Enforcing Fair Grid Energy Access for Controllable Distributed Solar Capacity

Noman Bashir
UMass Amherst
nbashir@umass.edu

David Irwin
UMass Amherst
deirwin@umass.edu

Prashant Shenoy
UMass Amherst
shenoy@cs.umass.edu

Jay Taneja
UMass Amherst
jtaneja@umass.edu

ABSTRACT

The rapid expansion of intermittent grid-tied solar capacity is making the job of balancing electricity's real-time supply and demand increasingly challenging. To address the problem, recent work proposes mechanisms for actively *controlling* solar power output to the grid by enabling software to cap it as a fraction of its time-varying maximum output. Utilities can use these mechanisms to dynamically share the grid's solar capacity by controlling the solar output at each site. However, while enforcing an equal fraction of each solar site's time-varying maximum output results in "fair" short-term contributions of solar power, it does not result in "fair" long-term contributions of solar energy. This discrepancy arises from fundamental differences in enforcing "fair" access to the grid to contribute solar energy, compared to analogous fair-sharing in networks and processors. In this paper, we present a centralized and distributed algorithm to enable control of distributed solar capacity that enforces fair grid energy access. We implement our algorithm and evaluate it on synthetic data and real data across 18 solar sites. We show that traditional rate allocation, which enforces equal rates, results in solar sites contributing up to 18.9% less energy than an algorithm that enforces fair grid energy access over a single month.

CCS CONCEPTS

•**Hardware → Smart grid;**

KEYWORDS

Solar, Fairness, Net Meter

ACM Reference format:
Noman Bashir, David Irwin, Prashant Shenoy, and Jay Taneja. 2017. Enforcing Fair Grid Energy Access for Controllable Distributed Solar Capacity. In *Proceedings of The 4th International Conference on Systems for Energy-Efficient Built Environments, Delft, The Netherlands, November 8–9, 2017 (BuildSys),* 11 pages.
DOI: 10.1145/3137133.3137147

BuildSys, Delft, The Netherlands

978-1-4503-5544-5/17/11…$15.00
DOI: 10.1145/3137133.3137147

1 INTRODUCTION

The amount of grid-tied solar power continues to grow at an exponential rate with capacity increasing by an average of 33% each year over the past six years [6]. This growth is driven by consistent drops in solar module prices, which have fallen 10% per-year on average over the past three decades. In many locations, the average cost of solar energy is now less than the cost of energy from fossil fuels. As a result, some estimates project that solar could contribute as much as 20% of global electricity consumption as early as 2030 [5].

Unfortunately, the increasing penetration of solar energy in the grid complicates utility operations. In particular, utilities are responsible for balancing electricity's real-time supply and demand, requiring them to compensate for variations in solar output. However, compensating for large solar variations using mechanical generators is challenging, since generator ramp rates are less than solar ramp rates. As a result, governments generally *limit* the amount of grid-tied solar capacity that can feed energy into the grid. In the U.S., these limits vary widely by state, and often restrict both the percentage of users with grid-tied solar, and their aggregate solar power capacity. The rapid growth in solar power is now causing states to frequently hit these limits, triggering a complex political process (often taking many months) that is required to raise them. Since the limits, which are a form of admission control, are hard, once they are hit, additional users cannot install grid-tied solar until they are raised. For example, users in Hawaii were recently barred from installing grid-tied solar for two years [3, 10].

Importantly, the aggregate power limits above are static and based on the *rated installed capacity* of each solar site, and not the amount of power they actually generate. Standard Test Conditions (STC) for rating solar module capacity specifies an irradiance of 1kW/m^2 with an air mass of 1.5, no wind speed, and a cell temperature of 25C. These conditions approximate the generation of a south-facing solar module (tilted at the same angle as the Sun) at solar noon near the equinox on a clear sunny day in the U.S. with an ambient air temperature of 0C. Of course, weather conditions are rarely this "ideal:" the ambient air temperature at STC is unrealistic, roof lines dictate non-ideal orientations and tilts, and solar irradiance is usually much less than 1kW/m^2, e.g., during the morning, evening, over much of winter, and under cloudy skies.

Thus, the *actual* aggregate solar power generated is rarely, if ever, at (or even near) the rated capacity, and varies widely each day, over the year, and as the weather changes. For example, on cloudy days, the aggregate contribution of solar power across many distributed sites is much less than on sunny days. As a result, on a cloudy day,

the grid could potentially accept solar power from many sites that are currently forced off-grid without exceeding its capacity limit. To address the problem, recent work proposes mechanisms [13] and policies [8, 12] for actively controlling solar power output to the grid. This work enables software to cap the solar power injected to the grid as a configurable fraction of its time-varying maximum output [13], and then, inspired by similar problems in networks, designs rate allocation policies to limit the aggregate contribution of distributed solar subject to the grid's capacity [8, 12].

An important metric when determining how to dynamically limit each solar site's power output is preserving *fairness* between sites. Prior work co-opts the traditional notion of "fairness" from the networking literature, which computes it with respect to the instantaneous sending rates of flows, and not the cumulative amount of traffic they send over time. This makes sense in networking, as senders can potentially generate an arbitrary amount of traffic at any time. Thus, if one idle sender does not generate traffic for a long period, then i) other senders should be able to increase their rate to consume any excess bandwidth during this time, and ii) the idle sender should not be able to accumulate unlimited credit for their idleness, enabling them to monopolize the link once they resume sending. The former property ensures allocations are work-conserving, while latter property prevents starvation of senders. Analogously, prior work attempts to maintain "fair" grid rate allocations, such that each solar site contributes near the same fraction of their time-varying maximum instantaneous power output.

The problem is that this traditional notion of fairness in networks does not map well to the grid. Instead, we argue that the grid should express fairness in terms of the total fraction of *energy* users contribute over time (with respect to each other) rather than in terms of their instantaneous rates of *power*. Ultimately, users care about the amount of total solar energy they can feed into the grid (over some time window), as a fraction of the total solar energy they could possibly feed in, since this impacts both the cost of their system and the revenue it generates. In particular, users directly receive compensation for the energy they feed in, which decreases with the fraction of energy they can contribute. The expected fraction of energy users *cannot* feed into the grid may also necessitate additional system costs to store excess energy.

As we show, enforcing fair instantaneous rates, as in networking, may result in *unfair* contributions of total energy over time. Unlike in networking, solar sites can only generate "traffic" at certain times based the Sun's irradiance, which is a function of location, time, local weather, and physical installation characteristics. Importantly, solar sites *cannot control their location, the Sun, the weather, and often their physical characteristics*, and thus have no control over when and how much solar power they can generate. In contrast, network clients that are not generating traffic are doing so voluntarily, and *could generate traffic if desired.* Clearly, if network clients directly received compensation for the total amount of data they sent, they most certainly would generate traffic all the time, and the total amount of data they sent over time would be critically important.

