Author Notes

- ✔ The type has been replaced in your figures; please proof the figures carefully for any errors.
- ✔ Please confirm all artwork in this article is either original or that permission for use has been obtained and credit given where needed.
- ✔ The last page contains callouts that might be used in the article, depending on design and spacing considerations. Are these callouts appropriate? If not, please provide several alternatives.
- ✔ When the magazine is printed, you will receive courtesy copies of the issue to distribute to your coauthors. Please reply with your complete mailing address, including phone number.

By Srinivas Katipamula, Robert G. Lutes, Sen Huang, and Roshan Kini

Building Energy Systems as Behind-the-Meter Resources for Grid Services

Intelligent load control and transactive control and coordination.
<AU: Please check that the edited subtitle is acceptable.>

xxxxxx

To mitigate the impacts of climate change, significant reductions in emissions from all sectors of the economy are needed. The electricity generation sector has embarked on an ambitious plan to include renewable generation as part of its decarbonization efforts, and many cities and states are mandating all-electric buildings. While renewable resources will reduce emissions, they are not dispatchable, they vary temporally, and their generation is uncertain. Under these conditions, traditional approaches to managing grid reliability, where supply follows demand, will not be efficient and may not be cost-effective. There is a more efficient alternative for balancing the supply–demand imbalance and for absorbing variability and uncertainty of renewable energy using distributed energy resources (DERs) as opposed to reserve generation. Because buildings consume more than 75% of total U.S. annual electricity consumption, behind-the-meter (BTM) DERs have a load flexibility of 77 GW of power and 90 GWh of virtual energy storage capacity nationwide (Kalsi, 2017). Therefore, some portion of the supply–demand imbalance can be met by these DERs at a lower cost compared to business-as-usual solutions.

This article describes the nature and source of BTM DER flexibility in buildings and how to tap into the flexibility to support grid reliability using the two applications: 1) intelligent load control (ILC) and (2) transactive control and coordination (TCC). In the TCC approach, we create an electricity market within a building to coordinate and control the DERs. The building market can also be coordinated with other external hierarchical transactive markets. The ILC approach also coordinates BTM DERs but does so through a building-level peak demand target. This article introduces the essential concepts underlying these two applications and provides results obtained from simulation studies as well as tests from real buildings that show peak load reductions of 10% to 20% for 4–6 h without significant loss of service levels. Deeper peak reductions are possible, but that will affect service levels. These solutions will be suitable for deployment in a diverse set of commercial buildings with and without a building automation

Digital Object Identifier 10.1109/MELE.2022.3211018
Date of current version: xxxxxx

2325-5897/22©2022IEEE

system (BAS). Lessons learned from the case studies are also presented.

BTM DER Flexibility

Buildings account for more than 75% of the total U.S. electricity use. The composition of the end-use loads within commercial and residential buildings are shown in Figure 1(a) and (b), respectively. About half of the electricity consumption (space cooling, ventilation, lighting, refrigeration, and water heating) in commercial buildings can be controlled; however, with the current control infrastructure, only space cooling and ventilation end uses are easily controlled. The situation is similar in homes—about one third of the electricity consumed could be controlled (cooling, water heating, and space heating). Buildings should offer additional flexibility as the building sector is completely electrified, and controlling other DERs will be a possibility soon as controls infrastructure is deployed in buildings, especially connected lighting.

When nonthermostatically controlled systems, such as connected lighting, shed or modulate their load, it will result in reduced energy consumption during the event. For these systems, there is no recovery after the grid service event. When thermostatically controlled systems, such as heating, ventilation, and air conditioning (HVAC) systems or hot water heaters shed, shift, or modulate, reduced energy consumption results during a grid service event. Even so, this could cause an overall increase in energy consumption. A nationwide estimate of demand flexibility from residential air conditioners and heat pumps (HPs), commercial packaged HVAC rooftop units (RTUs), residential and commercial water heaters, and residential pool pumps is about 77 GW of power and 90 GWh of energy (Kalsi, 2017). This virtual battery capacity is significant if it can be tapped to provide grid services.

Automated Grid Services Using BTM Assets

For many commercial buildings that are subjected to traditional utility rate structures, utilities charge not only for the energy (kilowatthour) consumed but also for the peak power (kilowatt) demand. The peak is calculated as a rolling 15- or 30-min average during a billing cycle (typically 30 days). The electricity consumption in commercial buildings typically peaks for a short duration one or more times during a day. A portion of the peak can be avoided by managing BTM DERs without sacrificing service levels. However, to benefit from reduced demand charges, the peak must be managed every day of the billing cycle. The use of BTM DERs for utility demand response is not a new concept. In some regions of the United States, aggregators have used BTM DERs to provide capacity relief for more than two decades. However, many of these deployments are either direct load control or highly customized, making them difficult to scale. Large-scale participation of BTM DERs requires a highly scalable deployment. One of the attributes of scalable deployment is highly automated application.

ILC and TCC are two complementary automated technologies that can be used to manage BTM DERs. Although the goal of both technologies is to manage building electricity consumption to reduce cost and support electric grid reliability, the approaches used to achieve this goal are different. While ILC was designed to serve the current utility needs and serve as a bridge toward the future, TCC is more forward looking and is a natural choice in a transactive energy environment. ILC can be used to support the

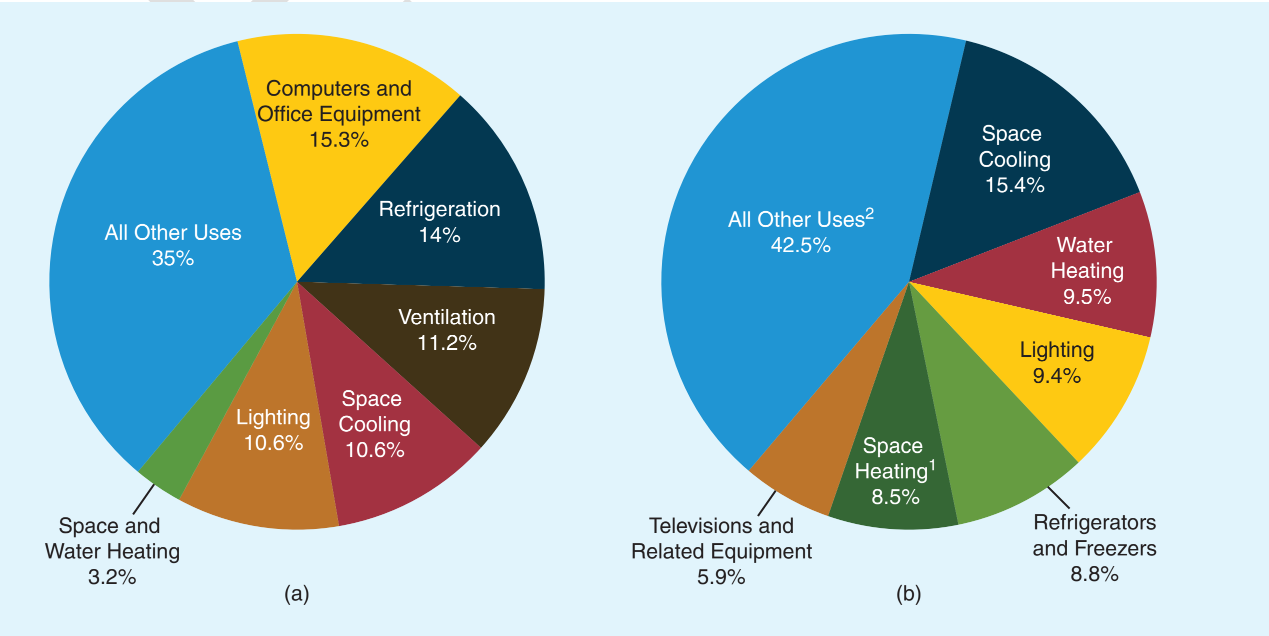

Figure 1. The opportunity of demand flexibility in buildings: the 2017 U.S. (a) commercial and (b) residential sector electricity consumption by major end uses (U.S. Energy Information Administration, 2018, Tables 4 and 5).

following grid service use cases: 1) peak load management (PLM), 2) time of use (TOU), 3) real-time pricing (RTP), 4) critical peak pricing (CPP), 5) capacity bidding, and 6) event-driven demand response.

