

vSolar: Virtualizing Community Solar and Storage for Energy Sharing

Stephen Lee, Prashant Shenoy, Krithi Ramamritham¹, David Irwin
University of Massachusetts, Amherst ¹Indian Institute of Technology, Bombay

ABSTRACT

Since many residential locations are unsuitable for solar deployments due to space constraints, community-owned solar arrays with energy storage that are collectively shared by a group of homes have emerged as a solution. However, such a group-owned system does not allow individual control over how the electricity generation from the solar array and energy stored in the battery is used for optimizing a home's electricity bill. To overcome this limitation, we propose vSolar, a technique that virtualizes community solar and battery arrays such that each virtual system can be independently controlled, regardless of others. Further, we present mechanisms and algorithms that allow homes with surplus energy to lend to homes with deficit energy.

ACM Reference Format:

Stephen Lee, Prashant Shenoy, Krithi Ramamritham¹, David Irwin. 2018. vSolar: Virtualizing Community Solar and Storage for Energy Sharing. In *e-Energy '18: The Ninth International Conference on Future Energy Systems, June 12–15, 2018, Karlsruhe, Germany*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3208903.3208932>

1 INTRODUCTION

In recent years, the penetration of renewable energy sources such as solar and wind in the electric grid has continued to grow. Due to continuous improvement in technology, the costs of deploying solar arrays have continued to drop, and the leveled cost of energy is now 12.5¢ per kWh [10], on par or less than traditional energy sources. This has fueled the growth of solar installations worldwide, and 70 gigawatts of solar capacity was deployed in 2016 alone [10].

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. *Community solar arrays (CSA)* have emerged as a solution to these challenges. CSA 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 array and is allocated a certain fraction of the solar array in proportion to their share.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

e-Energy '18, June 12–15, 2018, Karlsruhe, Germany

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-5767-8/18/06...\$15.00
<https://doi.org/10.1145/3208903.3208932>

As power output of solar arrays is intermittent, community storage systems are used in conjunction with CSA to smooth out fluctuations. A community storage system consists 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 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 systems generate energy for all owners as a single aggregated system and do not permit individual control to the users. To overcome this limitation, we argue that community solar and storage should be virtualized to maximize its effectiveness. Similar to how virtual machines provide an individual server abstraction and are multiplexed onto a physical machine, a virtual solar and battery arrays provide an abstraction of individually-owned solar and battery arrays that are multiplexed onto the community-owned physical solar and battery array. Importantly, virtualization enables each owner to independently manage their solar generation and stored energy as if it were a dedicated 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. Prior works have discussed the benefits of a shared pool of energy storage [8, 9] and energy sharing [4, 5, 11]. In our paper, we present mechanisms to virtualize community solar and storage and also to enable flexible energy sharing algorithms in such systems.

2 BACKGROUND

A non-virtualized community-owned system does not allow independent control of the output in the shared system. Energy output from the solar and battery is used to meet the aggregate energy demand across all homes, but doesn't permit individual owners to independently utilize their share of energy. On the other hand, dedicated systems allow independent control but are more expensive and may be infeasible in many residential locations. In contrast, a virtual solar and battery arrays give each owner an illusion of ownership of the unit. Such local control allows each owner to make the "optimal" decision on how to utilize their virtual array and batteries independently of other owners.

Our work assumes a community solar and storage array that is collectively-owned by a group of residents. Each resident is assumed to own a certain fraction of the community solar and storage array. 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

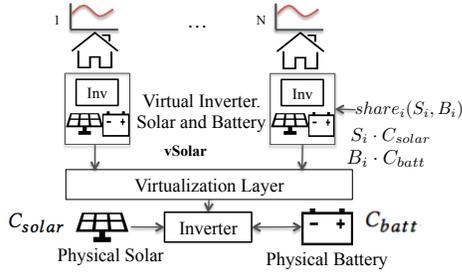


Figure 1: vSolar architecture diagram.

manner to perform energy optimization or reduce energy bills. 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).

