Grid Architecture for Buildings

Prepared for GMLC Grid Architecture 1.2.1

December 2019

B Nordman

Grid Architecture for Buildings

Prepared for GMLC Grid Architecture 1.2.1

B Nordman

December 2019

Abstract

Grid architecture traditionally has focused only on organizations—“grid entities”—whose only, or primary, function is to participate in the utility grid. Buildings, by contrast, only engage with the grid out of necessity, for electricity service. Traditionally, the electricity distribution infrastructure within buildings has been a simple extension of the utility grid, undifferentiated in any important way. Buildings have received electricity from the grid and paid for it, however, they have only rarely done so in ways that directly inform grid operation, other than how buildings consume energy in the aggregate. Until recently, buildings did not commonly receive direct control signals from the grid, and they have rarely had local generation, local storage, or the ability to operate independently of the utility grid.

With all this, buildings were particularly uninteresting from the perspective of grid architecture, and so received little attention. They lacked coordination with the utility grid and no other organization was involved in dynamic decision-making for building operation. Due to these and other reasons, grid architecture diagrams generally show no internal structure to buildings. However, this overall situation is beginning to change, with the growth of Distributed Energy Resources (DER) and the ability to use them for grid benefit. With technology advancement and economics favorable to DER, the pace of change will only increase. Public goals for integration of more renewable power, and concerns for cybersecurity, privacy, and resiliency, all suggest that more attention is needed to grid architecture from the buildings perspective.

The core issues for Grid Architecture for Buildings largely fall into two distinct areas:

- How buildings relate to the utility grid
- How buildings are organized internally

The combination of these elements determines how DER are coordinated with the grid. This paper does not address any relationships between grid entities. The utility grid may be organized internally differently across locations or across time at the same location, but so long as the interface at the meter does not change, the building need not be aware of or concerned with this diversity.

This paper describes the usual grid architecture for buildings today and assess possible future options. It does not aim to cover architectural decisions for the larger grid, but some implications are unavoidable.

The paper covers key terms and concepts, reviews current building / grid relationships, the structure and operation of price-based grid management (one alternative architecture), potential future building/grid relationships, and implications for grid architecture generally for the buildings perspective.

Summary

Grid architecture traditionally has focused only on organizations—"grid entities"—whose only, or primary, function is to participate in the utility grid. Buildings, by contrast, only engage with the grid out of necessity, for electricity service. Traditionally, the electricity distribution infrastructure within buildings has been a simple extension of the utility grid, undifferentiated in any important way. Buildings have received electricity from the grid and paid for it, however, they have only rarely done so in ways that directly inform grid operation, other than how buildings consume energy in the aggregate. Until recently, buildings did not commonly receive direct control signals from the grid, and they have rarely had local generation, local storage, or the ability to operate independently of the utility grid.

With all this, buildings were particularly uninteresting from the perspective of grid architecture, and so received little attention. They lacked coordination with the utility grid and no other organization was involved in dynamic decision-making for building operation. Due to these and other reasons, grid architecture diagrams generally show no internal structure to buildings. However, this overall situation is beginning to change, with the growth of Distributed Energy Resources (DER) and the ability to use them for grid benefit. With technology advancement and economics favorable to DER, the pace of change will only increase. Public goals for integration of more renewable power, and concerns for cybersecurity, privacy, and resiliency, all suggest that more attention is needed to grid architecture from the buildings perspective.

The core issues for Grid Architecture for Buildings largely fall into two distinct areas:

- How buildings relate to the utility grid
- How buildings are organized internally

The combination of these elements determines how DER are coordinated with the grid. This paper does not address any relationships between grid entities. The utility grid may be organized internally differently across locations or across time at the same location, but so long as the interface at the meter does not change, the building need not be aware of or concerned with this diversity.

This paper describes the usual grid architecture for buildings today and assess possible future options. It does not aim to cover architectural decisions for the larger grid, but some implications are unavoidable.

The paper covers key terms and concepts, reviews current building / grid relationships, the structure and operation of price-based grid management (one alternative architecture), potential future building/grid relationships, and implications for grid architecture generally for the buildings perspective.

Acknowledgments

This paper was not possible without the interest, encouragement, and support of the GMLC 1.2.1 project leader, Jeff Taft, whose creation of the discipline of Grid Architecture was a fundamental inspiration for this work. Also thanks to the other staff who supported the project, the other Grid Architecture team members, and the DOE staff who supervised the project.

Acronyms and Abbreviations

AC	Alternating Current
BEMS	Building Energy Management System
CPP	Critical Peak Price
DC	Direct Current
DER	Distributed Energy Resource
DERMS	Distributed Energy Resource Management System
DLC	Direct Load Control
ESI	Energy Services Interface
ESG	Energy Services Gateway
EV	Electric Vehicle
GMLC	Grid Modernization Laboratory Consortium
HVAC	Heating, Ventilation, and Air Conditioning
IP	Internet Protocol
kWh	kilowatt-hour
LPD	Local Power Distribution
NPI	Network Power Integration
OSI	Open Systems Interconnection
PSI	Power Systems Interface
PV	Photovoltaic
TOU	Time Of Use

Contents

Figures

Use F9 to update this List of Figures.

Tables

Use F9 to Update this List of Tables

1.0 Introduction

Grid architecture traditionally has focused only on organizations—"grid entities"—whose only, or primary, function is to participate in the utility grid. Buildings, by contrast, only engage with the grid out of necessity, for electricity service. Traditionally, the electricity distribution infrastructure within buildings[1] has been a simple extension of the utility grid, undifferentiated in any important way.[2] Buildings have received electricity from the grid and paid for it, however, they have only rarely done so in ways that directly inform grid operation, other than how buildings consume energy in the aggregate. Until recently, buildings did not commonly receive direct control signals from the grid, and they have rarely had local generation, local storage, or the ability to operate independently of the utility grid.

With all this, buildings were particularly uninteresting from the perspective of grid architecture, and so received little attention. They lacked coordination with the utility grid and no other organization was involved in dynamic decision-making for building operation. Due to these and other reasons, grid architecture diagrams generally show no internal structure to buildings. However, this overall situation is beginning to change, with the growth of Distributed Energy Resources (DER)[3] and the ability to use them for grid benefit. With technology advancement and economics favorable to DER, the pace of change will only increase. Public goals for integration of more renewable power, and concerns for cybersecurity, privacy, and resiliency, all suggest that more attention is needed to grid architecture *from the buildings perspective.*

In the past, buildings only consumed energy, so that they bore no resemblance to grid entities. However, the ability to add generation to buildings is becoming ever easier (less costly and less burdensome), and an increasing number of DER in buildings are participating in wholesale markets. The primary purpose of grid entities (e.g., bulk generation or bulk storage) is to participate in the grid.[4] Buildings, by contrast, have other primary purposes (e.g., housing, office space, etc.). Thus, they can always be readily distinguished.

The core issues for Grid Architecture for Buildings largely fall into two distinct areas:

- How buildings relate to the utility grid
- How buildings are organized internally

The combination of these elements determines how DER are coordinated with the grid. This paper does *not* address any relationships between grid entities. The utility grid may be organized internally differently across locations or across time at the same location, but so long as the interface at the meter does not change, the building need not be aware of or concerned with this diversity.

---

[1] The terms "facility," "building," "premise," and "customer" are all treated as equivalent for this discussion. While most buildings are residential or commercial, in this paper the term also includes isolated vehicle charging, public infrastructure (e.g., street lights), and industrial facilities. The term building is used as that is the term usually used by the occupants (end-users) of buildings. The GMLC project 1.2.2 on Interoperability, in its Strategic Vision document (Widergren et al., 2018), uses the term "DER Facility".

[2] There are of course exceptions, but they are at the margin and have not affected utility grids particularly. These include diesel backup generators, and uninterruptible power supplies for critical loads (e.g., communications).

[3] DER in this paper includes any device in a building that can coordinate with the utility grid including generation, storage, and flexible loads. It does not include any devices on the utility side of the meter.

[4] Even if some buildings are net producers of electricity, the value of that production is tiny in comparison to the building's primary value.

This paper only covers *energy*[1]—value that can be assessed by any utility meter that tracks energy flows on a time periodic basis. Grid services related to *power* and *power quality* (e.g., inverter control and reactive power) are substantially different from those concerned with energy, and these are not directly addressed in this paper.

This paper aims to describe the usual grid architecture for buildings today and assess possible future options. It does *not* aim to cover architectural decisions for the larger grid, but some implications are unavoidable. In grid architecture an *architectural view* is "A view of a System from the standpoint of a particular stakeholder" (from *Grid Architecture* glossary [REF]). As building owners and operators are stakeholders, this paper is an architectural view. It also presents one architectural approach in detail: a price-centric one. The paper is intended to be fair, but *not* neutral. It begins by viewing the building as central, and is based on a set of system architectural principles. Other architectures are explored in contrast to this base case. The recommended principles and architecture also take inspiration from the organizing principles of Internet architecture.

