


RESEARCH

Open Access



AutoShare: Virtual community solar and storage for energy sharing

Stephen Lee^{1*} , Prashant Shenoy², Krithi Ramamritham³ and David Irwin²

*Correspondence:

stephen.lee@pitt.edu

¹Department of Computer Science,
University of Pittsburgh, Pittsburgh,
USA

Full list of author information is
available at the end of the article

Abstract

Residential solar installations are becoming increasingly popular among homeowners. However, renters and homeowners living in shared buildings cannot go solar as they do not own the shared spaces. Community-owned solar arrays and energy storage have emerged as a solution, which enables ownership even when they do not own the property or roof. However, such community-owned systems do not allow individuals to control their share for optimizing a home's electricity bill. To overcome this limitation, inspired by the concept of virtualization in operating systems, we propose virtual community-owned solar and storage—a logical abstraction to allow individuals to independently control their share of the system. We argue that such individual control can benefit all owners and reduce their reliance on grid power. We present mechanisms and algorithms to provide a virtual solar and battery abstraction to users and understand their cost benefits. In doing so, our comparison with a traditional community-owned system shows that our AutoShare approach can achieve the same global savings of 43% while providing independent control of the virtual system. Further, we show that independent energy sharing through virtualization provides an additional 8% increase in savings to individual owners.

Keywords: Energy systems, Storage, Solar, Energy sharing

Introduction

Advances in technology have led to a decline in the installation costs of solar panels, and the prices are expected to drop further. The levelized cost of energy (LCOE) is now 12¢ per kWh (When Will Rooftop Solar Be Cheaper Than the Grid? 2016), on par or less than traditional energy sources. This has fueled the growth of solar installations worldwide with 512 gigawatts of solar capacity deployed globally in 2018 alone (Growth of Photovoltaics 2019).

A key challenge arising from the increased penetration of renewable sources is their intermittent nature. The output of solar arrays varies based on the changing position of the sun throughout the day as well as effects such as shade from trees and cloud cover in the sky. Such fluctuations in output, as well as lack of solar production during the night, complicates management of the grid where supply and demand must be continuously

balanced. Energy storage in the form of batteries has been proposed as a potential solution for dealing with this intermittency (McPherson and Tahseen 2018; Irwin et al. 2017; Zhu et al. 2013). Battery-based energy storage can smooth out fluctuations in output from solar arrays, while also providing a range of benefits such as exploiting electricity price differentials (Daryanian et al. 1989), reducing peak usage (Mishra et al. 2012), providing UPS backups (Ramamritham et al. 2017) and supporting demand-response (Wang et al. 2013). Until recently, the high cost of batteries has been a barrier to large-scale energy storage deployments. However, this is beginning to change with the development of new battery technologies and falling prices. Today, battery-based systems such as Tesla Powerwall and others are increasingly deployed in conjunction with solar array deployments.

While large-scale solar array deployments continue to grow rapidly, the majority of solar installations in North America and Europe continue to be small-scale rooftop systems, primarily in residential homes. However, not every type of residential building is a suitable candidate for rooftop installations. Consider, for example, apartment buildings that are a common form of housing in an urban city. In scenarios where apartment homes are rented, the residents who do not own the property cannot install solar array themselves. Even when the apartment homes are individually owned, the building roof is collectively owned by the residents precluding individual installations. Similarly, a neighborhood of homes surrounded by significant tree foliage may not be a suitable candidate for rooftop solar arrays.

Community solar arrays (CSA) have emerged as a solution to these challenges (Augustine and McGavisk 2016; Chan et al. 2017). A community solar array is an array that is collectively owned by a group of individuals and is deployed in a common location. Each owner leases or purchases a share of the solar array and is allocated a certain fraction of the solar array in proportion to their share. The apartment property owners can deploy CSA and lease it to apartment renters. CSAs can also be installed by a neighborhood of homes on the suitably located open ground, away from the tree foliage, and shared by all households in the community. Similarly, community storage systems consist of an array of energy storage batteries that are collectively owned by a group, with a fraction of the storage capacity allocated to each owner. While both community solar and community storage are nascent technologies, the combination of the two opens up new opportunities for increasing solar penetration and performing various energy optimization.

Community solar and storage (CSS) systems generate energy for all owners as a single aggregated system and, unlike dedicated single owner systems, do not permit individual control of how the output should be optimized based on an individual's local demand. In this work, we draw inspiration from how such a problem has been addressed in other areas of computer science, such as virtualization in cloud computing (Tanenbaum and Van Steen 2007). In cloud computing, a cloud server is shared by multiple customers of the cloud platform. Each customer is allocated a virtual server that resembles a physical server in all respects; multiple virtual servers, each implemented using virtualization, share the same cloud server, while allowing each customer to independently control their virtual server (Tanenbaum and Van Steen 2007). Although solar arrays are quite different from a cloud server, the notion of a virtual resource is applicable to community solar and storage even though the details of implementing such abstraction will vary. Inspired by past work on cloud and OS virtualization, including our own (Wood et al. 2007; Sharma et al. 2016), we propose the notion of providing a virtual solar and battery abstraction to

each owner of a community solar and storage (CSS) systems. A virtual solar and battery array will provide the illusion of a dedicated single owner solar array, and importantly, will enable independent control of the energy resources. Like in cloud server virtualization, virtual solar arrays and batteries belonging to multiple customers are mapped onto a single physical solar array and battery owned by the community. Thus, our virtual abstraction each user to independently manage their share of solar generation and stored energy as if the system were a dedicated single owner system.

A second key benefit of virtualization of a community-owned system is that it enables *sharing* of electricity generated or stored in batteries by each virtual system. Such energy sharing, which is not possible in dedicated independently deployed systems, allows a resident to temporarily borrow electricity from one or more neighbor's shares to provide capital and operational savings. Thus, an individual can plan how much to share based on their current and future needs. In this paper, we present virtual abstractions for community solar and storage and show it can enable flexible sharing algorithms to reduce overall energy costs. In designing AutoShare, our paper makes the following contributions.

- We propose the virtual abstractions for a physical solar and battery array to support multiple virtual solar and battery arrays. We design a set of software-defined mechanisms to enable such virtual abstractions and permit the independent control of each virtual solar and battery system.
- We present an energy sharing algorithm that enables homes to share surplus solar capacity and excess stored energy in a virtual battery with households experiencing energy deficits and discuss monetary incentives for borrowers and lenders to incentivize such sharing.
- We demonstrate the efficiency of our approach using electricity demand from real homes and solar generation traces from actual installations. We compare AutoShare to non-virtualized community-owned solar that implements a single global policy for all homes and to dedicated individual installations with no sharing capabilities. Our results show the benefits of allowing local control as well as energy sharing. Specifically, AutoShare can reduce solar array size by 14.6% and battery capacity by 75% compared to a dedicated system. Further, we show that energy sharing gives an additional 8% increase in savings while providing monetary benefits to borrowers and lenders.

Background and overview

Our work assumes a community solar and battery array that is collectively owned by a group of residents. Each resident is assumed to own a certain fraction of solar and a certain battery share of the community solar and storage system. Consequently, the corresponding share of solar output and stored energy is assigned to each owner. We assume that each community owner can use their portion of the solar array and battery in any manner to perform energy optimization or reduce energy bills.

This is done by exposing a virtual solar and battery abstraction to each community owner. The capacity of the virtual solar and battery array is determined by the fractional share of the physical solar array and battery assigned to each owner. Each virtual solar array and battery appear, in all regards, as a physical array and battery, and all virtual arrays are mapped onto (multiplexed onto) the underlying physical solar and battery array.

As noted earlier, this is analogous to server virtualization in cloud computing, where virtual machines belonging to multiple cloud users are multiplexed onto a physical server and resources of the physical server are sliced (i.e. partitioned) across co-resident virtual machines.

A non-virtualized community-owned system does not allow independent control of the output in the shared system (see Fig. 1(a)). Energy output from the solar and battery is used to meet the aggregate energy demand across all homes, prohibiting individual owners to decide on how to use their share of energy. On the other hand, dedicated systems allow independent control but are more expensive or may be infeasible in many residential locations (see Fig. 1(b)). In contrast, the notion of virtual solar and battery arrays provide multiple benefits. First, such abstractions can provide the illusion of a dedicated physical array and battery that can be utilized and managed independently of other virtual solar arrays and batteries. This allows each owner to make the “optimal” decision of how to utilize their virtual solar array and batteries independently of what other owners decide at each instant. For example, one owner may be using their share of solar output to power their loads, while another owner net-meter their share to the grid. Such local control is not possible in a non-virtualized community system, which is forced to implement a single global policy on behalf of all homes.