ILC

The ILC application is used to provide grid services by managing BTM-controllable DERs while mitigating service-level excursions (e.g., occupant comfort, lighting comfort, minimizing equipment on/off cycling) by dynamically prioritizing available DERs for curtailment using both quantitative (thedeviation of zone conditions from the set point) and qualitative criteria (e.g., the type of zone). This application uses a business decision-making process to prioritize DERs to generate a numerical score to prioritize each alternative load based on associated decision criteria.

Although the ILC process supports several grid service use cases, all use cases result in the generation of a building-level peak demand target (or goal) for ILC to manage. ILC decomposes the problems into a hierarchy of the elements influencing a system by incorporating three levels, as shown in Figure 2: the goal, criteria, and alternatives of a decision. The ILC process prioritizes a set of criteria used to rank the alternatives of a decision and distinguish, in general, the more important factors from the less important factors. Pairwise comparison judgments are made with respect to the attributes of one hierarchy level given the attribute of the next level up, from the main criteria to the subcriteria.

To illustrate the ILC approach for managing the peak electricity demand, we use a building that has a set of RTUs, as shown in Figure 2. Load management options are determined by comparing alternatives with respect to a set of criteria. The goal is to generate the dynamic load curtailment priority of individual RTUs for managing building electricity consumption to a target level. In this example, five decision criteria are used to manage building electricity consumption without significantly affecting occupant comfort. Additional criteria can be easily added, or existing criteria can be modified or removed. For a decision criterion to be effective, it must be able to capture important characteristics that have a direct impact on control. In this example, four quantitative criteria and one qualitative criterion (the type of zone) are used.

First, a pairwise comparison is conducted to qualitatively determine which criteria are more important and then assign a weight to each criterion. The last layer in Figure 2 consists of different decision alternatives, which are multiple RTUs that can be controlled to manage the building energy consumption to the desired target. The ILC process selects RTUs with the highest priority level that can be curtailed for the longest duration of time without a comfort penalty during the event period. A more detailed description of the ILC process is provided by Kim and Katipamula (2017). It was deployed for testing and validation on buildings with PLM, RTP, and capacity bidding use cases. In these tests, ILC was able to successfully control several DERs, such as RTUs, variable air volume (VAV)

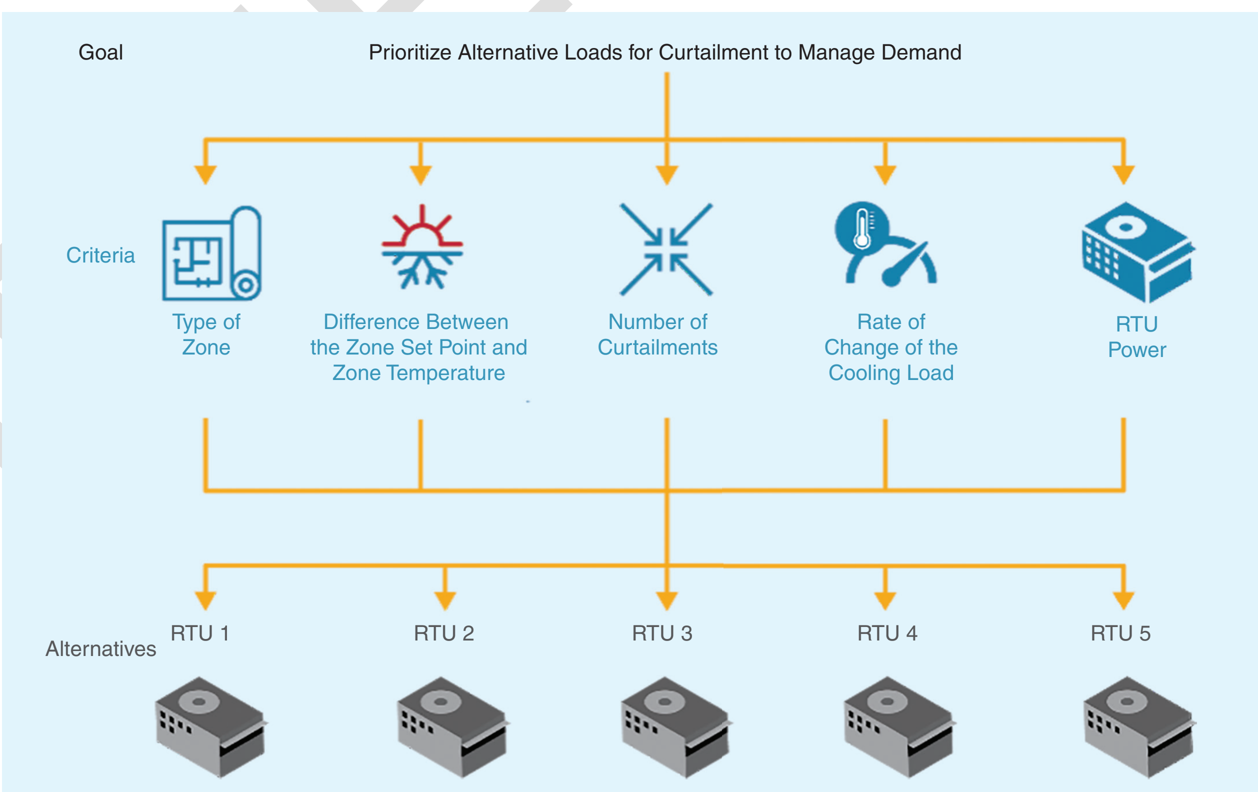

Figure 2. An example of the AHP ***<AU: Kindly spell out AHP.>*** process that uses RTU loads to manage the building peak demand.

boxes, dimmable lighting fixtures, etc. to achieve the necessary objectives.

TCC

Dynamic utility rates (the TOU, CPP, and RTP) and demand-response programs (e.g., capacity bidding) are designed on an implicit assumption of the aggregate load flexibility (or response) from buildings. Therefore, communication between the utility and buildings is one way, i.e., prices to devices, with no feedback from the building loads to establish the optimum price. Unlike "prices to devices," TCC is a two-way process; i.e., it includes feedback between the price and quantity. Because every flexible load in a TCC process plays a role in balancing the electricity supply and demand, the price is a function of electricity quantity in the coordination function. Other objectives besides the balancing objective might affect electricity price as well. On the consumption or demand side, quantity is a function of the price and status of the controlled devices within a building as well as other local conditions.

For illustrating the TCC concept, a distributed hierarchical market structure is used (Figure 3). In this approach, there is a market at the building level (shown by the oval), the neighborhood level, and the distribution level. In each of the markets, supply-and-demand curves are balanced, and other constraints and objectives are imposed. More details on the implementation of a hierarchical market approach can be found in Katipamula et al. (2019). In the TCC approach, the first step is to generate the price-capacity (or flexibility) curve for each participating DER. The price-capacity curve expresses the inherent flexibility of a device as a function of price, as shown in Figure 4(b). The coordination of hierarchical markets is critical for reliable implementation, including connectivity to external sources. Because this article is focused on the application of TCC to buildings, we discuss only how the demand curves are created and aggregated in a building and used to support the TCC process.