1.1 Grid Architecture Context

Figure 1 shows a good, representative architectural view from *Grid Architecture* (Taft, 2015). The range of lines shown in such diagrams vary modestly from figure to figure. Figure 1 includes power flows (green), grid services (black), control signals (red), retailing (light green), and regulation (blue, purple). Note that buildings occupy only a small portion of the figure, at the very bottom, reflecting their lack of current complexity with respect to utility grid interaction and internally for power distribution.

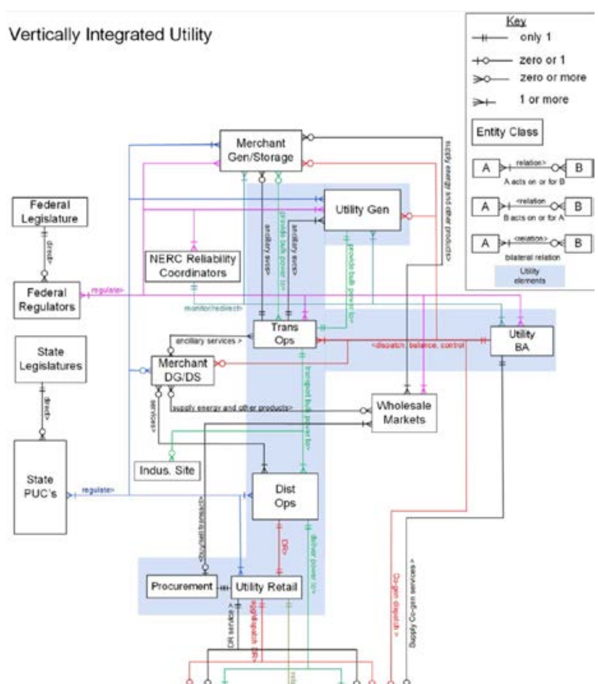

Figure 1. Vertically Integrated Utility Industry Structure (Taft, 2015; Figure 4.5)

Focusing on buildings specifically (Figure 2), the only difference between residential and commercial/industrial buildings is for the latter to have potential cogeneration coordinated with the utility.

[1] It actually only covers electricity but some of the principles, particularly about metering and pricing, also apply so other forms of energy, including intra-building distribution of hot water, chilled water, or steam.

This is primarily applicable to commercial buildings today but in future could be common in residential buildings as well. More importantly, photovoltaic panels (PV) plus storage is essentially equivalent to cogeneration for grid purposes. For grid architecture, it is preferable to not need to distinguish among building types. Thus, the relationship between buildings and the utility grid can be summarized as:

- Electricity flowing to (and sometimes from) the building.
- The utility retailing electricity to the building (metering and billing).
- The building providing grid services to the utility.
- The utility (and/or a third party) coordinating those grid services.

Figure 1 also shows the regulatory context that allows the grid and grid entities to exist, establishes minimum behaviors, and bounds their activities. This influences the building/grid relationship and what buildings are allowed to do. However, the regulatory context does not actively influence the architectural analysis.

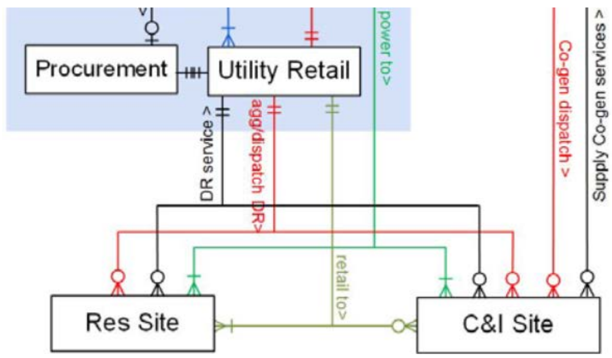

Figure 2. Vertically Integrated Utility Industry Structure (Taft, 2015; Figure 4.5 excerpt)

Building/grid relationships are defined by:

- The tariff the building faces for buying electricity from (or selling it to) the grid,
- The grid services delivered by the building to the grid, and
- Any financial and control arrangements in addition to the tariff.

This is essentially the same scope as covered by the Energy Services Interface (see below).

While grid architecture diagrams usually distinguish among many types of grid entities, this paper is concerned with grid entities only for their characteristics that affect buildings. This simplifies the issue since the same signal may be sent to buildings by different grid entities depending on the location. Therefore, the type of grid entity is immaterial; only the content of the information and its financial implications are important. For other grid architecture purposes, the differences between grid entities can be central.

1.2 Building Internal Grid Architecture

For over a century we have had a "unitary grid" in which most buildings did not have a distinct electrical identity separate from the utility grid to which they were connected. The reliability and quality of the power in the building and on the grid were nearly the same, and excepting any capacity constraints in the building (rarely an issue in the U.S.), the supply/demand balance in the building was the same as on the grid. In this context, the idea of grid architecture within a building makes no sense.

The defining characteristic of a microgrid is its ability to island from a utility grid. If the microgrid is a single customer, then it is not interacting with the grid during these times, so the fact that a building has microgrid capability is not relevant to grid architecture. Similarly, if the microgrid is a section of the utility grid, when it islands, the building is not affected, so the microgrid is still not relevant. Thus, while microgrids are reviewed in Section 6 of this paper, they ultimately do not affect building grid architecture, except to the degree that a utility grid and microgrid might voluntarily seek to separate even when both are operating.

Most modern microgrids also have sophisticated management of DER. As this can be deployed in buildings without the ability to island, this capability is distinct from the building being a microgrid.

In general, grid architecture is concerned with relationships among system domains (e.g., generation, transmission, and distribution) and with the economic or direct control signals being sent between grid entities to coordinate operation. Thus, one element of grid architecture in buildings is concerned with disaggregating the building into multiple electrical domains (if there is more than the utility grid domain) and into entities that can communicate for purposes of affecting local generation, local storage, or end-use device consumption.

1.3 Architecture Alternatives

Many approaches to exploring Grid Architecture for Buildings are possible. The field's openness has generated many diverse proposals, and the highly fragmented nature of electric utilities has meant that many paths have already been taken or at least explored. However, the alternatives vary widely in many respects, including the cost and complexity of creating and deploying the necessary technology, the effectiveness of the proposed solution, and interactions with related topics of privacy, cybersecurity, and reliability.

This study's goal was to explore the field of grid architecture for buildings in more depth, to identify a unified approach. A first step was to identify the most promising directions and clearly describe their assumptions and principles of operation. These were then evaluated and compared for important distinctions. There are many benefits of identifying a single approach for the country, and globally, such as easing production of end-use products and reducing burdens on installers and building owners. Continued diversity of coordination approaches is a great impediment to using DER most effectively.

The paper is organized as follows: Section 2 outlines several key terms and concepts; Section 3 reviews current building / grid relationships; Section 4 presents the structure and operation of price-based grid management (one alternative architecture); Section 5 considers potential future building/grid relationships; and Section 6 lays out implications for grid architecture generally for the buildings perspective.

2.0 Key Concepts

At the building/grid interface, a number of concepts are intertwined; some longstanding, some relatively new, and some only emerging. These include the following:

- Utility Meter: records energy flows and so can account for some grid services[1]
- Energy Services Interface (ESI): the communication path through which many (or all) energy coordination signals pass
- Grid Service: what the utility grid can request or encourage
- Distributed Energy Resource (DER): devices in buildings that can coordinate with grid entities for energy
- Distributed Energy Resource Management System (DERMS): the software that determines how the utility grid should orchestrate interactions with DER (through the ESI or otherwise).

These and related terms are not always used with the same meaning in different documents. The following discussions outline the current status of each of these concepts and where they may be headed.

2.1 Utility Meter

Grid architecture diagrams typically omit the meter, because it has generally played no important role in system structure. In this paper's architectural view (see Figure 6), the meter at a building (customer) site is explicitly represented, to clarify what mechanisms are accomplished using the meter and what are done separately. Although control signals in the form of time-varying pricing may not literally flow through the meter, they may be shown doing so in some diagrams to reflect that the building's response to the signal is assessed by the utility through the total consumption communicated by the meter. Other control signals, such as those for traditional command/control demand response (DR), are assessed through different means.

The range of grid services that a meter can cover depends on the time interval of the measurement reports. Common intervals are one or more months, 15 minutes, and 5 minutes. The smaller the interval, the more responsive a building can be. However, with more storage being deployed to the grid, the value of such smaller periods is dropping.

2.2 Energy Services Interface (ESI)

The concept of an Energy Services Interface (ESI) dates back to at least 2011, and likely earlier. It was then described as follows:

> "An ESI is a bi-directional logical interface that supports the secure communication of information between entities inside and entities outside of a customer boundary to facilitate various energy interactions between electrical loads, storage and generation within customer facilities and external entities." (Hardin et al., 2011)

This definition is general enough to cover a variety of architectures for delivering grid services and the control mechanisms needed to support them. Some key questions are:

- Is there a single ESI for each building, or can individual DER or groups of DER have separate ESIs?
- If there are multiple grid entities (e.g., the utility grid and one or more third-parties) involved in coordinating DER, do they all use a single ESI or each have their own?