This paper identifies this fundamental difference between fair rate allocation in networks and fair grid energy access for solar, and discusses how and why it arises. We then design a rate allocation algorithm to enforce weighted fair grid energy access and evaluate its tradeoffs. In doing so, we make the following contributions.

Fairness Definition. While preserving fairness is a first-class concern when sharing processors and networks, it has generally not been a metric of interest in electric grids. We introduce and define the notion of distributed solar fairness (DSF), and discuss how it differs from similar notions of fairness in computer systems and networking. We also discuss how unfairness arises among distributed solar sites with limits on their aggregate solar output.
Fair Energy Allocation Algorithm. We propose a simple energy allocation algorithm to enforce fair grid energy access among distributed solar sites. While this algorithm allocates rates to different solar "flows" over time, as in computer systems and networks, it varies these rates to ensure users contribute the same fraction of their actual solar energy capacity. The algorithm exposes tradeoffs in its convergence speed, fidelity to the aggregate limit it enforces, and robustness, i.e., the interval over which it must exchange data.
Implementation and Evaluation. We implement our algorithm above and evaluate it on both synthetic data and real data from 18 solar sites. We show that traditional equal rate allocation results in solar sites contributing up to 18.9% less energy over a single month than our algorithm that enforces fair grid energy access.

2 FAIRNESS IN THE ELECTRIC GRID

In this paper, we consider grid-tied solar arrays capable of "net metering." Currently, the grid allows a net metered grid-tied solar array to feed any amount of power into the grid, up to its maximum installed capacity, with no restrictions. Thus, the "admission control" decision of whether to allow a solar array to net meter must be made at installation time. As discussed earlier, this limits the number of solar installations the grid can permit, since policies must plan for the worst-case scenario, i.e., where all solar arrays concurrently feed in their maximum capacity, even though this scenario is highly unlikely (if not impossible), and can only occur one time per year.

Thus, enforcing such limits at "run time" has the potential to enable a much larger number of grid-tied solar arrays, while still limiting the total net metered power to a pre-specified capacity. In the future, we expect the grid to rate control the amount of power that can be injected by a grid-tied solar array. The ability to rate control solar arrays at the time-scale of minutes or hours has many benefits. For example, rate control can enable any number of solar arrays to connect to the grid, while limiting their stochasticity. In addition, rate control can also incentivize energy storage to save surplus solar power that the grid cannot accept due to the limits.

Given such a scenario, we examine the problem of how the grid should assign rates to different solar arrays, while maintaining an aggregate limit on solar output and fairness across users. Prior work has used an analogy to the rate allocation problem in computer networks and has taken the notion of fairness from networking. Specifically, prior work uses analytical models of TCP's rate control algorithm and weighted versions of this rate allocation problem to model the problem [2, 8, 12]. However, with solar, owners directly receive compensation for the solar energy they contribute and thus are incentivized to always produce as much power as possible.

Thus, rather than using a notion of fairness from networking, we instead propose a new fairness metric for rate-controlled solar arrays called *distributed solar fairness (DSF)* that is based on net metered compensation. Let $E_i^{actual}(t_2 - t_1)$ denote the actual energy

net metered by a solar array i over a duration $[t_1, t_2)$ in the presence of rate control, and $E_i^{max}(t_2 - t_1)$ denote the maximum amount of energy it *could have produced* in this time period with no rate control, e.g., using standard techniques such as maximum power point tracking (MPPT). Note that a site's maximum generation potential varies over time based on a site's unique location, weather, and physical characteristics. Since rate control reduces the total energy that can be produced, the reduction in net metered revenues over the interval $[t_1, t_2)$, which we term as $loss_i(t_2 - t_1)$, is $1 - \frac{E_i^{actual}(t_2-t_1)}{E_i^{max}(t_2-t_1)}$. This can be viewed as a direct monetary loss incurred by solar array i over the specified time interval due to rate control.

To be fair across users, we require that the percentage loss is the same for all arrays over any time interval $[t_1, t_2)$. Thus our notion of fairness requires that for any two arrays i and j,

$$|loss_i(t_2 - t_1) - loss_j(t_2 - t_1)| < \epsilon \quad (1)$$

While our ideal definition of fairness requires that this condition be true over any arbitrary time interval, in practice, achieving fairness over very short time scales may be infeasible. For example, if the sun has risen at the location of array i but it has yet to rise at the location of array j, it is not possible to guarantee fairness over a small time scale, since array j is unable to produce any power. In the next section, we describe a number of factors that complicate enforcing fairness at short time scales. However, it is both acceptable and feasible to enforce fairness over the much longer time scale of hours, days, or even at the time scale of a monthly billing cycle. In general, consumers' primary concern is whether their monetary percentage loss is fairly distributed across all arrays over these longer time scales. Thus, in practice, the grid only needs to ensure fairness over these longer intervals $[t_1, t_2)$.

In the case of networks, fairness guarantees are provided only when the network flows are *backlogged*, which requires that the flows can continuously send data when network capacity is available. In our case, providing fairness over very short time scales also requires that the solar arrays be capable of producing enough power to use their allocated rates. However, over longer time scales, it is possible for an array to not use its instantaneous allocation, since it is unable to produce sufficient power, and yet "catch" up later by injecting power at higher rates than other arrays.

Even when enforcing fairness over longer time scales, the problem of allocating rates to each array is complicated by many factors. For instance, a simple approach that allocates identical rates to two arrays of identical size can yield unfair results. This is because arrays of identical size can still produce vastly different power output at any instant due to local differences in weather, as well as factors such as tilt, orientation, and location. Ignoring these differences can cause the fairness measure to diverge for various arrays. Thus, a fair rate allocation algorithm must consider several factors: assuming identical weather conditions, two arrays at two different locations will have slightly different sunrise, sunset and solar noon times, yielding solar output curves that are time-shifted with respect to one another. In the networking case, this is analogous enforcing fairness for *time shifted* flows, where to two identical flows are time-shifted and start transmitting data with different start times. Similarly, two solar arrays that are in proximity to one another may also produce different output due to micro climates, different shading effects, etc. Finally, different arrays may have vastly different capacities and thus rates must be computed to equalize the percentage loss for such heterogeneous size arrays.

