

EcoBlock: Grid Impacts, Scaling, and Resilience

Aminy Ostfeld, Alexandra von Meier, Michael Whitmeyer, and Mohini Bariya
Berkeley Energy and Climate Institute
and Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

December 1, 2017

Introduction

Widespread deployment of EcoBlocks has the potential to transform today's electricity system into one that is more resilient, flexible, efficient and sustainable. In this vision, the system will consist of self-sufficient, renewable-powered, block-scale entities that can deliberately adjust their net power exchange and can optimize performance, maintain stability, support each other, or disconnect entirely from the grid as needed. This vision requires not only effective design of the EcoBlocks themselves to realize these capabilities, but also substantial changes in the operation of the grid itself. Ideally, the transition would be a win-win for EcoBlocks and the grid, and grid operators and planners would be proactive participants in the redesign. However, depending on the design and management of the EcoBlocks, the changes could also be detrimental to the grid and other participants, causing utilities and regulators to resist EcoBlock deployment. Therefore, for this model to succeed, it is necessary to anticipate the range of possible interactions between EcoBlocks and the grid as their number and penetration level scales up. The system architecture must be designed with the final vision in mind to ensure the means of managing EcoBlocks are fully scalable. At the same time, compatibility with today's planning and operational conventions is necessary to facilitate the entry of the first EcoBlocks into the power system.

This report is intended as an independent analysis of the potential relationships, both constructive and adverse, between EcoBlocks and the grid. We thoroughly survey the possible impacts of EcoBlocks and the distributed energy resources (DERs) embedded in them, initially making no assumptions about how the resources are controlled. Then, we evaluate possible strategies for managing both the resources within EcoBlocks and the EcoBlocks themselves, with specific focus on supporting positive impacts on the power system as the EcoBlock penetration level scales.

Of the many potential capabilities of EcoBlocks, we expect adaptive islanding to be particularly transformative. The ability of individual blocks to disconnect from the grid and self-supply their loads during an emergency, then seamlessly re-connect when the problem is solved, is a huge advance in power system resilience. While this is a major shift from how the grid operates today, it is in utilities' best interest to support efforts to improve resilience. Grid resilience is vital for safety and security, public health, and the economy, and is only increasing in importance as the occurrence of extreme weather events increases due to climate change. We therefore give particular attention to the themes of adaptive islanding, resilience and self-sufficiency throughout this report.

Section 1 will outline the diversity of potential impacts, positive and negative, that EcoBlocks could have in the context of today's grid. Section 2 discusses the range of design choices for EcoBlock hardware and operations, as well as the framework for their interactions with the main grid, to shift the balance toward positive impact. Section 3 focuses on the thresholds of penetration level at which specific changes in impact and grid operations may occur. Section 4 summarizes the changes in the form of paradigm shifts in power system functions. Section 5 is a quantitative analysis of resilience under the adaptive islanding paradigm, probabilistically estimating self-sufficiency of islanded EcoBlocks under a range of scenarios. Finally, section 6 presents suggestions for future work.

Contents

| | | |
|----------|--|-----------|
| 1 | Overview of potential grid impacts | 3 |
| 1.1 | Loading | 3 |
| 1.2 | Resilience | 5 |
| 1.3 | Flexibility | 6 |
| 1.4 | Situational awareness | 8 |
| 1.5 | Planning | 9 |
| 1.6 | Summary | 10 |
| 2 | EcoBlock design for positive impact | 11 |
| 2.1 | Installed assets | 12 |
| 2.2 | EcoBlock operating strategies | 15 |
| 2.3 | EcoBlock-grid interaction | 16 |
| 3 | EcoBlock scaling | 18 |
| 3.1 | Distribution power flow thresholds | 18 |
| 3.2 | Topological thresholds | 20 |
| 3.3 | Beyond the distribution level | 20 |
| 3.4 | Significance of the block scale | 21 |
| 4 | Summary of paradigm shifts | 21 |
| 4.1 | Power flows | 22 |
| 4.2 | Flexibility | 22 |
| 4.3 | Resilience | 22 |
| 5 | Resilience analysis | 23 |
| 5.1 | Simulation Procedures and Data | 23 |
| 5.2 | Results and discussion | 25 |
| 5.3 | Conclusions | 32 |
| 6 | Suggestions for future work | 32 |
| 7 | References | 34 |

1 Overview of potential grid impacts

The potential impacts of EcoBlocks on the electric grid are wide-ranging and depend on the EcoBlocks' physical characteristics, how they operate, and what knowledge and control grid operators have of their behavior. Analyzing the potential impacts of distributed energy resources—including photovoltaics (PV), inverters, electric vehicles (EVs), energy storage devices, and demand response (DR)—is already a significant area of research. Since EcoBlocks are a combination of all of these resources, much of the literature in this area could be relevant for EcoBlocks. The range of potential impacts on loading, resilience, flexibility, situational awareness, and planning are discussed in the following sections and summarized in Tables 1-5. For each of these categories, there are numerous ways in which DERs, and therefore EcoBlocks, can have positive effects or negative ones. Importantly, if the EcoBlock aggregates these resources into a single unit with a single point of connection to the grid, it offers the additional possibility of controlling their behavior in a holistic way to shift the balance toward the positive.

1.1 Loading

Today, there is concern that uncontrolled DERs will negatively impact the loading on distribution systems. For example, if large numbers of electric vehicles are allowed to charge simultaneously or during high load conditions, the increased peak current can exceed the thermal ampacity ratings of lines or transformers [1]. Other significant new electric loads could have a similar effect, for example if gas-powered heaters and water heaters are replaced in large numbers with inefficient electric counterparts. Similarly, during periods of low load and high PV generation, reverse power flow exceeding thermal limits can occur [2]. Upgraded protection devices may be required if the peak current due to new generation or loads exceeds the rating of the existing device. It is also possible for losses to increase as a result of increased real power flow or improper control of inverter power factor, for example if inverters absorb reactive power to regulate the voltage during reverse power flow conditions [3], [4]. Increased peak load and reverse power flow due to EVs and PV also increase the likelihood of undervoltage and overvoltage conditions, respectively, particularly on secondary circuits and long, high-impedance feeders that are most susceptible to voltage problems [1], [2], [5]. If not remedied, voltages out of the permissible range could damage sensitive equipment, including appliances belonging to neighboring customers. Additional or upgraded voltage regulation equipment would therefore be required. Furthermore, when variability in PV generation and electric vehicle charging load cause voltage to fluctuate, voltage regulation equipment such as tap changers may operate more frequently and require maintenance or replacement sooner [5], [6]. EcoBlocks with inadequate coordination among their generation and load resources and/or inadequate energy storage could potentially cause these same problems.

On the other hand, appropriate control of DERs can not only mitigate these negative impacts but even improve conditions relative to the baseline (pre-DER) scenario. These positive impacts arise from the combination of local generation and energy efficiency retrofits leading to low *average* load, plus scheduling of storage and DR to minimize *peaks* in load. First, reduced power flow in lines and transformers would reduce losses, decreasing generation requirements and cost. Optimizing inverter power factor can minimize reactive power flow, further reducing losses [3]. Generally, as electricity demand and distributed generation increase, it is expected that some equipment will become overloaded during times of peak load or generation and require upgrades. However, if EcoBlocks are managed to reduce peak load, the investment in these upgrades could be deferred, allowing a greater renewable energy penetration to be achieved at lower cost to the utility. Optimal scheduling of distributed energy storage has been shown to be effective at reducing peak load on distribution circuits [7], [8]. With high enough penetration of EcoBlocks performing peak load shaving, congestion at the transmission level could also be reduced. Additionally, if the power flow on distribution lines can be lowered, the voltage drop is also lowered. As a result, the voltage at the feeder head can often be reduced while still maintaining voltages within the allowable range (within 5% of the nominal value in the United States, according to the ANSI standard C84.1). This technique, known as conservation voltage reduction, has been proven effective at reducing energy demand on many US circuits

Table 1: Potential benefits and caveats of EcoBlocks related to loading.

| Benefits | Caveats |
|--|--|
| <ul style="list-style-type: none"> • Reduced losses due to reduced real power flow • Reduced losses due to reactive power control • Reduced peak load • Deferred equipment upgrades due to reduced loading • Reduced transmission congestion • Reduced voltage drop on distribution circuits • Increased potential for conservation voltage reduction • Reduced wear on legacy voltage regulation equipment • Three-phase balancing | <ul style="list-style-type: none"> • Potential to cause or worsen loading and voltage problems if EcoBlocks are ineffectively controlled • Benefits highly dependent on EcoBlock assets and controls, and feeder characteristics; loss reduction and voltage improvements may be negligible • Reduced load factor • Potential need for equipment upgrades (e.g. for protection coordination); these may outweigh other upgrade deferral benefits |

[9]. Reduced variation in load over time would also reduce variation in voltage. This in turn could reduce wear on voltage regulation equipment such as tap changers or capacitor banks, delaying the need for maintenance or replacement of the equipment. Finally, EcoBlocks could play a proactive role in three-phase balancing, which is not a major concern for utilities at present but is expected to become a bigger issue with the increasing amount of PV and EVs on single-phase connections. The impact of EcoBlock topology on three-phase balancing capability will be discussed further in Section 2.1.

EcoBlocks could in theory provide these benefits with minimal interaction with the grid operator, particularly if most of the customers on a circuit are EcoBlocks, each working independently to flatten their net load profiles at their respective points of common coupling. However, when there are other significant players on a circuit—for example, large or numerous customers not in EcoBlocks, or a large variable generation resource—nearby EcoBlocks could also be recruited to compensate for these and reduce the power flow upstream. EcoBlocks with excess generation could also provide power to their neighbors, relaxing the constraints on net load for individual EcoBlocks while still flattening the load profile of the aggregate.

A few caveats with these potential benefits should be noted. First, the benefit of reducing losses and voltage variation is highly dependent on the particular system being considered. Losses and voltage problems are most common on long, rural distribution feeders and improvements through DER could have substantial value, but it is not obvious how the EcoBlock concept will apply to areas with low load density. However, in an urban distribution system where voltage drops and losses already tend to be small, the improvement due to EcoBlocks could be negligible.

Second, EcoBlocks may reduce load factor (i.e. the ratio of average to peak demand) as seen by the grid. Even if they reduce loading in the typical case, they may still occasionally rely on the grid for providing peak load if their local energy resources are insufficient. This would require the grid operator to continue maintaining infrastructure sized for peak load, but with reduced revenue due to the reduction in kWh consumed. This challenge, and potential solutions such as alternative tariff structures for high EcoBlock penetration areas, will be discussed further throughout this report. Finally, the topology of the EcoBlock plays an important role. For example, if the EcoBlock interfaces with the main grid at a single point of common coupling (PCC), then minimizing net load translates directly to minimizing power flowing through that point. The situation would be quite different if the loads connect to the main grid at a different point from the energy sources (PV and storage)—for example if all the loads are AC and interface with the grid through existing service transformers—especially if EVs are included. In this case, the service transformers would need to support all the power flowing from the sources to the loads, and

could therefore be at risk of overload even if the *net* load of the EcoBlock is small.

In sum, the potential impacts of EcoBlock adoption on loading and the related issues of voltage, losses and equipment upgrades may vary widely, depending on the capacity of the resources within the EcoBlocks, how they connect to the main grid, and what objectives are used in their control.

1.2 Resilience

The National Infrastructure Advisory Council defines resilience as “the ability to reduce the magnitude and/or duration of disruptive events. The effectiveness of a resilient infrastructure or enterprise depends upon its ability to anticipate, absorb, adapt to, and/or rapidly recover from a potentially disruptive event” [10]. One major component of grid resilience is the ability to continue serving customers even if parts of the infrastructure fail. The traditional way to achieve this type of resilience is through redundancy, particularly at the transmission level. However, redundant infrastructure requires significant capital investment, and may still be insufficient in the case of major catastrophes. While systems are often designed with N–1 or N–2 redundancy—meaning they can survive a loss of 1 or 2 components—2012’s Hurricane Sandy, for example, was an N–90 contingency [11]. At the distribution level, where 90% of outages originate [12], circuits are typically radial, meaning a single failure will often cause an outage for all the customers downstream.

EcoBlocks, positioned at the distribution level and having the capability to operate as islands, have the potential to revolutionize grid resilience. As long as they have the necessary hardware and controls to disconnect and re-connect to the main grid, and sufficient generation and storage to meet demand while islanded, they can continue to serve loads regardless of what has occurred upstream on the main grid. If the investment required to make EcoBlocks self-sufficient is less than that required for an equivalent level of traditional infrastructure redundancy, then EcoBlocks are a favorable business proposition for utility planners. Adaptive islanding would also be, of course, advantageous for residential customers. In the event of major disasters, where the main grid is unavailable for a long time, benefits include improving public safety by keeping lights on, keeping food fresh in refrigerators, and keeping phones and other vital electronics charged. Many critical loads such as hospitals and data centers already have backup power systems, typically in the form of diesel generators. EcoBlock technology would be an upgrade for these customers as well, since solar PV systems do not cause local pollution and do not rely on fuel supply chains.

In addition to serving the customers within the EcoBlock, the presence of a few energized areas can also benefit other people in the region during a wide-area outage. For example, neighbors without power could go to EcoBlocks to charge critical electronics and vehicles that they may need in an emergency. In this way, EcoBlocks could help society cope with disasters even at low penetration levels.

Another way that the adaptive islanding behavior of EcoBlocks may contribute to grid resilience is through avoidance of cascading failures. Cascading failures occur when the protective mechanisms to disconnect grid components during excursions in voltage or frequency inadvertently worsen the grid condition and cause other components to disconnect or fail. For example, inverters tripping offline in response to low voltage can further depress the voltage and thus exacerbate the condition. If EcoBlocks can detect these events, or respond to a command from the grid operator, and automatically disconnect from the grid, they can reduce the loading on the system and potentially stop the cascading failure from continuing. EcoBlocks controlled for disturbance rejection could even prevent a cascading failure scenario from starting at all.