---

[1] This assumes an interval meter with time granularity of hourly or less. It is likely that five minutes may be the smallest time granularity needed.

- Might there be different ESIs for different grid services?
- Must an ESI be on-site or can it be implemented elsewhere (e.g., in the Internet cloud)?
- Does the ESI expose the individual existence, identity, and status of DER?
- Does the ESI cover all possible grid services?

The ESI serves the role for communications that the meter does for power—being the point of demarcation between the utility grid infrastructure and the building infrastructure. More formally, the ESI is "a logical demarcation point at a facility asset ownership or operations support boundary" (Holmberg et al., 2014). As an ESI is an interface, its building side must be hosted by some device.

If a building has an ESI for each DER then the ESI function can be hosted by each DER. Otherwise there needs to be some central entity in the building to perform the function of receiving data through the ESI and relaying it to each DER; in this document that entity is called the "ESI Host."[1] Today, ESI hosts exist in individual DER, centrally in buildings, and in the Internet cloud. Since a key function of the ESI Host is to do protocol translation, it might be usefully referred to as the Energy Services Gateway (ESG).

Grid service information received by the ESI Host may be relayed unchanged to DER, relayed with modification, or implemented as functional control signals sent to devices (e.g., turning devices on or off or changing operational parameters). Usually the ESI host will be a device with some other primary purpose (perhaps several). In larger buildings, a Building Energy Management System (BEMS) could take on this role, to use price information in its own decision-making, as well as pass prices on to other devices. In smaller buildings, a device that is well-networked and always on is a suitable choice, e.g., a Wi-Fi router.

In some cases the device hosting the ESI must obtain data from the meter (as when there are differential buy/sell rates at the meter, to know the marginal value of power at a particular point in time). The meter may need to get price signals, and could get these from the ESI host device.[2]

2.3 Grid Services

This paper is focused on the quantity of net electricity use (or production for buildings that can export power to the grid), which is what utility meters measure. Grid services that cannot be assessed (measured or otherwise validated) through ordinary interval meter data are out of scope. Lists of grid services are many and diverse. Most of them combine services for energy and power, which is unhelpful. Some grid services exist in only a limited number of utility markets, and some listings include services not in use anywhere. Naming among lists is often inconsistent, and definitions are scant.

Few of the lists clarify the relative magnitude of the need for such services currently, or likely in future. For example, the frequency regulation market is orders of magnitude smaller than the value of energy that is usefully shifted in time, yet these two services are usually listed as being equivalent on these lists. All frequency regulation needs can be satisfied by a small fraction of buildings (or by grid entities alone), but nearly all buildings can usefully contribute to shifting load.

---

[1] The GMLC project 1.2.2 on Interoperability, in its Strategic Vision document (Widergren et al., 2018), uses the term "DER Facility Management Function" for the overall entity that includes DERs internally.

[2] Conventionally, utility prices change less often than the meter reports energy consumption values, or in a few cases, as often. If prices changed more frequently than the meter reports, this could be accommodated by having the meter internally integrate and accumulate price and quantity with the price changing at any rate or no fixed rate, and then report the total value (in $) consumed at a less granular level.

2.3.1 Power Grid Services

Some grid services cannot be measured by ordinary interval utility meters but can still be usefully provided by buildings directly to a utility, or indirectly through an aggregator. Some grid service capabilities can be required of buildings by law or as a condition of connecting to the grid (as with inverter capabilities). These other grid services can be considered completely separately from control concerned with the aggregate quantity of electricity over time. Provision of some of these ancillary services generally requires exposing details of building components and their operation to the grid in ways not required for managing energy flows.

With this in mind, it is useful to separate grid services from buildings into two groups for purposes of grid architecture (and in general): (a) that portion concerned with aggregate energy demand (as the meter can assess), and (b) the portion concerned with all other grid services. In this paper these two are referenced as "energy grid services" and "power grid services," respectively. Each domain will have distinct flows of financial relationship and control and technology infrastructures, along with a different arrangement of these within the building. As increasing numbers of buildings and DER can provide power-related grid services, the price they can obtain from the utility is almost certain to drop, and many buildings may choose to provide only what is required. The total dollar value of these services is much smaller than the value of shifts in energy use, so there is a low limit on the value that buildings as a whole can tap into.

2.3.2 Energy Grid Services

Some grid services can be measured by ordinary interval utility meters. The simplest and most common way to deliver energy services is through a time-varying rate. Today this occurs mostly with Critical Peak Prices (CPP) and Time-Of-Use (TOU) rates.[1] The amount and type of energy grid services that can be delivered through prices alone is a function of how prices vary (in time granularity and magnitude) and meter capability. Some residential customers have prices that don't change at all over the year; for others, the prices is different every hour of every day. When electricity retailers can change prices in real time, they can use the price to increase, decrease, or shift demand, in ways beneficial to the grid. Customers simply change their consumption in response to the price changes for their own benefit, and with the correct price at the meter, the interests of the grid and the building are aligned.

When energy is coordinated through a mechanism other than price at the meter, then inevitably some of the electrons in the building are valued differently from other electrons being consumed at the same time. This also requires a second mechanism to track energy flows, to distinguish between these two groups of electrons. This can be done through submeters or through statistical analysis of meter data, or an aggregator-specific verification process.

With electrons all being the same, the idea that they should be valued differently at the same place and time is peculiar. The box below illustrates how this condition can be created with traditional demand response.

[1] Demand charges exist for some customers, mainly commercial and industrial, and while they do encourage some shifting of load, they do so with respect to the building peak. This often does not match the utility peak, and even that is often not the relevant concern. It is the relative balance of supply and demand, and costs, which should determine shifts. Thus, demand charges are not considered part of demand response generally or in this paper.

How much is a kilowatt-hour worth?

Consider a building that pays for electricity according to TOU rates of $0.15 cents/kilowatt-hour (kWh) noon–5 pm, $0.25 from 5 pm–9 pm, and $0.10 after 9 pm. It also is part of a traditional demand response program that is called up to 10 times per year, which requires reducing consumption by 25 kW for four hours and results in a payment to the customer of $400/year. This comes to $0.40/kWh for the energy that is shifted.

An event is called at 3 pm, and the four hours of 25 kW of load is shifted to start instead at 7 pm. The event time period spans two TOU rate periods, and the replacement time period (7–11 pm) also spans two TOU periods. For each time period, there is the rate charged at the meter, and the premium given for participating in the DR program ($0.40/kWh), which summed gives the initial benefit to the customer of the DR reduction. However, there is a cost to consume power later to deliver the service that was missed. This replacement cost reduces the value to the customer, to give a final net value for each period of time.

Table 1 shows important values for the two halves of the DR event period. At these times, the same electrons in the same place are involved, but they are valued at more than one price. The electrons that are moved from the DR time are valued much higher than the rest (the Net DR value in contrast to the Basic rate).

Table 1. Electricity Values for Energy for a Customer Participating in a Traditional DR Program Event

Time period	Basic rate	Gross DR value	Replacement cost	Net DR value
3–5 pm	0.15	0.55	0.25	0.30
5–7 pm	0.25	0.65	0.10	0.55

The net DR value is what the customer receives and what the utility pays, since other customers will need to cover those costs. In this example, the average cost to accomplish the shift is $0.425/kWh, which is much more than the TOU rate differential.

An alternative is to value all electrons the same price at any place and time. By bringing in all buildings and all devices, a much smaller price increase could likely accomplish all the load shifting, and so be less costly for both the utility and the average customer. The higher price in the event period could be the only change, or prices could be depressed at later (or earlier) times to shape where the shifted energy is instead consumed (This example is for a typical thermal load, like a water heater, in which there is no significant reduction in net overall consumption; for effects as for dimming lighting, the energy is not consumed later so the value to the customer is the gross value listed in the table).

2.3.3 Distributed Energy Resources (DER)

Distributed energy resources generally are defined to include flexible loads, battery storage, local generation, and vehicle charging (or discharging). The term in general includes equipment in buildings and in the distribution grid. In this paper it only references those within buildings—on the customer side of the meter. A key distinction between loads and vehicle charging versus storage and generation, is that the former deliver building services to people, and the latter do not. Loads that cannot coordinate with the grid (even manually) are not DER.

2.3.4 Distributed Energy Resource Management System (DERMS)

A distributed energy resource management system is generally defined as a system that knows the identity, location, status, and control capabilities of all DER in a utility service area (or a portion of it). The DERMS sends commands to these DER to meet utility needs to keep the grid within capacity limits (generation, transmission, or distribution) or for economic benefit to the utility (and consequently customers).