Price-Capacity Curves

In a price-capacity curve, the x-axis represents the capacity, and the y-axis represents the price. The capacity bounds (q_- and q_+) are constructed as follows: the system that this curve represents requires a certain minimum power or energy to meet the service levels and is willing to pay any price for it. For example, in a VAV air-handling unit (AHU), this could represent the minimum ventilation that the system must always provide to meet the service levels. Likewise, there is a certain maximum energy or power the system can consume, irrespective of the price (even if the price is negative), to meet the service levels. The minimum and maximum capacities bound the capacities curve. Generally, it is sufficient to linearly interpolate between the minimum and maximum capacities, as shown in Figure 4. However, for systems that are nonlinear, other points on the curve can also be estimated.

On the price side, the mean price is estimated using historical prices, such as prices from the last 24 h. The minimum price is the difference between the mean and the product of k and the standard deviation (σ) of the price, again based on the prices of the last 24 h. The maximum price is the sum of the mean price and the product

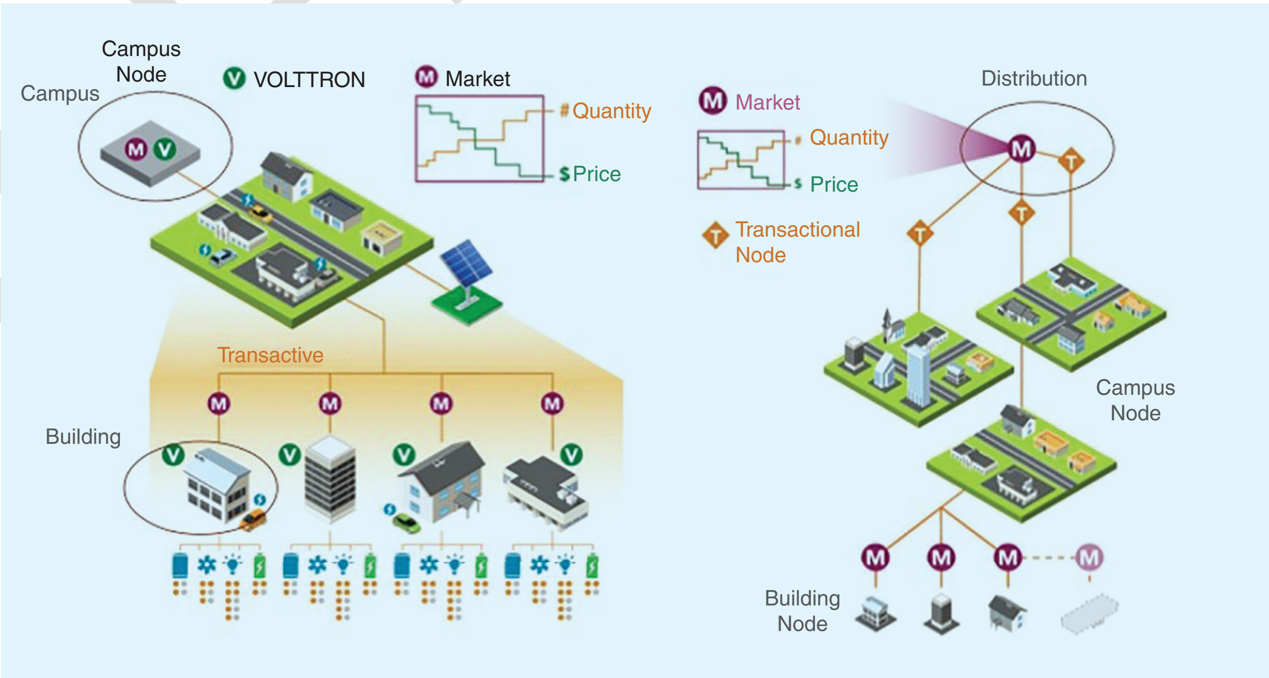

Figure 3. The distributed hierarchical market.

of k and σ. Note that this choice of estimating the minimum and maximum prices is somewhat arbitrary. The building manager or occupants determine the value of the scalar "k." For example, a value of k = 0 maximizes the demand reduction or energy savings, and, as k increases, the weighting of comfort increases relative to the weighting of demand (or energy) reduction. Because the flexibility of loads is dynamic, these price-capacity curves must be estimated for each market interval. If the market is based on a single interval, only one curve is required to represent the next market interval. For markets that are based on day-ahead pricing, flexibility must be estimated for the next 24 h.

The TCC Process in a Building

To illustrate this process, we consider a building with a VAV AHU. The chilled water and hot water to this system are provided by a chiller and a boiler, respectively. The TCC