EcoBlocks could also facilitate recovery from outage events, which is currently a difficult process that requires careful balancing of generation and load units that are re-energized step by step. If too much load is added all at once, the generators may not be able to support it and the system will collapse again. EcoBlocks, on the other hand, could be controlled to maintain zero net power exchange while the physical re-connection is in progress, then begin to inject or consume power gradually after ensuring the system is stable.

The adaptive islanding capability of EcoBlocks thus represents a true paradigm shift in grid resilience, which would improve as the number of EcoBlocks scales up. The rollout of EcoBlocks would be desirable both for customers, especially those with critical loads, and for grid operators whose bottom line is to keep

Table 2: Potential benefits and caveats of EcoBlocks related to resilience.

| Benefits | Caveats |
|---|--|
| <ul style="list-style-type: none"> • Ability to continue serving loads despite outage of main grid • Cleaner and more sustainable power source than diesel backup generators • Ability to help non-EcoBlock neighbors by providing pockets of service during wide-area outages • Prevention of cascading failures • Smoother and more reliable recovery from outage events | <ul style="list-style-type: none"> • High reliability and large energy resources required for long-term islanding • Loss of protection coordination due to reverse power flow • Potential to increase cascading failure risk by disconnecting at wrong time • Potential to accelerate grid defection • Potentially worse reliability for non-EcoBlock customers, due to reduced utility budget for infrastructure maintenance |

the lights on. However, this vision is not without its challenges. First of all, new resources that include their own control systems and generation have the potential to disrupt existing procedures for protection and reliability. Just as disconnecting at the right time could mitigate disturbances, disconnecting at the wrong time could exacerbate disturbances and increase the risk of cascading failures. Also, reverse power flow on distribution networks designed for unidirectional flow can lead to a loss of protection coordination. The operating strategy of the EcoBlock must therefore be well-designed and synergistic with that of the main grid so that these problems are avoided.

The usefulness of the EcoBlock as a backup power system is limited by the self-sufficiency and reliability of the EcoBlock infrastructure itself. If faults occur frequently on the EcoBlock itself, while the adaptive islanding paradigm could prevent them from impacting other customers, it would not help the reliability of service within the EcoBlock. Also, to function in islanded mode for long periods of time (for example, in response to an extreme weather event that destroys grid infrastructure) the EcoBlock must have enough generation, storage, and demand response capacity to be self-sufficient in the long term.

Another challenge with the scaling of EcoBlocks is that as fewer loads and even fewer critical loads remain on the main grid, there would also likely be less motivation and less revenue for grid operators to maintain reliable infrastructure. This could lead to an increase in service interruptions for non-EcoBlock customers. If the cost to customers of upgrading to EcoBlock infrastructure is significant, low-income customers would likely participate later or not at all, and would therefore suffer a disproportionate number of interruptions. It will be necessary to plan in detail which aspects of infrastructure maintenance are the responsibility of the grid operator and which are the responsibility of the EcoBlocks, as well as to design a tariff system that gives each party enough revenue to perform its responsibilities.

The grid defection problem is not limited to EcoBlocks; similar concerns exist for individual customers with local generation and storage. However, as will be shown in Section 5, block-level aggregation is more reliable than individual home self-consumption. Therefore, the option to participate in EcoBlocks may accelerate grid defection. On the other hand, if the EcoBlock design is presented to utilities along with a viable market solution to the grid defection problem, it may be highly appealing.

1.3 Flexibility

In addition to the ability to adapt and recover after a disturbance, grid resilience also requires flexibility—the ability to balance supply and demand dynamically at various timescales. Flexibility services are characterized by the magnitude and direction of the power adjustment (increase or decrease), the starting time and duration of the adjustment, and the location of the resource [13]. Today, multiple markets for flexibility exist at different time scales, from frequency regulation (which relies on automated control with a new setpoint every few seconds) to hour-ahead balancing markets to day-ahead energy markets [14].

Table 3: Potential benefits and caveats of EcoBlocks related to flexibility.

| Benefits | Caveats |
|---|--|
| <ul style="list-style-type: none"> • Increased operating reserves • Increased resource for voltage support, through reactive power control • Increased interruptible load resource • Increased range of locations where flexibility resources are available | <ul style="list-style-type: none"> • Potential to cause local problems while providing higher-level services • Burden of optimally dispatching and communicating with large number of blocks |

Since EcoBlocks include a diversity of resources that can increase or decrease net power, and some with flexible time of operation (such as energy storage), they have the potential to perform a range of services while connected to the grid. For example, energy storage can be treated as operating reserves and dispatched quickly to meet demand. If vehicle-to-grid power is possible, electric vehicle batteries can also act as a form of reserves. Inverters can be recruited to provide voltage support through reactive power compensation. PV generation can be curtailed if it is causing unfavorable conditions. In addition, since EcoBlocks can island and reduce consumption, they can act as flexible or interruptible loads in order to help the grid operator manage situations such as supply shortages and voltage constraints. Increased penetration of EcoBlocks would be expected to improve flexibility (assuming they all choose to offer flexibility services), since it would not only increase the overall capacity of the flexibility resource, but also the number of locations where the resources are available. One concern with DERs participating in flexibility markets is that in the process of performing one service they may cause other problems. For example, distributed resources performing services for the transmission level such as frequency regulation could cause voltage problems locally. Widespread deployment of EcoBlocks could mitigate this issue by allowing grid operators to select the optimal EcoBlocks to dispatch, although the success of this approach would depend greatly on the quality of distribution-level circuit models and the grid operator’s knowledge of the EcoBlocks’ characteristics. EcoBlocks will be a particularly useful resource if they can perform the services at a lower cost—for example, if the cost of interrupting EcoBlocks is lower than that of other interruptible loads—or if the existing resource is already fully exhausted.

To enable these services, the grid operator must not only know what capacity is available in the form of EcoBlocks, but also be able to influence EcoBlock operation, either through direct control or through price signals. The practicality of recruiting EcoBlocks for flexibility will be limited if the grid operator needs to negotiate with each EcoBlock as an individual resource, particularly as the number of EcoBlocks increases. Needing to compute the optimal dispatch of a very large number of EcoBlocks, each having relatively small capacity on its own, could also be a burden. Today, virtual power plants (VPPs) exist to facilitate recruitment of diverse resources, which may or may not be geographic neighbors, for flexibility services [15]. An aggregation of EcoBlocks could offer similar services as the VPP, plus the additional option of islanding one or more of the blocks. In this model, the grid operator would communicate with an aggregator representing a group of EcoBlocks and request a service (such as providing a particular amount of power or interrupting a particular amount of load, at a certain time). The aggregator would then decide how to distribute the commands among its EcoBlocks to produce the desired outcome. Alternatively, a location-specific dynamic pricing scheme could be implemented to incentivize EcoBlocks to behave in desirable ways. Communication would be straightforward, if the grid operator simply broadcasts price signals without negotiation, but significant computation would still be required in order to determine the prices that would cause the desired behavior.

Table 4: Potential benefits and caveats of EcoBlocks related to situational awareness.

| Benefits | Caveats |
|--|---|
| <ul style="list-style-type: none"> • Reduced uncertainty due to statistical aggregation of resources • Reduced uncertainty due to active compensation for generation and load fluctuations • Opportunity to install sensing infrastructure and improve visibility | <ul style="list-style-type: none"> • Increased uncertainty if generation is masked by load • Increased uncertainty if EcoBlocks island without proper notification • Potential to overwhelm grid operator with non-standard communications and large data size |

1.4 Situational awareness

Situational awareness refers to knowledge of conditions in the grid including physically installed grid infrastructure, past and present states of the network (e.g. voltages, power flows, and switch status), and the health and security of the network and its components. This knowledge helps grid operators make decisions that are safe, efficient and economical. Historically, the focus has been on observability of transmission systems; monitoring distribution systems, with their smaller signals of interest and much greater numbers of nodes, has generally not been considered worth the investment. However, with the increasing penetration of DERs, particularly variable and unpredictable generation such as PV, visibility at the distribution tier is becoming increasingly important as well.

Variable and unpredictable generation can present a significant challenge for situational awareness. In particular, when PV generation is behind a net meter, the grid operator is only able to see the net load; the total amount of load is masked by the generation. As a result, net load may fluctuate significantly when clouds pass over, but it is challenging to predict the magnitude of these changes. The grid operator is also unaware of the cold load pickup requirement should an outage occur. It is possible that EcoBlocks could increase these forms of uncertainty, thereby increasing exposure to contingencies. This would occur if the EcoBlocks allow the net load at their PCC to fluctuate freely (i.e. if the storage does not smooth the net load curve adequately) and if the grid operator lacks information on their generation and consumption. Uncertainty would be further increased if the EcoBlocks disconnect and re-connect from the main grid at their convenience without communicating with the grid operator. In this case, the grid would be exposed to sudden increases or decreases in net load as EcoBlocks connect and disconnect, without awareness of the total number of EcoBlocks connected at each moment in time.

On the other hand, uncertainty could also be reduced due to the aggregation of the generation, load, and storage resources within EcoBlocks. Although the generation and load may fluctuate on their own, storage or demand response with sufficiently fast response time could compensate for these fluctuations and maintain the net load within a reasonable range. In this case, EcoBlocks would greatly reduce uncertainty compared to PV systems without integrated storage. Furthermore, variability of PV generation tends to decrease with spatial distribution of the PV systems, because passing clouds impact only some of the panels at a time [16], [17]. As a result, for a given PV penetration on a circuit, variability would be reduced if the PV panels are distributed across many homes in many EcoBlocks rather than concentrated in one large array. These reductions in variability are not only beneficial in terms of uncertainty, but would also reduce voltage volatility and related issues such as excessive operation of tap changers.

The design and retrofit process for the EcoBlocks also presents an opportunity to install monitoring and communication infrastructure to improve situational awareness. For example, a device such as a micro-phasor measurement unit (μ PMU) [18] could be installed at the PCC of each EcoBlock to measure AC current and voltage phasors. Additional sensors on the EcoBlock’s DC system could monitor behavior of components inside the EcoBlock, such as PV generation and charging rates of the storage and EVs. This data could be transmitted to the grid operator to disaggregate the net load observed at the PCC and aid in

the prediction of future conditions. The potential challenge of the EcoBlock connecting and disconnecting at will could also be remedied with communication. For example, the EcoBlock could simply notify the grid operator when it connects or disconnects, or it could request confirmation from the grid operator beforehand. If the switch to disconnect the EcoBlock from the grid is on the AC side of the inverter, a μ PMU could independently validate the switch status by measuring the voltage angle difference between the EcoBlock and the main grid. With all this information, EcoBlocks could not only avoid worsening situational awareness, but in fact increase visibility at the distribution level from what it is today. The ability to observe previously undetected problems on the grid, coupled with controllable resources located at the same site as the sensors, has great potential to improve power quality, reliability and efficiency.

It should be noted that since grid operators are already exposed to large amounts of data and frequent alerts, more information is not always better for situational awareness. With a large number of EcoBlocks in a system, visualizing information from every one of them could prove overwhelming for the grid operator, complicating their decision making or obscuring more urgent problems. Therefore, some sort of aggregation would be essential to condense the data from a large number of EcoBlocks into a smaller number of messages that are useful and actionable for the grid operator. What exactly these messages should be, and how to distill them from the data, is an active area of research [19]. This wealth of information could also support automated operations, including control of the EcoBlocks themselves, in situations not requiring human involvement. For automated operations and human operators alike, the use of non-standard communication schemes for EcoBlocks could complicate interactions and compromise situational awareness. Therefore, the communication protocols and data formats should be standardized across all EcoBlocks so that the new data sources can easily be integrated into a unified monitoring framework.

1.5 Planning

Today, utilities are responsible for calculating the PV hosting capacity of their feeders for compliance with power quality and network constraints. This process requires an accurate model for each feeder in question, since the impact of PV depends on the characteristics of the circuit, and a large number of simulations to represent the range of possible installation locations and sizes. If the hosting capacity on a circuit is low or unknown, impact studies for proposed PV installations are required even at low penetration. The utility may also choose to mitigate the expected PV impacts by upgrading conductors or transformers, or installing new voltage regulation equipment [20]. Both of these options increase costs and disincentivize PV adoption. On the other hand, suppose the designer of an EcoBlock guarantees the EcoBlock will have certain behavior—e.g. limits on power and its rate of change—regardless of the actual amount of PV present. Then, not only can the permissible PV penetration levels increase, but the burden of calculating them can be shifted away from the utility. Additionally, as discussed in section 1.1, equipment upgrades can be deferred if EcoBlocks reduce peak load and voltage range compared to an uncontrolled DER scenario. Of course, it is essential for these calculations to be dependably correct: if they are not, the planning burden for the utility would not decrease, or could even increase, if they must both check the EcoBlock designer’s work and redo some of the calculations.

Although EcoBlocks may have positive impacts on planning for normal operation scenarios, they may require additional upgrades to ensure proper protection coordination. Several protection issues can occur when there is significant PV penetration on a circuit. For example, excessive reverse power flow could trip protection devices in non-fault conditions, or additional fault current due to PV could exceed the interruption rating of existing protection devices. If the substation current is very low due to large amounts of generation downstream, a fault might not cause the substation breaker to trip. Atypical power flow conditions could also cause the wrong protection devices to operate: for example, if EcoBlocks on one feeder are providing power to another feeder, and a fault occurs on the second feeder, the protection on the first feeder might operate first, which is not the correct response. These problems would become more likely with higher penetration of EcoBlocks, requiring the utility to perform protection studies and potentially upgrade devices. It is worth evaluating whether the EcoBlock designer can participate in the design of new protection schemes, for example by incorporating certain protection devices into the EcoBlock to

Table 5: Potential benefits and caveats of EcoBlocks related to planning.

| Benefits | Caveats |
|---|---|
| <ul style="list-style-type: none"> • Increased permissible PV penetration levels • Improved ease of PV hosting capacity determination and other planning studies if EcoBlocks participate • Equipment upgrade deferral • Potentially improved power quality and reliability in the neighborhood, e.g. through harmonic cancellation and voltage support | <ul style="list-style-type: none"> • Increased planning burden for utility if EcoBlock calculations are unreliable • Need for protection studies and possibly upgrades • Potentially worse power quality and reliability in the neighborhood due to grid defection or inadequate EcoBlock management |

ensure that it does not cause certain protection coordination issues. This way, the burden for the utility of designing (and installing) protection schemes for circuits with EcoBlocks could be reduced. Again, of course, accuracy and reliability of the EcoBlock designer’s contribution to protection design are vital to ensure that EcoBlocks do not worsen or cause dangerous situations.