For power services, this concept makes sense. However, for energy services, if all (or most) DER management is done through pricing, then the role of the DERMS system shifts dramatically to doing fine-grained management of prices and forecasts, including locationally. In this case, the DERMS does not need to know the identity, location, or status of any DER, but infers their existence and performance based on measured observations from utility meters at customer sites and elsewhere. These observations show how price changes affected consumption in previous times that are similar—in time of day, time of week, and weather (which affects both production and consumption). For a price-based system, if a price change does not have the expected energy effect (or too much effect), the price can simply be changed again.

Figure 3 below shows a recent snapshot of the outlook page of the California Independent System Operator. The y-axis does not start at zero, so the size of the differences between actual consumption and forecast are smaller than they appear. The day-ahead forecast is spot-on for many hours; the worst hour is just over 5 percent off. The hour-ahead values are much closer. The ISO is very capable of understanding consumption (and in this case unmetered in-building PV production), so it should be similarly capable in understanding the effect of price changes on consumption. Prices are more useful to customers if they include a forecast, and as the ISO can forecast consumption, it should be similarly able to forecast prices (these are the retail prices, which do not have to correspond exactly to wholesale prices, though being correlated is helpful).

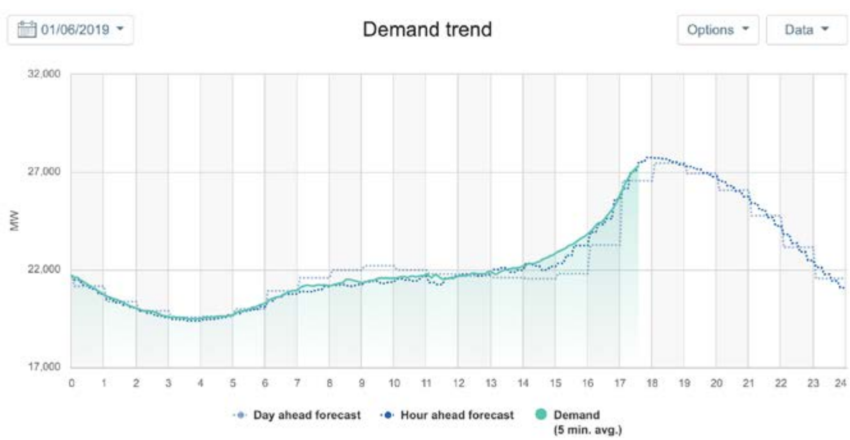

Figure 3. A snapshot of the California Independent System Operator Total Demand and Forecasts

3.0 Current Building/Grid Relationships

Relationships between buildings (utility grid customers) and the utility grid cover the types of relationships in grid architecture generally: electricity flows, financial relationships, and control signals. Control signals can be *direct*, as with commands to an individual DER device (or the building as a whole), or can be *indirect*, as with dynamic changes in price. Different grid services can have different control mechanisms.

3.1 Conventional Grid

Traditionally buildings have only had one relationship with the grid—with their electric utility—and that relationship has been mediated by the meter. Because rates have been largely fixed (acknowledging limited use of demand charges and time-of-use rates), they provided no meaningful control signal to coordinate building energy use with grid goals. Generally there is no direct communication between the building and the grid.

Figures 1 and 2 (earlier in the report) show this relationship in the format used in the *Grid Architecture* reports (Taft et al., 2015; Taft, 2016). Electricity flows into the building from the grid. Information flows from the meter to the grid. Money flows from the building to the grid.

Figure 4 shows a traditional Smart Grid view of the relationship; buildings are just like any other part of the grid and communicate directly with four other domains of the utility grid. There are no special characteristics that might identify the building as fundamentally different from other grid entities. This diagram is not reflective of current communication flows with buildings in general, but rather is aspirational to potential communications.

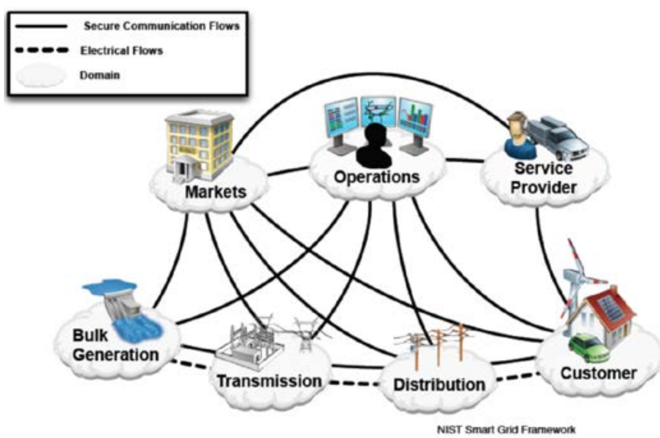

Figure 4. An Architectural View of a Traditional Building Relationship to the Utility Grid (NIST, 2012, Figure 3-1. Interaction of Actors in Different Smart Grid Domains through Secure Communication).

3.2 Traditional Demand Response

In the United States it has been difficult and slow to evolve the rate offerings made to customers, so that even as utilities recognize real opportunities for new coordination between the grid and buildings, they have found it difficult to use the meter relationship to mediate these or to engage and enlist customers to participate.[1] Thus, other grid entities (e.g., aggregators) have been enlisted to work around the meter to accomplish that coordination.[2] Grid operators have often recognized that the marginal value of power at a peak time is many times the price charged at the meter and so found it well worthwhile to pay customers to reduce or defer consumption for economic advantage, to ensure that available supply can even meet demand, or address capacity constraints in the transmission and/or distribution systems. In addition, before digital smart meters, it was much more difficult and expensive to offer time-varying prices than it is today.

One inherent result of this situation is that different electricity use in the same building at the same time is valued differently. In this case, the basic financial relationship with the utility is unchanged, but there is a second (or multiple) set of financial and control relationships with these other entities (or the utility if it is the aggregator).

Two primary mechanisms have been utilized in Traditional DR: Direct Load Control (DLC) and Event-based DR.

In the DLC model, the aggregator (and by extension the utility) needs to know the internal equipment characteristics of the building to know what agreement could be made. Some mechanism for verification of control operation needs to be in place for the utility to be confident that the DR action is actually occurring. In many cases the verification of action being taken is to statistically differentiate between "normal" days and "event" days. With multiple grid services layered on top of each other, with multiple DER in the same building on different programs and the possibility of wanting to shape electricity use on *all* days of the year rather than just a few, the feasibility of continuing to take this "baseline" approach is questionable. Common examples of DLC include cycling of air conditioners or water heaters. With event-based DR, a prior agreement is made that on receipt of a signal from the grid, a building will reduce its power level by a specified amount for a specified duration.

Figure 5 shows this as an architectural view. Electricity can flow in either direction between the utility and grid (though usually net to the building, at least on an annual basis). Demand response signals flow directly to individual devices or systems in the building, and a payment flows to the building distinct from that for the utility bill from the meter. A convention used in this paper for such views is that when a box in the figure is open, that means that the rest of the system architecture is not described. This is done because the rest of the system is not the focus of this discussion, and this relationship could apply to many different grid architectures. If the utility implements the aggregation function, then the two upper boxes are merged.

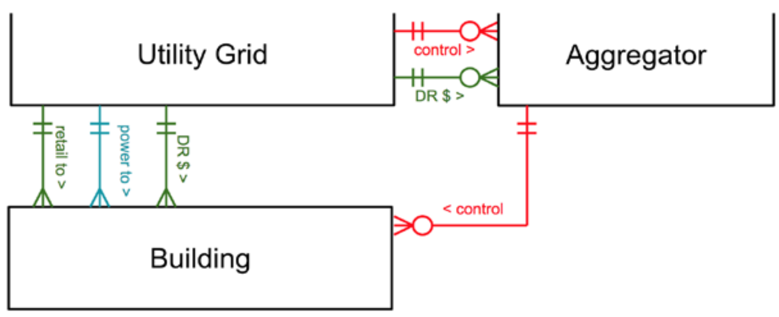

Figure 5. An Architectural View of a Building's Relationship to the Utility Grid with DR Aggregators

[1] Critical Peak Pricing (CPP) and Time-Of-Use (TOU) rates are an exception to this.
[2] Utilities sometimes do the aggregation function themselves.

3.3 Time-Varying Prices

A third common approach to the building/grid relationship is to use only prices as the principal coordination mechanism. Today this is most commonly CPP or TOU prices, though there are utilities that offer hourly prices that are different each hour of the year (more on this in the next section). Decisions about how to change prices are informed by general knowledge of building types and by past experience with the relation between prices, price changes, and building consumption; they can also be based on an existing price stream, such as wholesale spot prices.

When only prices are used, the utility does not need to know anything about what equipment is in the building, or its state of operation. In addition, only the existing meter (assuming it is already a smart meter) is needed to track the value of energy consumed and energy shed or shifted.