3 VSOLAR DESIGN

3.1 vSolar 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 Figure 1). 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. To implement this abstraction, vSolar exposes a set of virtualization primitives ($charge_i(t)$, $discharge_i(t)$, $send_to_grid_i(t)$, $draw_from_grid_i(t)$) that can be controlled by software algorithms in each virtual controller (see Appendix A.1).

3.2 vSolar Virtualization Algorithm

We now present the vSolar virtualization algorithm that uses the above primitives to implement software control of the virtual solar and battery system within the virtual controller. Since the solar output of virtual array i is $solar_i(t)$ and demand of home i is $demand_i(t)$; the vSolar algorithm can determine if the current solar output is adequate to satisfy the demand. If so, the net surplus is $surplus_i(t) = \max(solar_i(t) - demand_i(t), 0)$. If not, the net deficit is $deficit_i(t) = \max(demand_i(t) - solar_i(t), 0)$.

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:

$$charge_i(t) = \min(max_charge_rate, surplus_i(t)) \quad (1)$$

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:

$$send_to_grid_i(t) = solar_i(t) - demand_i(t) - charge_i(t) \quad (2)$$

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:

$$draw_from_grid_i(t) = deficit_i(t) \quad (3)$$

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

$$discharge_i(t) = \min(deficit_i(t), max_discharge_rate) \quad (4)$$

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.

$$draw_from_grid_i(t) = \max(demand_i(t) - solar_i(t) - discharge_i(t), 0) \quad (5)$$

Thus, the vSolar algorithm within each virtual controller can make independent decisions based on the solar output, battery level and demand of each home.

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

$$charge(t) = \sum_{i=1}^N charge_i(t) - \sum_{i=1}^N discharge_i(t) \quad (6)$$

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

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

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:

$$send_to_grid(t) = \sum_{i=1}^N (send_to_grid_i(t) - draw_from_grid_i(t)) \quad (8)$$

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.

3.3 vSolar Energy Sharing Algorithm

vSolar's virtualization algorithm allows each owner to operate their virtual solar and battery array independently of others. Since all virtual arrays are multiplexed onto a common physical array, there are opportunities for the virtual systems to *collaborate* with one another. One form of collaboration is *energy sharing* where virtual systems with surplus solar generation or 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.

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 temporarily increase its capacity by borrowing from surplus homes. 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 energy sharing, vSolar virtual inverters need two additional virtualization primitives: $borrow_i(t)$ and $lend_i(source, t)$. These primitives enable a virtual controller to implement any energy sharing algorithm that is best suited to its needs.

We design the vSolar energy sharing algorithm based on the vSolar virtualization algorithm described above. Our energy sharing algorithm is as follows. 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:

$$lend_i(solar, t) = solar_i(t) - demand_i(t) - charge_i(t) \quad (9)$$

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:

$$lend_i(battery, t) = \min(max_discharge_rate - discharge_i(t), 0) \quad (10)$$

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:

$$borrow_i(t) = demand_i(t) - solar_i(t) - discharge_i(t) \quad (11)$$

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 battery 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:

$$borrow(t) = \sum_{i=1}^N borrow_i(t) \quad (12)$$

The solar capacity available for lending is computed as:

$$lend(solar, t) = \sum_{i=1}^N lend_i(solar, t) \quad (13)$$

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:

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

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

$$draw_from_grid_i(t) = borrow_i(t) - borrowed_power_i(t) \quad (15)$$

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

$$send_to_grid_i(t) = lend_i(solar, t) - lent_solar_i(t) \quad (16)$$

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.

To determine the amount of energy a home can lend requires predicting future demand and solar output. We build a demand model using Support Vector Machines (SVM) to predict the electricity needs of the home for the following day. Separately, we use the technique discussed in [7] to predict the solar output for the next day. Combining the future solar output, electricity demands generated from the model, and using the current battery capacity a candidate home can estimate the amount of energy to lend from the virtual battery. Specifically, the surplus energy is defined as

$$surplus_energy_i(t) = \max(future_solar_i(t) + charge_i(t) - future_demand_i(t), 0) \quad (17)$$

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

$$lend_i(battery, t) = \min(surplus_energy_i(t), max_discharge_rate_i - discharge_i(t)) \quad (18)$$

Our energy sharing algorithm is summarized as a flowchart in Appendix A.2.