Unless a circuit consists exclusively of EcoBlocks, the potential positive or negative impacts on power quality and reliability for non-EcoBlock customers must be considered. The planning process must therefore also include establishing guidelines and accountability for these impacts. EcoBlocks could improve power quality in their vicinity if their inverters are used for harmonic cancellation and their DERs provide voltage support (such as actively reducing voltage volatility and canceling flicker). On the other hand, they could worsen voltage volatility and harmonics. EV charging in particular can cause significant harmonic distortion, which can accelerate transformer aging and damage sensitive electronic loads [21]. As discussed previously, EcoBlocks could also mitigate disturbances or exacerbate them, including preventing or contributing to cascading outages, depending on how they are controlled. The non-EcoBlock customers on the circuit are exposed to these same risks, but without the safety net of local backup power and islanding capability that the EcoBlocks have. To avoid these issues, particularly as the penetration of EcoBlocks increases, EcoBlocks should be held to a high standard in terms of power quality and reliability impacts. For example, they could promise to keep their harmonics injection within certain limits, offer control of their resources for voltage support, and design their islanding schemes to minimize the risk of exacerbating disturbances. If the EcoBlocks themselves are not responsible for maintaining an acceptable level of power quality and reliability, then the utility must take up the responsibility. Since this would require additional investment on the part of the utility, and charging the near-zero-net-energy EcoBlocks by the kWh would produce insufficient revenue at high penetration levels, a new tariff structure would be required to keep rates from increasing for non-EcoBlock customers.

1.6 Summary

Since EcoBlocks are at heart a collection of PV generators, energy storage, electric vehicles and other loads, they could potentially cause all the problems associated with uncoordinated operation of these components. But with appropriate design and control of the resources within the EcoBlocks, plus communication with the grid operator, it should be possible not only to avoid these negative impacts but to add positively to grid performance and resilience. In an ideal scenario, EcoBlocks would present a smooth net load profile to the grid, minimizing losses and voltage problems and reducing uncertainty; revolutionize resilience and flexibility through adaptive islanding; improve situational awareness through the addition of new sensors and communication channels; assist in planning and delay infrastructure upgrade needs. It is crucial to understand the “control knobs” in the design of the EcoBlock that will avoid the negative outcomes and facilitate the positive ones. These design choices center around the specific assets in the EcoBlocks, their operating strategies, and how they interact with the grid operator. The next section will elaborate on these

choices and how they can enable the positive scenario, particularly as the number of EcoBlocks scales up.

2 EcoBlock design for positive impact

From the potential impacts discussed above, a set of functional requirements for EcoBlocks and the system surrounding them can be outlined:

- While connected to the main grid, the EcoBlock should be able to control the net power it exchanges with the grid. At a minimum, this requires control of energy storage charging and discharging. The ability to modify the consumption or timing of some loads and to curtail generation would further expand the range of net power available to the EcoBlock. This control capability will allow optimization of the EcoBlock net power, for example to flatten the net-load curve or to minimize costs, or at least ensure no constraints are violated. Ability to tune reactive power, e.g. through the inverter power factor, would offer an additional control knob for further optimization.
- Some communication is needed to indicate what the EcoBlock's net power should be. This could take the form of a direct command sent by the grid operator or an aggregator, a price signal, or even peer-to-peer communication with neighboring EcoBlocks. Any entity making decisions requires information (data or models, preferably both) for the resources it controls as well as other relevant network and load characteristics.
- While connected to the grid, the EcoBlock should not adversely affect power quality for its neighbors, for example by injecting excessive harmonics or causing voltage flicker.
- While disconnected from the main grid, the EcoBlock should be self-sufficient. If the islanded portion consists of DC power distribution only, this translates to balancing power supply and demand, while if AC is included, both real and reactive power must be balanced. Therefore, the total rated power of controllable resources (energy storage real power, inverter reactive power, and any interruptible loads) must be sufficient to ensure demand is met without additional input from the grid. The energy capacity of storage resources must also be sufficient to meet demand, at least for a certain pre-determined amount of time.
- To ensure stability while islanded, there must be resources in the EcoBlock with sufficiently fast response time to adjust their output dynamically as load and generation change. Like for longer-term self-sufficiency in the point above, stability here includes real power balancing plus reactive power balancing if AC devices are present. These resources must be controlled to maintain voltage (and frequency, in the AC case) within allowable bounds.
- The components in the EcoBlock should have sufficient reliability and/or redundancy so that if there is a failure on the main grid, the EcoBlock is highly likely to remain operational.
- The islandable portion of the EcoBlock should be able to connect and disconnect from the main grid through a single point of common coupling. If only the DC system can island, the inverter is the natural choice for this point, while if the AC system can island, AC switchgear is required. Ideally, seamless resynchronization would be performed in the AC cases by monitoring the frequency and voltage phasor difference across the switch, then controlling resources within the EcoBlock to match the phasors before closing the switch.
- Any faults occurring in the EcoBlock must be immediately isolated from energy sources within the EcoBlock (PV, flywheel, and any EVs having vehicle-to-grid capability), as well as from the main grid. If a fault occurs on the main grid upstream of the EcoBlock, the EcoBlock must island to avoid energizing the fault. However, if a disturbance occurs that is not a fault, for example a voltage sag due to a momentary load increase on the grid, EcoBlocks can help to remedy the issue by supplying power and should therefore *not* island.




Figure 1: Power distribution scenarios under consideration for the EcoBlock. Scenarios are: (1) standard PV and storage at home scale; (2) community DC collection, distribution, and storage at block scale; and (3) community DC collection, distribution, and storage at block scale with home DC loads. Credit: Integral Group, Inc.

There are several factors in and around the EcoBlocks that will collectively determine their impact. These factors are: installed assets and their capabilities; operating strategies of the EcoBlock itself; the grid operator’s prior knowledge of the EcoBlock operating strategy; the grid operator’s ability to influence the operating strategy or exercise direct control; and the tariff structure. This section will outline design choices for these factors that could meet the requirements for positive grid impact.

2.1 Installed assets

As illustrated in Figure 1, three power distribution scenarios are considered:

- **Scenario 1: Standard PV and Storage at Home Scale.** Each home has its own PV and energy storage installation. Residents’ AC loads are powered by these devices (through inverters located at each home) and by the utility grid. The entire system can be islanded from the grid and the homes have the ability to share power with each other.
- **Scenario 2: Community DC Collection, Distribution, and Storage at Block Scale.** The rooftop PV systems connect to a shared DC bus and charge a shared energy storage system. Residents’ AC loads are powered from the DC bus (through inverters located at each home) and by the utility grid. The entire system can be islanded from the grid.
- **Scenario 3: Community DC Collection, Distribution, and Storage at Block Scale with Home DC Loads.** The rooftop PV systems connect to a shared DC bus and charge a shared energy storage system. Both DC and AC power are distributed to homes, which have a combination of AC and DC loads. A central inverter allows for power exchange between the DC bus and the utility grid. Only the DC bus stays energized when the system islands.

The hardware requirements for the EcoBlock to perform the desired functions include some requirements that are common to all three scenarios and some that depend on the scenario. These requirements are listed in Table 6. While additional hardware will be required depending on DC voltage levels, EV charging needs, demand response capabilities and other design choices, this discussion is limited to the requirements for desired grid impacts and general resilience.

Table 6: EcoBlock hardware and important characteristics. For characteristics specific to a particular power distribution scenario, the scenario number is given in parentheses.

| Asset | Important Characteristics |
|--|--|
| PV systems | Located on all feasible rooftops. Size for net-zero energy with expected energy efficiency retrofits and full EV deployment. |
| Energy storage | Located at each house (1) or central (2,3). Size based on duration of islanding that is desired (see section 5). |
| Inverter | One per home (1,2) or per EcoBlock (3). Size based on each customer's generation (1), each customer's load (2), or total EcoBlock generation + storage power rating (3). Bidirectional (3). Should allow control of power factor. |
| DC charge controller | One per home (1) or per EcoBlock (2,3). Sufficiently fast response time to maintain system stability while islanded. |
| Energy manager | Master device that optimizes power injections and extractions of controllable components in EcoBlock (see section 2.2). Could be combined with charge controller (2,3), or could be separate entity that sends commands to charge controller. |
| Cables within EcoBlock | Underground, for resilience against weather events and vegetation. Should be sized to support anticipated EV charging loads. |
| Service transformers | If the PCC is upstream of the transformers, they should serve only EcoBlock customers so that only EcoBlock customers participate in islanding. |
| Switchgear | Either AC (1,2) or DC (3). If AC, must be able to communicate with resources in the EcoBlock to match phasors for resynchronization after islanding (1,2). |
| AC sensor (μ PMU) | At PCC. If on EcoBlock side of switch, can be used to match phasor to rest of feeder for resynchronization (1,2). |
| Other meters | Measure power flows and voltages within EcoBlock, including PV generation, storage charge/discharge rate, and power exchanged with grid. Some of these may already be integrated into other components such as inverter or charge controller. Will be AC or DC in different places depending on the scenario. Should also include meters for utility billing purposes. |
| Real-time communication channel with grid operator | Send data to grid operator and receive signals for desired EcoBlock behavior (see section 2.3). |
| Protection devices at individual homes | Interrupt fault currents and isolate faults at homes from rest of EcoBlock. AC (1) or both AC and DC (2,3). |
| Protection devices for entire EcoBlock | Interrupt fault currents and isolate faults on EcoBlock from generators, storage and rest of grid. AC (1) or both AC and DC (2,3). |

The choice of AC or DC power distribution, or both, has many important implications for the management of individual EcoBlocks as well as their relation to the wider grid ecosystem. Protection is more straightforward for AC systems than for DC, since current can be interrupted easily at the zero crossing that occurs twice per AC cycle. On the other hand, DC power islands are easier to manage than AC because they do not require frequency regulation or reactive power balancing, and because they do not

need to be synchronized with the main grid before re-connecting. If large amounts of load remain connected to the grid when the EcoBlock islands, as would be the case in scenario 3 if most load remains on the AC system, then present-day restoration problems such as cold load pickup may continue to exist. Considering these tradeoffs, as well as the efficiency benefit of DC power for many loads, the best option in the long term may be to choose scenario 3 and incentivize as many loads as possible to migrate to the DC system—relying on continued innovation and commercialization of DC protection devices.

The choice of AC vs. DC is also expected to affect the rates and impacts of EcoBlock scaling. Since AC equipment and loads are more developed technologies than their DC counterparts, the retrofit process would be relatively straightforward and inexpensive, leaving more resources available to accelerate scaling at the early stages. Deployment of DC EcoBlocks would likely take longer to reach the same penetration level due to the increased effort and cost required to retrofit each EcoBlock. However, once large-scale deployment has been achieved, the benefits of DC would become more evident. The reduction in power flows due to higher-efficiency DC appliances, negligible at low penetration, would scale roughly proportionally to the number of EcoBlocks. In addition to facilitating re-connection to the main grid after islanding, the choice of DC could also enable entirely new strategies for managing regions saturated with EcoBlocks. For example, EcoBlocks could connect to each other via a DC link, again avoiding the requirements of synchronization, balancing reactive power and frequency. While we show in Section 5 that connecting two already self-sufficient blocks together does not further improve self-sufficiency significantly, this approach *could* be beneficial in other scenarios, for example if it is more economical to have a large energy storage system shared by several blocks than to design each block for complete self-sufficiency.

This same strategy could even be extended beyond the few-block scale. DC is advantageous for power transmission over long distances due to the lack of reactive power (which reduces losses), reduced capacitive losses in underground and underwater cables, and the lack of phase angle stability concerns. One can envision a scenario where entire power systems—transmission, distribution, and EcoBlocks—are converted to DC simultaneously, improving efficiency and stability across the board and avoiding the need for AC-DC conversion. However, this vision would require a complete overhaul of utility infrastructure and operations, and is unlikely to occur in the near future. Therefore, it is essential for EcoBlocks to be able to participate in the AC grid of today even if their main internal power distribution system is DC.

One way in which EcoBlocks could assist the present-day AC grid is through phase balancing, which is expected to become a more important concern in the near future, as PV systems and EVs may be connected in unequal numbers to the three phases. We expect most EcoBlocks in urban areas to have access to three-phase connections. The balanced three-phase central inverter in scenario 3 would ensure that the installed PV generation does not worsen phase imbalance. The greater the fraction of loads—particularly EV chargers—that are on the DC side, the more balanced the EcoBlock would be. In scenario 1, where there is no DC system, the EV chargers could be connected to each AC phase in equal numbers to improve balancing. In all cases, the EcoBlock construction process would provide an opportunity to determine the homes’ phase connectivity and change it if needed. Demand response (and storage, in scenario 1) at individual homes could also be leveraged for dynamic phase balancing, although it would not be fair to count on these resources alone to correct systemic imbalances. If only a single phase is available for the EcoBlock, as in more rural areas, the design could be modified to use a single-phase inverter. The impact of single-phase EcoBlocks on phase balancing requires further evaluation, but in general, deployment of EcoBlocks at the same rate on all phases is likely to be an appropriate strategy.