Figure 6 shows an architectural view for the distributed approach. An arrow is shown below the meter to reflect that prices are communicated to the building as a whole, and the same prices used to calculate the bill. The building box is shown as open on the bottom, as the building internal architecture is not shown.

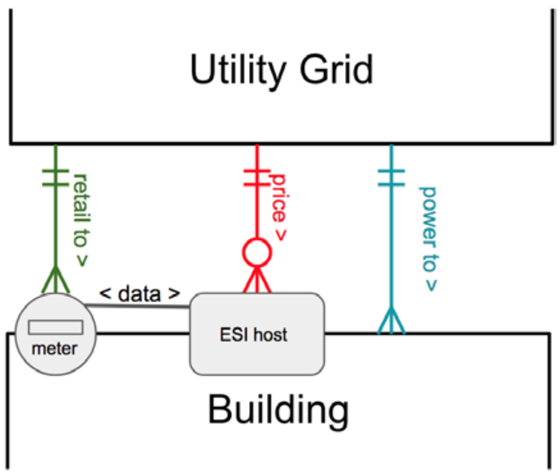

Figure 6. An Architectural View of a Building's Relationship to the Utility Grid with Distributed Control

4.0 Price-Based Grid Management

This section illustrates how grid architecture for buildings can be applied by describing one approach in detail. It would be helpful if several other approaches were also described in similar detail, to clarify their assumptions. Such descriptions are best crafted by people who have been deeply involved in their development and testing.

Most discussions of how to coordinate buildings and the DER within them are based on unstated assumptions about what is possible and desirable, how effective or efficient different systems might be, and what coordination is needed and where it should be provided. Usually it is difficult to determine what is an assumption brought into the analysis and what is a conclusion of it, as the policy and technology choices in scope are mutually determinative. The following description of the assumptions about this topic are made by the author, and they frame the evaluation of alternative schemes that follow. The price-based perspective derives from a buildings-centric view of electricity systems.

While pricing has always been part of utility grids, the concept described here is derived much more from Internet technology; in particular, not tightly constraining how future systems are structured to their traditional management model. To make large leaps forward, it is generally required to make great changes in technology and system architecture. That said, it is also important to understand how the system could plausibly evolve from its present state to the (much) better future one.

This approach intends to support a variety of features, including the following:

- Enabling local generation and storage to be "plug-and-play," as IT devices are
- Enabling those devices to automatically perform efficiently and effectively
- Greatly reducing the cost of having single-customer microgrids; including, for example, by having the control mechanism used when operating as a microgrid identical to the mechanism used when grid-connected
- Minimizing risks to privacy and security by minimizing the amount of information that must be exchanged between buildings and the utility grid, and by minimizing the number of entities on each side required to interact
- Taking into account conversion losses, including AC/DC conversions, to minimize these

The section begins with design principles drawn from Internet architecture; reviews how IT networks and electricity distribution systems are similar and different; presents a detailed description of the elements of a price-centric model; and finishes with some discussion.

4.1 Design Principles

A good discussion of design principles underlying Internet technology can be found in "Architectural Principles of the Internet" (Carpenter, 1996). The following are selected principles that seem applicable to our future electricity systems, presented in order of appearance; quoted text from the document is indented. There are important differences between communications and electricity, so adaptation of technology and principles must be done carefully.

<u>Constant Change</u>

> Principles that seemed inviolable a few years ago are deprecated today. Principles that seem sacred today will be deprecated tomorrow. The principle of constant change is perhaps the only principle of the Internet that should survive indefinitely.

Internet architecture is an empirical result of considering various solutions and choosing the one best suited to current needs and technology. Just because a mechanism has been used in the past does not mean it is necessarily a good choice for the future.

The End-to-End Principle

> The basic argument is that, as a first principle, certain required end-to-end functions can only be performed correctly by the end-systems themselves.
> … Therefore, providing that questioned function as a feature of the communication system itself is not possible.

This concept is further described in Saltzer et al. (1984), which states "End-to-end arguments are a kind of 'Occam's razor'" (see below). This argues for a simpler network and more sophisticated devices at the edge of the network. A simpler network makes it more resilient to disruption, change, segmenting, and security challenges.

Heterogeneity[1]

> ... is inevitable and must be supported by design.

There will inevitably be multiple physical layer and application layer protocols, to serve different purposes.

"If there are several ways of doing the same thing, choose one"
Well said.

Scaling

> All designs must scale readily to very many nodes per site and to many millions of sites.

It often has been noted that Internet technology is independent of scale—of the speed of network links, the complexity of the network, etc. Technologies that can easily scale up or down in electrical capacity or network size are preferred.

Performance

> ... and cost must be considered as well as functionality.

When considering solutions for the electricity system, many different systems likely would work to coordinate entities in the electricity system, but some will perform better, and some will cost less.
A system that can do anything may have cost and complexity burdens which make it less preferred to a simpler one that does somewhat less but well-covers the most useful capabilities.

Keep it simple

> When in doubt during design, choose the simplest solution.

Complexity costs money and is generally worse for privacy and security than simpler systems.

Modularity is good

> If you can keep things separate, do so.

Practicality

---

[1] Fourteen "General Design Issues" are listed, but the last seven are too specific to Internet technology to be useful here.

In many cases it is better to adopt an almost complete solution now, rather than to wait until a perfect solution can be found.

The perfect can be the enemy of the good. Even if we do not know all details of a solution, if general directions are sound, we can head towards them.

Prefer unpatented technology
Open standards are best.

Privacy
Privacy and related goals should be always addressed by end systems, whether or not the network also provides this.

4.2 Communications and Power Systems

While the electricity grid is arguably the most important technology achievement of the 20th century, Internet technology similarly can be seen as the most important one of the last 25 years.[1] It is important to recognize that it replaced an earlier achievement, the traditional phone system, and is based on a number of fundamentally different technology assumptions than the phone system: digital rather than analog, distributed and not central, packet-switched rather than circuit-switched, and highly dynamic in functionality rather than being largely static. For technology, the Internet has highly functional end devices (in the phone system phones had minimum functionality), and uses data storage everywhere (in the phone system data went from one end of the system to the other without any storage). There are other parallels as well, but fundamentally, the phone system was a "central" and "unitary," one and the Internet is "distributed" and "networked." This all suggests that we should be open to the possibility of a markedly different technology basis for electricity systems, and that if one emerges, that the network metaphor may be central to it.

Central to Internet technology and architecture is the Open Systems Interconnection (OSI) model (ISO, 1989) of layering of functionality, with the Internet Protocol (IP) at the "narrow waist" of interoperability. The architectural principles reviewed above also included that "... there should be one, and only one, protocol at the Internet level." The Internet Protocol is the "fundamental mechanism" that drives the Internet—everything else is hung off of it, whether it is a physical layer protocol below or an application layer protocol above. For the electricity system design presented here, price (for a quantity of electricity) is the fundamental mechanism.

Figure 7 shows the OSI model[2] (for communications) alongside the Network Power Integration (NPI) model (for energy). In the OSI model, the application-layer protocols (layers 4 and 5) do not need to know or care what physical layers they transit, and the physical-layer protocols (layers 1 and 2) do not need to know anything about the content of the data they carry; there is complete separation. For energy in buildings, we can have one set of technologies for managing power flows and a separate set for device functionality.

[1] Internet technology is of course older than this but it only saw widespread use in the early 1990s.
[2] Strictly speaking, the OSI model has seven layers, though the distinctions between layers 5–7 are much less clear than between other layers, and they do not contribute to this discussion. The five-layer version is shown here.

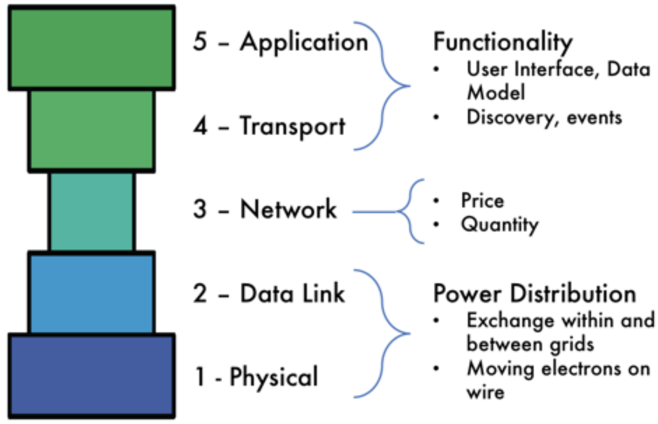

Figure 7. OSI model and Network Power Integration (Christensen and Nordman, 2013)

The NPI model only applies to end-use devices that provide building services to people (or the structure or local environment); it does not apply to power infrastructure, including storage and generation devices that may be within a building, as they do not provide such services. Since grid entities never provide services to people, NPI only applies within buildings. NPI does not in any way replace the traditional OSI model, and communication physical layers still carry application-layer data. The two contexts (power distribution and building services) are integrated internally within each end-use device (or system of devices).