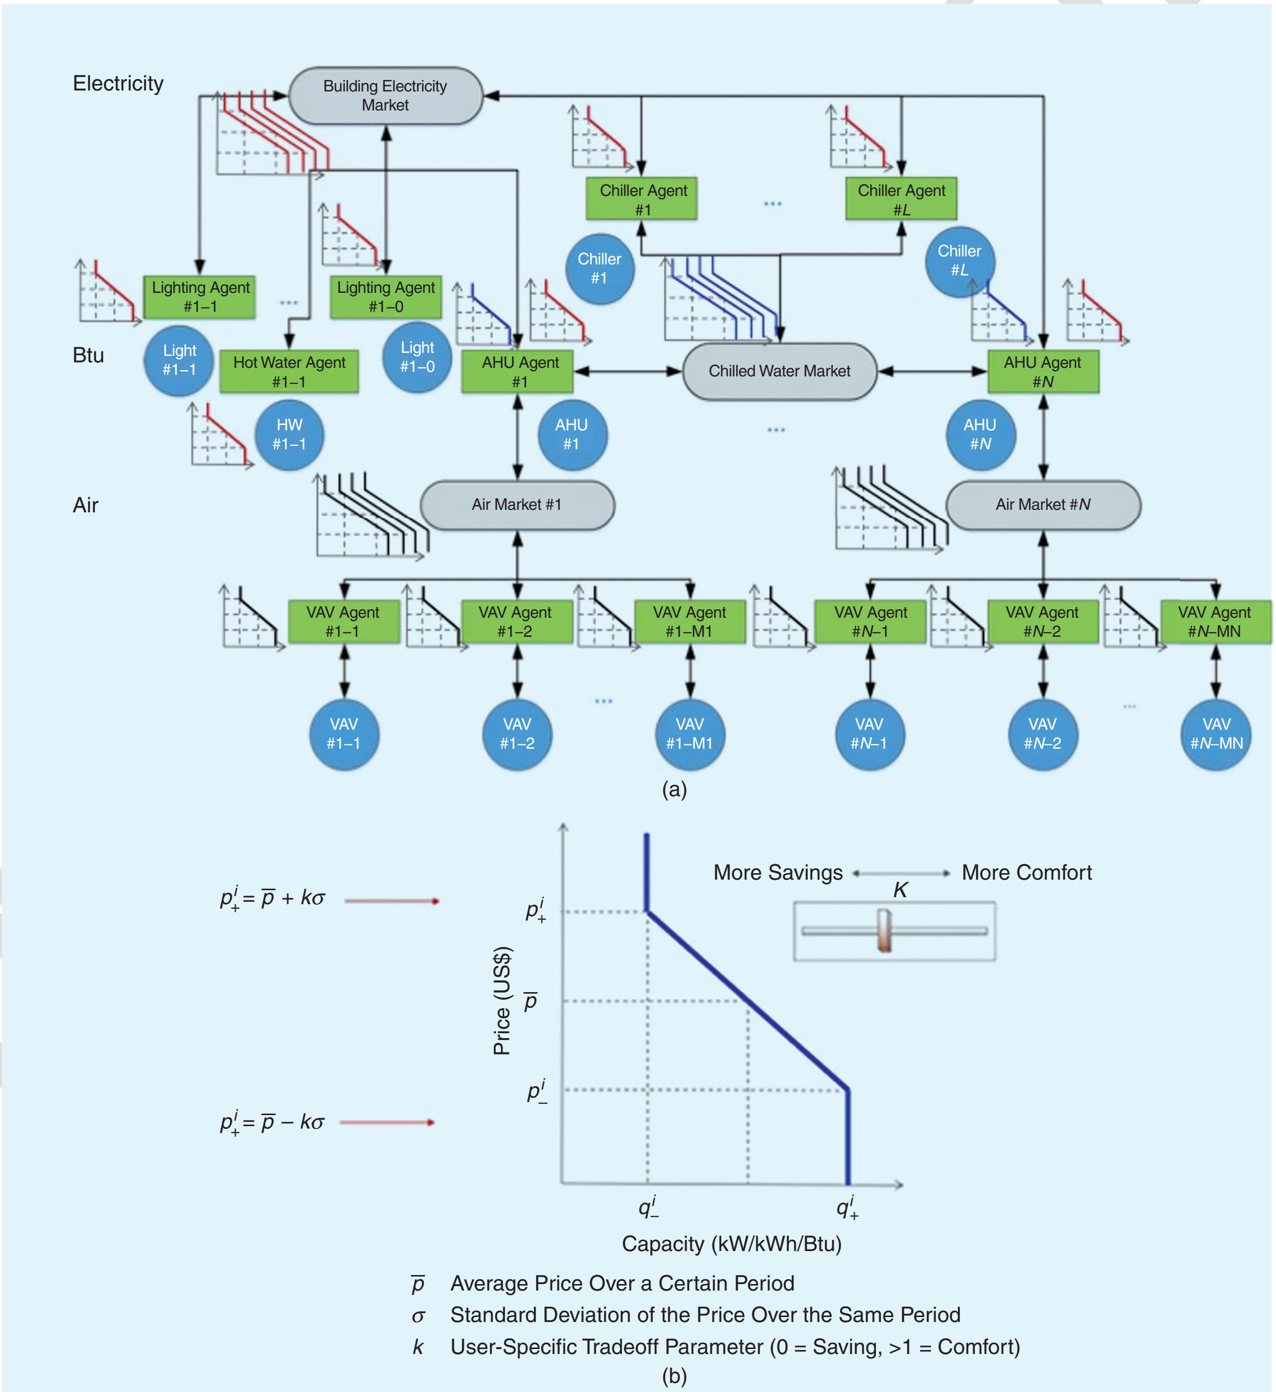

Figure 4. (a) Multiple commodity markets inside a commercial building and (b) a generic price-capacity curve. ***<AU: Please check if k, q_-, and q_+ should be set in roman or italic fonts.>***

process for this use case includes the following five different commodities: 1) conditioned air, 2) chilled water, 3) hot water, 4) local electricity, and 5) local natural gas. Each commodity requires a separate market. However, because grid services are only for electricity and because conditioned air and chilled water are generated using electricity, these commodities are eventually converted into equivalent electricity at the building level and become part of the local (inside the building) electricity market. Therefore, the local electricity market constitutes the electricity consumption associated with conditioning air, the generation of chilled water, other flexible electricity loads (such as connected lighting, electric hot water tanks, etc.), and all nontransactive or inflexible loads (such as plug loads, emergency lighting, etc.).

Although storage devices (thermal and battery) are flexible loads, for simplicity, these devices are not considered in this use case. Because hot water for reheating is generated using natural gas, hot water use is not converted into electricity. Therefore, the TCC process includes only three markets: the "air" market for buying the conditioned air flow that is necessary to maintain the required zone temperature; the "Btu" market for buying chilled water for the AHUs, which is purchased from the chillers; and the local "electricity" market that represents all electricity loads that participate in the TCC market as well as the nontransactive/noncontrollable electric loads, as illustrated in Figure 4. In Figure 4, the devices and controllable loads are shown in blue, agents that represent the DERs are in green, and markets are in the gray background.

The bidding or market cycle begins with the thermostat agent, which uses the price-temperature curve that essentially relates the price to the predetermined comfort expectation of the building occupant or manager. In this TCC process, an AHU obtains electricity at a certain cost directly from the building-level energy market and then sells its product—cool air—to zones within the building. The zones electronically "bid" on the cooling capacity based on price and desired occupant comfort levels. The curve influences AHUs to either reduce the power to balance cost and comfort objectives or, in cases of abundant economical electricity, perhaps increase consumption to perform tasks, such as precooling a building in anticipation of higher ambient temperatures that coincide with higher electricity prices. For more details on how the markets are cleared, refer to Katipamula et al. (2019).

Application of ILC and TCC in Buildings

In this section, we show how ILC and TCC applications can be used to manage BTM DERs to provide grid services.

PLM With ILC

PLM can be used to reduce the building peak electricity or when the utility system peaks, which may not be coincident with the building peak. To illustrate how ILC can be used for PLM, we use a 22,000 ft^2, **<AU: Kindly check that the expansion of "sf" is correct.>** single-story building (building 8) located in eastern Washington, which was built in 1980 with office, maintenance shop, and storage areas. The building has 11 HPs with electric backup heaters that can be controlled to manage the peak. To manage the building peak, the building peak consumption over the 30-day billing period is first forecasted. Next, a peak demand target is selected. Using the established target as the goal, ILC is used to control BTM DERs. For this use case, the selection of the target is somewhat arbitrary; however, tests have shown that reducing the peak between 10% and 20% will not result in a significant loss of service levels. Deeper reductions (>20%) are possible, but some compromise in services will occur.

To forecast the electricity consumption, a whole-building energy load forecast model is used with historical electricity consumption data and a set of independent variables (e.g., the forecasted outdoor air temperature, day of the week, and hour of the day). Using this model, the expected energy consumption is computed using future weather forecast data. To manage the building peak when the utility system peaks, a signal for the utility will be required. A building target is selected using the signal and the incentive associated with the signal.

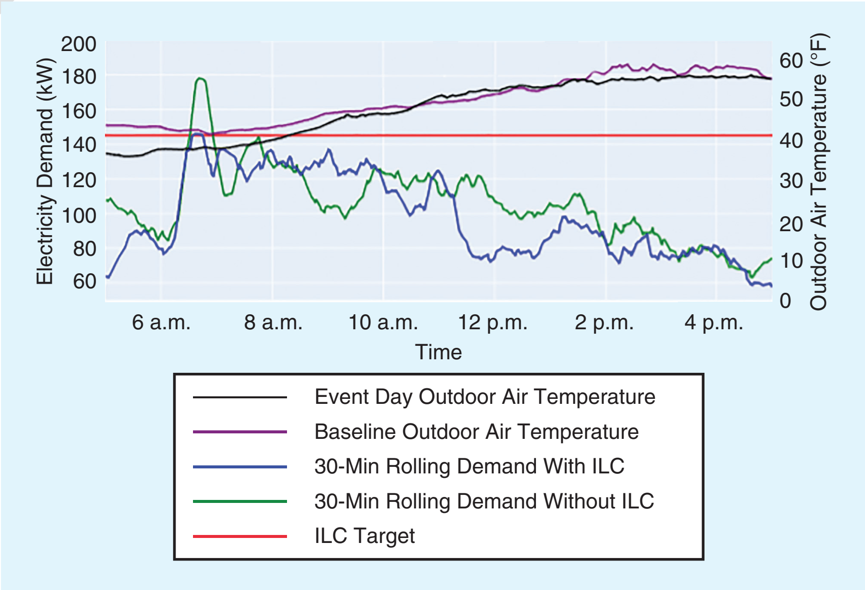

Figure 5. The 30-min rolling average electricity demand (green and blue), outdoor temperature (purple and black on the right *y*-axis), and peak demand target (red) during the heating season in building 8.