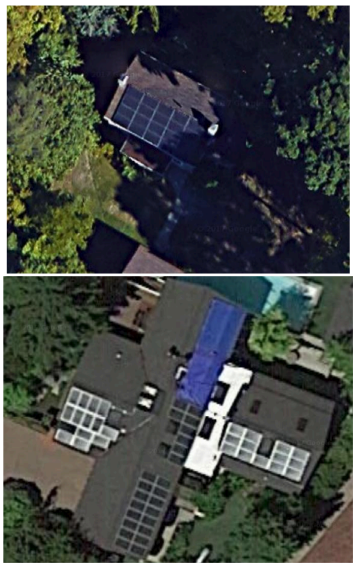

Figure 1: Illustrative examples of non-ideal solar sites.

Next, we describe fair rate allocation for solar arrays that achieves our notion of fairness while accounting for these factors.

3 FAIR GRID ENERGY ALLOCATION

The previous section compares the different notions of fairness in computer systems and networks, and in the grid. In this section, we first examine how unfairness arises from the differences in the shape of solar output across multiple sites. We describe the different types of effects that cause the "shape" of a solar curve to differ even across sites that are near each other. We then present our fair energy allocation algorithm, and its tradeoffs.

3.1 Solar Shape Diversity

Unfairness in solar energy access to the grid derives from the difference in output between solar sites, even when they are near each other. There are many reasons why solar output between solar sites differs. We describe the different reasons below.

Solar Potential. The Sun's position in the sky is unique at each location on Earth at each instant of time. The Sun's position in the sky affects the air mass light must travel through, which reduces the amount of irradiance that reaches the ground. As a result, even with clear skies, the maximum solar generating potential is different at every solar site at any moment.

Weather Effects. The weather also affects solar generation potential. In particular, solar power correlates with cloud cover, which is much more stochastic and localized than other weather metrics, such as temperature. For example, microclimates, such as those near large bodies water, can cause weather, and thus solar generation potential, to be significantly different at two nearby locations.

Physical Characteristics. Finally, the physical characteristics of a solar site also affect its solar output. These include the solar module's tilt and orientation, as well as any occlusions from surrounding buildings, trees, or mountains that may shade them. For example, an east-facing solar module will both start and stop generating power well before a west-facing one in the morning and evening, respectively. In general, rooftop solar deployments are complex and not ideal. Figure 1 illustrates typical rooftop solar deployments with multiple modules at different non-ideal tilts and orientations with significant shading from trees and other surroundings. In addition, soiling from debris can also cause solar generation to differ between two nearby sites with identical solar modules.

The differences above manifest themselves as differences in the shape of solar output at each site. We characterize these differences below, which are the root of unfairness in solar allocation.

Shifts. Shifts occur when a solar curve is shifted with respect to another solar curve, such that the first curve starts before or ends after another curve. Shifts occur either due to differences in the orientation of modules or differences in location. For example, east- and west-facing modules at the same location will be shifted with respect to each other. Similarly, a difference in longitude between two locations also results in a shift, since the sun rises and sets at different times (for the same daylength).

Squeeze. Squeezes occur when a solar curve is narrower with respect to another solar curve, such that the first curve starts before and ends after another curve. Squeezes occur either due to differences in the tilt of modules or differences in location. For example, a south-facing vertically tilted module will be squeezed with respect to a horizontally flat tilted one. Similarly, a difference in latitude between two locations also results in a squeeze, since the length of a day changes with latitude.

Dips and Cuts. Dips occur when the solar output drops below the power level seen when the sky is clear. Dips may be caused by clouds, shade from trees, or nearby buildings and reduce the amount of sunlight seen by an array. The amount of the power dip depends on the magnitude of the reduction in the sunlight seen by at array. Similarly, cuts occur when a solar curve's power is cut-off (or blocked) with respect to another solar curve, such that the first curve generates power normally while the second curve generates nothing. Cuts typically occur in the morning and evening, since these blockages are more prevalent when the Sun is low in the sky.

Note that each solar deployment can exhibit an arbitrary combination of the characteristics above. These characteristics are also static, since they are purely a function of a site's location, physical characteristics, and surroundings. As a result, if a solar site experiences a shift, squeeze, or dip relative to another solar site one day, it will often experience it every day (although the extent of it may change). In addition, different weather conditions between sites also create differences in the solar curves.

Figure 2 illustrates how two nearby homes can exhibit different solar output over a day. In this case, Home 2, is more east-facing, as in Figure 1 (bottom), than Home 1, and thus its power generation is shifted with respect to Home 1 on this day. However, Home 2 has a cut near the end of the day, indicating a blockage in solar output that causes its output to drop to zero, as in Figure 1 (top), which has trees on its west-side that block sunlight near the end of the day. In this case, imposing a limit on the aggregate power from

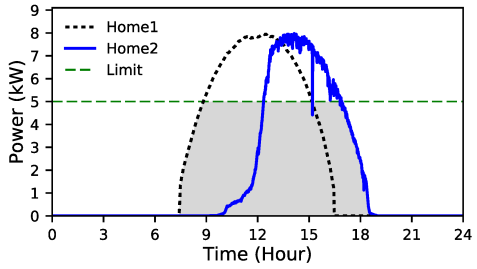

Figure 2: Profile of solar output for two homes 80km apart.

the two homes, and then satisfying this limit by allocating equal rates of solar power output between the two homes results in an unequal solar energy contribution at the end of the day.

This occurs because at the beginning of the day Home 2 is generating no power, and thus Home 1 is able to contribute a high fraction of its generation up to the limit. Due to the cut in power, once Home 2 starts generating power it must share the grid with Home 1 by contributing an equal fraction of its time-varying maximum power potential up to the limit, even though Home 1 has already a contributed a significant amount of energy to the grid. Thus, even though Home 2 contributes the same fraction of power as Home 1 at all times, its *fraction of energy* always remains less than Home 1, since it is never able to catch up.

3.2 Fair Energy Allocation Algorithm

We assume a mechanism exists to remotely control the time-varying fraction of maximum power a solar deployment contributes to the grid, as described in recent work [13]. We also assume that a grid balancing authority sets limits on the aggregate solar energy output across all solar sites by controlling this mechanism at each individual site. We assume that the grid's transformers and feeders are well-provisioned to handle the maximum solar generation, such that the transformers never exceed their capacity and feeders do not reverse their power flow. These assumptions are likely true for the foreseeable future, as transformers and feeders are generally over-provisioned for energy consumption, and grid-tied solar power actually reduces the energy consumption. As a result, we need not consider the impact of the grid's topology or the capacity of its distribution infrastructure in determining rate allocations.

Instead, the grid balancing authority sets aggregate limits on the distributed solar output based solely on net metering regulations. However, in our case, we assume these limits are dynamic and based on actual solar generation, rather than static and based on the rated capacity of solar sites as is the case today. The balancing authority may also alter the limit to improve operations, such as increasing it during times of peak demand to allow more solar energy to flow into the grid. In this case, the curtailed solar power operates like high-quality reserve capacity or a demand response resource.