This model applies most clearly to distributed systems in which every end-use device is autonomous. Some devices today and in the future have limited or no internal decision-making and so delegate that to a central control system; this is particularly prevalent in large HVAC systems.

Technology layering is well known to improve interoperability and better enable technology evolution, as layers can change independently. Devices that functionally coordinate may be powered differently, and those that are on the same power distribution circuit may have no functional relationship.

4.3 Pricing Applied

Prices are used extensively between grid entities, and are determined through many different mechanisms, such as auctions and long-term contracts. The important point is that there are prices. Grid operators have long known that the value of electricity varies over space and time. In California for example, there are over 4,000 nodes at which a Locational Marginal Price is calculated, and these are updated at five-minute intervals.

Prices at customer meters have been primarily static over the course of the day and year, often with just two different seasonal rates. This has been due to lack of adequate metering technology, lack of perceived need, lack of ways for customers to automatically use dynamic prices, and the simplicity of explaining flat rates. All but the last of these are now rapidly disappearing, and we are seeing increased use of TOU rates with three or four time periods per day, and differences between weekends and weekdays, and seasonally. In Illinois, two utilities have offered hourly dynamic rates for years (Ameren et al., 2016; ComEd et al., 2016), and tens of thousands of residential customers save millions of dollars and are generally happy (the rates are opt-in). Two retailers in Texas offer plans that directly pass through wholesale market price values, one at 15-minute intervals. The question is not whether time-varying rates will spread and become increasingly sophisticated, but how fast.

Pricing can be used within buildings as well. The utility meter price can be broadcast to end-use devices for them to use in decision-making (“prices-to-devices”; EPRI, 2006), but that is only the beginning.

Once a building has local storage and/or generation (and in some cases for any building), the availability of electricity within the building can be different from what it is on the utility grid on the other side of the meter. As price is reflecting availability, this indicates that there should be a different price on the building side of the meter—a "local price." The following examples illustrate how local prices arise:

- Differential buy/sell prices. More places are making use of prices that change depending on the direction of power flow at the meter, to encourage more utilization of local (usually PV) generation (e.g., a utility sells power at $0.20/kWh and buys it at $0.10/kWh). The marginal value of power then depends on whether the building is in buy or sell mode. If no power is flowing at all—readily attainable with a local battery—the value is somewhere in between.
- Carbon valuation. While utilities rarely include the cost of carbon pollution in rates they charge, most customers understand that damage from carbon pollution is real. At least one service, WattTime, provides data and mechanisms for people to use marginal carbon emissions values in building and device operation (Mandel and Dyson, 2017). Customers can readily add these real pollution costs to the metered price to ensure better DER operation.
- DC circuits. While long used for niche applications in buildings, particularly for communications, there is recently more use of DC distribution in buildings for other end uses. Local generation and storage are DC, as are most end-use devices internally, and there are energy efficiency advantages to keeping power in the DC domain (Gerber et al., 2017). If the net flow of power is from AC to DC, then the DC price should be set to above the AC price by at least the conversion loss. For example, if the AC/DC conversion loss is 10 percent, and the AC domain price is $0.20/kWh, then the DC price should be at least $0.22/kWh. However, when the power flows in the other direction, then the price in the DC domain should be lower than the AC domain. While usually small, wire losses should also be included in price adjustments.
- Peer-to-peer (P2P) exchange. In recent years, many people have promoted the idea of financial exchanges between customers that give the appearance of excess electricity from one customer being purchased by another customer.[1] However, true peer-to-peer exchange would be over a direct wired connection between two customers, not through the utility grid; this is particularly easy to visualize if the interconnection is a DC circuit. Peer-to-peer exchanges are contemplated in California as the "over-the-fence rule" (Lindl, 2009); this allows exchange with up to two adjacent properties so long as public right-of-way is not involved. With differential buy/sell prices, neighbors who are in different modes will benefit from a P2P exchange, with a price perhaps halfway between the utility grid buy and sell prices. While many jurisdictions do not allow for this (and perhaps none do when a public right-of-way is involved), doing so can be similar to sharing a Wi-Fi password with a neighbor—no one is likely to know or interfere.
- Microgrids. When the utility grid goes down, customers in general would prefer that at least part of their electrical infrastructure would continue to operate—to have a microgrid. When islanded, it is even more important than under normal conditions to carefully balance supply and demand, and as there is no grid price, the price must be generated locally. This can be usefully combined with the peer-to-peer exchange use case.
- Capacity constraints. An electrical domain in a building (perhaps the whole building) may have more potential demand than the supply wiring or other infrastructure (e.g., a transformer) can supply. Conventionally, systems are oversized to avoid circuit breakers tripping. With communications, such oversizing can be avoided, and a price signal (or some other mechanism) used instead to ensure that demand does not exceed supply, by increasing the price as the capacity constraint is neared. A reverse condition can occur when there is too much local production, in which case reducing the local price is the solution. Demand charges have a similar effect as a

[1] The utility has to be involved since the electrons flow through the utility meters, over its wires. In addition, utility bills would need to be adjusted accordingly, and it is not the same electrons. In reality, it is two coordinated transactions, each between one customer and the grid.

capacity constraint in that a building management system will seek to keep interval energy consumption below a particular demand limit.

- Battery management. If a local battery has a state of charge sufficiently different from how its controller thinks it should be, then adjusting the local price up or down can make power more or less available in the power domain to enable charging or discharging. Battery losses should also be considered in making battery charge/discharge decisions.

While there are likely more, the seven reasons above should be sufficient to conclude that prices need to be local to be correct.

At a minimum, a price is only instantaneous value. Utility meters commonly have fixed time periods in which prices change (from minutes to hours long), but in principle there is no need for fixed periodicity at utility meters. Local prices will likely change at times when the meter price is constant. Prices are more useful when future prices are included, whether they are guaranteed, or just a "best-effort"[1] forecast. Whether the future prices are guaranteed or not is not so important, as DER should seek to optimize on the basis of the best information they have.

The concept of a local price is part of Local Power Distribution (LPD), a "network model of power" (Nordman, 2017). It was developed to facilitate electrical systems which can be readily operated off-grid (whether always or only occasionally), and to manage DC distribution systems. However, as seen above with local prices, LPD also applies to AC-only buildings. Pricing can be applied to local generation when fuel is used, to correctly account for marginal costs and generation efficiency; PV marginal costs are generally zero.

4.3.1 Discussion

A core tenet of the price-based approach is that pricing can do the great majority of what we need to do to manage flows of energy over space and time. What remains—what utility needs that pricing alone cannot satisfy—may only be determined after considerable experience with wide use of prices that vary over space and time. This is particularly the case with what is needed on the grid with much higher levels of renewables production, as well as the deployment of other new technologies such as large amounts of stationary storage and electric vehicle charging.

A price for electricity is always a control signal. This is in addition to being a mechanism for collecting revenue or distributing payments. Even when prices don't vary over the course of a year, their absolute value determines how building owners make efficiency investments and trade-offs with other forms of energy. When prices vary over time of day, that encourages shift in demand of grid resources.

Pricing is for quantities of energy. It does not directly address electrical capacity constraints, though does indirectly with raising prices when power limits are neared. Some scheme(s) for negotiating capacity will be needed for anomalous consumption (or production), principally from vehicle charging. These capacity management mechanisms can be separate from prices for the underlying electricity.

Price is a unit of measurement. While it is usually used in conjunction with cash exchange (between different parties), it also can be readily used within a building where both devices exchanging power are owned by the same individual or organization. A consumer may own multiple bottles of wine or boxes of candy with very different costs per unit of consumption, and will take the value of the item into account in making consumption decisions, even though they don't "buy" the item from the shelf in their kitchen.

---

[1] Internet Protocol packet delivery is also "best-effort." It is not guaranteed, but on many networks it is surprisingly reliable.

These prices are also local. The relative scarcity or plenitude of an object within a building can shift the value to the owner well above the original purchase price. If there are lots of pieces of fruit and they are going to go bad soon, the value will likely be less than the original purchase cost. If there is limited milk for breakfast, and there is no store nearby, it will have higher value.

4.4 Summary

Prices may not solve all problems with grid/building coordination (for in-building DER), but they may solve the vast majority. There is a role for technology to manage capacity constraints, to ensure that power limits are not exceeded and circuit breakers not tripped; this can occur within the utility grid (vehicle charging will raise this issue in many places), at the meter, or at some internal building interconnection.

Issues of power quality and inverter control are not addressed well or at all by pricing and require a different solution. The concept of a PSI (Power Services Interface) as a complement to the ESI can cover these concerns.

Pricing as a fundamental mechanism, and the solution outlined above, has a high degree of resonance with the Internet rchitecture principles above. Pricing is simple, universal, and scalable. Prices can be determined entirely locally (when not grid-connected) or strongly tied to a utility grid price. Pricing can completely isolate the domains of power distribution and functional control from each other, for best use of layering.

