prices (LMPs) throughout the day hovered around \$40/MWh, ranging from \$25 per MWh to \$60/MWh throughout the day. As shown, the power set point inversely tracks the LMP, increasing power consumed when prices are low and decreasing when prices are high.

Fig. 7. Illustrative results from the LoadShaper service dispatching the virtual battery (upper figure) based on the trajectory of locational marginal prices (lower figure).

When operating in PeakCrusher mode, the PVB can reduce costs associated with annual and monthly peak loads. From

May 1st to September 1st, the utility scheduled 36 peak events ranging from 2 to 6 hours in duration. Figure 8 illustrates the performance of the PVB with 81 devices deployed and a peak scheduled from 15:00 to 21:00 on a weekday in August. Assuming that the peaks are accurately scheduled, this type of peak reduction can substantially reduce utility costs for regional transmission and generation capacity, which are assessed based on contributions to annual peak loads.