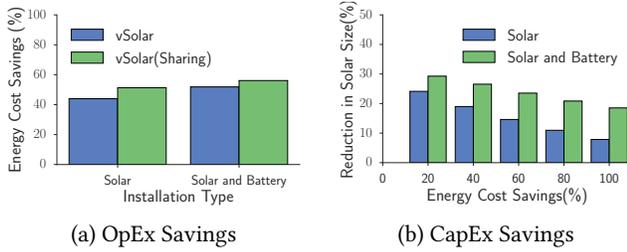


Figure 2: (a) Median cost savings using 40 kW solar and 80 kWh battery. (b) Reduction in solar size through energy sharing in comparison to a dedicated solar and battery array.

4 EVALUATION

We focus on evaluating the potential benefit of vSolar using trace-driven simulations. To do so, we use real electricity load dataset from 50 homes over a two year period between 2014 to 2015. The electricity dataset was gathered from the New England region of United States and consists of energy consumption information at a resolution of 30 minutes [6]. 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 greater than one. In contrast, *night profile* homes use energy mostly during the off-peak pricing period. Next, we randomly select homes 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 [3]. Typically, wholesale prices are 30% to 50% of the retail price [2]. We assume the wholesale electricity rates to be 40% of the retail price and the apartments share energy to others at prices between wholesale and retail 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 [1].

4.1 Experimental Results

vSolar Benefits: Figure 2(a) shows the median energy cost savings of a home for both vSolar and the energy sharing scenario. With 40 kW of solar array, vSolar achieves 43% of energy cost savings and yields 8.8% higher savings when coupled with an 80 kWh battery. The energy cost savings further increases when energy is shared. Compared to vSolar, 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. Figure 2(b) shows the reduction in solar array size (CapEx) of a home with a dedicated system can achieve when the system is virtualized, and energy is shared with others. We note that a dedicated system can reduce its solar array size by 14.6%, through virtualization and energy sharing, to achieve 60% energy

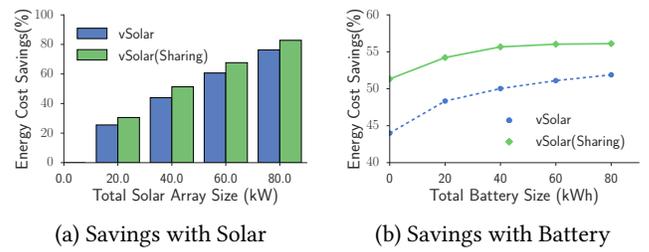


Figure 3: Median energy cost savings with varying sizes.

cost savings, which is 8.6 kW reduction in absolute values. Using a battery capacity of 80 kWh, which is roughly 4 kWh per home, we note a dedicated system can achieve 23.5% reduction in solar array size, which is 13.8 kW reduction in absolute values.

Impact of Solar and Battery Arrays. Figure 3(a) shows the median energy cost savings across homes using vSolar’s virtualization for different solar array size and a battery size of 80kWh. The cost savings is computed by comparing the energy cost of a home using vSolar to its original energy cost. Clearly, the energy cost savings is higher with energy sharing than the no sharing scenario. 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 vSolar. This is because 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.

Figure 3(b) shows the median energy cost savings for different battery sizes and a solar array size of 40 kW. We observe that the cost savings increase from 43% to 50% with a 40 kWh battery size. Moreover, the cost savings increases an additional 5.6% with energy sharing. 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 does not draw energy from batteries at all times, owners can allow energy discharge from their share to fulfill part or all of the local demand to reduce energy costs.