Pricing *is* “transactive energy,” since exchanges of electricity based on a price are the most basic type of transaction we have.

5.0 Future Building/Grid Relationships

Existing and proposed mechanisms for buildings to coordinate with the grid are diverse, along many dimensions, some of which are interdependent. Determining how best to summarize and organize these is challenging. It is possible that no single organization is best.

5.1 Essentials of Building/Grid Coordination

A variety of elements are essential to ensuring successful coordination between a building and the grid:

A **utility grid** determines either its current state of relative power availability (with a retail price) or a desired change in that availability (for DER to increase or decrease their consumption). This has various names, including "grid signal," "demand response signal," or for a subset, "price."

A **DER device** can change its functional/operational state to increase or decrease consumption. This can be expressed directly (e.g., to change power up or down) or indirectly via a setpoint or level change or command (e.g., "shed" or "load up").

Some entity needs to **translate** the grid indication into a functional control or command. This could be a grid entity (utility or aggregator), a central device in a building, an intermediate device, or the end-use device itself. Figure 8 shows these alternatives graphically.

There needs to be **communication** from the grid to the DER device, possibly through one or more intermediate entities (in-building devices or grid entities).

There must be a **financial reward** for customers to have their DER participate, either through price-based or event-based approaches.

Price-based systems use interval utility meter data to directly reward changes in consumption according to the existing billing relationship. They only require one-way communication from the grid. The translation to functional control can occur centrally in a building, in an intermediate device, or in the end-use device.

Event-based systems have separate financial arrangements (with the utility or a third party) with the customer. They usually have two-way communication (many such architecture require two-way communications) and sometimes involve comparing "event" and "non-event" (baseline) days to verify DER operation statistically. Event-based systems can include direct coordination with individual DER or a command to a building as a whole (e.g., reduce load by X kW) that is translated to a functional command in the building.

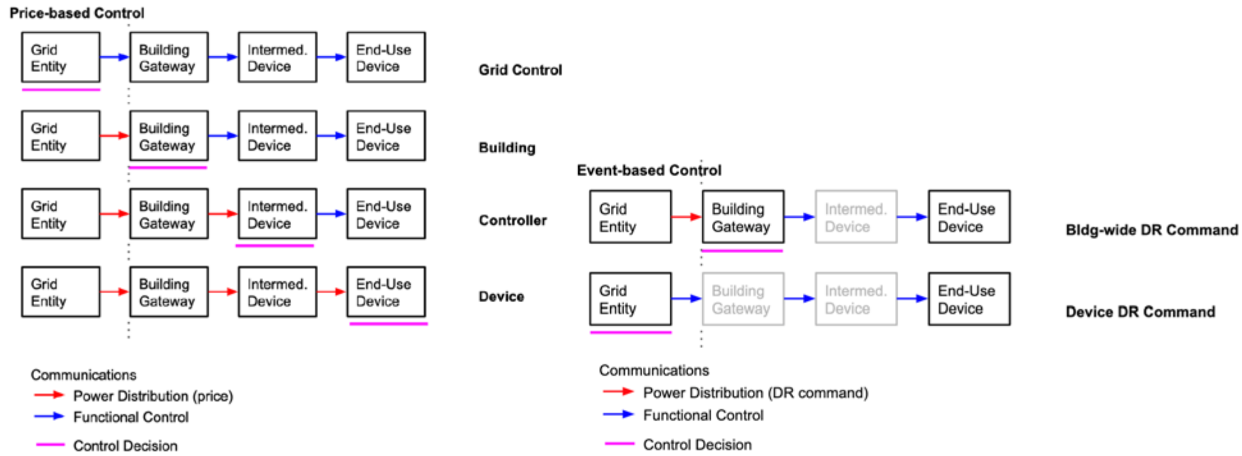

Figure 8. Price-based Control and Event-based Control Approaches

This collection of mechanisms also raises other ways they are distinguished.

5.2 Basic DR Types

Below is a list of basic types of demand response that are different in some structural sense. The terms are for convenience. This disaggregation covers coordination mechanisms, and while closely related to the grid services delivered, is distinct.

A Distributed Energy Resource (DER) can be a flexible load, battery, electric vehicle (EV), or dispatchable local generation. With Direct Load Control (DLC), a grid entity sends functional control commands to a DER—usually directly, but the commands could be relayed by intermediate devices. A grid entity can be a utility or third-party aggregator.

5.2.1 DR Types

Blind DLC: A grid entity sends a one-way command to the DER. Common examples include on/off or setpoint commands to water heaters or air conditioners.

Traditional DR: A grid entity sends a command to a building to reduce load (the amount and duration of which has been previously agreed upon). This is commonly implemented with a single DER, but it doesn't have to be. Generally the rest of the building is uncoordinated with the grid.

Informed DLC: A grid entity knows the operational state of DER and uses that knowledge in sending custom control commands. Most aggregators use Informed DLC.

Price-based DR: Any mechanism in which the price at the meter varies over the course of a day. This can be CPP, TOU, or continuously varying pricing.

DR Bidding: The building tells the grid entity its willingness to use more or less than that and the payment required to make this change (versus a self-forecast consumption or a baseline). This is what enthusiasts of the term *Transactive Energy* usually want.

Ad hoc DR: Alerting customers to an upcoming peak cost time and rewarding them to the degree that they comply as assessed by the meter. This is similar to a Critical Peak Price but the financial relationship is more indirect. A California company, OhmConnect, does this.

Table 2. Comparison of characteristics of DR types

DR Type	Direct Load Control	Event-based	Price	Bidding	Communications	Locus of Decision	Scope
Blind DLC	X	X			one-way	Grid	DER
Traditional DR		X			two-way	Bldg/Dev	DER
Informed DLC	X	X			two-way	Grid	DER
Price-based DR			direct		one-way	Device	Building
DR Bidding			negotiated	X	two-way	Bldg/Dev	Building
Ad hoc DR		X	indirect		one-way	Bldg/Dev	Building

Event requires advance agreement to participate. Price uses meter price directly; DR Bidding "Decision" is where the decision is made about translating grid goals to device control. Traditional DR may be day-ahead or day-of notification.

5.2.2 Details

One helpful organization approach was developed by Kok and Widergren (2017), who identified two dimensions of "energy management": (1) location of decision-making and (2) whether communications are one-way or two-way. Their table is worth reproducing (Table 3).

Table 3. Smart Grid Energy Management Categories (Kok and Widergren, 2016; emphases added)

	One-Way Communications	**Two-Way Communications**
Decisions on Local Issues Made Locally	Price Reaction + Full Use of Response Potential – Uncertain System Reaction – Market Inefficiency + No Privacy Issues	Transactive Control + Full Use of Response Potential + Certain System Reaction + Efficiency Market + No Privacy Issues
Decisions on Local Issues Made Centrally	Top-Down Switching – Partial Use of Response Potential – Uncertain System Reaction – Autonomy Issues	Centralized Optimization + Full Use of Response Potential + Certain System Reaction – Privacy and Autonomy Issues – Low Scalability

Several assumptions seem embedded in Table 3. "System Reaction" presumably refers to the sum of the behaviors of a collection of buildings—perhaps a modest number if only a single electrical feeder is present, or a large number if the scope is an entire substation or region. While the reaction to price changes of an individual building may be hard to predict, as one aggregates large numbers of them, average responses become much more regular. The greatest challenge for DER is to address variability of renewable supply, which is principally a regional or systemwide issue. The assertion of "Market

Inefficiency" for price reaction seems to presume that the grid is unable to send a price that balances supply and demand well.

Role of the meter
The meter can be used as the sole mechanism for financial calculations and coordination, or it can be used in addition to other mechanisms. Bills can be calculated solely based on a price for each metered time period, or they can add additional elements, with a monthly "demand charge" being the most common.

Building as a unitary or disaggregated entity
The building can be seen as a single entity, with all coordination and valuation done with the building as a whole. Or, one or more parts of the building can have separate relationships for control, measurement, or valuation. The latter approach can be achieved with the utility grid only or with one or more third-party aggregators.

5.3 Other Topics

Some issues introduce complexities that don't well map onto the categorization above.

Advance Purchase
Several architectures require buildings to forecast their future demand, then engage in processes to buy that energy in advance. These vary in detail, but some involve complex market structures with thousands or millions of active participants. The initial coordination from the building might be via a central entity or multiple entities. It is not the purpose of this discussion to address how they work internally, but rather on how the building operates and interacts with outside entities.

Usually there are provisions to account for the fact that a building will likely not consume precisely what it expects to, and the difference in consumption will be valued at a rate different from the advance purchase rate (in TEMIX the "spot" rate). As managing these variations from expected demand is burdensome for the utility, the deviation rates are generally higher, so consuming close to the expected amount is advantageous for the building.