A final characteristic of scenario 3 worth mentioning is that its design with a central energy storage system and inverter takes greatest advantage of economies of scale. However, these components also represent single points of failure and may limit the reliability of the system. The motivation for deploying EcoBlocks is greatly reduced if they are less reliable than the main grid or have greater maintenance requirements. It may therefore be worth considering investing in redundant infrastructure such as a second inverter. Since it consists of several smaller flywheels in parallel, the flywheel system has some redundancy built in already; the system for managing the flywheels should support an operation mode with one or more flywheels out of service.

2.2 EcoBlock operating strategies

The operating strategy of the EcoBlock should include the following functions:

1. Dynamic power balancing (and reactive power balancing, if applicable) to maintain stability, particularly while islanded
2. Longer-term energy balancing to maintain a reasonable reserve of stored energy
3. Further optimization of power flows, for example to minimize cost, losses, peak load, or power exchanged with the main grid
4. Connection and disconnection from the main grid at appropriate times and without causing major disturbances to either the grid or the EcoBlock

There are a number of possible strategies that could be used to control the hardware listed in the previous section in order to meet these objectives. In centralized control, a single controller computes a target set of conditions (for example, the power flows in or out of each resource) that satisfy some criteria (for example, minimizing losses while maintaining voltage within certain bounds). The central controller then sends signals to each resource, causing the resources to behave in the desired way. This approach is most effective when the central controller has an accurate model of the system, the controllable resources and their constraints, so that optimal behavior can be achieved with minimum iteration. Central control is well-suited for the long-term energy balancing and optimization functions of the EcoBlock (functions 2 and 3) because it provides a single optimal snapshot of the system that every resource agrees upon. It is also appropriate for resynchronization of AC islands (function 4), since the resources in the EcoBlock must be collectively controlled to match the voltage phasors on both sides of the PCC before closing the switch.

However, central control requires communication with every resource and the optimization may be computationally intensive, particularly for systems with many resources and constraints. In topology scenario 1, for example, optimizing the charge and discharge of the energy storage systems located at every home to balance house-level and block-level objectives may be challenging. When a very fast response time is needed, such as for function 1, a decentralized control strategy is preferable. In this case, a resource able to perform decentralized control would immediately adjust its power according to a local control curve: for example, increasing power output if it observes that voltage has gotten too low. If the resources only have local information, though, optimizing their collective behavior for the other three functions is not possible. Therefore, the ideal approach is likely a combination of local control for stability plus some sort of centralized control for optimality [22], [23].

For the centralized layer, there are multiple options for the type of signal the central energy manager sends to its resources. The energy manager may simply command each resource to produce or consume the particular amount of power that is optimal at each point in time. Or, it may send a signal that implies a certain behavior is preferable. For example, in a dynamic pricing scheme, the energy manager would broadcast a price based on present energy availability, and each entity in the EcoBlock would decide how much power to buy or sell. In the ARDA DC microgrid design [22], the energy manager regulates the DC voltage as a control signal. Increasing the voltage signifies that energy is abundant and consumption is encouraged, while decreasing the voltage signifies the opposite. Devices in the system have their own internal algorithms to automate their responses to the voltage signal. The voltage- and price-signaling approaches give each resource more freedom to determine its individual behavior, compared to the command-and-control approach. However, for the energy manager to achieve a particular deterministic outcome with these approaches—for example when performing resynchronization—it either needs accurate models of each resource’s behavior, or must iterate after observing the responses to the original signal. There should be an additional mechanism in place to prevent individual actors from doing something detrimental to the group (for example, consuming an unfair share of power in islanded mode when the EcoBlock is at risk of running out of energy.)

2.3 EcoBlock-grid interaction

The potential strategies for managing groups of EcoBlocks on the grid mirror the options for managing the resources in a single EcoBlock. Architectures may be central or distributed, or some hybrid of the two, and control strategies may be direct or indirect. A fundamental question to answer in the design of the EcoBlock-grid interaction is to what degree the grid operator should be able to influence EcoBlock behavior. On one extreme, the grid operator could view the EcoBlocks as fully controllable resources and command them to have a particular net power at each point in time. This approach would ensure the grid operator’s objectives are met, but at the expense of optimality for the EcoBlocks. It might not even be possible for EcoBlocks to provide the requested net power at times, depending on their internal generation, load, and storage state of charge. On the other extreme, the EcoBlocks could be allowed to behave however they choose, similar to traditional energy customers today. This would place great responsibility on the utility to ensure grid infrastructure—including utility-scale generation—is sized appropriately to support the range of EcoBlock behaviors. Unless strict limitations are placed on EcoBlock behavior, this approach would only be acceptable for the utility up to a certain hosting capacity for EcoBlocks, above which infrastructure upgrades would be required.

The ideal approach is surely somewhere between these two extremes. The EcoBlocks should make at least a few guarantees to be good citizens on the grid; these include rules for islanding and re-connecting at appropriate times only, and limits on power to avoid constraint violations. Once these rules are set, several options remain for balancing the benefits to EcoBlocks and the grid [8]. One possibility is for each EcoBlock to offer the amount of power it prefers to inject or consume from the grid, or multiple possible amounts with different associated prices. Then, similar to today’s energy markets, the grid operator would optimize cost across the offers and communicate the dispatch to the EcoBlocks. Since this option requires bidirectional communication with all EcoBlocks and optimization on the part of the grid operator, it is not well suited to short timescales, and scalability to large numbers of EcoBlocks is limited.

Alternatively, EcoBlocks could offer to provide flexibility services at times when it is beneficial for them. During those times, the grid operator would be able to exercise direct control over them, at least within the range of power injection specified in the offer. EcoBlocks not contracted to provide flexibility at those times would be allowed to operate freely. This approach could also provide a solution to the challenge of maintaining sufficient revenue for the utility when most customers are consuming near net-zero energy. Instead of paying or receiving a rate for each kilowatt-hour they consume or produce, each customer would pay a fixed amount to the utility for the grid infrastructure they rely upon. Then, those that offer flexibility services would receive compensation based on the level of flexibility they offer (regardless of the net energy they end up providing to the grid). The exact amounts for the infrastructure charges and flexibility tariffs would have to be optimized to ensure that the utility can afford to maintain its infrastructure and that there is enough flexibility at all times. It may be challenging for the EcoBlocks to be financially viable with this infrastructure fee model, compared to today’s net metering model where customers with self-generation are effectively subsidized. Scaling up the penetration of EcoBlocks would increase the number of variables in the optimization but, importantly, also increase the amount of flexibility available to make this approach possible.

A hierarchical control and communication architecture could facilitate grid-EcoBlock interactions in the high-penetration scenario. Distribution feeders typically serve about 1000 customers each, which would translate to about 30 EcoBlocks per feeder in the 100% EcoBlock penetration limit (assuming roughly 30 customers per EcoBlock). Monitoring every EcoBlock would be unmanageable for distribution operators, who currently monitor only about 5 or 6 signals per feeder. Therefore, aggregating the EcoBlocks into groups of 5 or 6 would be preferred for maintaining a constant level of complexity in distribution operations. The grid operator (and automated distribution management systems) would communicate with this intermediate aggregation level, which would in turn communicate with the individual EcoBlocks in its jurisdiction. Each aggregator should represent a set of EcoBlocks that are geographical and electrical neighbors so that there is a clear point differentiating the aggregator’s jurisdiction (downstream) from the grid operator’s (upstream).

Recent advances in distribution synchrophasor technology enable a new paradigm in distributed control, in which resources are controlled to track a target voltage phasor rather than a particular power injection [19]. In this model a supervisory controller (grid operator) periodically performs an optimization to determine target phasors for various nodes in the network, then communicates those targets to local controllers, which in turn leverage their local resources to meet the targets. In the EcoBlock case, the local controllers could be the EcoBlock energy managers themselves (in which case the phasor targets would be set at the PCC of each EcoBlock), or there could be an intermediate level of aggregators between the supervisory controller and EcoBlocks, depending on the complexity of the system. There are several potential advantages to this approach. Phasor-based control is uniquely suited to adaptive islanding because tracking a target voltage phasor is precisely what must happen during resynchronization—specifically, matching the phasors on both sides of the switch to be closed. Framing the control objective as a voltage phasor inherently avoids voltage constraint violations and other voltage problems. This approach also promises stability: if something unexpected occurs, for example an EcoBlock disconnects suddenly, the voltage phasors in that region will deviate from their targets and the local controller will immediately know to take action.

Finally, rather than controlling the physical variables of voltage and power, a dynamic pricing scheme could be implemented to incentivize EcoBlock behavior. Here, the grid operator would compute the optimal energy price to offer each EcoBlock at each point in time and broadcast that information; each EcoBlock would then perform a local optimization to determine how much power it should inject or extract from the grid given that price. Compared to options that require negotiation, dynamic pricing should be faster. On the other hand, it is likely to be slower than direct control because at every time step the grid operator must convert from physical variables to prices, then the EcoBlock must convert from price back to its own physical variables. Dynamic pricing may circumvent one of the challenges with direct control: the possibility that the target set by the grid operator is undesirable or even impossible for the EcoBlock. With dynamic pricing, the EcoBlock would always have the option of paying more for a certain amount of power if the benefit is sufficient. However, there is the challenge of how exactly the grid operator should decide the prices. If too many EcoBlocks are offered the same price, there is a risk that they will all behave the same way, reducing diversity and potentially causing new problems. For example, if prices are increased due to a shortage of generation, a large number of EcoBlocks may choose to inject power or reduce demand, turning the supply shortage into an excess. To avoid this, the price would need a high granularity with respect to location. However, excessively high locational granularity could lead to equity issues, for example if one block is consistently charged higher rates than its neighbor because it is at the end of a long feeder with voltage problems.

While challenging, designing the framework for EcoBlock-grid interaction is essential to enable an efficient, resilient and cost-effective scenario with high EcoBlock penetration. We add a few final comments on additional constraints and possibilities.

- The accuracy of the grid operator’s (or aggregator’s) optimization is highly dependent on the accuracy and completeness of the information they have. This information should include models of the grid infrastructure, or at least the parts of it in their control domain, including connectivity, impedances, and the control schemes of any automated regulation equipment present. It should also include ampacity limits of the conductors and other equipment; voltage constraints; and limits on real and reactive power, rate of change of power, and energy of the controllable DERs. Measurements of physical variables on the network are also important, to provide real-time feedback and to support offline analysis.
- If EcoBlocks are given the freedom to decide their net power, there must be some backup plan (such as automated control, or an option for the grid operator to exercise control) to be used in case of emergency.
- To facilitate interactions between EcoBlocks and the grid, in both planning and operations stages, a third party company (perhaps a spinoff of existing utility companies) could manage the EcoBlocks.

Similar to PV designers and financiers today, this company could work to ensure that plans and real-time operating strategies are favorable for both EcoBlocks and the grid.

3 EcoBlock scaling

As we have discussed throughout the previous sections, scaling the penetration of EcoBlocks from 0 to 100% represents a paradigm shift along multiple axes of power system operation, spanning both normal and abnormal conditions. Fundamentally, the transition to full deployment of EcoBlocks is from a centralized architecture to a decentralized one. Today, the grid and its utility-scale generators are by far the dominant source of power for customers, and in most places still the only source of power. In the high-penetration limit, they will become secondary, relied upon to supplement the EcoBlocks' individual sources only when it is necessary or economical. At the same time, the responsibilities of the EcoBlocks must increase, as they become more and more capable of impacting conditions on the grid.

With the large number of possible impacts, and presumably different deployment strategies in different regions, EcoBlock scaling will likely be a continuous transformation. However, there are a few specific penetration thresholds at which qualitative changes in EcoBlock impacts and capabilities begin to take place. A few of these changes are discussed here, along with the approximate number of EcoBlocks per distribution feeder at which they may first occur. Impacts on the distribution level are expected to occur first: with a lower overall power flow compared to the transmission system, the power injection of each EcoBlock is more significant. Also, impacts such as reverse power flow may require new equipment and operating strategies on distribution circuits designed for unidirectional power flow, while on transmission systems, bidirectional power flow is already expected.

It is important to distinguish between the penetration level at which EcoBlocks become *capable* of producing a certain impact and the level at which they are *likely* to have that impact. This discussion focuses on thresholds of capability because grid operators need to understand and prepare for the range of possible operating conditions, both normal and abnormal. Furthermore, at what penetration level an impact becomes likely to occur depends on the specific operating strategies, which have yet to be fully decided. Awareness of these capability thresholds should help the EcoBlock designer create a plan that appropriately synchronizes the construction of EcoBlocks themselves with the deployment of systems for EcoBlock management.

3.1 Distribution power flow thresholds

The first threshold where EcoBlocks may begin to have impact is when their aggregate load surpasses the pre-existing level of load variability on a distribution feeder. At this point, their behavior becomes statistically observable. To establish this noise floor, the standard deviation of current at the feeder head has been quantified for a few feeders in previous studies with μ PMUs. For example, for a particular 12 kV line-to-line residential feeder in the southern United States, the standard deviation of current taken in 10-minute windows is typically between 2 and 6 A per phase (depending on the day and time), or about 0.5-2% of the mean.¹ For a small 12 kV primarily commercial feeder in the western United States, the typical variation is between 1 and 3 A per phase, or 1-3% of the mean, while for a larger feeder in the same region, it is 4-7 A or 2-3%. Under a 1% noise level, for example, a change in power as low as 27 kW per phase on the southern feeder, and 8 and 19 kW per phase respectively on the two western feeders, could be observable. The significance of this result is that if the average load of an EcoBlock is about 40 kW (as expected for the Oakland EcoBlock after the energy efficiency retrofits—not including generation), then even a single EcoBlock transitioning from zero net load (or islanded state) to average load could be

¹This result was obtained by computing the mean and standard deviation of the current magnitude data stream from a μ PMU with 100-millisecond resolution, at 10-minute intervals, for three nonconsecutive weeks spanning different seasons. The middle 80% of the 10-minute chunks were considered the “typical” range of current variation. Major events were excluded from the analysis.