5 CONCLUSION

In this paper, we proposed vSolar — a mechanism to virtualize community-owned solar and battery, wherein each virtual solar and battery can be assigned to an owner and controlled independently, regardless of other. Further, vSolar provides virtualization abstractions that allow each owner to implement their custom energy optimization policy. To show vSolar’s potential, we implemented an energy sharing algorithm that enables energy sharing to minimize their local electricity bill. We compared vSolar to a non-virtualized community-owned system and showed similar cost savings while providing local control of the virtualized system. Further, we showed that energy sharing using vSolar achieves both OpEx and CapEx savings over non-virtualized community-owned and dedicated solar and battery systems.

Acknowledgment: This research is supported by NSF grants IIP-1534080, CNS-1645952, CNS-1405826, CNS-1253063, CNS-1505422, and the Massachusetts Department of Energy Resources.

REFERENCES

- [1] 2017. *Weather dataset API*. <https://www.wunderground.com/>.
- [2] 2017. *Wholesale Electricity Price in the US*. <https://www.eia.gov/electricity/wholesale/>.
- [3] 2017. *Wisconsin Electric Rates*. <http://bit.ly/11RsroB>.
- [4] T Carpenter, S Singla, P Azimzadeh, and S Keshav. 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*.
- [5] Z Huang, T Zhu, Y Gu, and Y Li. 2016. Shepherd: sharing energy for privacy preserving in hybrid AC-DC microgrids. In *Proceedings of the Seventh International Conference on Future Energy Systems*. ACM.
- [6] S Iyengar, S Lee, D Irwin, and P Shenoy. 2016. Analyzing Energy Usage on a City-scale using Utility Smart Meters. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*. ACM.
- [7] S Iyengar, N Sharma, D Irwin, P Shenoy, and K Ramamritham. 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*.
- [8] J Liu, N Zhang, C Kang, D S Kirschen, and Q Xia. 2017. Decision-Making Models for the Participants in Cloud Energy Storage. *IEEE Transactions on Smart Grid* (2017).
- [9] A Mnatsakanyan and S W Kennedy. 2015. A novel demand response model with an application for a virtual power plant. *IEEE Transactions on Smart Grid* (2015).
- [10] When 2016. When Will Rooftop Solar Be Cheaper Than the Grid? <https://goo.gl/h1Ayy5>. (March 31st 2016).
- [11] C Wu, D Kalathil, K Poolla, and P Varaiya. 2016. Sharing electricity storage. In *55th Conference on Decision and Control (CDC)*. IEEE.

A APPENDIX

A.1 Virtualization API Abstractions

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 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 rate at which electricity should be drawn from the electric grid to meet a portion of $demand_i(t)$
- $borrow_i(t)$ which specifies the amount of power that home i wishes 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.

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

A.2 Energy Sharing Algorithm Flowchart

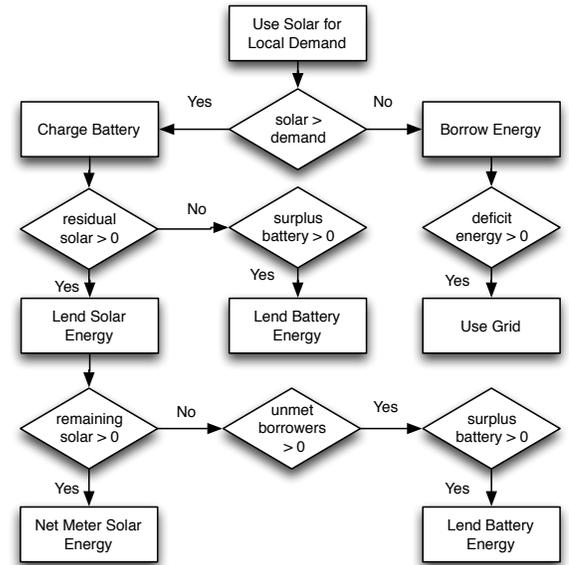


Figure 4: A flow chart of vSolar's energy sharing algorithm.