Second, the virtual abstraction gives each owner an illusion of ownership of the unit, even though the group collectively owns the physical resources. In some sense, the virtual solar and battery arrays act like N independent smaller array and battery installations — while being cheaper to install due to economies of scale of installing a single large solar and battery array over N smaller ones. Also, many components like inverters can be shared rather than duplicated.

Third, a virtual solar array and battery system can be utilized for many differing energy optimizations. In scenarios with time-of-use pricing with different pricing slabs, surplus solar production during off-peak or mid-peak price periods can be stored in the battery for later use, such as peak price periods to maximize cost savings. During peak periods,

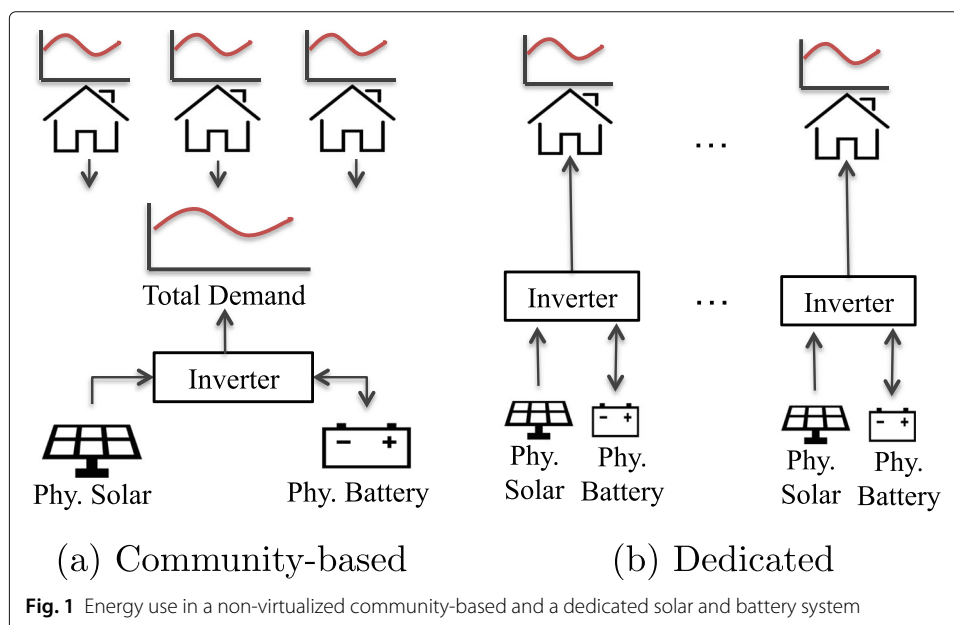


Fig. 1 Energy use in a non-virtualized community-based and a dedicated solar and battery system

the system prioritizes the use of solar production and stored energy in the virtual battery overdrawing power from the grid. Finally, if there is surplus solar production after serving local loads and charging the battery, this excess energy can be net-metered to the grid to earn revenues (which offset changes in the monthly electricity bill).

In our paper, we also assume that each owner can share electricity from their virtual solar or battery array with their neighbors. For example, rather than net-metering surplus electricity to the grid, it is also possible to sell (or lend) this surplus to a neighbor who has high current demand (and is drawing electricity from the grid). Similarly, it is feasible to sell or lend stored energy from a virtual battery to others, such sharing in the form of lending and borrowing provides both capital and operational cost benefits — it allows a community owner to provision a smaller virtual solar and battery system than dedicated setup and borrows from others during peak periods. It also provides additional operational benefits by increasing cost savings from a solar and battery system.

Virtual abstractions simplify the implementation of energy sharing. Analogous to how virtual machines can be resource multiplexed onto a physical machine and flexibly allocate resources, virtual solar and battery system can be used to dynamically allocate more electricity or more stored energy from the underlying community solar and battery system.

Our work also assumes an energy sharing pricing model. Currently, utilities purchase any surplus solar energy at retail prices from users (via net metering). However, when the wholesale price (cost to sell power to the utility) is same as the retail price (cost to purchase power from the utility), selling energy to the grid or the neighbor does not provide any additional cost benefits to a user. Cost benefits from sharing energy arise when the wholesale price is less than the retail price. Instead of selling electricity at wholesale price to the utility, a user can earn a profit by selling energy to its neighbors at a rate higher than the wholesale price. Similarly, in this scenario, borrowing energy at a rate lower than the retail price yields cost benefits to both the borrower and the lender. Increased solar penetration have impacted the grid (i.e., wholesale prices turned negative (California is getting so much power from solar that wholesale electricity prices are turning negative 2017), and with more solar adoption, utility companies will have to rethink how they purchase electricity from distributed sources. We assume that, in the future, utilities will purchase power at a lower rate than retail price, which will enable borrowing and lending of energy. Below, we present an algorithm and a system to virtualize a community solar and battery array. We also present “smart” algorithms that maximize the benefit of energy sharing with the overall goal of maximizing electricity bill savings due to the use of solar arrays and batteries.

AutoShare design

In this section, we first present key primitives to provide the abstraction of a virtual solar and battery array system and then show how these primitives can be used to implement algorithms to control the energy generated and stored in each virtual array.

AutoShare virtualization mechanisms

Consider a community solar array consisting of P panels with a capacity C_{solar} . The system also consists of a community battery array of B battery cells with a total capacity of C_{batt} . We assume the community-owned system is collectively-owned by N residents, and

allocate each user a virtual share of the solar and battery array. Suppose that the i^{th} owner is allocated a fraction S_i of the solar array and a fraction B_i of the battery array, where $\sum_{i=1}^N S_i = 1$ and $\sum_{i=1}^N B_i = 1$, $0 < S_i < 1$ and $0 < B_i < 1$. This implies that $S_i \cdot C_{solar}$ capacity of the aggregate solar array and $B_i \cdot C_{batt}$ capacity of the aggregate battery array is allocated to owner i .

From a virtualization standpoint, the system presents the illusion of N smaller solar and battery arrays of the corresponding size, each of which appears as a dedicated system to its owner (see Fig. 2). That is, owner i sees a virtual solar array of size $S_i \cdot C_{solar}$, a virtual battery of size $B_i \cdot C_{batt}$ and a virtual controller (e.g., a virtual inverter) to determine how the solar and battery array output is used at each instant. Virtualization allows each owner to make independent decisions on how to use the system, regardless of how others use their system. The N virtualized systems are “multiplexed” onto the underlying physical solar and battery array, and the overall behavior of the system at any instant represents the aggregate decisions made by each individual virtualized system.

To implement this abstraction, AutoShare exposes a set of virtualization primitives that can be controlled by software algorithms in each virtual controller. Let us assume that the array uses a virtual or physical sensor to monitor the solar output, the energy stored in the battery and the electricity demand of each owner. Let $solar_i(t)$, $battery_i(t)$ and $demand_i(t)$ represents the electricity output of virtual array i , energy stored in virtual battery i and electricity demand of home i at time instant t . To enable an owner to control their virtual system based on these monitored values independently, the physical controller exposes these following software-defined primitives to each virtual controller:

- $charge_i(t)$, which specifies the rate at which the virtual battery should be charged using the output of the virtual solar array at time t
- $discharge_i(t)$, which specifies the rate at which the virtual battery should be discharged to meet a portion of $demand_i(t)$
- $send_to_grid_i(t)$ which specifies the rate at which surplus solar electricity should be transmitted (net metered) to the electric grid at time t
- $draw_from_grid_i(t)$ which specifies the rate at which electricity should be drawn from the electric grid to meet a portion of $demand_i(t)$

Together these primitives enable each virtual controller to implement flexible software algorithms to control how the solar output and energy storage in the virtual solar and

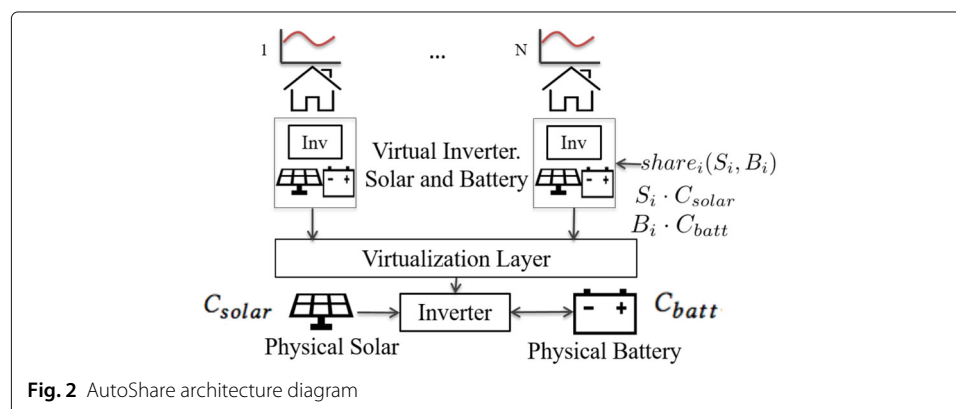


Fig. 2 AutoShare architecture diagram

battery array should be used. Each virtual controller can implement its own decisions regardless of how other owners behave.