Figure 2: Example of μ PMU current magnitude measurements on one phase of a typical 12kV residential distribution feeder for a spring day, to indicate magnitude and variability. Inset: half an hour of data at a higher zoom level, showing minute-to-minute variation. Magnitudes of the Oakland EcoBlock’s PV generation capacity, flywheel charge/discharge rate, and estimated average load are shown on the same scales, for comparison with feeder load current and variability.

observable, depending on whether it has a single-phase or three-phase connection. However, the 10-minute time window was intentionally chosen to capture the effects of individual loads while neglecting the effect of time of day. The change in load from off-peak to peak times of day on these feeders is about an order of magnitude greater than the standard deviation in the 10-minute windows. Therefore, while a major shift in power at a single EcoBlock may be observable compared to the noise, it falls well within the normal daily range of load and is unlikely to impact operations. Figure 2 provides a visual indication of how some possible “signals” due to EcoBlock behavior would compare to the “noise” of existing feeder load variability. The figure shows an example current magnitude trace for one of the feeders, measured by a μ PMU, to demonstrate the typical magnitudes and variability over the course of a day and from minute to minute. The current levels corresponding to certain EcoBlock characteristics (PV generation capacity, maximum flywheel charge/discharge rate, and average load) are compared on the same scales.

The second threshold is where there are enough EcoBlocks that their collective behavior could trigger a control action, either taken by a distribution operator or automatically. For example, a sufficiently large change in current would trigger an automatic tap change operation under a line-drop compensation scheme. This threshold is highly dependent on the feeder impedance characteristics, which determine how sensitive the voltage is to changes in load, and on the specifics of the line-drop compensation scheme. In one example, a change in current of 70 A per phase (on an 11kV line-to-line feeder) triggers a tap change operation [8]. This magnitude of current change is equivalent to, for example, the 200-kW flywheels in two EcoBlocks transitioning from zero to full charge or discharge rate. This is the threshold where the operating strategy of the EcoBlock begins to influence grid operations directly. If the EcoBlock operating strategy allows the power injection to fluctuate frequently over a large range, then at this penetration level it may have adverse effects such as accelerated wear on voltage regulation equipment. On the other hand, if the EcoBlocks manage their power injection, and infrastructure for grid interactions is already in place, EcoBlocks can avoid these adverse effects and may begin to counteract the impacts of other variable generators or loads on the grid.

Third, when the maximum power injection from EcoBlocks surpasses the minimum load on a feeder, reverse power flow becomes possible. The maximum possible power injection at midday is equal to the sum of the PV generation capacity and the energy storage maximum discharge rate; at night it is just the energy storage maximum discharge rate. The threshold at which EcoBlocks are likely to enable reverse power flow therefore depends not only on the minimum load itself but also at what time of day it occurs. For the prototype Oakland EcoBlock, we assume the maximum PV generation is 360 kW, the maximum discharge rate of the flywheel system is 200 kW, and the minimum load in the EcoBlock is small by comparison,

and compare to the daytime load of the three reference feeders analyzed above. For the residential feeder, the minimum daytime load is approximately 2 MW per phase. If 4 EcoBlocks per phase were deployed on this feeder, and both their PV generation and flywheel discharge rate were at maximum, they could cause reverse power flow. Similarly for the commercial feeders, under the same conditions, just one EcoBlock per phase on the smaller feeder and 3 on the larger feeder could cause reverse power flow. Today, reverse power flow is often considered a problem (for protection coordination and voltage regulation), but it can also be viewed from a much more positive standpoint. When EcoBlocks have the potential to cause reverse power flow, it signifies that they are capable of supporting the other loads on the feeder if it is necessary or economic, and they are crossing the threshold into being able to support other parts of the grid as well.

3.2 Topological thresholds

While the impacts above relate to the number of EcoBlocks on a feeder and are relatively independent of topology, there are a few additional impacts that do depend on topology and where the EcoBlocks are located in a system. Different strategies for EcoBlock deployment may make sense depending on what the priorities are in the particular system. For example, customers at the end of a feeder have the greatest impact on voltage because power flowing between them and the substation travels along the entire length of the feeder, causing the greatest voltage drop. Therefore, for feeders with pre-existing voltage problems, the ideal EcoBlock deployment strategy may be to prioritize locations at the end of the feeder, and manage the EcoBlock net power in order to regulate voltage. This strategy would mitigate the voltage problems as quickly as possible.

Often, feeders or sections thereof can be completely isolated from the rest of the grid by switches or breakers. A section entirely made up of EcoBlocks could potentially be operated as an island in the event of a problem on the main grid. Allowing EcoBlocks to share power could help overcome generation-load mismatches on individual blocks, improving resilience compared to the case where each block operates independently. This would be especially advantageous in long-term grid outage scenarios and if generation and load are highly variable on some of the EcoBlocks. While a section containing non-EcoBlock customers could also be islanded in this way, the self-sufficiency of the island would be reduced if they do not contribute generation, and managing both EcoBlock and non-EcoBlock customers in the same island might be challenging. Therefore, in areas that are at high risk of outages—for example, if they interface with the main grid through a small number of low-reliability connections—the ideal deployment strategy may be to fully populate islandable sections with EcoBlocks one at a time, before moving on to lower-risk areas.

Although sensing infrastructure can of course be deployed independently of EcoBlocks, system observability could be an additional benefit of high-penetration EcoBlock deployment. Assuming that each EcoBlock is accompanied by a monitoring device such as a μ PMU at its point of common coupling, there will be a threshold at which the penetration of monitoring devices affords full observability of the system. Full observability requires a large fraction of the nodes in a system to be monitored; the exact fraction of nodes needed is dependent on the system topology. Optimal PMU placement algorithms have been developed to enable observability of distribution systems with the minimum number of sensors [24].

3.3 Beyond the distribution level

A threshold for utility financial operations will occur when there are enough near-net-zero energy EcoBlocks (or other “prosumers” with significant DERs) in a system that traditional energy rate structures are no longer feasible to cover grid infrastructure maintenance costs. At this point, utilities should already have a new business model in place to avoid the so-called “death spiral”. Similarly, as the reliance on traditional power plants is reduced, there will be a point for each plant where the capacity factor and revenue are low enough compared to the maintenance costs that it is no longer economical to keep the plant in operation. Generators with the lowest costs and greatest flexibility are likely to remain in operation the longest.

The penetration levels at which these thresholds occur depend on the physical assets and structure of the network, as well as the existing market and rate structures, and are outside the scope of this report.

3.4 Significance of the block scale

The discussion above focuses on the impacts of scaling the penetration level of EcoBlocks given that the city-block size has already been selected. The block scale is certainly an appropriate choice from the construction and internal management perspectives, as it offers a convenient physical location for shared resources, including electric infrastructure (e.g. energy storage systems, inverters, electric vehicle chargers), water reclamation systems, and even food systems (gardens, composting). However, it is also important to consider explicitly other ways the grid could be partitioned, besides the block scale, and compare their benefits and downsides to those of EcoBlocks.

On the smallest end of the scale, each individual home could act as a microgrid, with control over its own energy resources and islanding capability. However, the high variability of power consumption in an individual home means that a comparatively larger amount of energy storage is required to achieve the same level of self-sufficiency while islanded, relative to a scenario where the homes are allowed to share power [25]. The scalability of this model is also limited, as managing a large number of home energy systems to achieve distribution-level benefits such as voltage regulation would be computationally and communication intensive.

Beyond the customer meter, the smallest electrical unit of aggregation that can be readily connected or disconnected is a lateral or section of a distribution feeder. But these separation points are mainly intended for protection, not active switching, and may not be equipped with remote sensing and control; they also may not have loads balanced across three phases. Therefore, the smallest operationally relevant unit for separation from the grid would typically be an entire distribution feeder, which can be easily disconnected at the substation. A key problem with the feeder scale is that a feeder encompasses a much larger number of customers (hundreds or even thousands), making it more difficult to coordinate generation and voluntary load curtailment (not to mention liability for power quality issues).

Block-scale microgrids, by contrast, are small enough for internal administration, while also having a point of common coupling that provides a safe and convenient transfer mechanism to disconnect and re-connect their aggregated loads and generation to the main grid. While a block consists of far fewer customers than a feeder, it is nevertheless sufficient to provide the benefits of statistically aggregating variable loads and generation, as will be demonstrated in Section 5. Furthermore, particularly in suburban and lower-density urban regions, distribution feeders often include overhead lines that are susceptible to weather- and vegetation-related damage. If these lines are damaged, a collection of independent EcoBlocks (especially with underground internal power distribution) could continue to be energized while a single feeder-level power island could not. In these ways, EcoBlocks would introduce operational flexibility to distribution systems at an intermediate scale, between individual home and entire feeder.

One notable challenge with the block scale arises from the frequent disconnect between the electrical and spatial layouts of distribution systems. In particular, customers on the same city block may be connected to different feeders, and service transformers may supply customers on more than one block. Changing the connectivity, reconfiguring or replacing utility equipment may therefore be necessary to create block-scale electrical units that connect to the grid at a single point of common coupling. It may be helpful to include these modifications in the EcoBlock development plan and budget so that supporting EcoBlock construction is not cost prohibitive for the utility.

4 Summary of paradigm shifts

In the preceding sections, we detailed the potential impacts of EcoBlocks on the grid, discussed the conditions under which they are likely to occur, and evaluated design choices from the standpoint of balancing benefits to EcoBlock participants and the grid. Here, as a form of summary, we review the three major

paradigm shifts in the nature of the power system that we expect to occur at high penetration levels of EcoBlocks.

4.1 Power flows

The architecture of today’s power grid is designed to transfer power from large generation facilities to consumers in different locations. As a result, conductors and other grid infrastructure consistently carry large power flows. At the distribution level in particular, protection and voltage regulation are designed assuming unidirectional power flow, a requirement limiting the permissible DG penetration. Since most customers consume more energy than they produce, pricing electricity by the kWh generates revenue for the grid operators to maintain the infrastructure.

With high penetration of EcoBlocks, on the other hand, distributed generation will be ubiquitous and co-located with loads. Customers’ consumption and production will differ on short time scales, but will be comparable in the long term, leading to near-zero net kWh consumption. While the grid infrastructure will be used less frequently and with reduced loading, it will likely still be relied upon at times to provide power from distant, low-cost bulk energy sources, and to facilitate power sharing between EcoBlocks. Maintenance needs will therefore be reduced but not disappear entirely. Protection and voltage regulation may need to be updated in many places to support bidirectional power flows. Pricing by the kWh will likely no longer be sufficient to finance the maintenance, necessitating a new market structure.

4.2 Flexibility

Today, resources that are flexible, such as energy storage and interruptible loads, represent a relatively small portion of the total power being exchanged on the grid. Due to their special capabilities, these resources may be recruited to solve problems such as supply-demand imbalance (due to the large amount of inflexible demand), thermal and voltage constraint violations, and emergencies.

In the new paradigm, by contrast, most of the participants in the grid will be flexible. There should no longer be concern about the supply adequacy of flexible resources, since the remaining inflexible demand will be small by comparison. The presence of a much greater number of resources to choose from will likely enable cost reductions, as well as new optimization objectives such as dispatching the resources geographically closest to the site of the problem. At the same time, a new challenge will emerge around optimizing the dispatch of the greatly increased number of resources and communicating with them in a timely manner.

4.3 Resilience

Today, when grid infrastructure goes out of service (for example, when power lines are damaged in a storm), customers in the area experience outages—particularly those downstream of a failed component. Since generation and load are geographically separate, multiple repairs and a long time may be required before service can be restored after a major event. Even where local generators such as rooftop solar installations are present, they are not allowed to energize the system under these conditions. The restoration process itself is risky and complex, since not all generators have black-start capability, and even a brief supply-demand imbalance during restoration can cause the system to collapse again. Therefore, generators and loads must be brought online gradually and in an appropriate order.

In the future, the deployment of EcoBlocks is expected both to reduce the severity of outages and to facilitate service restoration. With co-located generation, storage and loads, and the ability to disconnect from the main grid, EcoBlocks can power themselves if components in the main grid fail. Since the EcoBlocks by definition exchange zero net power with the grid while islanded, the risk and complexity of system restoration will be greatly reduced compared to the present-day scenario. The rules and processes for islanding, power balancing while islanded, and re-connection to the main grid must be well-designed and robust to ensure this adaptive islanding model is reliable and predictable.

5 Resilience analysis

In 2015, the average PG&E customer experienced about 150 minutes of power outages, and 36,000 customers experienced outages longer than one day due to weather events [26]. Outages frequently result from weather or vegetation damaging overhead conductors, and are worsened when the conductors are in remote locations where locating the fault is difficult. With the increasing occurrence of extreme weather events due to climate change—consider hurricanes Sandy, Harvey, Irma, and Maria, for example—improving grid resilience is a growing priority [27]. The adaptive islanding capabilities of EcoBlocks have great potential to improve resilience: since they can disconnect and self-supply, loss of a component in the transmission or distribution system will not necessarily cause an outage for EcoBlock customers. This is in contrast to today’s concepts of virtual power plants and DER aggregators, which may have similar benefits to EcoBlocks in terms of peak shaving and loss minimization, but do not allow islanding because the DERs and loads are not necessarily geographic or electrical neighbors. Since adaptive islanding for resilience is directly in line with utilities’ goal of keeping the lights on as much as possible, it is perhaps the strongest argument in favor of EcoBlock development. However, it is not guaranteed that all demand in an EcoBlock will be met during an islanding event, since generation, load, and the timing of the event are variable. Therefore, this section attempts to quantify probabilistically the self-sufficiency of EcoBlocks—defined as the fraction of load demand served—and its dependence on EcoBlock design parameters.