A general principle of these is that every building must forecast its future demand at multiple different price values. With complete information on the demand curve, the utility can combine the demand curve with its supply curve to find the "market clearing" price that should be most economically efficient.

Multiple organizations per building
Most buildings have just one organization involved in electricity use. Even if it is rented, the tenant usually pays for electricity. A question arises as to whether those buildings that do have multiple organizations introduce any important considerations for grid architecture. If the organization that controls energy use is not the one that pays for the energy, sub-optimal decisions will be made; this is well-known (the "principal agent problem"). However, this is unrelated to grid architecture. Tenant spaces, whether residential or commercial, can be considered separate buildings, and often have individual utility meters (along with one for common spaces). If there is common electrical infrastructure, such as PV on the roof —or even community PV not on the same site—this can be treated as a financial transaction separate from grid coordination.

Third-party DER
Some companies may find it advantageous to place DER at customer sites and manage those DER for the customer, or provide the management to customer-owned DER. Any control signals from the third-party should be through a separate interface from the grid ESI (though it could use the standard ESI technology

and could be based on prices from the grid). The financial arrangement between the customer and the grid would be unaffected. Thus, these arrangements can be invisible to the grid.

Grid-deployed DER
Some utilities offer to place batteries in customer premises for use by the utility in normal times but can be used as backup power for the customer when the utility grid is down. The utility presumably nets out the effect of battery charge, discharge, and losses from the customer bill.

Vehicle Charging
For energy purposes, the electricity that goes into vehicles can and should be valued the same as any other energy, but there may be a need to create mechanisms to manage electrical capacity since the power levels of a vehicle charger are so much greater than those of many buildings; particularly residential units. This requires two-way communications and likely a system to reserve this capacity in advance. Capacity constraints can now be managed with communications rather than relying only on circuit breakers and/or equipment oversizing.

Domains
Power domains within a building are generally created when there are circuits that provide power quality or reliability different from that on the grid; usually it is when a circuit has higher quality or reliability than the grid. Higher quality is most commonly provided with DC power; higher reliability with local storage and/or generation, also often with DC power. A capacity constraint within the building can create multiple domains if it leads to managing entities differently on either side of the constraint. Electrical domains can be entirely isolated from the utility grid (e.g., local generation coupled to a pumping function) or can be intermittently disconnected, most notably with vehicles. A building with microgrid capability may use that microgrid to cover only part of its power distribution infrastructure; resulting in one domain that is still part of the utility grid and a microgrid domain with the higher reliability. Many solar+storage systems sold today for residence-scale buildings provide for four circuits of reliable AC power.

One proposed terminology for such domains is “nanogrid” (SGIP, 2016), which is a “single domain of power”—for capacity, quality, reliability, administration, and price. While the concept was developed to enable rich functionality for DC nanogrids, many aspects of the approach also apply well to AC systems.

6.0 Implications for Grid Architecture

This section reviews other products of the GMLC 1.2.1 Grid Architecture project. Most of the GMLC 1.2.1 project results are encompassed in five “tracks” as follows.

- Advanced bulk energy systems. The issues covered by these systems are all distant from buildings, so there is no obvious overlap.
- Urban converged networks. Most of it this track is expected to be relationships between these converged networks in the infrastructure, and not in buildings, but to the degree that there are also connections between them in buildings, they are relevant. If all such networks used dynamic pricing as their primary control mechanism, there would be a single mechanism to cross energy forms, as well as cross space and time. (report not yet available at time of this report).
- High DA-DER-Storage Reference Architecture: This track is focused on community solar and storage, which generally has nothing to do with buildings, as the community resources are electrically separate from the building that may be owned by someone who has part ownership in them. It is possible to entangle the bills, and while this may be harmless, it does not serve a grid purpose. This does not mean that such approaches are without value; they are just a separate topic. There should be a High-DER architecture in which most of the DER are in buildings and not separate, particularly as this is how DER occur today.
- High Resilience Reference Architecture. The document has very little reference to buildings. That is not necessarily a problem as it is resilience of the grid rather than that of buildings that is in scope. Power distribution within buildings should ideally be resilient both when grid-connected and when the grid is unavailable.[1] Many of the “Core Resilience Architecture Principles” have good applicability to buildings. The laminar coordination concept can also be applied inside of buildings.

The five tracks are not mutually exclusive and do not necessarily need to be taken whole; parts of each and insight from other sources can be used. In that spirit, conclusions from this report can similarly be used in any of the tracks or by any user. This report is not incompatible with any of the tracks.

[1] Buildings that are never grid-connected are not a focus of this document, but technologies for microgrids in general will be useful there.

7.0 Summary

Buildings have been previously neglected in grid architecture, mostly appropriately so. While grid organization varies from place to place, and is necessarily complex, the organization within buildings should ideally be consistent, and explicable through more simple means. That said, we do need ways to properly integrate buildings into the discipline of grid architecture, and use that to draw insight.

There are two primary aspects to this challenge: the internal structure of buildings, and how they relate to the utility grid. The nature of the building/grid interface determines if the building can be a simple “black box” to the grid or if details of building components need to be exposed to the grid for status and control.

There are a variety of alternative ways to construct the building/grid interface that vary in communications required, financial and control relationships, and what they require of technology in buildings. There is a lack of quantitative evidence to show their relative performance and costs, so grid architecture is one of the few organized ways to compare them. This paper focuses primarily on one of these alternatives—the price-centric approach—but does compare it to others.

Better integration of buildings into grid architecture requires only modest adjustment to its methods and data. One is to clearly articulate what building/grid relationship(s) are covered by a particular analysis or diagram, as they do change some of the entities and lines. Examples could include the following:

- Distinguishing whether lines to buildings are indications of grid status (e.g., price or a desire to change net building load) or device functional control (e.g., changing a thermostat setpoint or turning on or off an individual device).
- Distinguish between energy grid services and power grid services, as the nature of the entities and their control and valuation relationships are different.
- Showing how this affects grid entities: for example, with highly effective dynamic pricing, there may be no need for third-party aggregators for energy grid services.

All of this will take time and much discussion to sort through. This paper is the first word on the topic, and by no means the last.

8.0 References

Ameren Corp., Elevate Energy, Illinois Power Smart Pricing, 2016 Annual Report, April 1, 2017

Carpenter, B., editor. 1996. “Architectural Principles of the Internet.” Request for Comments: 1958. June.

Elevate Energy, ComEd’s Hourly Pricing Program, 2016 Annual Report, April 25, 2017.

EPRI (Electric Power Research Institute), Advancing the Efficiency of Electricity Utilization: “Prices to Devices” Background Paper, 2006 EPRI Summer Seminar.

Kok, Koen, and Steve Widergren, *A Society of Devices*, IEEE Power and Energy Magazine, May/June 2016.

International Organization for Standardization, ISO/IEC 7498-4:1989 -- Information technology -- Open Systems Interconnection -- Basic Reference Model: Naming and addressing. 1989.

Lindl, Tim, Letting Solar Shine: An Argument to Temper the over-the-Fence Rule, 36 Ecology L.Q. (2009). Available at: http://scholarship.law.berkeley.edu/elq/vol36/iss4/3

Jamie Mandel, Jamie and Mark Dyson, WattTime Validation and Technology Primer, Rocky Mountain Institute, 2017. https://www.watttime.org/app/uploads/2019/01/Automated-Emissions-Reduction-Primer_RMI-Validation_June2017.pdf

Marnay, Chris, Microgrids and Heterogeneous Power Quality and Reliability: Matching the Quality of Delivered Electricity to End-Use Requirements, LBNL-63524, 2007.

Melton RB, S Pal, Problem Domain Reference Model: Functional Groups and Reference Architectures for the Electric Grid, GMLC, January 2019

Nordman, Bruce "The Fit Grid"-Principles for a better electric future, submission for the SEPA 51st State Process, Phase 2, March 2016.

Nordman, Bruce,, Local Grid Definitions, A Smart Grid Interoperability Panel (SGIP) Whitepaper, 2016.

Nordman, Bruce, and Ken Christensen, L*ocal Power Distribution with Nanogrids*, International Green Computing Conference, Arlington, VA, June 2013.

Nordman, Bruce, A Simulation of Local Power Distribution Control Strategies, Second International Conference on DC Microgrids, Nurnburg, Germany, June 2017.

J.H. Saltzer, D.P. Reed and D.D. Clark*, End-To-End Arguments In System Design, ACM Transactions in Computer Systems 2, 4, November, 1984, pages 277-288.

Taft, JD, A Becker-Dippmann, Grid Architecture (Final), PNNL-24044, January 2015.

Taft, JD, Grid Architecture 2, PNNL-24044 2, January 2016.

GRID
MODERNIZATION
LABORATORY
CONSORTIUM
U.S. Department of Energy

Address Line 1
Address Line 2
City, ST Zip
Phone Number

https://gridmod.labworks.org/