AutoShare virtualization algorithm

We now present the AutoShare virtualization algorithm that uses the above primitives to implement software control of the virtual solar and battery system within the virtual controller. For home i , let us assume the solar output of virtual array is $solar_i(t)$ and demand is $demand_i(t)$; the AutoShare algorithm can determine if the current solar output is adequate to satisfy the demand. If so, the net surplus is computed as:

$$surplus_i(t) = \max(solar_i(t) - demand_i(t), 0) \quad (1)$$

If not, the net deficit is computed as:

$$deficit_i(t) = \max(demand_i(t) - solar_i(t), 0) \quad (2)$$

In the event of a surplus, after first using the solar output to satisfy the entire demand, the controller needs to determine how to utilize the remaining surplus. In this case, if the virtual battery is not fully charged, the surplus is first used to charge the battery at the max charging rate as follows:

$$charge_i(t) = \min(max_charge_rate, surplus_i(t)) \quad (3)$$

If the battery is full, $charge_i(t)$ is set to zero. If there is additional solar output left after charging the battery at max rate, the rest is net metered to the grid as follows:

$$send_to_grid_i(t) = solar_i(t) - demand_i(t) - charge_i(t) \quad (4)$$

Conversely, in the event of a deficit, the controller must determine how to satisfy the portion of the demand not met by the virtual solar array. In this case, the decision will depend on the current electricity prices. If off-peak pricing is in effect at time t , then it is better to conserve battery energy for peak periods and satisfy the current deficit from the electric grid:

$$draw_from_grid_i(t) = deficit_i(t) \quad (5)$$

If peak prices are in effect and the battery is not empty, the controller first draws power from the battery i.e.

$$discharge_i(t) = \min(deficit_i(t), max_discharge_rate)$$

so long as $battery_i(t) > low_threshold$. If the stored energy in the battery is below the $low_threshold$, then $discharge_i(t)$ is set to zero. Any unsatisfied demand beyond the maximum discharge rate from the virtual battery is met from the grid.

$$draw_from_grid_i(t) = \max(demand_i(t) - solar_i(t) - discharge_i(t), 0) \quad (6)$$

Thus, the AutoShare algorithm within each virtual controller can make independent decisions based on the solar output, battery level and demand of each home. Further, using the virtualization primitives, AutoShare enable software-driven algorithmic control of the virtual system.

Mapping virtual controller decisions to a physical system

The physical solar and battery controller aggregates all of the decisions made by individual virtual controllers to implement physical control as follows. If the total charge rate of all virtual batteries is greater than the total discharge rate, then the physical battery is charged at a rate

$$charge(t) = \sum_{i=1}^N charge_i(t) - \sum_{i=1}^N discharge_i(t) \quad (7)$$

In contrast, if the total discharge rate across all virtual batteries is greater than the total charge rate, then the physical battery is discharged at the rate of

$$discharge(t) = \sum_{i=1}^N discharge_i(t) - \sum_{i=1}^N charge_i(t) \quad (8)$$

Similarly, if the total power transmitted to the grid by all virtual solar arrays is greater than the total power drawn from the grid, the physical solar array will perform overall net-metering at the following rate:

$$send_to_grid(t) = \sum_{i=1}^N (send_to_grid_i(t) - draw_from_grid_i(t)) \quad (9)$$

If the opposite is true, no power is net-metered, since all of the solar output is used to satisfy the local demands of all homes and to store energy in the battery.

Energy sharing in AutoShare

We now discuss how AutoShare's virtualization mechanisms can be employed to permit flexible energy sharing. The energy sharing algorithm aims to maximize the energy cost savings across homes while incentivizing borrowers and lenders.

AutoShare energy sharing algorithm

AutoShare's virtualization algorithm allows each owner to operate their virtual solar and battery array independently of others. In this case, each virtual system is isolated from others, and there is no direct interaction between them. However since all virtual arrays are multiplexed onto a common physical solar array, there are opportunities for the virtual systems to *collaborate* with one another. One form of collaboration is *energy sharing* where virtual systems with a surplus solar generation or surplus stored energy shares it with virtual systems that have a deficit. Such sharing further reduces reliance on the grid, since some or all of the demand of a home is met from other neighboring virtual systems with surplus capacity. In practice, opportunities for energy sharing arise since different homes have different demand profiles. Some homes with daytime occupants will see higher peak usage during day hours, while homes with working occupants will see low usage during day hours with peak solar generation. The latter homes can lend surplus electricity that would otherwise be net-metered to the grid to the former homes. Similarly, during evening peak periods, demand from homes may peak at different times (e.g., homes with evening peak versus those with late night usage). In such cases, virtual batteries with surplus stored energy can lend it to others if it is not being used locally for any reason.

Energy sharing makes economic sense only under certain types of electricity pricing schemes. In scenarios where the cost at which grid purchases electricity is the same as the retail cost of buying electricity from the grid, energy sharing provides no monetary benefit. A virtual system can then sell any surplus to the grid via net-metering and neighbors with deficit can buy it back from the grid at the same price, requiring no direct cooperation between virtualized systems. However, in scenarios where the grid purchases net-metered electricity at wholesale generation prices and sell it to homes at retail prices, direct lending without grid involvement provides monetary benefits. In this case, rather than purchasing electricity from the grid at a retail price, a virtual system can procure this electricity from a neighboring virtual system with surplus electricity and do so at a price that is *higher* than the wholesale price but *lower* than the retail prices. This incentivizes systems with a surplus since they can sell the surplus at a rate higher than the grid's wholesale prices, while homes with a deficit can purchase this surplus at a price that is lower than the grid's retail rate.

From a virtualization standpoint, energy sharing relaxes the assumption of strict isolation between virtualized systems. It allows a virtual solar array or a virtual battery to increase its capacity by borrowing from surplus homes temporarily. This is analogous to virtual machines that temporarily use unused physical CPU capacity that is allocated to other virtual machines but not currently used.


To implement such energy sharing, AutoShare virtual inverters need two additional virtualization primitives.

- $borrow_i(t)$ which specifies the amount of power that home i needs to borrow from any other virtual solar or battery system at time t
- $lend_i(source, t)$ which specifies the amount of surplus power that home i will lend from the specified source at time t . The source can be *solar*, in which case surplus power is lent from the virtual solar array, or *battery*, in which case power is drawn for energy stored in the virtual battery.

These primitives enable a virtual controller to implement any energy sharing algorithm that is best suited to its needs. For our current work, we design a AutoShare energy sharing algorithm that is directly based on the AutoShare virtualization algorithm presented in the previous section. Our energy sharing algorithm is an enhancement to the basic virtualization algorithm as follows (see Fig. 3). First, the algorithm determines if the current home should become a borrower, a lender, or neither, at time t . A home is a candidate for lending electricity if its virtual solar array has surplus power that it would have net-metered to the grid. In this case, all of this surplus power becomes available for lending to other homes rather than being net-metered. A home is also a candidate for lending electricity if its virtual battery has a high charge level (above a high watermark threshold) and is willing to share some of the stored energy with others. Specifically, if $solar_i(t) - demand_i(t) - charge_i(t) > 0$ then the home has surplus power it would have previously net-metered and the virtual controller indicates it is willing to lend this power:

$$lend_i(solar, t) = solar_i(t) - demand_i(t) - charge_i(t) \quad (10)$$

Further, if the battery has a high charge level indicated by $battery_i(t) > high_threshold$ and the battery power is not being consumed at the maximum discharge rate, the surplus can be drawn as follows:



$$lend_i(battery, t) = \min(max_discharge_rate - discharge_i(t), 0) \tag{11}$$

Conversely, a home becomes a candidate for borrowing electricity if it has a deficit that would normally require drawing power from the grid. In this case, the home can first request surplus power from other virtual systems, and only request grid power if its deficit cannot be fully met by other lenders. That is, if $demand_i(t) - solar_i(t) - discharge_i(t) > 0$ the home has a unmet deficit and it can make a borrow request as follows:

$$borrow_i(t) = demand_i(t) - solar_i(t) - discharge_i(t) \tag{12}$$

In all other cases, $lend_i$ and $borrow_i$ are set to zero. Note that it is possible for a home to neither be a lender nor a borrower at time t , a scenario that occurs if it has zero deficit (i.e., has no need to borrow) but can not lend either since all solar electricity is being directed to the virtual battery, which itself has a low charge level (and thus has no solar or batter capacity to lend).