5.1 Simulation Procedures and Data

One relevant design parameter is topology, particularly the three topology scenarios considered in Figure 1. In Scenario 3, only the loads connected to the DC system can be served in the islanded state, while in the other two, all loads can be served in the islanded state. This analysis focuses primarily on the case where all loads are present in the island, since that represents an upper bound on demand, but also compares cases representing Scenario 3 where all or only half of the candidate loads are migrated to DC. A second design parameter we evaluate is aggregation level. As an alternative to block-level aggregation, each household could maintain its own generation and storage and operate independently. Due to the variation in individual generation and demand, aggregating multiple homes into a microgrid has been shown to increase the fraction of total demand that can be met [25], [28]. This analysis assesses the value of statistical aggregation in the EcoBlock context, and further evaluates whether connecting multiple EcoBlocks into a larger island offers additional benefits. We further consider factors influencing self-sufficiency during a particular islanding event: conditions at the start of islanding (time of day and storage state of charge); duration of islanding; and generation and load, which depend on weather, time of year and other factors.

Since customer meter data for the pilot EcoBlock location was not available at the time of writing, this analysis was instead performed using residential load and PV generation data from Pecan Street, Inc, located in Austin, Texas [29]. A set of thirty homes with PV was selected from the Pecan Street data set to create a simulated EcoBlock. These homes were selected randomly from the subset of homes that had load and PV generation data for the entire year of 2014, and that had multiple disaggregated load time series corresponding to individual appliances and circuits within the home. This disaggregated load data is necessary to estimate the self-sufficiency of a DC system in which only certain loads participate. The Pecan Street dataset also includes households that have and have not participated in demand response programs. Since we expect the behavior of many, but not all, EcoBlock residents to be similar to those in demand response programs, we introduced a bias in the selection of the homes for the simulated EcoBlock so that 24 of them (80%) participated in the programs. A few key differences between the Pecan Street data and the expected load and generation at the Oakland EcoBlock site should be noted:

- Austin has higher summer temperatures and therefore greater air conditioning usage compared to Oakland.

- Since the Pecan Street PV installations were not designed for self-sufficient microgrid operation, the ratio of load to PV generation differs from the Oakland EcoBlock plans, and energy storage is not present.
- Since the homes differ in size and in number and type of appliances, the peak and average load values in Pecan Street differ from those in Oakland. In particular, the Pecan Street homes in this analysis are single-family homes, while the Oakland EcoBlock includes several multi-family residences.
- While a relatively smaller effect, we note for completeness that Austin is at about 30 degrees north latitude, whereas Oakland is at about 38 degrees. As a result, days are about 40 minutes longer in the summer and 40 minutes shorter in the winter in Oakland compared to Austin, likely making the seasonal differences in generation more extreme.

Despite these differences, there are several reasons justifying the use of Pecan Street data for this analysis. First, while the summer temperatures and air conditioning use in Austin differ from those of Oakland, they are more similar to conditions in parts of Southern California and the Central Valley. Therefore, analysis conducted for Austin would be more directly applicable to these regions, which are also of interest for future EcoBlock deployment. Second, the Pecan Street data offers unparalleled insight into the energy consumption of individual end-use loads as well as the statistical distributions of household load and generation over the course of multiple years. For resilience analysis, which by nature deals with system behavior under unlikely circumstances, data reflecting the variability of individual loads and PV generation is of utmost importance. The analysis framework demonstrated here with Pecan Street data could easily be applied to other locations such as Oakland when sufficient customer meter data becomes available. Finally, the Pecan Street PV generation was scaled so that the ratio of annual generation to annual energy demand would equal that of the Oakland EcoBlock. Additionally, the storage capacity in the simulated EcoBlock was chosen to give the same ratio of storage capacity to annual energy demand as in the Oakland EcoBlock.

At the time of writing, the most recent estimates for the Oakland EcoBlock were: 357,500 kWh of load per year, assuming conversion of all gas appliances to electric plus energy efficiency retrofits; 509,500 kWh of solar generation per year (1.425 times the annual load); and 200 kW or 800 kWh of storage capacity (1/447 times the annual load). The generation and storage in the simulated EcoBlock were scaled to give the same ratios compared to annual load: annual demand was 384,100 kWh, annual generation 541,500 kWh (prior to scaling it was 220,700 kWh), and the storage was sized at 215 kW or 860 kWh. The other storage parameters used in the simulations were based on the Amber Kinetics flywheel currently on the market [30]: charge and discharge efficiencies of 94%, giving a round-trip efficiency of 88%, and a constant self-discharge rate of 65 W per 8 kW of capacity. Since the Oakland EcoBlock design is intended to power streetlights in addition to customer loads, streetlights were added to the block-level demand in this analysis as well. Streetlight power consumption was assumed to be 900 W, based on the 10 streetlights currently installed within the EcoBlock boundaries and an assumed consumption of 90 W per light, the mean of the LED streetlights for residential locations compared in [31]. Streetlights were assumed to be on whenever the generation dropped below 1% of peak (encompassing nighttime hours but not cloudy days), and off otherwise. A very simple demand response scheme was assumed: at each time step, the EcoBlock served as much of the load as possible given the energy available from generation and storage. Any remaining load was recorded as unserved. This is equivalent to assuming that all loads have some sort of demand response capability, allowing an arbitrary fraction of load to be shed at each time step. The ability to shed noncritical loads is essential in order to maximize the power delivered to critical loads: in the absence of demand response, if power demand exceeds supply, the entire microgrid will experience an outage [25], [28].

The flow of the self-sufficiency simulation for the simulated EcoBlock is described below and in Figure 3:

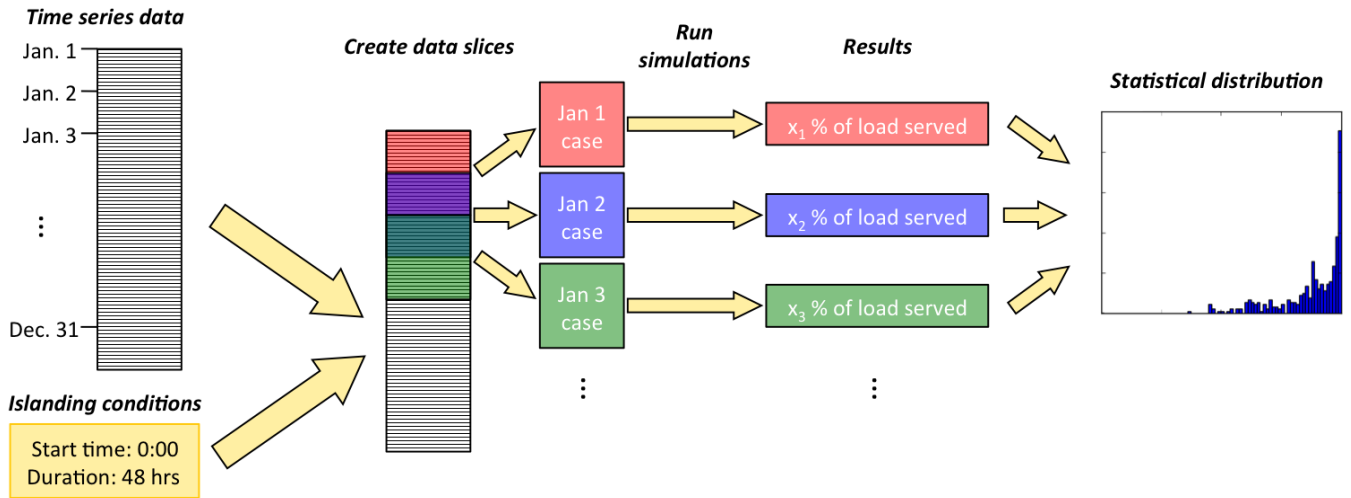


Figure 3: Self-sufficiency simulation process flow.

1. Data preparation
 1. Scale generation and storage to match Oakland EcoBlock ratios
 2. Select load: all loads, all DC-candidate loads, or half of DC-candidate loads
 3. Add streetlight load
 4. Outputs:
 1. Time series of generation E_G and load demand E_L , with time step dt (here, $dt = 1$ hour), for one year
 2. Storage parameters: capacity E_C , maximum charge/discharge rate P_{max} , charge and discharge efficiencies η_C and η_D , self-discharge rate P_{SD}
2. Self-sufficiency simulation
 1. Choose a start time t_{start} , duration T , and initial storage state of charge $E_{S,init}$
 2. Select all slices of generation and load data with that start time and duration (one slice for each day of the year)
 3. For each data slice:
 1. At each time step: compute energy delivered to loads $E_{L,served}$, energy stored E_S , and excess (spilled) generation $E_{G,excess}$; as a function of E_G , E_L , and $E_{S,init}$; within the storage rate and capacity constraints
 2. Calculate total demand, fraction of total demand served, and total excess generation
 4. Analyze the statistical distributions of total demand, fraction of total demand served, and total excess generation over the year

5.2 Results and discussion

Impacts of islanding start time and duration

We first focus on the case where the entire block islands as one—including all loads, both AC and DC—and consider a long-term unplanned islanding scenario. This type of scenario would most likely be the response to a fault on the main grid that would otherwise cause an outage, such as a weather event, component failure, or vegetation-related conductor damage. As a baseline case we assumed the energy storage was 50% charged at the start of islanding, and determined the effects of varying islanding start time and duration on self-sufficiency.

We determined the distributions of self-sufficiency, as a function of islanding duration, for start times of midnight, morning (8:00), and afternoon (16:00). The means of these distributions are plotted in Figure 4. For short islanding periods of just a few hours, the block only needs to rely on the initial energy stored

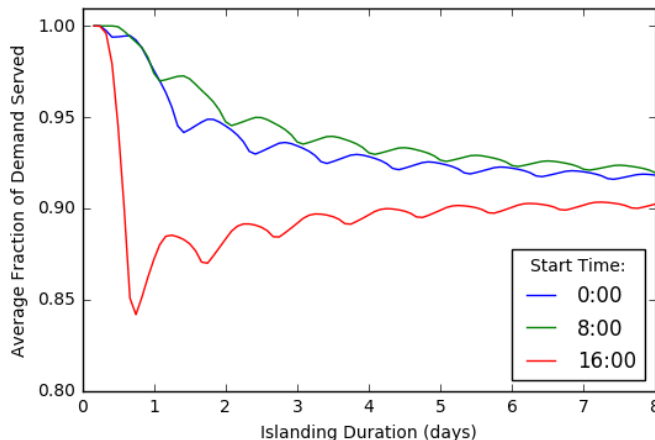


Figure 4: Dependence of the mean fraction of demand served during islanding on the start time and duration of islanding.

in the flywheel, and is therefore highly self-sufficient regardless of start time. For intermediate durations (half a day to about 2 days), the start time has a significant influence on self-sufficiency. In particular, the greatest amount of load has to be shed when islanding begins in the afternoon, because load is near peak at the same time that PV generation is ending for the day. When the islanding duration increases to several days, the effect of the starting conditions diminish, and the steady-state distributions of generation and demand take over. The ripples seen in Figure 4 result from the periodically changing fraction of the islanding event that includes daylight.

The distributions of self-sufficiency that result in these mean values are explored in more detail in Figure 5. Figures 5a-c show histograms of self-sufficiency for the three start times with durations of 1 day and 8 days. Figures 5d-f show scatter plots of self-sufficiency versus total demand for a few notable islanding scenarios, with results from the four seasons indicated in different colors. The histograms show that for most conditions, nearly 100% of load is served in the majority of islanding events, with a few low-self-sufficiency outlier events. As shown in Figures 5d-e, these outliers tend to occur in fall and winter. Although demand is modest during these times, overcast weather is most common, resulting in days of very low generation. Figure 6 shows generation and demand time series for two consecutive winter days, one overcast and one clear. Although demand is similar for the two days, the extreme difference in generation leads to very different self-sufficiency results. If the EcoBlock is islanded for the first day, only 35% of demand can be met, while if it is islanded on the second day, all demand can be met. For the 0:00 and 8:00 start times, Figures 4 and 5 show that the probability of low self-sufficiency increases for long durations of islanding. This results from the fact that the weather on one day affects the stored energy available on subsequent days. In particular, after a cloudy day, the stored energy is low, increasing the need for load shedding on the following day.