Although only single-day results from a heating season highlight how BTM DERs can be used for PLM, it can be applied during other seasons as well. To benefit from lower kilowatt charges over a billing period, ILC must run each day to manage the peak. Figure 5 shows the 30-min rolling average power consumption (baseline, without ILC in green, and the ILC-managed peak in blue), the peak demand target value (red), and the outdoor temperature (purple without ILC and black with ILC plotted on the right *y*-axis). The rolling 30-min average power consumption value is calculated using measured 1-min power consumption. The power consumption shown for 14 March (green), which represents business as usual, or normal building operation, peaks at 180 kW during the morning warm-up period between 6 a.m. and 8 a.m. The next day, on 15 March, ILC was deployed to manage the peak in the building (blue). The peak demand target was set at 145 kW, which is 20% lower than the peak that was established on 14 March. As the building electricity consumption was peaking early in the morning, ILC prioritized HP operations, shutting down some units while allowing others to run. Overall, the results show that the ILC was able to control the HPs and keep the actual consumption under the target demand level of 145 kW.

Using the 30-min moving average to prioritize the loads for curtailment will result in the peak demand overshooting the target value. Likewise, the use of instantaneous power measurement (1-min frequency) will result in an excessive curtailment of loads and a suboptimal result. Therefore, instead of using the 30-min average or instantaneous power consumption value to make curtailment decisions, ILC uses an exponential moving average (with a span of 15). The exponential moving average puts additional weights to values closer to the latest time, which gives ILC a good sense of when the 30-min average is likely to exceed the target value. ILC will start managing DERs whenever the exponential moving average exceeds the target value. Between 6 a.m. and 8 a.m., the exponential moving average exceeded or was equal to the target value on three different occasions.

While curtailing HPs, ILC also ensures that the service levels are within the established limits. ILC curtails HPs by lowering the heating set point between 1 °F and 2 °F. For most HPs, when the zone temperature is decreased more than 2 °F below the set point, both the heater and compressor(s) would turn on at the same time. To prevent this, the set point reduction for any HP is limited to 2 °F.

RTP With ILC

ILC can also be used to manage the building peak under dynamic rates, such as RTP. Under dynamic rates, the objective is to manage the electricity consumption in a building to minimize the energy cost. Utilities offer two types of RTP rates: RTP and day-ahead RTP. In the case of day-ahead RTP, as the name implies, the price of electricity for the next 24 h (e.g., midnight to midnight) is published a day ahead. Under the day-ahead RTP, there is an opportunity to optimize consumption over a 24-h period. For the other RTP rate, the price of electricity is published an hour ahead. Therefore, the energy cost cannot be easily optimized beyond the hour.

To illustrate how ILC can be used to manage building peak, we use a 20,000 ft^2, single-story office building (building 6) built in the 1980s and located in eastern Washington. The RTP signals for the California Independent System Operator **<AU: Please note that footnotes are not permitted as per magazine style. This footnote has been incorporated into the "For Further Reading" list. Please check that the placement is okay.>** are replayed. For ILC to manage the peak in a building, a target must be established. The target is established using the same concept as the price-capacity curve described previously. Figure 6 compares the 15-min rolling average whole-building electricity (WBE) for the baseline day (green), RTP control day (blue), target (red), and RTP (right *y*-axis). Because the building is occupied between 6 a.m. and 6 p.m., the target for ILC is estimated only during that period. For this use case, the consumption is not optimized over the day. In general, the target value is higher when the price is low and is relatively lower when the price is high.

Coordination of DERs in a Medium Office Using TCC

To illustrate the TCC process, we use results from one simulated and one real building. The simulated (referred to as *medium office*) and the real (referred to as *medium office 2*) buildings are approximately 23,000 ft^2, and each is served by two multizone VAV and two single-zone constant-volume systems. The two VAV AHUs serve 22 zones, and the two single-zone constant-volume systems each serve one zone. Note that, with these buildings, markets are created at the zone, system, and building levels as described previously. Markets are also created at the distribution and utility levels, but those are not discussed in this article. The price-capacity curves for the various commodities are automatically created and aggregated, and they are eventually aggregated to a single price-electricity capacity curve at the building level.

The electricity curve at the building level not only incorporates flexibility from VAV boxes, AHUs, and chillers, it also includes inflexible electricity loads. Although there could be multiple suppliers of electricity, in the tests we conducted, there is only one supplier of electricity that uses a two-tiered wholesale electricity price under which nearly all of the network's electricity is supplied. In addition to the two-tiered electricity price, there are also locational price differences that are small but meaningful and are intended to represent the impact of electricity losses. The wholesale electric rate is presumed to exist at a remote transmission circuit region from which the utility imports its electricity. Because losses are incurred while importing this electricity, the effective cost of electricity at the utility node is somewhat higher than at the remote transmission point from which the electricity is imported. Losses are also incurred as the distribution node imports its electricity from the utility node, so the distribution node's price is found to be a little higher than that of the utility node.

The real-building tests were conducted during the winter period, but the simulation tests did not have any restrictions. Therefore, we present the simulation results for a typical summer day and the real-building results for a typical winter day. Although the simulated building is identical to the real building, a direct comparison of results is not possible for many reasons. Figures 7 and 8 show the actual and predicted zone airflow rates for selected zones in the simulated medium office and real medium office 2 buildings and the cleared electricity price (black, right y-axis). The selection of zones to present was somewhat random; however, results from two interior (119 and 120) and two exterior zones (142 and 143) are presented in the figures.

The yellow shaded area in each plot represents the occupied period. Ideally, if the uncertainty in the airflow predictions is small, the actual zone airflow will be close to that predicted. It appears that the airflow rate prediction for the interior zones of the simulated building under summer conditions is more uncertain compared to exterior zones. Also, note that the actual airflow rate matches the predictions for both the interior and exterior zones in the real building very well. This occurs because there is not much demand flexibility during winter, and the airflow is mostly at the minimum flow rate, which is relatively easy to predict.