Our problem is to allocate the fraction of maximum power contributed by each site such that all sites contribute the same fraction of energy over each time window T. In general, we assume T is a long period, such as a week or a month, since it may be difficult or infeasible to ensure fairness over shorter time periods. The analogous rate allocation problem in networking, if we assume the

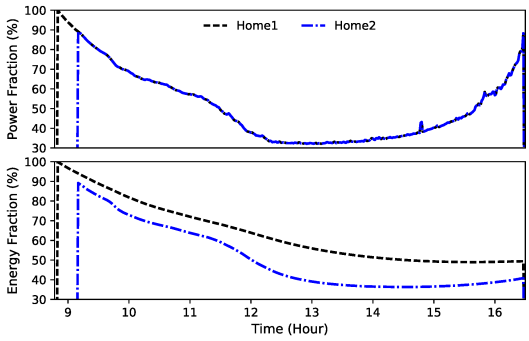

Figure 3: Divergence in the fraction of energy contributed by Homes 1 and 2 from Figure 2, even when the fraction of power they contribute is equal, assuming a 5kW limit.

Variable	Description
n	Number of solar sites
i	Index of sorted homes
T	Duration over which the fairness is enforced
$P_i(t)$	Maximum power that a site i can generate at time t
$P_i^{assigned}(t)$	Fraction of maximum power assigned to site i at time t.
Energy Fraction (EF)	Fraction of solar energy fed into the grid over interval T for a given site i.
Fair Energy Fraction (FEF)	Fair fraction of solar energy over interval T.
$L(t)$	Aggregate limit on solar capacity at time t.
$P^{avail}(t)$	Difference between aggregate limit and assigned power at sites at time t.
$P_{est}^{agg}(t)$	Estimated aggregate power
K	Correction gain

Table 1: Variable definitions for Algorithms 1-3.

grid's transformers and feeders are well-provisioned, is to simply enable all sites to contribute the same fraction (or rate) of their time-varying maximum power at all times. Thus, to enforce an aggregate limit, the balancing authority might enforce that all sites contribute only 50% of their maximum power. Note that, we assume the grid balancing authority specifies the aggregate limit as an absolute power (as in current net metering policies), and thus it will have to adjust the equal fraction of power contributed by each site over time as it varies to maintain the limit. In this case, we can compute this *equal rate* across all sites as simply the aggregate limit (L) divided by the sum of the current power output (P) of each of the n sites at any time t. We can augment this approach to include a weight, as in weighted fairness [4], such that the allocated rates are in proportion to each site's weights, rather than being equal.

$$Rate(t) = min(\frac{L(t)}{\sum_{i=1}^{n} P_i(t)}, 1) \quad (2)$$

However, as discussed above, this does not result in an equal (or weighted) contribution of energy over time. Figure 3 illustrates this behavior for Homes 1 and 2 in Figure 2. While the rate, expressed as a fraction of each site's maximum generation potential, is always equal (top), the fraction of energy each contributes diverges (bottom). In this case, the aggregate limit is set to 5kW throughout the day. Since Home 2 does not generate any power early in the day, Home 1 is able to feed a disproportionate amount of energy into the grid. Then, once Home 1 starts generating power, Home 1 and Home 2 each feed power in with equal rates. However, as the bottom graph indicates, the initial generation early in the day enabled Home 1 to feed in more energy (as a fraction of its total energy generation potential) relative to Home 2. In this case, Home 1 fed in 10% more energy than Home 2 in only a single day.

To address this problem, we design a rate allocation algorithm that enforces fair energy access to the grid. We first discuss a centralized version of this algorithm, assuming a tightly-coupled system, and then present a distributed version. In both cases, the algorithms first start by computing the equal rates above, and then determine which and how much sites can deviate from this equal rate based on the their current cumulative fraction of energy. We use the equal rate allocation as a starting point, since we need some basis for assigning initial rates to users. Equal rate allocation represents a good starting point, since under ideal conditions, i.e., where sites have exactly the same solar profile at all times, setting equal rates above *will* result in equal long-term energy contributions. Only when the solar profiles diverge does the equal rate allocation also diverge from a fair long-term energy allocation.

Centralized Algorithm. Algorithm 1 shows the pseudocode for our centralized algorithm, which we label as fast centralized allocation. Table 1 defines the algorithm's variables. In the centralized case, we assume that each solar site knows the fraction of solar energy each other site has fed into the grid over the current time window T, e.g., a month, which we call the Energy Fraction (EF). The algorithm then simply sorts each solar site by their EF, and assigns rates based on a solar site's position in the list. In particular, lower-ranked solar sites get allocated higher rates than higher-ranked solar sites to allow them to "catch up." The algorithm enables sites to catch up fast, since it allocates rates to 100% of solar power in sorted order, starting with the lowest-ranked site, until it reaches the aggregate power limit or it reaches a site that has an energy fraction equal to the mean across all sites, which we call the Fair Energy Fraction (FEF). At this point, the algorithm sets the rates of sites with energy fractions above the FEF based on the fair rate allocation algorithm above, but where the limit $L(t)$ is the remaining power after setting rates for the low-ranked sites. Thus, the algorithm is work-conserving in that it does not penalize sites that have contributed more than their fair energy by not allowing them to feed solar into the grid. As above, we can also apply a weight to each site, such that the fraction of energy they feed in should be in proportion to their weight.

One problem with the algorithm above is that it has the potential to starve out solar sites if other sites are not able to feed in solar for a long period. For example, after a snowstorm, the snow may melt off solar modules at different rates, enabling large differences in their maximum power. As a result, some solar site may not be able to feed power into the grid, and will thus "get behind" in terms of its energy contribution. Once the snow melts from this solar site, the algorithm above would set its rate to 100% until it catches up, which would reduce the rates of other flows. To mitigate the starvation problem, we can limit the catch-up rates for sites that are behind. In this case, rather than set these sites to 100%

Algorithm 1 Centralized Energy Allocation (Fast)

Require: $P_i(t)$ *and* $P_i^{assigned}(t)$ *for all homes over time* T, $L(t)$

1: *Compute* $EF_i = \frac{\sum_{t=0}^{T} P_i^{assigned}(t)}{\sum_{t=0}^{T} P_i(t)}, \ \forall\ i$
2: *Compute* $FEF = \frac{\sum_{i=0}^{n}\sum_{t=0}^{T} P_i^{assigned}(t)}{\sum_{i=1}^{n}\sum_{t=0}^{T} P_i(t)}$
3: *Sort & index homes in ascending order of EF* $(\uparrow_{i=1}^{n})$
4: $P^{avail}(t) = L(t) - \sum_{i=1}^{n} P_i^{assigned}(t)$
5: **while** $(P^{avail}(t) > 0)$ **do**
6: **if** $(EF_i < FEF)$ **then**
7: $P_i^{assigned}(t) = P_i(t)$
8: *Update* $P^{avail}(t),\ i++$
9: **else**
10: *break*
11: $Rate(t) = \frac{P^{avail}(t)}{\sum_{i=1}^{n} P_i(t)}$ *for homes above FEF*
12: *Update EF for all homes*