Mapping virtual sharing requests onto the physical system

Both $lend_i$ and $borrow_i$ indicate the maximum amount of power that each virtual controller i wishes to lend or borrow based on its current generation and demand. The actual amount of power that is lent or borrowed must then be computed by the physical controller by matching borrowers and lenders. To do so, the physical controller first computes the total borrowing needs as;

$$borrow(t) = \sum_{i=1}^N borrow_i(t) \tag{13}$$

The solar capacity available for lending is the:

$$lend(solar, t) = \sum_{i=1}^N lend_i(solar, t) \tag{14}$$

If the solar lending capacity $lend(solar, t)$ exceeds the borrowing demand $borrow(t)$, then all of the borrowing needs can be met from the surplus solar capacity that is available. Each lender can lend an equal amount to meet the total borrowing need or lend in proportion to its solar share S_i . If the total borrowing demand exceeds the total solar capacity, any unmet borrowing need can be lent from stored battery energy that can be lent. The maximum battery power that can be lent is:

$$lend(battery, t) = \sum_{i=1}^N lend_i(battery, t) \quad (15)$$

Finally, if the borrowing need is still not satisfied by the lending solar and battery capacity (that is, $borrow(t) > lend(solar, t) + lend(battery, t)$), the rest must be drawn from the grid.

$$draw_from_grid_i(t) = borrow_i(t) - borrowed_power_i(t) \quad (16)$$

Conversely, if all of the borrowing needs are met by surplus solar energy, any remaining solar can be net metered to the grid as follows

$$send_to_grid_i(t) = lend_i(solar, t) - lent_solar_i(t) \quad (17)$$

where $send_to_grid_i(t)$ is zero if there are no surplus solar energy, and $borrowed_power_i(t)$ and $lent_solar_i(t)$ are amount of power borrowed or lent by home respectively,

Surplus energy prediction

The above algorithm assumes a simple threshold-based approach to determine when stored energy can be shared with others. A fixed threshold-based approach to determine the amount of energy to lend does not consider the owner's future electricity demands. In practice, the threshold should not be a fixed value. In scenarios where the owners are away from home, the threshold should be set lower to enable more energy sharing and increase their benefit. Similarly, if there is a higher demand the following day, the threshold should be set higher to reduce energy sharing to draw from the grid. Accurate future knowledge will help in deciding how much excess energy is available to share. To handle such dynamic mechanisms, we design a machine learning-based approach to predict future energy demands and estimate the surplus energy that can be lent from the battery. For simplicity, our work assumes each owner uses a greedy policy to share surplus energy. That is, it lends its surplus energy with others only when future demand is also met.


To determine the surplus energy to lend, we first build a demand model of the home to learn the energy usage behavior. The model is then used to predict the electricity needs of the home for the following day. We explore two techniques: Autoregressive Integrated Moving Average (ARIMA) and Support Vector Machines (SVM) with different kernel functions. ARIMA is a popular model used for time series prediction. It uses the Autoregressive (AR), Integrated (I) and Moving Average (MA) components to predict future points. The autoregressive part is a linear combination of their own lagged values. The moving average part captures the regression error as a linear combination of the past error term. Finally, the integrated part indicates whether the data is differenced to make data stationary. Together, it forms a linear equation that uses past values and errors to determine future points.

In contrast, SVM is a supervised regression technique that can use exogenous information such as weather, along with past energy usage to predict energy usage for the following day. We use multiple features as inputs to the SVM model: day of the week, day of the year, outside temperature, humidity, heating and cooling degree days, past power consumption, weekday or weekend information, and holidays. Since the day of the week is cyclical, i.e., repeats every seven days, we encode it by transforming the day into an angle (in steps of $2\pi/7$) and use the sine and cosine values as input feature vectors. We use a similar approach to encode the day of the year. Further, we use one-hot encoding, a standard machine learning technique for encoding categorical labels, to encode weekday or weekend data, before using it as input features. For other data points, we use raw values as inputs. In summary, the above features are used as input to solve a regression problem, where we predict the future energy demand at a given resolution.

To train our models, we use the dataset described above that has energy demand of homes at a 30-minute resolution. We use one year of the dataset in 2014 for training and the next year (2015) for testing. To train the ARIMA model, we perform a grid search on the parameters and select the parameters with the lowest mean absolute percentage error (MAPE). Figure 4 shows the median MAPE value of different prediction models across all homes. As seen, the SVM-rbf model has a median MAPE value of 31%, and is the lowest among all the models used. Since SVR-rbf performs better than other techniques, we use it to predict the future surplus energy available for sharing.

Determining the surplus energy to lend also requires predicting the solar output for the following day. We use the technique discussed in Iyengar et al. (2014) to predict the solar output for the next day. The method takes into account irradiance and other weather parameters as features to accurately predict future solar production. By combining the future solar output, future electricity demands generated from the model, and the current battery capacity, a candidate home can estimate the amount of energy to borrow from the virtual battery. Specifically, the surplus energy is defined as

$$\begin{aligned} \text{surplus_energy}_i(t) = & \max(\text{future_solar}_i(t) + \text{charge}_i(t) \\ & - \text{future_demand}_i(t), 0) \end{aligned} \quad (18)$$



where $future_solar_i(t)$ and $future_demand_i(t)$ are the future predictions for the next 24 hours. Amount of energy to lend, when $surplus_energy_i(t) > 0$, is defined as

$$lend_i(battery, t) = \min(surplus_energy_i(t), \quad (19)$$

$$max_discharge_rate_i - discharge_i(t))$$

Thus, each home independently calculates how much to lend based on predicted future demand, generation and available battery capacity.

Evaluation

We focus on evaluating the potential benefit of AutoShare using trace-driven simulations. To do so, we use real electricity load dataset from 50 homes over two years between 2014 to 2015. The electricity dataset was gathered from the New England region of the United States and consists of energy consumption information at a resolution of 30 minutes (Iyengar et al. 2016). We construct different demand profile mixes of 20 homes each from these 50 homes to generate the diversity of homes in a building. To construct the demand profile mixes, we separate the homes into *day* and *night demand profiles*. We define *day profile* homes as homes that have most of their energy demand during the peak pricing hours, i.e., peak to off-peak energy usage is higher than one. In contrast, *night profile* homes use energy mostly during the off-peak pricing period. Next, we randomly select houses belonging to either of the demand profiles proportionately and ran our experiment multiple times to report the overall savings.

We use Wisconsin electric's time-of-use (TOU) prices as a representative pricing model (Wisconsin Electric Rates 2017). However, we also use other pricing models, which we present in our results. Wisconsin's peak pricing periods are between 7 a.m to 7 p.m, while the off-peak periods are from 7 p.m. to 7 a.m. Typically, wholesale prices are 30% to 50% of the retail price (Wholesale Electricity Price in the US 2017). We assume the wholesale electricity rates to be 40% of the retail price, but also evaluate the effect of other wholesale prices on cost savings. For simplicity, we assume that the apartments use the average of wholesale and retail price to sell their energy to others but sell the electricity to the grid at wholesale price. Note that the residents share only their portion of the community solar or battery energy. The share of solar and battery for each home is determined based on their energy consumption in the previous year, i.e., we assign a solar and battery proportionate to their overall yearly load.

We also require weather data for predicting future electricity usage of a home and its solar generation output. The weather data is available at a one-hour granularity, which we gather from the Weather Underground website (Weather dataset API 2017). We resample the weather dataset to 30-minute resolution, before running the prediction algorithms. The decision on when and the amount of energy to share is determined using the prediction output, as described in "Energy sharing in AutoShare" section.

Experimental results

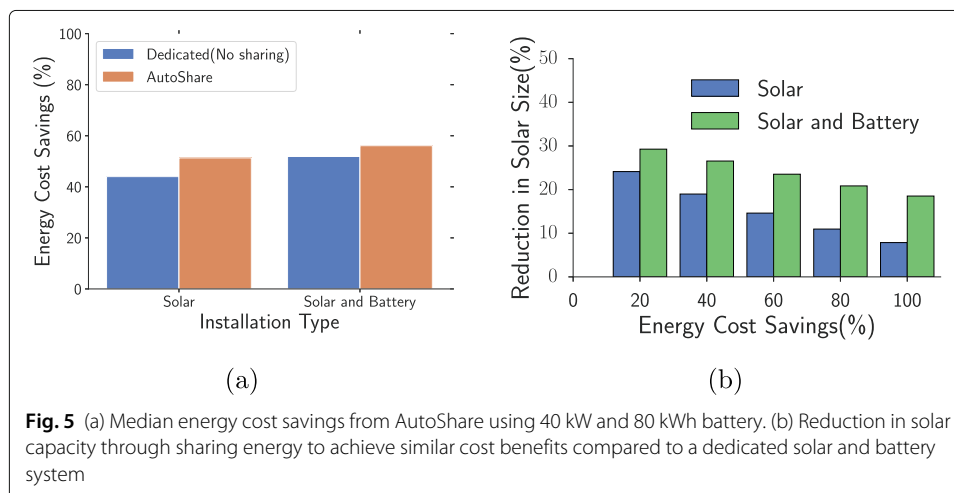
In this section, we describe the effects of energy sharing within a community sharing solar energy and battery.