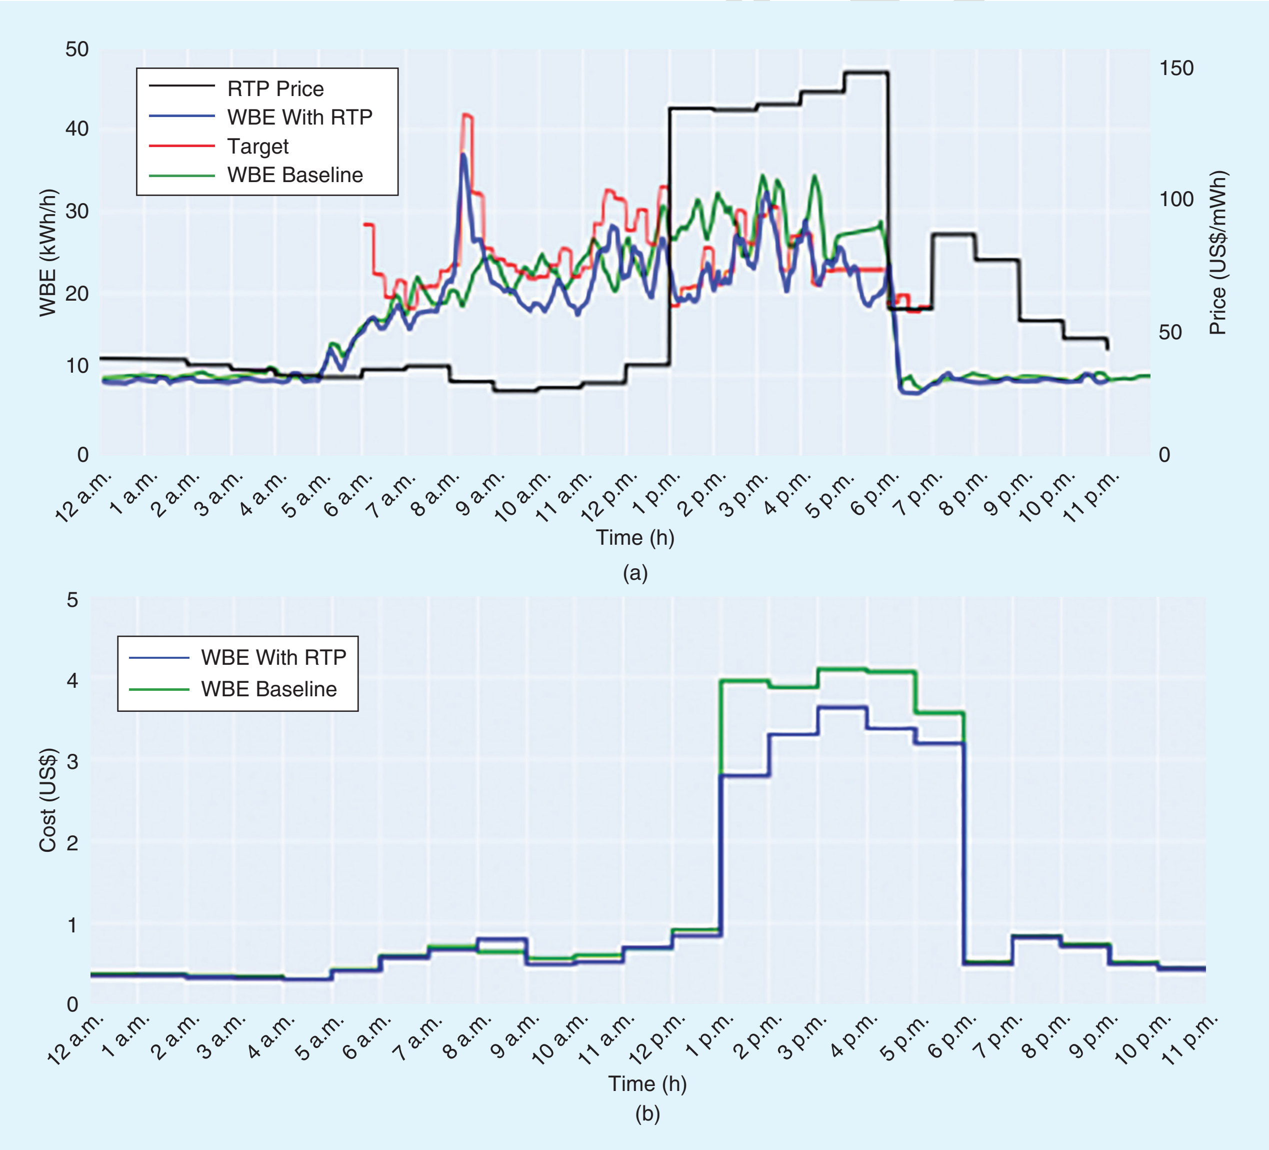

Figure 6. (a) A comparison of the 15-min rolling average electricity demand on baseline day (green: 7 October), RTP control day (blue: 9 October) and target (red), and RTP (black on the right *y*-axis) during the cooling season in building 6. (b) A comparison of the electricity cost between baseline day (green: 7 October) and RTP control day (blue: 9 October). WBE: whole-building electricity.

Figure 9(a) shows the WBE consumption without TCC (green) and with TCC (blue) as well as the cleared electricity price (black) for a summer day in the medium office building. The high-price period is between 6 a.m. and 10 p.m., and the low-price period is the rest of the day. The electricity prices are slightly higher midday than during the early morning and late afternoon because electricity transport losses are proportional to the square of electrical current, which increases near midday. The actual consumption (blue) with TCC is lower than the baseline (the operation of the building without TCC or transactive markets, i.e., business as usual) without TCC (green) when the prices are high during the daytime, which results in reduced electricity consumption when TCC is controlling the VAV boxes. Because it is a summer day, the flexibility is high in the building.

Figure 9(b) shows the actual WBE consumption (blue) and the cleared electricity price (black) for a winter day in the medium office 2 building. Because the building is heated with hot water generated from a natural gas boiler, there is negligible electric demand flexibility during winter. The other HVAC systems also have limited flexibility during heating operations because AHU supply fans operate near minimum speeds when spaces require heating.

Challenges and Lessons Learned From Using Buildings as a Grid Asset

Buildings can be a significant asset to the grid because they can support increasing the hosting capacity of distributed renewable generation and electrification of the built environment. However, to operationalize the deployment of these grid services and realize the demand flexibility in buildings, we need to overcome several challenges:

- More than 85% of commercial buildings and almost all homes do not have proper control infrastructure to