If islanding begins in the afternoon, while 100% self-sufficiency is still possible, significant load shedding is more likely regardless of the season, as shown in Figures 5c and e. At this time of day, generation is ending at the same time as load is nearing its peak. Since the storage is only half charged at this time, there is often not enough energy to last through the night. While winter peak load is lower than summer in this climate, generation is also lower and ends earlier in the day, leading to similar overall ranges of self-sufficiency throughout the year. A well-managed EcoBlock would try to have the storage fully charged by the afternoon to avoid this scenario. If that were the case, the self-sufficiency of the EcoBlock would be much greater, both on average and in the worst case, as shown in Table 7. Table 7 also shows that while a greater initial state of charge is advantageous for short-term islanding, the effect is reduced for long-term islanding. A complete analysis would incorporate a statistical distribution of the storage state of charge that depends on the time of day. Deriving this statistical distribution would require knowledge

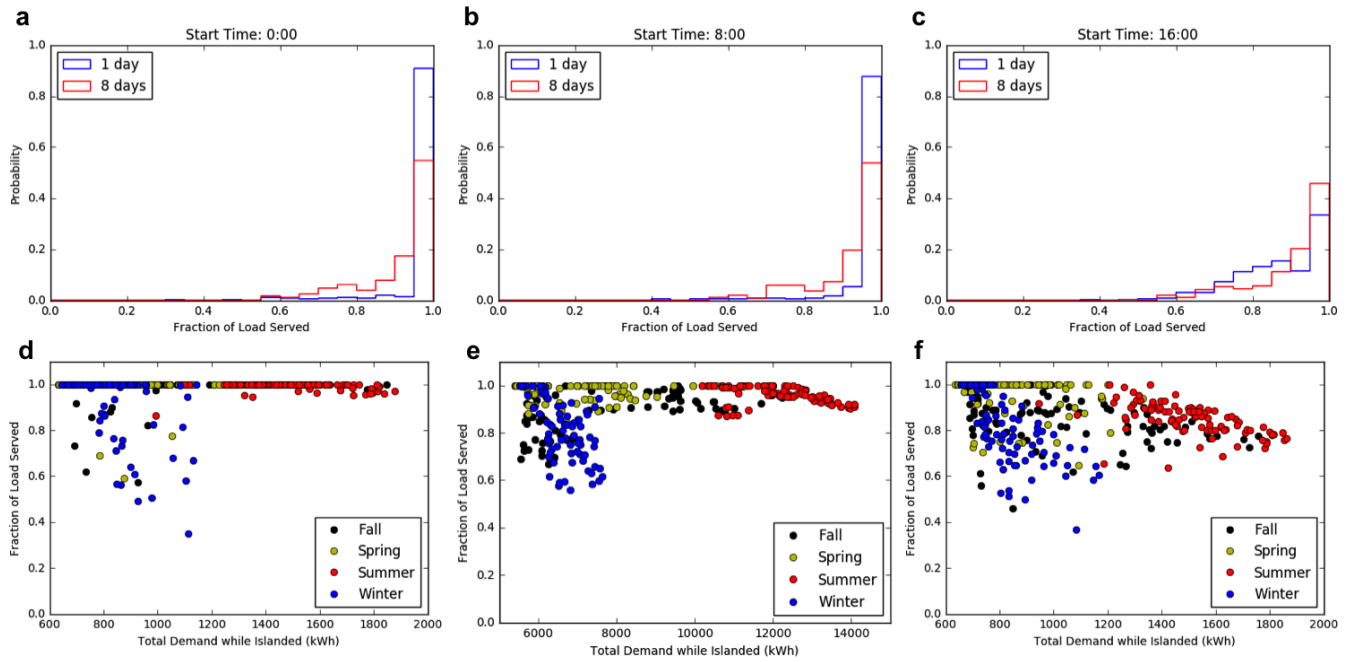


Figure 5: (a-c) Histograms of the fraction of demand served with islanding durations of 1 and 8 days, with start times of 0:00 (a), 8:00 (b), and 16:00 (c). (d-f) Scatter plots of the fraction of demand served versus total demand while islanded for each season of the year. Start times and durations are 0:00 and 1 day (d), 0:00 and 8 days (e), and 16:00 and 1 day (f).

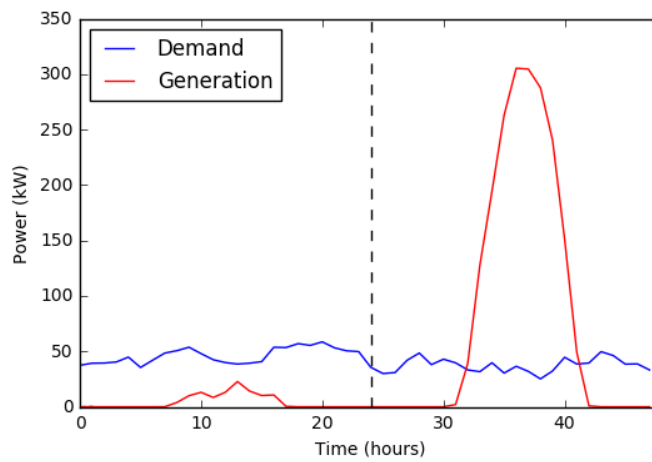


Figure 6: Generation and demand time series for two consecutive winter days, one overcast and one sunny.

Table 7: Mean and minimum self-sufficiency results for various islanding scenarios, comparing 50% and 100% state of charge for the energy storage at the start of islanding.

| Duration | Start time | Self-sufficiency: 50% initial charge | | Self-sufficiency: 100% initial charge | |
|----------|------------|--------------------------------------|------|---------------------------------------|------|
| | | Mean | Min | Mean | Min |
| 1 day | 0:00 | 0.97 | 0.35 | 0.999 | 0.79 |
| 1 day | 8:00 | 0.97 | 0.44 | 0.99 | 0.80 |
| 1 day | 16:00 | 0.87 | 0.37 | 0.99 | 0.72 |
| 8 days | 0:00 | 0.92 | 0.56 | 0.93 | 0.59 |
| 8 days | 16:00 | 0.90 | 0.50 | 0.92 | 0.57 |




Figure 7: Cumulative density function of excess generation for islanding duration of 1 day and start times of 0:00, 8:00, and 16:00.

of the EcoBlock operating strategy while connected to the grid.

In addition to the fraction of demand met, we quantify the excess generation, which is the generation that cannot be stored in the flywheel because the flywheel is already fully charged. About half of the homes in the simulated EcoBlock already have electric vehicles; the excess generation can be considered an estimate of the energy that is available to charge additional EVs without needing to shed any additional loads or install additional power sources. For the excess generation to be useful for this purpose, it is of course necessary for EVs to be present and needing charge at the times the generation is available. Figure 7 shows cumulative density functions of excess generation with an islanding duration of 1 day. For the best case, the 8:00 start time, about half of the days have over 600 kWh/day of excess generation, and 75% of days have at least 220 kWh/day. Considering the 30 kWh battery of the Nissan Leaf [32], for example, this means that seven Nissan Leafs could be fully charged, while islanded, 75% of the time. It should be noted, though, that a significant fraction of the days do not have excess generation (seen in the figure as a y-intercept greater than zero). The EV charging capability could be investigated further given more information on customers' EV usage and constraints. Constraints include the number of miles driven by each EV customer and the frequency that they drive, times of day that the EVs are plugged in to charge, and how critical it is for each EV to be charged. With this information, it would be possible to determine a distribution of the critical EV load that could be served and the amount of other, noncritical load that would need to be shed in the process.

Impacts of topology and aggregation level

Up to this point it has been assumed that all loads will be present on the islanded EcoBlock. However, in topology scenario 3, only loads connected to the DC system would be powered during islanding. The




Figure 8: Mean fraction of demand served during islanding vs. islanding duration, for a start time of 16:00, if the island includes all loads (“AC”), all DC-candidate loads (“DC”), or half of DC-candidate loads (“DC partial”).

main candidate loads to migrate to the DC system are expected to be lighting, electronics, and thermal loads (heating, refrigeration, air conditioning). These loads are likely to be more efficient when powered by DC, since they are inherently DC (in the case of LED lighting and electronics) or in the case of thermal loads, could be replaced by high-efficiency DC heat pump devices. These loads are also generally the most important to keep powered during emergencies. The shared electric vehicle charging stations may also be connected to the DC network. Therefore, we simulated the DC scenario by adding up the disaggregated load data for each home corresponding to EVs, lighting, plug loads, heating and water heating, air conditioning, and refrigeration. Since we expect that not all residents would upgrade to DC appliances immediately, we also modeled a scenario where only half of the homes’ DC-candidate loads (randomly selected) are actually migrated to the DC system. Differences in power consumption between the actual loads and their DC counterparts were neglected. Figure 8 shows that, as expected, including only a subset of loads in the island greatly increases the fraction of that load that can be powered. In fact, self-sufficiency is 100% in the scenario where only half of DC-candidate loads are on the island. The significance of this result is that if fewer loads are able to connect to the islandable circuit in the first place, fewer of them need to have demand response capabilities to ensure self-sufficiency.

We next evaluated the benefits of sharing power and storage resources at the block scale. This was accomplished by comparing the amount of load served at each of the 30 homes, assuming they operate as independent islands, against that of the aggregated EcoBlock. The generation at each home was scaled by the same multiplier so that the *total* generation of the block would be the same as the block-level aggregation case, while preserving the differences in generation-load ratio between homes. This accounts for the fact that homes have differing amounts of viable roof space for PV. The storage sizing at each home followed the same rule as the entire block: 1 kWh of storage for every 447 kWh of annual energy demand. Self-sufficiency in the independent home case was then calculated by dividing the total load served at all of the homes by the total demand. As shown in Figure 9, aggregating the resources offers a significant advantage compared to independent home self-consumption. The fraction of load served in the aggregated block is not only greater than the average of the independent homes; it is greater than the fraction of load served in even the best-performing individual home. This is a similar overall trend to that observed in [25], [28], but with greater overall self-sufficiency due to the larger sizes of generation and storage relative to load.

A major reason for this result is that for individual homes, variation in net load is very high; even turning on or off a single appliance can change a home’s load significantly. That same appliance is insignificant on the scale of a 30-home block. While generation tends to be highly correlated across homes in the

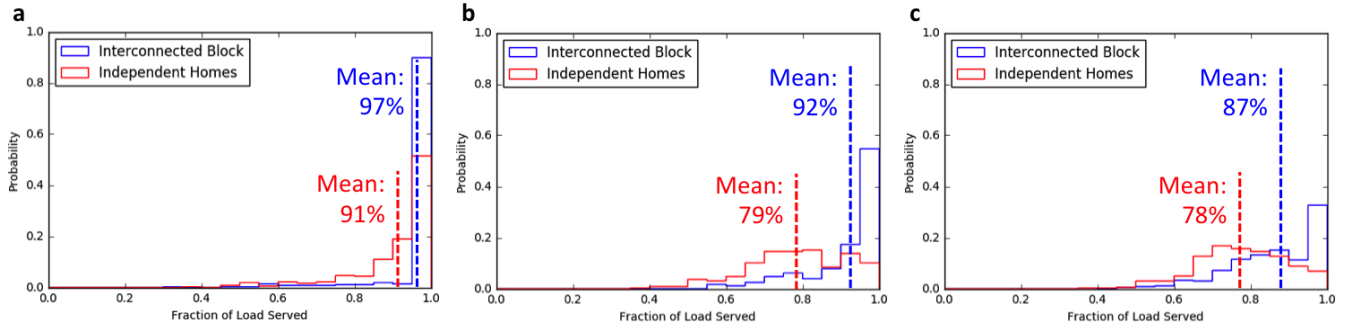


Figure 9: Distribution of fraction of load served for various islanding conditions, if each home operates as its own independent island (red) or if all 30 homes share power and storage resources (blue). Start time and duration of islanding are 0:00 and 1 day (a), 0:00 and 8 days (b), and 16:00 and 1 day (c). The mean of each distribution is labeled with a dotted line.

same neighborhood, load spikes at individual homes are not as well synchronized. As a result, block-scale aggregation reduces variability of net load, thereby reducing the likelihood of running up against storage rate or capacity constraints. For the individual homes, the net load surpassed the storage charge or discharge rate limit an average of 820 hours of the year, and as much as 2100 hours of the year for one of the homes. This was reduced to 410 hours of the year for the aggregated block.

We must recall that this analysis was performed under the assumptions that all loads on the island are of equal priority and can be disconnected at any moment if necessary. With a more intelligent demand management scheme, we expect a similar fraction of load to be served over the course of an islanding event, but with a different distribution in time and with the most critical loads taking priority. While the mean fraction of load served is quite high even for independent homes, the minimum fractions served during the worst-case islanding events are much lower. Of the 8-day events starting at midnight, for example, 56% of load is served in the interconnected block in the worst-case event. During that same event, only 40% of load would be served in the independent home scenario, and for 5 of the homes less than 30% would be served. Under these conditions it is likely that even relatively important loads would need to be shed at some homes. Therefore, aside from increasing the probability of meeting 100% of its load, an EcoBlock’s electrical infrastructure would ease the selective disconnection of non-critical loads, rather than risking a complete outage in case local resources are not adequate. The ability to intelligently adapt loads to available resources would substantially mitigate the impact on customers. Furthermore, block-level aggregation could help ensure that the distribution of energy among residents is fair, by equalizing the availability and cost of power among households that may have differing available roof space for PV.

To determine whether the statistical aggregation benefits extend beyond the block level, we created a second simulated EcoBlock using a second set of 30 homes located in the same city. The generation and storage were scaled to have the same ratios relative to load as for the first EcoBlock. The self-sufficiency of these two blocks were then compared against an aggregated “double block” in which the two were allowed to share power and storage resources as needed. As seen in Figure 10, the two blocks perform similarly—as expected given that the same PV and storage sizing process was used for both—and only a very small increase in self-sufficiency is observed when they are connected. Aggregation of geographically distributed PV systems has been shown to reduce variability on short timescales because cloud transients affect only a few systems at once [16], [17]. Here, though, we consider longer timescales (the hourly time step) because the large energy storage resource can mitigate the minute-to-minute variability. The greater correlation across geographic areas at this longer timescale, plus the great reduction in load variability that has already occurred at an aggregation level of 30 homes, combine to give minimal statistical aggregation benefit at the two-block level. Since connecting multiple blocks into a larger island is challenging from the perspectives of safety and legality—it requires energizing utility infrastructure between blocks, for example—we conclude that it is not likely to be worthwhile for urban residential blocks. We note that in some settings, islands

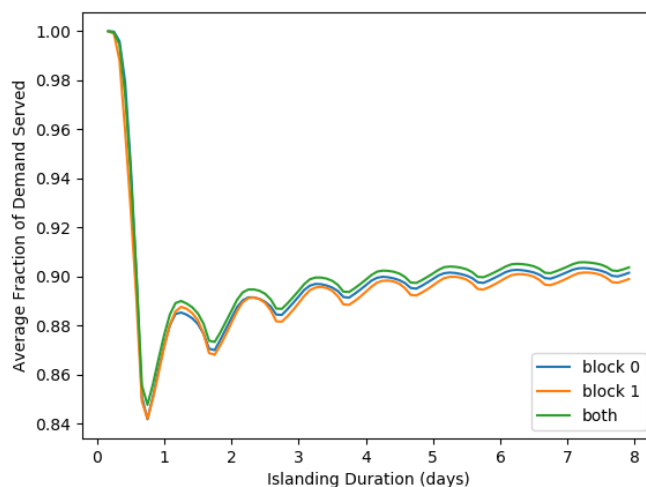


Figure 10: Mean fraction of demand served during islanding vs. islanding duration, for a start time of 16:00, for two simulated EcoBlocks and for the aggregated double block where both are connected together.

with dimensions larger than a city block may offer greater advantage. These may include commercial or industrial settings, where individual customers' load profiles may differ greatly, rural areas with large distance between customers, or very high population density areas where there is not enough space for large energy storage or rooftop PV installations on every block. It may be of interest to determine the marginal benefit in self-sufficiency as a function of the number of customers or geographic size of the island and how that depends on the load type and location.