Energy sharing benefits

We first analyze the cost benefits of using dedicated community solar and battery system and compare it to the energy sharing scenario using AutoShare. Figure 5(a) shows the median energy cost savings of a home for both dedicated and the energy sharing scenario. With 40 kW solar array, a dedicated system achieves 43% of energy cost savings and yields 8.8% higher savings when coupled with an 80 kWh battery. We see higher energy cost savings when we share energy. In particular, energy sharing provides an additional 8% increase in cost savings. This is because, rather than net metering to the grid at wholesale prices, users can sell surplus electricity at a higher rate to its neighbors, benefiting both the borrower and the lender.

We now examine the capital expenditure (CapEx) savings achieved through energy sharing. Clearly, dedicated solar arrays and batteries for each home cost more, assuming it is feasible to install one in a community area, simply because we cannot get *economies of scale*. On the other hand, a virtual community-owned solar and battery arrays do not require additional inverters, separate wirings, and thus cost less. In our analysis, we use the reduction in solar arrays and batteries as a proxy for the reduction in CapEx costs. Figure 5(b) shows the reduction in solar array size (CapEx) of a home if we want to achieve higher energy cost savings. The figure shows that to achieve higher cost savings will require bigger solar panels. In other words, as energy cost savings increase, the percentage reduction in solar array size decreases. This is because, with larger solar installations, most homes will have surplus energy to lend, which will reduce the need to borrow energy from others. We note that a virtual community-based system can reduce its solar array size by 14.6%, through energy sharing, to achieve 60% energy cost savings compared to a dedicated system. This is equivalent to 8.6 kW reduction in absolute values. We observe that a higher reduction in solar arrays can be achieved when batteries are installed. Using a battery capacity of 80 kWh, which is roughly 4 kWh per home, we note a virtual system can achieve a 23.5% reduction in solar array size, which is 13.8 kW reduction in absolute values.

Summary: AutoShare achieves 43% energy cost savings while providing each user with independent control over the virtual system. Moreover, a virtual community system can reduce the size by 14.6% compared to a dedicated system through energy sharing.



Impact of solar arrays

We analyze the impact of different solar array sizes on energy cost savings using AutoShare. Intuitively, larger solar arrays generate more solar energy, which in turn reduces the reliance on the grid and minimizes energy costs. However, it is not evident a priori how much cost benefits sharing energy provides. Figure 6(a) shows the median energy cost savings across homes using a 40kW solar array. As expected, the graph shows that with an increase in solar array size, the median energy cost savings of a home increase. Moreover, the energy cost savings is higher with energy sharing than the no sharing scenario. Since some occupants during the day may not use their share of solar energy, instead of net metering, the surplus solar energy can be lent to other homes to achieve higher cost savings. Thus, the variations in demand profile among homes allow the sharing of energy. In particular, we observe that the median energy cost savings for a home is 43% with a 40 kW solar array size and increases to 51% when energy is shared. Moreover, an 80 kW solar panel can yield energy cost savings of up to 85% using AutoShare.

Summary: Sharing energy increases the energy cost savings from 43% to 51% for a 40 kW solar installation.

Impact of energy storage

We now study the benefits of employing energy storage. The cost savings from batteries arise due to two primary reasons. First, batteries reduce the amount of energy net metered by storing surplus energy. The energy stored can then be used during peak pricing hours to increase cost savings. Second, sharing any surplus stored energy with others, especially during peak pricing periods, also increases cost savings. Figure 6(b) shows the median energy cost savings for different battery sizes and solar array size of 40 kW. The graph shows that as battery capacity increases, the energy cost savings increase. With a small battery size of 40 kWh, which is roughly equivalent to 2 kWh of battery per home, AutoShare can increase energy cost savings from 43% to 50%. The energy savings further increases by an additional 5.6% when energy stored in the battery is shared with others. This is because the discharge rate of any battery is limited. Even if the battery has sufficient energy to meet local demands, it may not be possible for a battery to fulfill all of the local energy needs as its maximum discharge rate may limit how much local demand it can satisfy. Since homes may not need to draw energy from batteries at all times, owners can allow energy discharge from their share to fulfill part or all of the local demand, thereby

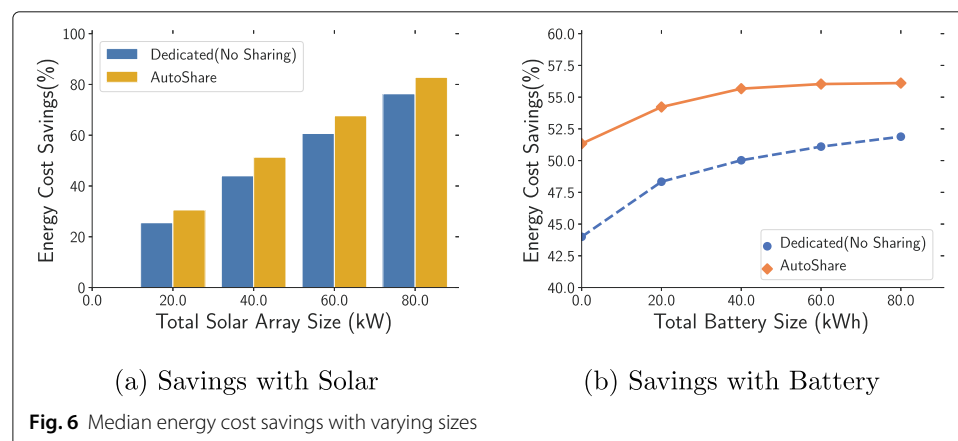


Fig. 6 Median energy cost savings with varying sizes


reducing energy costs. Unsurprisingly, we observe diminishing returns with an increase in battery capacity. Since the solar output is finite, there are fewer price differentials to exploit.

We observe that sharing energy also provides battery CapEx savings, i.e., a dedicated system will require a smaller battery size to achieve similar savings compared to the sharing scenario. As seen in Fig. 6(b), a battery capacity of 80 kWh is required to achieve 51.8% of cost savings. On the other hand, when energy stored in the battery is shared, AutoShare achieves 54.2% energy cost savings with a battery capacity of 20 kWh — a 75% reduction in battery capacity yielding significant CapEx savings. Since AutoShare allows energy sharing, it achieves higher cost savings from a solar and battery system compared to a dedicated system for each home.

Summary: Energy cost savings increase an additional 5.6% with energy sharing compared to a dedicated system. Moreover, sharing provides CapEx benefits and can reduce the battery capacity of a dedicated system by 75%.

Effect of energy pricing models

As mentioned earlier, energy cost savings is sensitive to the pricing model. Pricing models such as TOU pricing allow smart algorithms to exploit the price differential between peak and off-peak periods to achieve higher savings. So far, we used Wisconsin’s time-of-use pricing model, wherein the off-peak to peak ratio is roughly 1:2. We now analyze how a change in different off-peak to peak ratio impact energy cost savings. Figure 7(a) shows the impact on energy cost savings for varying off-peak to peak ratio from a 40 kW solar array size. Clearly, when peak and off-peak prices are similar, there will be no price differential to exploit, and the energy savings will be small. With higher off-peak to peak ratio, we will see more cost savings. As expected, the graph shows that as the difference between peak and off-peak price increases so does the cost savings. While an off-peak to peak ratio of 1:2 gives 43% in energy cost savings. Further, as off-peak to peak ratio decreases, the energy cost savings increases to 50.8%. Similarly, in the sharing scenario, the energy cost savings increases from 51% to 58.7%. Separately, using 80 kWh battery, the energy cost savings increases from 56.1% to 62.6% with energy sharing. Since batteries play the role of shifting the solar energy to generate cost savings, it has more potential to exploit the price differential in the TOU pricing model.