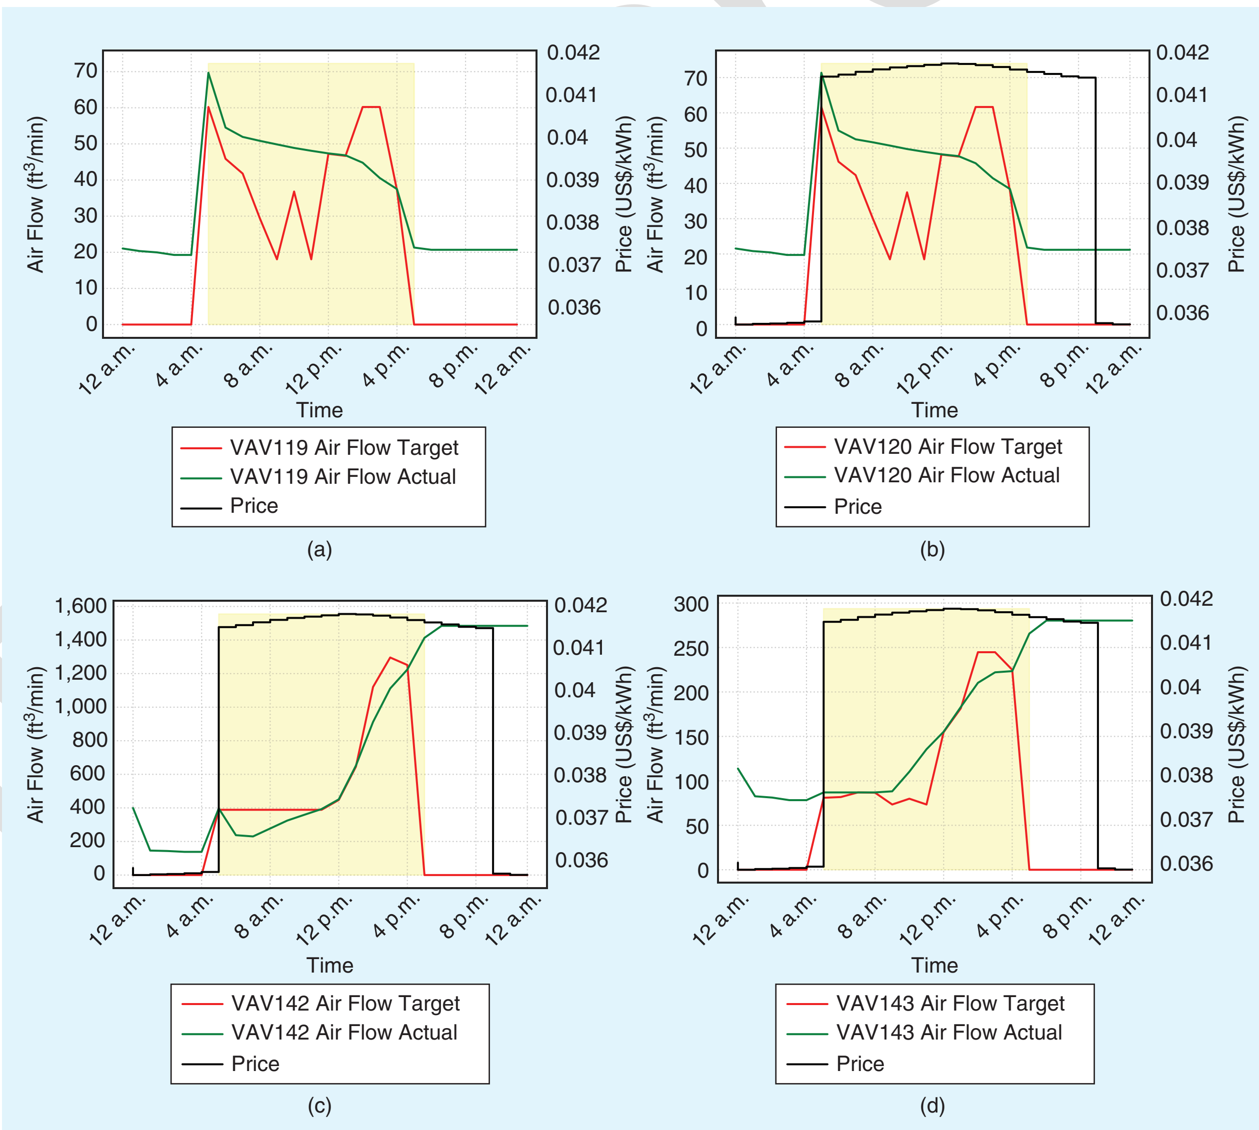

Figure 7. The actual (green) and target (red) zone airflow rates for zones (a) 119, (b) 120, (c) 142, and (d) 143 and the clearing price (black, right *y*-axis) in the simulated medium office building for one summer day.

enable building–grid integration. Therefore, there is a need for low-cost control infrastructure.

- Most demand-response programs that utilities offer are vertically integrated; i.e., they focus on controlling a large number of similar devices, for example, connected thermostats or connected electric water heaters. Many of these programs are direct load control programs with no ability for the customer to opt out. This is not sustainable in the long run.
- Most building–grid integration efforts that coordinate multiple DERs are one-off custom deployments. This approach is neither cost-effective nor scalable. To scale deployment, we need fully automated grid services applications, and we need to accelerate the deployment of these services.
- The decarbonization of electricity generation and electrification of buildings by switching fossil fuel-based heating equipment to electric HPs is a significant step to mitigate climate change. However, it could put a severe strain on the electric grid on a cold winter morning. During this period, HPs will not have enough heating capacity, so they will use backup electric heat, which could create significant winter morning peaks that need to be mitigated. The variable nature of renewable generation as well as significantly lower power generation for solar in the winter months means that additional measures must be taken to mitigate these winter peaks with very high heating demand. Applications such as ILC and TCC can mitigate these effects, but they must be deployed at scale.
- Transactive signals or dynamic rates can be used to exercise demand flexibility in buildings; however, most utilities do not offer such rates. Utilities and

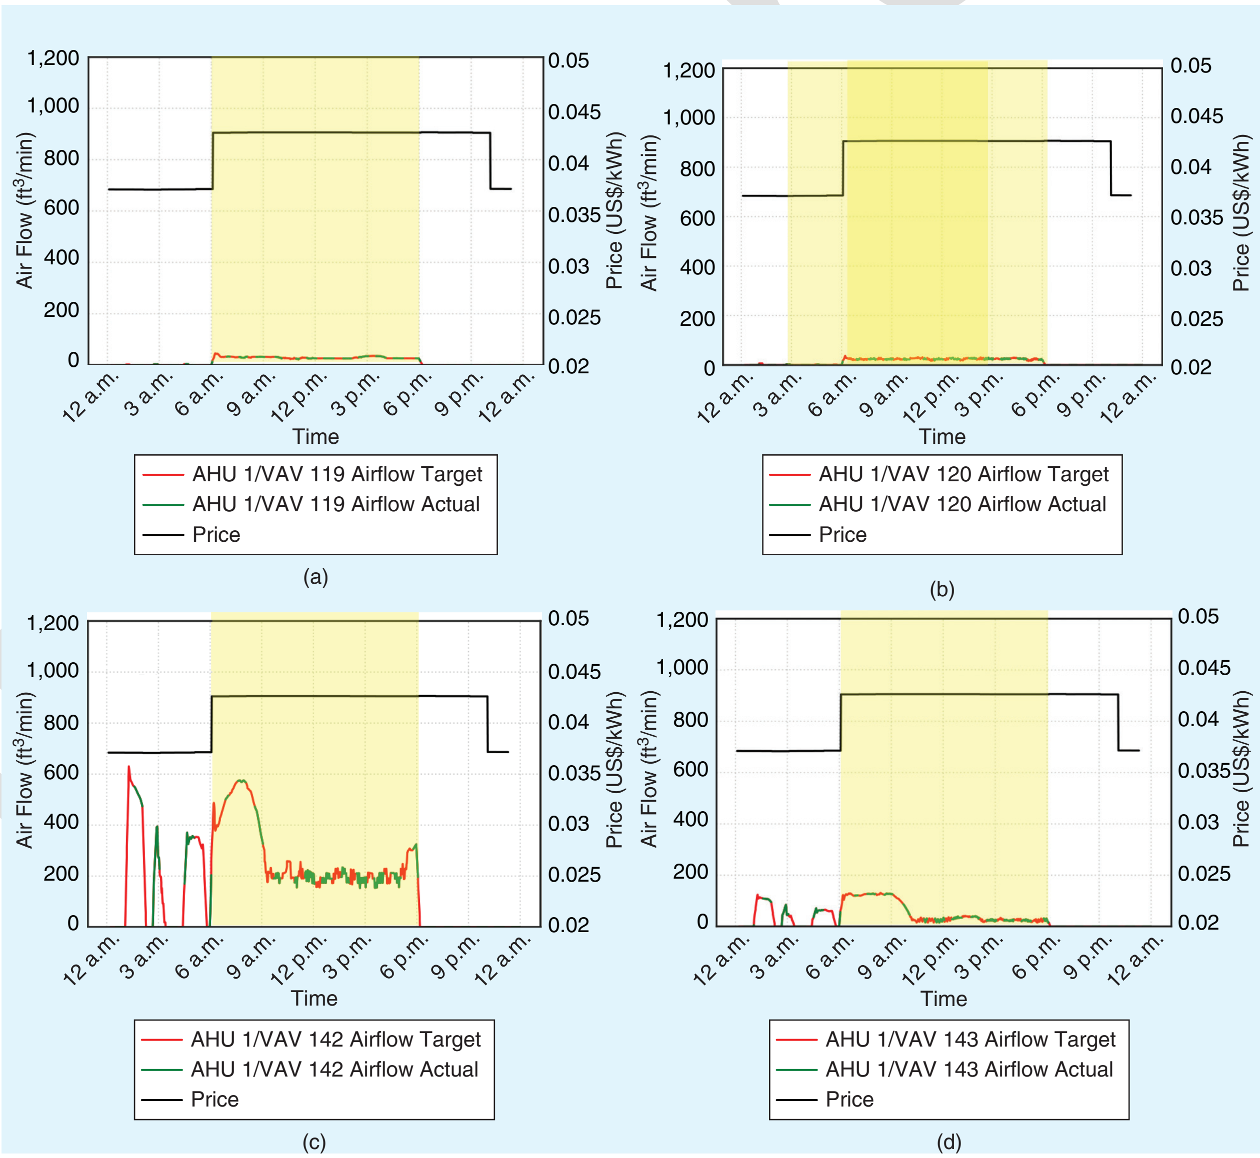

Figure 8. The actual (green) and target (red) zone airflow rates for zones (a) 119, (b) 120, (c) 142, and (d) 143 and the clearing price (black, right *y*-axis) in the real medium office 2 building for one winter day.