Algorithm 2 Centralized Energy Allocation (Slow)

Require: $P_i(t)$ *and* $P_i^{assigned}(t)$ *for all homes over time* T, $L(t)$

1: *Compute* EF_i *as in Algorithm* 1
2: *Compute FEF as in Algorithm* 1
3: *Sort & index homes in ascending order of EF* $(\uparrow_{i=1}^{n})$
4: *Compute fair rate* $Rate(t) = \frac{L(t)}{\sum_{i=1}^{n} P_i(t)}$
5: $P^{avail}(t) = L(t) - \sum_{i=1}^{n} P_i^{assigned}(t)$
6: **while** $(P^{avail}(t) > 0)$ **do**
7: **if** $(EF_i < FEF)$ **then**
8: $P_i^{assigned}(t) = (1 + (FEF - EF_i)) \times Rate(t)$
9: *Update* $P^{avail}(t),\ i++$
10: **else**
11: *break*
12: $Rate(t) = \frac{P^{avail}(t)}{\sum_{i=1}^{n} P_i(t)}$ *for homes above FEF*
13: *Update EF for all homes*

of their maximum power, we can set a limit between the equal rates computed in Equation 2 and 100%. In our algorithm, we apply proportional control to set these rates, such that the more behind a solar site, the faster it is able to catch up. In particular, we increase the rate in Equation 2 by the same proportion the site is behind in energy. Thus, if a solar site has 20% less than their "fair" fraction of global energy, we allow it to increase its rate in Equation 2 by 20%. Algorithm 2 shows the pseudocode for this algorithm, which we label as slow centralized allocation, where line 7 applies the proportional adjustment to the rate.

Distributed Algorithm. The centralized algorithms above assume accurate generation information is available from all solar sites in real-time, and that it is able to instantaneously set the rates of all solar sites without any delay. This implies that solar sites form a tightly-coupled system with utilities, where they stream generation data to utilities in real-time and utilities are able to instantaneously control their rates. Implementing such a tightly-coupled system is not realistic today. Most smart meters communicate wirelessly over the cell network and thus have limited bandwidth and periodic connectivity issues. In addition, a centralized approach represents a single point of failure and is not robust to network failures. Thus, we also present a distributed algorithm that uses incomplete information propagated at lower rates, e.g., minutes to hours.

Algorithm 3 Distributed Energy Allocation

Require: $L(t), P_{agg}^{est}(t)$, *and* $Rate(t)$ *over time* T

1: *Estimate aggregate power* $P_{agg}^{est}(t)$ *using gossip protocol*
2: *Compute* EF_i *as in Algorithm* 1
3: *Estimate* $FEF = \frac{\sum_{t=0}^{T}(P_{agg}^{est}(t) \times Rate(t))}{\sum_{t=0}^{T} P_{agg}^{est}(t)}$
4: *Compute fair rate* $Rate(t) = \frac{L(t)}{P_{est}^{agg}(t)}$
5: $P_i^{assigned}(t) = (1 + K(FEF - EF_i)) \times Rate(t)$

In this case, individual sites do not know the specific power and energy generation of other sites, and thus cannot compute precise rates that satisfy the aggregate limit and correctly apportion fair rates across sites. Individual sites can only increase or decrease their rate relative to the equal rates in Equation 2 and based on the difference between the globally fair energy fraction and their local fraction of energy. Thus, in our distributed algorithm, sites that are both above and below the globally fair energy fraction decrease and increase, respectively, the rate in Equation 2 by the same proportion that the site is ahead or behind in energy.

Algorithm 3 shows the pseudocode for this algorithm, which we label as distributed energy allocation. Each solar site independently runs the distributed algorithm at a specified interval to determine their solar rate. The length of this interval represents the expected time period between disseminating new generation information to other solar sites. While each solar site can broadcast to all other solar sites, full mesh communication has the same issues as the tightly-coupled centralized approach. Instead, similar to prior work on distributed rate limiting in networks [11], we can use a more robust push-sum gossip protocol that periodically disseminates recent generation information to a random set of N other sites each interval [7]. This push-sum gossip protocol may take a few intervals to converge, such that each site has an accurate estimate of the "fair" fraction of global energy and the global equal rate from Equation 2. We also add a multiplicative gain factor, K, as a configurable parameter to adjust how fast sites catch up in the distributed algorithm, similar to Algorithm 2.

3.3 Fidelity of Control

Both the centralized and distributed algorithms must make decisions based on stale information, as solar power changes continuously. In the centralized case, even though this time period may be small, e.g., one minute, solar output can fluctuate significantly even over these short time periods. Since large fluctuations can have a negative impact on electronics, the fidelity of the control, i.e., how close the algorithm is able to maintain the aggregate limit that is set, is an important performance metric. In addition, large fluctuations in the rates from the algorithm can also have a negative impact on the electronics that control solar output, and thus are also undesirable. As we show, the centralized algorithm with a

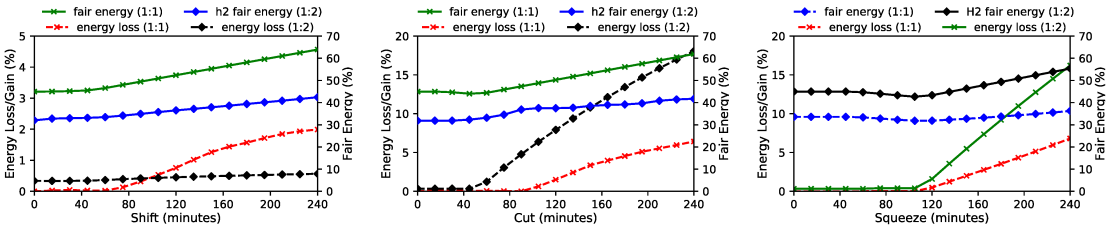

(a) Shifts (b) Dips and Cuts (c) Squeeze

Figure 4: Impact on energy fairness as a function of the magnitude of shifts (a), cuts (b), and squeezes (c) for two solar sites.