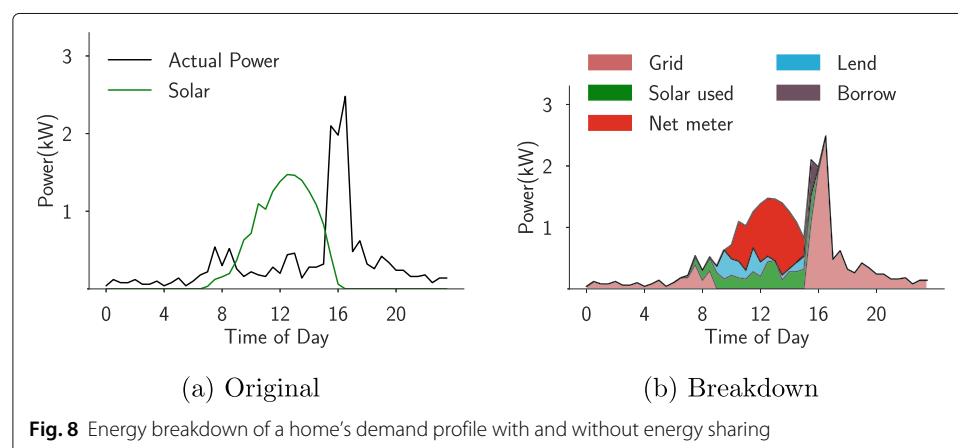



Next, we examine the impact of wholesale prices (i.e., the price at which electricity is sold to the grid) on energy cost savings. Intuitively, if the wholesale price is the same as the retail price, it is more beneficial to net meter the surplus energy to the grid than share. This is because storing energy for later use may result in loss of power due to battery inefficiency, thereby impacting cost savings. However, we observe that sharing is beneficial even when there is battery inefficiency as long as there is some price differential to exploit between wholesale and retail prices. Figure 7(b) illustrates the energy cost saving with different wholesale to retail price ratio from a solar array size of 40 kW. As wholesale price increase and approaches the retail price, the energy cost savings increase. This is because of the increase in profit by selling electricity to the grid at a higher rate. However, in the sharing scenario, this increase in profit is marginal as the surplus energy is already sold at a higher rate than wholesale prices to others. As expected, the energy cost savings from AutoShare equals the sharing scenario as the wholesale price approaches retail price.

Summary: As off-peak to peak ratio decreases, energy cost savings in the sharing scenario increases from 51% to 58%.

Effect on demand profiles

We now examine the effect of our energy sharing algorithm on different demand profiles. Figure 8 illustrates the energy distribution of a home for a day and compares energy sharing to a non-virtualized setup. The figure shows that without virtualization, any surplus energy is net-metered to the grid. However, with energy sharing, surplus energy is first lent, and the remaining energy is net metered. Further, when energy demand rises late afternoon, and local solar output is insufficient, energy is borrowed from others. Figure 9 depicts the normalized cost distribution of two homes with day and night demand profiles, with median energy cost savings in their respective cohort. We normalize the cost with its final cost. Note that the energy costs from the grid are higher for night profiles than day profiles. Since solar energy is only available during the day, most of the solar output is either net-metered or lent, resulting in a higher profit than day profile homes, and higher grid costs. In contrast, day profile homes tend to net meter less energy to the grid, with lower net meter cost, but also have to borrow more power during the daytime, resulting in higher borrow costs. It is important to note that while all homes use the same

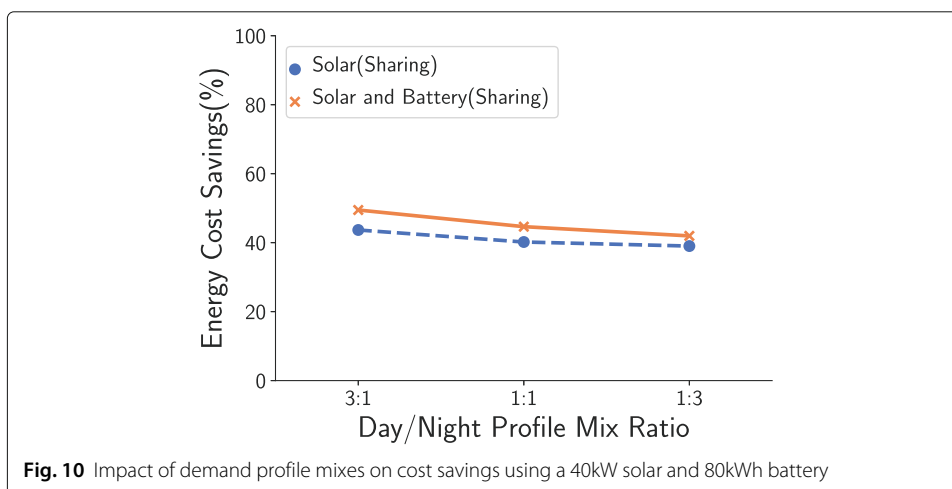




energy sharing policy, a virtual controller can have its own sharing policy, which may impact the energy and cost distribution of a home.

Finally, we examine the different mix of demand profiles in a building and their impact on the overall cost savings. The overall cost savings is computed by summing the final energy cost of all homes and the original cost. As discussed in “Evaluation” section, the day profile homes have high energy demands during peak hour periods, while night profile homes have high energy demands during off-peak periods. Figure 10 illustrates the impact on cost savings as the demand profile mix varies for a 40 kW solar array size. We compare the overall cost savings with solar only versus having both solar and battery. The graph shows that the overall cost savings are higher for profile mixes where the number of day profile homes is more than the night profile homes. Since the peak pricing period occurs during the daytime, most day homes can exploit the price differential and benefit from it. In contrast, homes with night demand profile, do not have high energy usage during the day to exploit the price difference. In particular, overall cost savings vary from 43.6% to 39% as the number of day demand profile decreases compared to night demand profiles. As expected, an addition of an 80 kWh battery increases the overall cost savings as it shifts the surplus solar energy to other periods, and provides an additional potential to exploit price differential.

Summary: An increase in the day profile homes increases cost savings as more homes can leverage the price differential.



Discussion and future work

There are several design considerations for realizing a virtual abstraction for energy sharing. First, our approach assumes homes have smart electric meters to monitor local demand (*demand_i*) at an appropriate time granularity (e.g., every few minutes). Second, our approach assumes that in addition to paying monthly electricity bills for drawing power from the grid, homes will make or receive micro-payments for borrowing or lending electricity to or from other community homeowners. The information on the amount of electricity borrowed or lent must be tracked by the physical controller (i.e., inverter) and periodically “settled” via actual payments. Third, modern inverters for non-virtualized community solar and battery arrays (e.g., Schneider inverters (Schneider Solar Hybrid Inverter Systems 2017) provide configuration controls on how much to net-meter and how much or when to charge the battery. Such an inverter can be easily enhanced to support AutoShare’s virtual abstraction. AutoShare can be implemented as a software layer on top of these configuration controls. Specifically, the software layer would expose a controller with AutoShare’s abstraction primitives to each homeowner. This provides the illusion of independent control of each system. However, in practice, the software abstraction layer takes the decisions from the virtual controllers, aggregates them and directly exercises the aggregate decision on the physical configurations exposed by the inverters. This makes it appear as though each user owns a dedicated system making independent judgments. Although, in practice, the decisions are “multiplexed” onto the physical solar and battery. As explained in the previous section, the abstraction layer exposes the API described in Table 1 to each virtual controller to support such energy sharing mechanism.

Billing and reconciliation

We also assume that we can net meter and share energy with others. In practice, borrowing and lending of energy require a bill reconciliation infrastructure to account for the micro-payments between homes. Smart meters available today are capable of reporting the local energy consumption of a home. A billing infrastructure can be implemented using smart meters and AutoShare primitives, wherein energy consumption information of each home can be logged along with energy borrowed or lent between homes. Separately, we require a billing agreement among homes to determine the costs of energy borrowed or lent.

Other benefits

There are two key benefits of our approach over multiple dedicated individual-owned or a non-virtualized community-owned system. A dedicated solar and battery systems may be

Table 1 Summary of the virtual abstraction API primitives AutoShare exposes to a virtual controller

Name	Description
<i>charge(t)</i>	charge rate of the virtual battery
<i>discharge(t)</i>	discharge rate of the virtual battery
<i>send_to_grid(t)</i>	amount of energy virtual inverter sends to the grid
<i>draw_from_grid(t)</i>	amount of energy virtual inverter draws from the grid
<i>borrow(t)</i>	amount of energy to borrow from others
<i>lend(source, t)</i>	amount of energy to lend to others from <i>solar or battery</i>

infeasible in most cases. Even if the installation was feasible, sharing is not possible as each installation will be on separate circuits. On the other hand, a non-virtualized community-owned installation allows sharing but do not provide individual control of the system. So the total bill across all homes is reduced, but it does not minimize the bill of each house. In contrast, AutoShare allows each home to minimize their local bills rather than the overall bill across all households. Another key benefit of AutoShare is users can design their own optimization policies. The virtualization primitives AutoShare enables powerful enough to implement other algorithms. For example, in our energy sharing algorithm, users first use surplus energy locally (i.e., satisfy local demand and charge batteries) and then share remaining energy or net meter. However, a range of algorithms is possible. A user can first share surplus energy, and then use the remaining power to store in batteries or net meter. Separately, instead of minimizing energy bills, other policy objectives can also be implemented. A user may have shiftable loads that they can directly control, and thus might define a policy for their virtual battery in conjunction with their controllable background loads.