Planned islanding

Besides coping with a fault on the main grid, another use case for adaptive islanding is to keep customers powered during planned maintenance of grid infrastructure. In the planned islanding case, we can assume the storage can be charged completely beforehand since the start time is known in advance. The objective is then to find the optimal start time, given a particular duration of maintenance, that minimizes the fraction of load shed. This analysis was performed by running the self-sufficiency simulations for all possible start times, with a given islanding duration corresponding to a particular type of maintenance.

For simple maintenance activities such as replacing a distribution power pole, outages of around 4 hours are typical. Since the storage capacity of the EcoBlock was designed to provide 4 hours of power, 100% of demand in the simulated EcoBlock was met during a 4-hour planned islanding scenario, for any start time, on every day of the year. More substantial maintenance activities such as replacing substation equipment may require up to 16 hours of work at a time. Therefore, the self-sufficiency of the simulated EcoBlock was also evaluated under a 16-hour planned islanding scenario. In this case, the EcoBlock was still fully self-sufficient with many of the start times. As shown in Figure 11, only for start times in the afternoon (1:00 pm to 9:00 pm) is the EcoBlock at risk of load shedding. The particular days that are not self-sufficient with afternoon start times are almost all in the late summer or early fall, when peak load is high due to high temperatures but generation does not extend long into the evening.

These results imply a new set of guidelines for scheduling maintenance on circuits with high penetration of EcoBlocks. Today, maintenance activities are often scheduled to take place in the middle of the night to minimize the nuisance for customers. These results show that maintenance on EcoBlock circuits could instead take place during daylight hours without causing outages, potentially improving safety and convenience for lineworkers. Additionally, forecasts for irradiance (affecting generation) and temperature (affecting load) should be taken into account when scheduling maintenance. Maintenance starting in the afternoon should be avoided whenever possible, especially for more work-intensive jobs and on days where low generation and significant air-conditioning load are expected.




Figure 11: Mean and minimum fraction of load served during a 16-hour planned islanding scenario as a function of islanding start time.

5.3 Conclusions

The conclusions drawn from the simulated EcoBlock, summarized here, can be applied to the proposed EcoBlocks in California. First, aggregating homes at the block level offers significant improvements in self-sufficiency relative to each home operating as an independent island. Second, while letting only a subset of loads participate in a DC island increases the fraction of load that can be served reliably, self-sufficiency is excellent even if all loads participate in the island, due to the large generation and storage resources. The simulated EcoBlock was fully self-sufficient during the majority of simulated islanding events, even with durations as long as 8 days. Furthermore, in addition to the 14 electric vehicles already on the simulated EcoBlock, there was enough excess generation to charge several more on most days if they are plugged in during daylight hours. Although load was the highest in the summer due to the air conditioning demand, overcast winter days made up most of the events requiring significant load shedding. We expect the same for the Oakland EcoBlock, where air conditioning needs are far less than in Austin. Since major outage events most often result from storms [26], the need to island may coincide with cloud cover a disproportionate amount of the time. Therefore, although the results show that little load shedding is needed on average, it is worth having the ability to shed large amounts of load during storm events. Finally, since the starting amount of stored energy matters greatly, especially for short islanding events, we recommend controlling the flywheel to maintain a minimum stored energy while connected to the grid, in case islanding needs to occur. This minimum energy should depend on the time of day, demand response capabilities of the EcoBlock loads, and perhaps weather forecasts. Overall, our analysis shows that, with real data for residential household loads and PV generation, the EcoBlock design principles—large generation and storage capacity, coupled with demand response—produce a highly self-sufficient islandable microgrid.

6 Suggestions for future work

Our work has focused on evaluating the potential design choices for EcoBlocks from the perspectives of grid impacts and scalability. To move from the hypothetical evaluation stage to deployment, decisions and progress need to be made in many areas of the design, including:

- Selection of the EcoBlock topology from among the scenarios in Figure 1, and selection of corresponding hardware
- Design of the EcoBlock’s protection system, which will be particularly challenging if DC

- Selection of the control architecture, objectives and signals that will be used to control devices within the EcoBlock, both for short-term stability and long-term energy balancing
- Selection of the control architecture, objectives and signals that will be used to influence the behavior of EcoBlocks within the grid, particularly control over net power while grid-connected and rules for islanding
- Development of a data analytics and visualization framework to improve grid operators' situational awareness that is scalable to large numbers of EcoBlocks
- Deeper investigation of the economic, legal and regulatory changes that would enable the most effective operation and scaling of EcoBlocks.

Quantitative evaluation of EcoBlock performance and impacts will be more accurate once these choices have been made. At the same time, simulation results can help inform the design choices. This suggests the need for a holistic, iterative approach where designs are adjusted based on simulation results and vice versa. For example, the self-sufficiency analysis in Section 5 could be made more realistic for the prototype EcoBlock through the use of:

- Oakland EcoBlock customer meter data
- A more realistic demand response model, such as one that includes generation and load forecasts and prioritization of loads
- Models for the control of EV charging as an additional form of demand response
- A time-varying statistical distribution of the storage state of charge at the start of islanding, derived from the operating strategy of the EcoBlock while connected to the grid.

We plan to develop this analysis further as a tool for optimally sizing energy storage resources according to certain self-sufficiency criteria. Furthermore, the potential impacts of grid-connected EcoBlocks on power system performance characteristics—such as voltage variation, load factor, overloading, and losses—can be determined quantitatively using power flow simulation tools such as GridLAB-D. This analysis would run time-series simulations representing each EcoBlock as a single time-varying net power injection at its PCC, and could be used to estimate the dependence of these impacts on:

- Number and location of EcoBlocks
- Feeder characteristics such as voltage, size, impedance, level of imbalance: there are a number of real and realistic feeder models available for such simulations
- EcoBlock operating strategy: could validate the effectiveness of a particular strategy or help identify best-performing strategies.

7 References

- [1] G. Montes and F. Katiraei, “Laboratory testing and field measurement of plug-in electric vehicle grid impacts,” California Energy Commission, Tech. Rep. CEC-500-2015-093, 2015.
- [2] P. P. Barker and R. W. D. Mello, “Determining the impact of distributed generation on power systems. I. radial distribution systems,” in *2000 Power Engineering Society Summer Meeting*, vol. 3, 2000, 1645–1656 vol. 3.
- [3] Y. Liu, J. Bebic, B. Kroposki, J. de Bedout, and W. Ren, “Distribution system voltage performance analysis for high-penetration PV,” in *2008 IEEE Energy 2030 Conference*, Nov. 2008, pp. 1–8.
- [4] F. Katiraei and J. R. Aguero, “Solar PV integration challenges,” *IEEE Power and Energy Magazine*, vol. 9, no. 3, pp. 62–71, May 2011.
- [5] EPRI, “Understanding the grid impacts of plug-in electric vehicles (PEV): Phase 1 study - distribution impact case studies,” EPRI, Palo Alto, CA, Tech. Rep. 1024101, 2012.
- [6] M. Lave, M. J. Reno, and R. J. Broderick, “Characterizing local high-frequency solar variability and its impact to distribution studies,” *Solar Energy*, vol. 118, pp. 327–337, 2015.
- [7] A. Oudalov, R. Cherkaoui, and A. Beguin, “Sizing and optimal operation of battery energy storage system for peak shaving application,” in *2007 IEEE Lausanne Power Tech*, Jul. 2007, pp. 621–625.
- [8] E. L. Ratnam and S. R. Weller, “Receding horizon optimization-based approaches to managing supply voltages and power flows in a distribution grid with battery storage co-located with solar PV,” *Applied Energy*, 2017.
- [9] K. P. Schneider, J. C. Fuller, F. K. Tuffner, and R. Singh, “Evaluation of conservation voltage reduction (CVR) on a national level,” Pacific Northwest National Laboratory (PNNL), Richland, WA (US), Tech. Rep. PNNL-19596, 2010.
- [10] NIAC, “Critical infrastructure resilience final report and recommendations,” National Infrastructure Advisory Council (NIAC), Tech. Rep., 2009.
- [11] A. Gholami, F. Aminifar, and M. Shahidehpour, “Front lines against the darkness: Enhancing the resilience of the electricity grid through microgrid facilities,” *IEEE Electrification Magazine*, vol. 4, no. 1, pp. 18–24, Mar. 2016.
- [12] R. Arghandeh, M. Brown, A. D. Rosso, G. Ghatikar, E. Stewart, A. Vojdani, and A. von Meier, “The local team: Leveraging distributed resources to improve resilience,” *IEEE Power and Energy Magazine*, vol. 12, no. 5, pp. 76–83, Sep. 2014.
- [13] C. Eid, P. Codani, Y. Perez, J. Reneses, and R. Hakvoort, “Managing electric flexibility from distributed energy resources: A review of incentives for market design,” *Renewable and Sustainable Energy Reviews*, vol. 64, pp. 237–247, 2016.
- [14] B. Kirby, “Ancillary services: Technical and commercial insights,” 2007.
- [15] S. M. Nosratabadi, R.-A. Hooshmand, and E. Gholipour, “A comprehensive review on microgrid and virtual power plant concepts employed for distributed energy resources scheduling in power systems,” *Renewable and Sustainable Energy Reviews*, vol. 67, no. Supplement C, pp. 341–363, 2017.
- [16] T. E. Hoff and R. Perez, “Quantifying PV power output variability,” *Solar Energy*, vol. 84, no. 10, pp. 1782–1793, 2010.
- [17] A. Nguyen, M. Velay, J. Schoene, V. Zheglov, B. Kurtz, K. Murray, B. Torre, and J. Kleissl, “High PV penetration impacts on five local distribution networks using high resolution solar resource assessment with sky imager and quasi-steady state distribution system simulations,” *Solar Energy*, vol. 132, pp. 221–235, 2016.

- [18] A. von Meier, D. Culler, A. McEachern, and R. Arghandeh, “Micro-synchrophasors for distribution systems,” in *Innovative Smart Grid Technologies Conference (ISGT), 2014 IEEE PES*, Feb. 2014, pp. 1–5.
- [19] A. von Meier, E. Stewart, A. McEachern, M. Andersen, and L. Mehrmanesh, “Precision micro-synchrophasors for distribution systems: A summary of applications,” *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 2926–2936, Nov. 2017.
- [20] R. Sherick, S. Gupta, W. Marin, A. Guerra, A. Gebeyehu, S. Collins, K. Schneider, J. Fuller, E. Mayhorn, S. Sridhar, A. Dinkel, L. Young, and B. Fitzsimons, “Advanced distribution analytics enabling high penetration solar PV,” Southern California Edison, Pacific Northwest National Laboratory, and Qado Energy Systems, Tech. Rep., 2016.
- [21] R. Liu, L. Dow, and E. Liu, “A survey of PEV impacts on electric utilities,” in *ISGT 2011*, Jan. 2011, pp. 1–8.
- [22] L. E. Zubietta, “Are microgrids the future of energy?: DC microgrids from concept to demonstration to deployment,” *IEEE Electrification Magazine*, vol. 4, no. 2, pp. 37–44, Jun. 2016.
- [23] X. Han, K. Heussen, O. Gehrke, H. W. Bindner, and B. Kroposki, “Taxonomy for evaluation of distributed control strategies for distributed energy resources,” *IEEE Transactions on Smart Grid*, vol. PP, no. 99, pp. 1–1, 2017.
- [24] A. A. G. Mabaning, J. R. C. Orillaza, and A. von Meier, “Optimal PMU placement for distribution networks,” in *To be presented at IEEE ISGT Asia Conference*, Dec. 2017.
- [25] A. H. Habib, V. R. Disfani, J. Kleissl, and R. A. de Callafon, “Optimal energy storage sizing and residential load scheduling to improve reliability in islanded operation of distribution grids,” in *2017 American Control Conference (ACC)*, May 2017, pp. 3974–3979.
- [26] PG&E, “2015 annual electric reliability report,” Pacific Gas and Electric Company (PG&E), 2016.
- [27] D. T. Ton and W. T. P. Wang, “A more resilient grid: The U.S. department of energy joins with stakeholders in an R&D plan,” *IEEE Power and Energy Magazine*, vol. 13, no. 3, pp. 26–34, May 2015.
- [28] A. H. Habib, E. L. Ratnam, V. R. Disfani, J. Kleissl, and R. A. de Callafon, “Optimization-based residential load scheduling to improve reliability in the distribution grid,” in *2016 IEEE 55th Conference on Decision and Control (CDC)*, Dec. 2016, pp. 2419–2424.
- [29] Pecan Street, Inc. (2017). Pecan Street Dataport, [Online]. Available: <https://dataport.cloud/> (visited on 11/13/2017).
- [30] Amber Kinetics. (2017). Amber Kinetics Model 32 specifications, [Online]. Available: <http://amberkinetics.com/products-2/> (visited on 11/13/2017).
- [31] L. Radetsky, “Streetlights for local roads,” *National Lighting Product Information Program: Specifier Reports*, vol. 14, no. 1, Feb. 2011.
- [32] Nissan. (2017). Nissan Leaf specs, [Online]. Available: <https://www.nissanusa.com/electric-cars/leaf/versions-specs/> (visited on 11/13/2017).