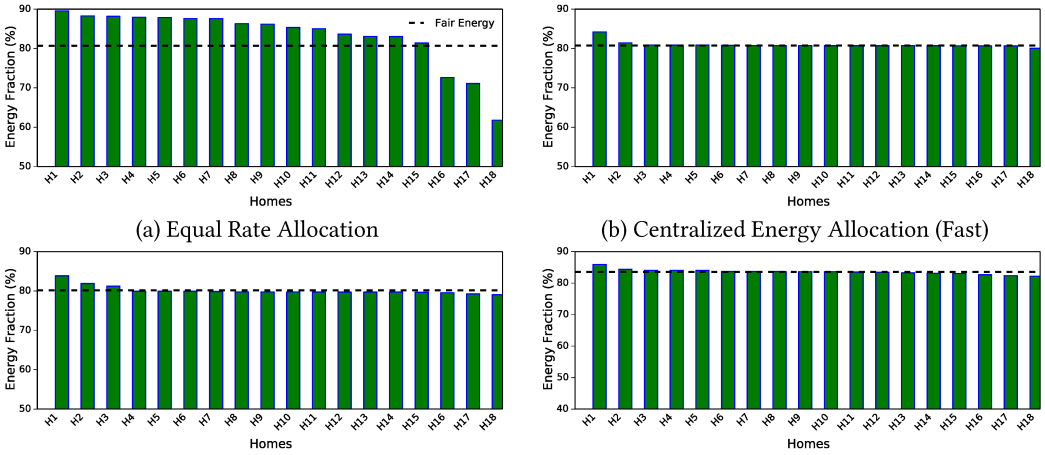

(a) Equal Rate Allocation (b) Centralized Energy Allocation (Fast)

(c) Centralized Energy Allocation (Slow) (d) Distributed Energy Allocation

Figure 5: Distribution of energy allocation (relative to maximum energy potential) under a limit of 60kW for the equal rate (a), centralized fair energy (fast) (b), centralized fair energy (slow) (c), and distributed fair energy algorithms (d).

fast catch-up suffers from increased fluctuations as it periodically focuses solar allocation on a few sites by increasing their rates to 100% of maximum output, and thus causes large changes in allocated rates. Of course, the distributed algorithm may also take more time to propagate information, causing it to diverge more from the aggregate limit. We evaluate the fidelity of control and fairness of this algorithm under different conditions in §5.

4 IMPLEMENTATION

We evaluate our centralized and distributed algorithms from the previous section in simulation using both real and synthetic solar traces. We derive our synthetic solar traces from clear sky solar irradiance models implemented in the Pysolar Python library [1]. The resolution of this synthetic solar data is one minute, and we convert the irradiance into power assuming a typical solar module efficiency of 18%. We then vary the maximum solar capacity of different sites from 1-20kW, and also vary the orientation and tilt angles of the simulated modules. For our real solar sites, we use data from 18 solar sites in the Western part of the U.S. We implement our simulator in Python and vary the simulated interval by which each site propagates its generation information.

5 EVALUATION

We evaluate both the impact of diversity in solar output on fairness using the equal rate allocation algorithm, as well using the different variants of our fair energy allocation algorithm. In addition, we also evaluate the tradeoff between fairness and the fidelity of the algorithm to maintain an aggregate limit. We quantify the fidelity using Mean Absolute Percentage Error (MAPE) between the limit and the actual aggregate generation, as below.

$$MAPE = \frac{100}{T}\sum_{t=0}^{T}\left|\frac{L(t) - \sum_{i=0}^{n} P_i^{assigned}(t)}{L(t)}\right| \quad (3)$$

Note that we only compute the MAPE for all t where we enforce the aggregate limit, i.e., $\sum_{i=0}^{n} P_i(t) > L(t)$. Quantifying fairness is more challenging than accuracy, since average fairness metrics, such as Jain's fairness index, can obscure highly unfair behavior between any two sites by averaging over many sites. For example, if there are many flows, Jain's fairness index can be close to 1 (indicating a fair allocation) even though some set of solar sites (or solar "flows") may experience highly unfair allocations. Since energy fed into the grid directly correlates with money, unfairness even among a few users is problematic. Thus, we avoid aggregate measures of fairness across many sites, and instead quantify fairness by examining the distribution of energy allocations across sites.

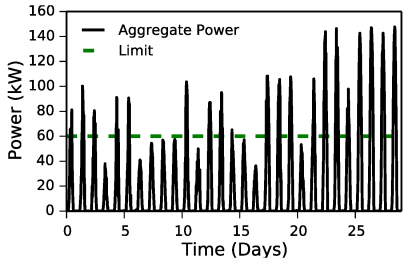

Figure 6: Aggregate power of the 18 homes over 30 days.

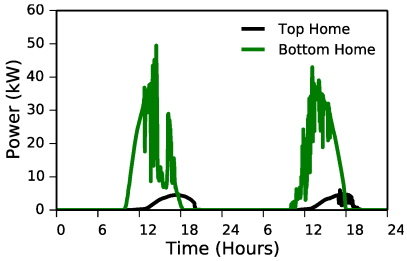

Figure 7: Daily profiles of home H1 and H18 in Figure 5(a).

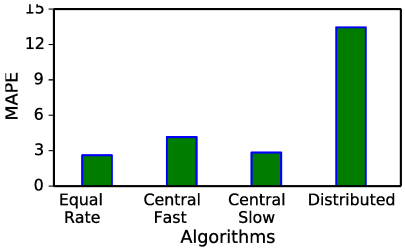

Figure 8: Fidelity of each algorithm at enforcing a limit with a one-minute communication interval.

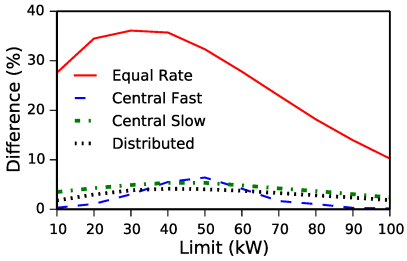

Figure 9: Energy difference between H1 and H18 in Figure 5(a), as a function of the aggregate limit.

5.1 Microbenchmarks: Shape Diversity

Figure 4 first looks at the impact on fairness between two solar sites for different magnitude shifts, dips and cuts, and squeezes. We use the equal rate allocation algorithm, which always satisfies the aggregate limit by setting rates equal to each other. For this experiment, we use synthetic data based on clear sky generation for two sites at the same location, and then alter one site's generation to shift it, cut it, or squeeze it by a certain amount of time. Thus, these results do not include other effects that could impact energy fairness, such as weather, location, or tilts. The results are also a function of the aggregate limit, which we set to 14kW in this case, where the maximum power of the sites is 10kW (or 20kW total). These experiments quantify the effects over an entire year, and include two scenarios: one where the weights are equal (where each site should contribute the same fraction of their maximum solar energy potential) and one where the weights are in a 1:2 ratio.