Future work

For simplicity, our pricing model assumes the wholesale price (i.e., selling price of energy) is at a fixed percentage of the retail price. That is, if the retail price fluctuates, the wholesale price will also vary as a fixed percentage. However, it is possible that the wholesale prices may be very different from the retail prices. In this case, to maximize user profit, energy sharing will need to be based on the wholesale price signals. Users can take advantage of price differentials in wholesale price and sell their surplus energy (from solar or by discharging the battery) when the prices are higher to maximize expected revenue. Since many factors affect market pricing (Yu et al. 2019), to maximize expected revenue, uncertainties in load, solar, and market pricing need to be modeled and can be explored as part of future work.

Related work

Our work is related to these areas: energy systems, economics, and energy sharing in smart grids.

Energy systems

Recent work has looked at providing abstractions to physical resources to control and improve energy efficiency in buildings (Agarwal et al. 2011; Dawson-Haggerty et al. 2013; Lam et al. 2014; Keshav 2016; Karmakar et al. 2015; Wei et al. 2016; Pan et al. 2015). BOSS (Dawson-Haggerty et al. 2013) and Microsoft HomeOS (Dixon et al. 2012) provides service abstractions of shared physical resources. These abstractions connect different physical systems and allow the development of applications that enable energy optimization, such as HVAC or electrical lighting control. Also, there are studies on developing systems to control HVAC systems in buildings (Agarwal et al. 2011; Weng and Agarwal 2012; Karmakar et al. 2015). There has also been designing metadata schema for representing building applications and capture relationships between sensors and varying subsystems within a building (Balaji et al. 2018; 2016). However, prior works do not propose abstractions for solar or battery arrays for designing energy-efficient systems.

Energy economics

Numerous studies have focused on the use of renewable energy and its impact on the grid (Mishra et al. 2013; Carpenter et al. 2012; Narayanaswamy et al. 2012; Pan et al. 2013; Keck et al. 2019). Since electricity from renewable sources such as solar is highly intermittent and unreliable, intelligent use of batteries can minimize reliance on grid electricity without impacting daily usage patterns (Mishra et al. 2013). Studies show the benefits of energy storage with renewable for different market paradigms and power grid configurations (McPherson and Tahseen 2018). Also, energy cost savings achieved with batteries may be sensitive to certain pricing schemes (Carpenter et al. 2012), batteries are profitable when combined with demand response programs (Exarchakos et al. 2009). A significant amount of work has also focused on designing algorithms to reduce energy consumption in buildings (Doukas et al. 2007; Mishra et al. 2013). There are also studies on designing grid networks using solar and storage to manage their local power needs (He et al. 2008). Our work is built on previous work, where we leverage the variation in demand profiles of homes to optimize the use of solar and batteries.

Energy Sharing

There has been work in pricing and incentivizing energy trading in a microgrid scenario (Zhong et al. 2014; Thakur et al. 2014; Wang et al. 2015; Gregoratti and Matamoros 2015; Huang et al. 2016; Liu et al. 2017; Liu et al. 2017). For example, (Zhong et al. 2014) proposes a novel pricing model to incentivize energy sharing. Further, energy trading between microgrids is shown in Gregoratti and Matamoros (2015), while the energy sharing model with a price-based demand response is proposed in Liu et al. (2017). Other incentives include energy sharing to mitigate privacy leakage (Huang et al. 2016). In contrast, our work is complementary and focuses on the systems issues of virtualizing solar and battery sharing. Some of the pricing incentives discussed in prior work can be used in conjunction with our approach. Recent work has also studied bill reconciliation in energy trading (Mnatsakanyan and Kennedy 2015; PowerLedger: Peer to peer energy trading 2018; LO3 Energy: Brooklyn micro-grid 2018). PowerLedger, an energy startup company, uses a distributed ledger to enable a bill reconciliation platform for energy trading (PowerLedger: Peer to peer energy trading 2018). Again, their work is complementary to ours as they provide a platform for tracking energy usage, and do not provide mechanisms for controlling solar or battery arrays. Previous approaches have also studied energy sharing in microgrids (Han et al. 2016; Zhu et al. 2013; Wu et al. 2016). Zhu et al. designed an energy matching algorithm to minimize energy loss via energy sharing (Zhu et al. 2013). Much of the work on community shared solar on designing programs to improve affordability and access (Chan et al. 2017; Awad and Gül 2018; Hoffman and High-Pippert 2015; Chwastyk and Sterling 2015; Augustine and McGavisk 2016). These studies focus on understanding the barriers and implementing community solar and storage (CSS) programs to improve solar adoption among consumers. Since policies may differ from one state to another, it is essential to understand the driving factors in CSS participation. Recently, studies focused on solar energy trading and the benefits of CSS over individual ownership (Kasaei et al. 2017; Yu et al. 2019; Hafiz et al. 2019). For example, (Gai et al. 2019) proposes a privacy-preserving algorithm to prevent data leakage on energy traded and stored on a blockchain. Separately, fairness issues have also been in solar energy trading, arising from varying energy generation rates across different solar sites (Bashir et al.

2017; Feng et al. 2020). However, our work is complementary as we look into the systems aspects and focuses on providing control to a shared community solar and storage system.

Conclusion

In this paper, we proposed AutoShare — an energy sharing mechanism for community-owned solar and battery, wherein each virtual solar and battery can be assigned to an owner and controlled independently, regardless of others. To show AutoShare's potential, we implemented an energy sharing algorithm that enables energy sharing to minimize their local electricity bill. We compared AutoShare to a dedicated community-owned system and showed energy savings could increase the energy cost savings. Moreover, we showed that energy sharing also provides CapEx savings and requires a smaller solar array and battery compared to dedicated systems.

Abbreviations

LCOE: Levelized cost of energy; CSA: Community solar arrays; CSS: Community solar and storage; ARIMA: Autoregressive Integrated Moving Average; SVM: Support vector machines; MAPE: Mean absolute percentage error; TOU: Time-of-use; CapEx: Capital Expenditure

Acknowledgements

Not applicable.

Authors' contributions

SL designed and carried out the studies and drafted the manuscript. PS, KR and DI participated in its design and coordination and drafted the manuscript. All authors read and approved the final manuscript.

Funding

This research is supported by NSF grants IIP1534080, CNS-1645952, CNS-1405826, CNS-1253063, CNS-1505422, #2021693, #2020888, and the Massachusetts Department of Energy Resources.

Availability of data and materials

The datasets generated and/or analysed during the current study are not publicly available due to the nondisclosure agreement with the utility company but are available from the corresponding author on reasonable request.

Declarations

Competing interests

The authors declare that they have no competing interests.

Author details

¹Department of Computer Science, University of Pittsburgh, Pittsburgh, USA. ²College of Information and Computer Sciences, University of Massachusetts, Amherst, USA. ³Indian Institute of Technology, Bombay, Mumbai, India.

Received: 19 March 2021 Accepted: 15 June 2021

Published online: 12 July 2021

References

- Agarwal Y, Balaji B, Dutta S, Gupta RK, Weng T (2011) Duty-cycling buildings aggressively: The next frontier in HVAC control. In: Proceedings of the 10th International Conference on Information Processing in Sensor Networks
- Augustine P, McGavisk E (2016) The next big thing in renewable energy: Shared solar. *Electr J* 29(4):36–42
- Awad H, Gül M (2018) Optimisation of community shared solar application in energy efficient communities. *Sustain Cities Soc* 43:221–237
- Balaji B, Bhattacharya A, Fierro G, Gao J, Gluck J, Hong D, Johansen A, Koh J, Ploennigs J, Agarwal Y, et al (2016) Brick: Towards a unified metadata schema for buildings. In: Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments. pp 41–50
- Balaji B, Bhattacharya A, Fierro G, Gao J, Gluck J, Hong D, Johansen A, Koh J, Ploennigs J, Agarwal Y, et al (2018) Brick: Metadata schema for portable smart building applications. *Appl Energy* 226:1273–1292
- Bashir N, Irwin D, Shenoy P, Taneja J (2017) Enforcing fair grid energy access for controllable distributed solar capacity. In: Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments. pp 1–10
- California is getting so much power from solar that wholesale electricity prices are turning negative (2017). <http://bit.ly/2pdJKg2>. Accessed 05 May 2021
- Carpenter T, Singla S, Azimzadeh P, Keshav S (2012) The impact of electricity pricing schemes on storage adoption in Ontario. In: Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet

- Chan G, Evans I, Grimley M, Ihde B, Mazumder P (2017) Design choices and equity implications of community shared solar. *Electr J* 30(9):37–41
- Chwastyk D, Sterling J (2015) Community solar: Program design models. *Solar Electr Power Assoc*
- Daryanian B, Bohn RE, Tabors RD (1989) Optimal demand-side response to electricity spot prices for storage-type customers. *IEEE Trans Power Syst* 4(3):897–903
- Dawson-Haggerty S, Krioukov A, Taneja J, Karandikar S, Fierro G, Kitaev N, Culler DE (2013) BOSS: Building Operating System Services. In: *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation*
- Dixon C, Mahajan R, Agarwal S, Brush A, Lee B, Saroiu S, Bahl P (2012) An operating system for the home. In: *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*
- Doukas H, Patlitzianas KD, Iatropoulos K, Psarras J (2007) Intelligent building energy management system using rule sets. *Build Environ* 42(10):3562–3569
- Exarchakos L, Leach M, Exarchakos G (2009) Modelling electricity storage systems management under the influence of demand-side management programmes. *Int J Energy Res* 33(1):62–76
- Feng Y, Li Q, Chen D, Rangaswami R (2020) Solartrader: Enabling distributed solar energy trading in residential virtual power plants. In: *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. pp 60–69
- Gai K, Wu Y, Zhu L, Qiu M, Shen M (2019) Privacy-preserving energy trading using consortium blockchain in smart grid. *IEEE Trans Ind Inform* 15(6):3548–3558
- Gregoratti D, Matamoros J (2015) Distributed energy trading: The multiple-microgrid case. *IEEE Trans Ind Electron* 62(4):2551–2559
- Growth of Photovoltaics (2019). https://en.wikipedia.org/wiki/Growth_of_photovoltaics. Accessed 05 May 2021
- Hafiz F, de Queiroz AR, Fajri P, Husain I (2019) Energy management and optimal storage sizing for a shared community: A multi-stage stochastic programming approach. *Appl Energy* 236:42–54
- Han J, Choi M, Lee I, Kim S (2016) Photovoltaic energy sharing system in a multifamily residential house to reduce total energy costs. In: *IEEE International Conference on Consumer Electronics*
- He M, Reutzel E, Jiang X, Katz R, Sanders S, Culler D, Lutz K (2008) An architecture for local energy generation, distribution, and sharing. In: *Energy 2030 Conference*
- Hoffman SM, High-Pippert A (2015) Community solar programs and the democratization of the energy system. In: *Proceedings of the European Consortium for Political Research*. Montreal, QC, Canada. pp 26–29
- Huang Z, Zhu T, Gu Y, Li Y (2016) Shepherd: sharing energy for privacy preserving in hybrid AC-DC microgrids. In: *e-Energy*
- Irwin D, Iyengar S, Lee S, Mishra A, Shenoy P, Xu Y (2017) Enabling distributed energy storage by incentivizing small load shifts. *ACM Trans Cyber-Phys Syst* 1(2):1–30
- Iyengar S, Lee S, Irwin DE, Shenoy PJ (2016) Analyzing energy usage on a city-scale using utility smart meters. In: *BuildSys*
- Iyengar S, Sharma N, Irwin D, Shenoy P, Ramamritham K (2014) SolarCast: a cloud-based black box solar predictor for smart homes. In: *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*
- Karmakar G, Kabra A, Ramamritham K (2015) Maintaining thermal comfort in buildings: feasibility, algorithms, implementation, evaluation. In: *Proceedings of the Real-Time Systems Symposium*
- Kasaei MJ, Gandomkar M, Nikoukar J (2017) Optimal management of renewable energy sources by virtual power plant. *Renew Energy* 114:1180–1188
- Keck F, Lenzen M, Vassallo A, Li M (2019) The impact of battery energy storage for renewable energy power grids in Australia. *Energy* 173:647–657
- Keshav S (2016) Technical perspective: the chemistry of software-defined batteries. *Commun ACM* 59(12):110–110
- Lam AH, Yuan Y, Wang D (2014) An occupant-participatory approach for thermal comfort enhancement and energy conservation in buildings. In: *Proceedings of the 5th International Conference on Future Energy Systems*
- Liu N, Yu X, Wang C, Li C, Ma L, Lei J (2017) An energy sharing model with price-based demand response for microgrids of peer-to-peer prosumers. *IEEE Trans Power Syst* 32(5):3569–3583
- Liu J, Zhang N, Kang C, Kirschen D, Xia Q (2017) Cloud energy storage for residential and small commercial consumers: A business case study. *Appl Energy* 188:226–236
- LO3 Energy: Brooklyn micro-grid (2018). <https://lo3energy.com/>. Accessed 05 May 2021
- McPherson M, Tahseen S (2018) Deploying storage assets to facilitate variable renewable energy integration: The impacts of grid flexibility, renewable penetration, and market structure. *Energy* 145:856–870
- Mishra A, Irwin D, Shenoy P, Kurose J, Zhu T (2012) Smartcharge: Cutting the electricity bill in smart homes with energy storage. In: *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*
- Mishra, A, Irwin D, Shenoy P, Kurose J, Zhu T (2013) Greencharge: Managing renewable energy in smart buildings. *IEEE J Sel Areas Commun* 31(7):1281–1293
- Mnatsakanyan A, Kennedy SW (2015) A novel demand response model with an application for a virtual power plant. *IEEE Trans Smart Grid* 6(1):230–237
- Narayanawamy B, Garg VK, Jayram T (2012) Online optimization for the smart (micro) grid. In: *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*
- Pan J, Jain R, Paul S, Vu T, Saifullah A, Sha M (2015) An internet of things framework for smart energy in buildings: designs, prototype, and experiments. *IEEE Internet Things J* 2(6):527–537
- Pan D, Wang D, Cao J, Peng Y, Peng X (2013) Minimizing building electricity costs in a dynamic power market: Algorithms and impact on energy conservation. In: *Proceedings of the 34th IEEE Real-Time Systems Symposium*. IEEE
- PowerLedger: Peer to peer energy trading (2018). <https://powerledger.io/>. Accessed 05 May 2021
- Ramamritham K, Karmakar G, Shenoy P (2017) Smart energy management: A computational approach. In: Krishna RP, Ashish S, Sharma C, Subhash B (eds). *Big Data Analytics*. Springer International Publishing, Cham. pp 3–14
- Schneider Solar Hybrid Inverter Systems (2017). <https://solar.schneider-electric.com/>. Accessed 05 May 2021
- Sharma P, Chaufournier L, Shenoy P, Tay Y (2016) Containers and virtual machines at scale: A comparative study. In: *Proceedings of the 17th International Middleware Conference*

- Tanenbaum AS, Van Steen M (2007) *Distributed Systems: Principles and Paradigms*. 2nd ed.. Prentice-Hall, Upper Saddle River. <https://www.bibsonomy.org/bibtex/2aeb78dfa5b91b66dce027f86ecdeeb81/flint63>
- Thakur S, Saha M, Singh A, Agarwal Y (2014) Wattshare: Detailed energy apportionment in shared living spaces within commercial buildings. In: *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*
- Wang Z, Gu C, Li F, Bale P, Sun H (2013) Active demand response using shared energy storage for household energy management. *IEEE Trans Smart Grid* 4(4):1888–1897
- Wang H, Zhang JX, Li F (2015) Incentive mechanisms to enable fair renewable energy trade in smart grids. In: *Sixth International Green Computing Conference and Sustainable Computing Conference (IGSC)*
- Weather dataset API (2017). <https://www.wunderground.com/>. Accessed 05 May 2021
- Wei P, Chen X, Chandrasekaran R, Song F, Jiang X (2016) Adaptive and Personalized Energy Saving Suggestions for Occupants in Smart Buildings. In: *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*
- Weng T, Agarwal Y (2012) From buildings to smart buildings-sensing and actuation to improve energy efficiency. *IEEE Des Test Comput* 29(4):36–44
- When Will Rooftop Solar Be Cheaper Than the Grid? (2016). <https://goo.gl/h1Ayy5>. Accessed 05 May 2021
- Wholesale Electricity Price in the US (2017). <https://www.eia.gov/electricity/wholesale/>. Accessed 05 May 2021
- Wisconsin Electric Rates (2017). <http://bit.ly/11RsroB>. Accessed 05 May 2021
- Wood T, Shenoy PJ, Venkataramani A, Yousif MS, et al (2007) Black-box and gray-box strategies for virtual machine migration. In: *NSDI*
- Wu C, Kalathil D, Poolla K, Varaiya P (2016) Sharing electricity storage. In: *Proceedings of the 55th Conference on Decision and Control*. IEEE
- Yu S, Fang F, Liu Y, Liu J (2019) Uncertainties of virtual power plant: Problems and countermeasures. *Appl Energy* 239:454–470
- Zhong W, Huang Z, Zhu T, Gu Y, Zhang Q, Yi P, Jiang D, Xiao S (2014) ides: Incentive-driven distributed energy sharing in sustainable microgrids. In: *Proceedings of the International Green Computing Conference*
- Zhu T, Huang Z, Sharma A, Su J, Irwin D, Mishra A, Menasche D, Shenoy P (2013) Sharing renewable energy in smart microgrids. In: *ACM/IEEE International Conference on Cyber-Physical Systems*

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