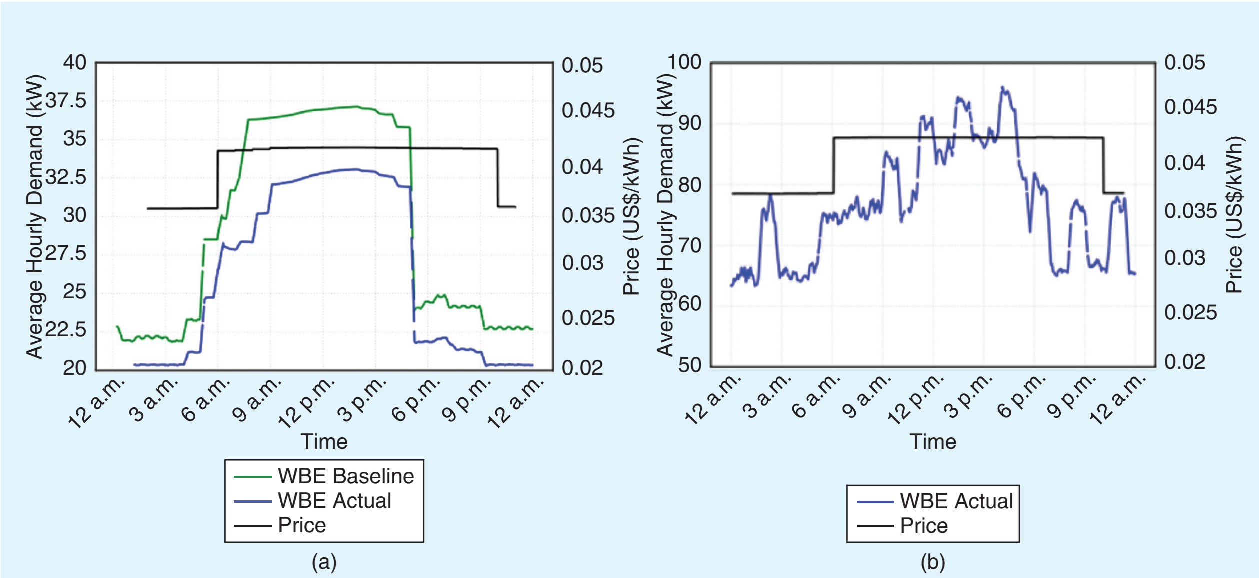

Figure 9. (a) The simulated building electricity consumption with TCC (WBE actual) and without TCC (WBD baseline) and cleared electricity price for the medium office building for one summer day. (b) The real physical building electricity consumption with TCC (WBE actual) and the cleared electricity price of one winter day for the medium office 2 building. Consumption is shown on the left *y*-axis and price on the right *y*-axis.

public service commissions must introduce dynamic rates to incentivize BTM DERs to support grid reliability.

If these challenges are overcome, which is possible, BTM DERs can be reliable resources for providing grid services. There are several lessons learned from testing and validating grid services in commercial buildings:

- Deploying grid services in large commercial buildings with BACnet-based BASs is possible. However, the effort to coordinate different BTM DERs is still labor intensive. The lack of a standard naming convention for the various sensors and control points in a BAS and lack of meta information (units, the association of the sensor with a system, etc.) make the deployment of grid service applications labor intensive. Efforts such as the American Society of Heating, Refrigeration and Air Conditioning's 223P Designation and Classification of Semantic Tags for Building Data should help.
- In small commercial buildings and homes, interoperability among various connected devices, including connected thermostats, is lacking, making it difficult to deploy a single solution that works with all devices.
- Although building–grid integration has the potential to benefit both the grid and building owners/managers, many building owners/managers are not aware of the benefits.
- As noted previously, many existing demand-response programs directly control a DER, and the customer generally does not have an easy way to opt out. To make grid services ubiquitous in buildings, the deployment must be usercentric.
- Deploying forward-looking TCC in buildings and the distribution network will result in a more optimal price of electricity. However, significant automation within a building is required to make it a scalable process.

Conclusions

Given the growing desire to mitigate climate change, utilities are increasing generation from renewable sources, and many cities and states are mandating all-electric buildings. The electrification of the building sector—which already consumes more than 75% of the total electricity generated in the United States—will increase electricity consumption. Because renewable generation is variable and not dispatchable, it will create significant supply–demand imbalance. Mitigating the imbalance using generation reserves will be more expensive and not efficient. It will be less expensive and more efficient to use BTM DERs to mitigate some of the supply–demand imbalance.

Buildings have more than 77 GWs of demand flexibility potential. As we have shown in this article, the use of grid service applications in commercial buildings will result in peak electricity reduction between 10% and 20% for 4–6 h without significantly compromising the service levels. A deeper reduction in electricity consumption is possible, but it will affect the service levels. ILC was developed to support current utility demand-response programs, such as PLM and capacity bidding, but it can also be used for TOU and CPP. TCC is forward looking and needs markets at each transactive node, including multiple commodity markets within a building.

Acknowledgment

We would like to acknowledge the Building Technologies Office of the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy, the U.S. Department of Energy's Office of Electricity, and the Washington State

Department of Commerce Clean Energy Fund for supporting this research and development effort.

For Further Reading

K. Kalsi. (2017). "Virtual batteries." Presented at the U.S. Dept. Energy Building Technol. Office Peer Rev. [Online]. Available: https://www.energy.gov/

S. Katipamula, R. G. Lutes, S. Huang, J. Lian, H. Ngo, and D. Hammerstrom, "Coordination of behind-the-meter distributed energy resources for transactive grid services: multi-building," Pacific Northwest Nat. Lab., Richland, WA, USA, PNNL-28430, 2019.

W. Kim and S. Katipamula, "Development and validation of an intelligent load control algorithm," *Energy Buildings*, vol. 135, pp. 62–73, Jan. 2017, doi: 10.1016/j.enbuild.2016.11.040.

"California Independent System Operator." [Online]. Available: http://www.caiso.com **<AU: Please provide the date of access.>**

"Annual Energy Outlook 2018," U.S. Energy Information Administration, Washington, DC, USA, Feb. 2018. [Online]. Available: https://www.eia.gov/pressroom/presentations/Capuano_02052018.pdf

Biographies

Srinivas Katipamula (srinivas.Katipamula@pnnl.gov) is with Pacific Northwest National Laboratory, Richland, WA, USA. **<AU: Please provide the postal code for the current affiliation.>**

Robert G. Lutes (robert.lutes@pnnl.gov) is with Pacific Northwest National Laboratory, Richland, WA, USA. **<AU: Please provide the postal code for the current affiliation.>**.

Sen Huang (huangs@ornl.gov) is with Pacific Northwest National Laboratory, Richland, WA, USA. **<AU: Please provide the postal code for the current affiliation.>**.

Roshan L. Kini (roshan.kini@pnnl.gov) is with Pacific Northwest National Laboratory, Richland, WA, USA. **<AU: Please provide the postal code for the current affiliation.>**.

While renewable resources will reduce emissions, they are not dispatchable, they vary temporally, and their generation is uncertain.

Buildings account for more than 75% of the total U.S. electricity use.

While ILC was designed to serve the current utility needs and serve as a bridge toward the future, TCC is more forward looking and is a natural choice in a transactive energy environment.

The ILC application uses a business decision-making process to prioritize DERs to generate a numerical score to prioritize each alternative load based on associated decision criteria.

Unlike "prices to devices," TCC is a two-way process; i.e., it includes feedback between the price and quantity.

PLM can be used to reduce the building peak electricity or when the utility system peaks, which may not be coincident with the building peak.

Buildings can be a significant asset to the grid because they can support increasing the hosting capacity of distributed renewable generation and electrification of the built environment.

Although building–grid integration has the potential to benefit both the grid and building owners/managers, many building owners/managers are not aware of the benefits.

Utilities and public service commissions must introduce dynamic rates to incentivize BTM DERs to support grid reliability.

ILC was developed to support current utility demand-response programs, such as PLM and capacity bidding, but it can also be used for TOU and CPP.