Figure 4(a) shows the effect of a shift, where the x-axis indicates the duration of the shift, the right y-axis is the percentage of energy lost due to unfairness in the allocation, and the left y-axis is the mean fraction of energy the solar site should have fed into the grid. The figure shows that the energy loss is only modestly impacted by shifts (1%-2%), in large part because they cancel each other out, such that a shift increases one site's allocation at the beginning of each day, but decreases it at the end of each day. As also illustrated in Figure 2, cuts (in Figure 4(b)) have a much larger impact on the energy loss, causing one site to lose nearly 10% of its energy relative to a fair allocation in the case of equal weights, and nearly 20% when weights are in a 1:2 ratio. The unequal weights increase the relative loss, since it exacerbates the amount of solar power one site is able to feed into the grid when another site is unable to generate power. This effect is similar for squeeze with losses near 10% and 20%, respectively, with equal and weighted rates.

5.2 Fair Energy Allocation

The previous subsection demonstrated the relative difference in fairness between two ideal synthetic homes with different shifts, cuts, and squeezes. We also experiment with controlling a small group of 18 homes in the western U.S. to get a sense of the differences in energy allocation across many homes with real solar power. In this case, we experiment with the equal rate allocation algorithm, as well as the three different variants of our fair energy allocation algorithm, including the extreme centralized algorithm with fast catchup, the centralized algorithm with slow (proportional) catchup, and our distributed algorithm. For these experiments, we assume all the rates are equal, and set the limit to 60kW. Figure 6 shows the aggregate power across all the homes over a month-long period, as well as the 60kW limit. While we maintain a fixed limit, note that, in practice, a balancing authority may vary this limit over time.

Figure 5 then shows the distribution of the energy gain/loss relative to the fair energy in each case over a one month period, which corresponds to a typical billing cycle. Note that this percentage directly translates into the fraction of money gained and lost from net metering. In the equal rate allocation case (a), the largest difference is over 27%, such that one home gets 27.8% less than another home and 18.9% less than their fair energy allocation. For each of the other algorithms, the percentage drops to near 0%, since they explicitly attempt to maintain a fair energy allocation over time. Figure 7 then shows a sample sunny day for both the most advantaged and disadvantaged solar site with equal rate allocation; we can see from this graph the impact of shifts, cuts, and squeezes on fairness, as these two homes have significantly different solar curves. As the figure shows, these sites have significantly different capacities, with one site having a capacity near 50kW and the other

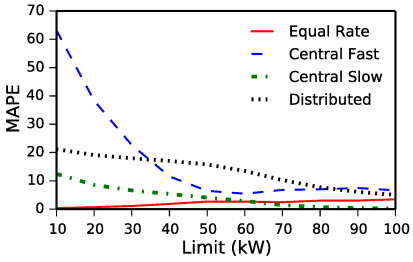

Figure 10: Fidelity of maintaining the aggregate limit as a function of its magnitude for the different variants.

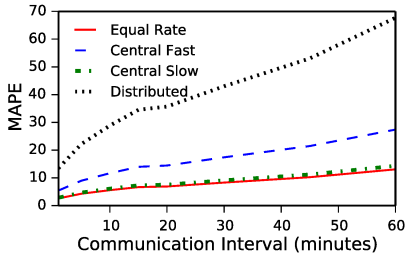

Figure 11: As the propagation delay increases the fidelity of control for the distributed algorithm decreases.

having a capacity of only 7kW. Note that a goal of our fair energy access algorithm is to enable both of these sites to contribute the same fraction of their maximum generation potential, which is relative to their capacity. In contrast, despite these differences, in all variants of the fair energy allocation algorithms, we see this difference narrowing significantly, with all having a difference of less than 1% in terms of grid energy access over the month.

In all of the algorithms above, we assume a one-minute update interval, such that the rate is updated once every minute based on data from the previous minute. Figure 8 shows the fidelity of each algorithm in maintaining the limit with this update interval. We see that the equal rate allocation has the highest fidelity (corresponding to the lowest MAPE), since it adjusts rates instantaneously. The small divergence here is due to the minute-to-minute changes in solar power, as the algorithm can only adjust rates after it senses that solar output has changed (which takes 1 minute in this case). The centralized algorithm with the fast catchup has a lower fidelity, which is also exacerbated due to the stochasticity in solar at minute-levels. This algorithm results in highly imbalanced rates during its catch-up phase, where some solar sites are contributing 100% of their energy generation. As a result, if these sites change their output significantly within a minute (before the rates are updated), the aggregate solar power will diverge from the limit. In the equal rate allocation case, the likelihood of such aggregate changes is low because it requires all sites to suddenly change their output in unison. However, when a small number of sites are catching up and have a disproportionate share of grid energy, it increases the likelihood that these sites will alter their generation within a minute. The centralized algorithm that uses a slower proportional catch-up mitigates this effect and has a MAPE near that of the equal

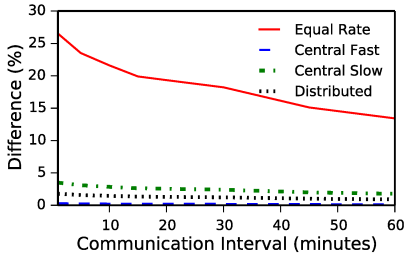

Figure 12: As the propagation delay increases, the fairness for the distributed algorithm also increases.

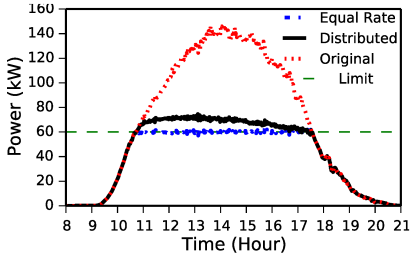

Figure 13: Maintaining the limit at a 1-minute interval.

rate algorithm. Finally, the distributed algorithm has significantly lower fidelity than the others due to its long propagation delays.

The deviation above changes with the limit as shown in Figure 9. For the equal rate algorithm, the unfairness decreases as the limit increases, since it mitigates the effect of differences in the solar curve between sites. However, the difference between the different variants of our fair energy algorithms remain largely constant and generally under 5%. However, Figure 10 shows that the equal rate algorithm has the highest fidelity across all aggregate limits. For the fair energy algorithms, the lower the limit, the worse the fidelity at maintaining the aggregate limit. This impact of low limits is particularly severe for the centralized algorithm with fast catch-up, since at low limits it is subject to increasingly more extreme versions of the effects described above.

5.3 Distributed Algorithm

Finally, we explore the impact of information propagation delay in the distributed algorithm. Figure 11 shows this delay on the x-axis, while the y-axis shows the resulting MAPE relative to the limit. The graph demonstrates that, as expected, the fidelity of the control decreases (yielding a higher MAPE), as the propagation delay increases. This increase is faster for the distributed algorithm, since it takes some time for the rates to converge. However, in contrast, fairness actually *improves* as the delay increases. Figure 12 shows the percentage maximum difference in the percentage of energy gain/loss between any two homes (in this case, H1 and H18 from Figure 5(a)). The graph shows that as the propagation delay increases this percentage trends towards 0%. Of course, the equal rate algorithm is unfair and thus takes longer to converge. With longer propagation delays, solar sites operate at the same fraction of power for longer windows of time. As a result, the amount of energy they contribute to the grid relative to each other converges.

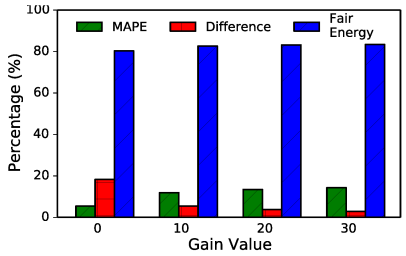

Figure 14: Impact of accelerating the "catch up" of sites that are behind in their fair energy allocation by a multiplicative gain factor in the distributed algorithm.

Thus, our fair energy access algorithms enable a tradeoff between propagation delay, fidelity of control, and fairness.

Figure 13 illustrates the fidelity of maintaining an aggregate 60kW limit for the distributed algorithm over a representative sunny day with a communication interval of one-minute. The graph shows that the centralized equal rate algorithm is able to maintain the 60kW limit precisely, while the distributed algorithm maintains a limit that is slightly above the 60kW threshold. Finally, Figure 14 shows how we mind the gap between fidelity and fairness by accelerating the catch-up amount in the distributed algorithm. In this case, we specify a gain value, which is a multiplicative factor applied to the typical rate computed by the distributed algorithm (which enables sites to increase their rate in proportion to the amount of energy they are behind). Here, a gain of 0 indicates no additional increase, while a gain of 10 increases the rate by a factor of 10. The graph illustrates the tradeoff between fairness and fidelity: as we increase the gain value (to accelerate catching up sites that are behind in their energy allocation), the MAPE of the aggregate limit increases (reducing the fidelity of control), while the fairness increases (as specified by the decrease in the largest difference in energy allocation between two sites). For comparison, we also plot the fair energy fraction for the distributed algorithm, which increases slightly, as more power is fed into the grid (as a result of overshooting the limit as seen in Figure 13).

6 RELATED WORK

There is a large body of work on fair rate allocation and scheduling in networking. This work differs from our work in that it maintains instantaneous fairness when flows are backlogged, and not fairness over long periods. Recently, there have been adaptations of this work to the grid to dynamically manage solar energy [8] and electric vehicles [2]. We show that this work can lead to *unfair* energy access. Similarly, iPlug [12] proposes a policy for decentralized dispatch of solar power based on congestion-aware network protocols. iPlug differs from this work in that solar sites backoff after sensing grid congestion, e.g., due to a deviation in grid voltage and frequency. One issue with this approach is that it requires degrading grid power quality to send feedback signals. Balancing authorities are unlikely to allow such degradation. Thus, we adopt an approach that directly communicates generation via the network to maintain a fair energy allocation over time. Finally, iPlug's approach is not fair, since different users sense different voltages and frequencies depending on their position in the grid. For example, a user at the end of the distribution line will have lower voltages, and thus backoff more than a user further up the line.

Finally, prior work in the power systems community explores different strategies for curtailing solar power. However, these approaches have largely focused on preserving the reliability of the grid, and responding to over-voltage situations [9, 14–17]. Instead, our work focuses on enabling fair control of distributed solar capacity, which has not been a metric of interest in prior work.

7 CONCLUSIONS

This paper highlights that enforcing fairness based on the relative amount of energy injected into the grid over time is more important than enforcing instantaneous rates. This discrepancy arises from fundamental differences in enforcing "fair" energy access to the grid, compared to analogous fair-sharing in networks and processors. To address the problem, we present both a centralized and distributed algorithm to enable control of distributed solar capacity, while enforcing fair grid energy access. We implement our algorithm and evaluate it on both synthetic data and real data from 18 solar sites. We show that traditional rate allocation that enforces equal rates results in solar sites contributing up to 18.9% less energy than an algorithm that enforces fair grid energy access.

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

REFERENCES

[1] 2017. PySolar. http://pysolar.org/. (2017).
[2] O. Ardakanian, C. Rosenberg, and S. Keshav. 2013. Distributed Control of Electric Vehicle Charging. In *e-Energy*.
[3] D. Cardwell. 2015. New York Times, Solar Power Battle Puts Hawaii at Forefront of Worldwide Changes. (April 18th 2015).
[4] A. Demers, S. Keshav, and S. Shenker. 1989. Analysis and Simulation of a Fair Queuing Algorithm. In *SIGCOMM*.
[5] J. Farmer and F. Lafond. 2016. How Predictable is Technological Progress? *Research Policy* 45 (2016), 647–665. Issue 3.
[6] J. Hill. 2017. GTM Forecasting More Than 85 Gigawatts of Solar PV To Be Installed in 2017. In *CleanTechnica*.
[7] D. Kempe, A. Dobra, and J. Gehrke. 2003. Gossip-based Computation of Aggregate Information. In *FOCS*.
[8] S. Lee, S. Iyengar, D. Irwin, and P. Shenoy. 2017. Distributed Rate Control for Smart Solar-powered Systems. In *e-Energy*.
[9] D. Lew, L. Bird, M. Milligan, B. Speer, X. Wang, E. Carlini, A. Estanqueiro, D. Flynn, E. Gomez-Lazaro, N. Menemenlis, A. Orths, I. Pineda, J. Smith, L. Soder, P. Sorensen, A. Altiparmakis, and Y. Yoh. 2013. *Wind and Solar Curtailment.* Technical Report. National Renewable Energy Laboratory.
[10] A. Mulkern. 2013. Scientific American, A Solar Boom So Successful, It's Been Halted. (December 20th 2013).
[11] B. Raghavan, K. Vishwanath, S. Ramabhadran, K. Yocum, and A. Snoeren. 2007. Cloud Control with Distributed Rate Limiting. In *SIGCOMM*.
[12] S. Rongali, T. Ganuy, M. Padmanabhan, V. Arya, S. Kalyanaraman, and M. Petra. 2016. iPlug: Decentralised Dispatch of Distributed Generation. In *COMSNETS*.
[13] A. Singh, S. Lee, D. Irwin, and P. Shenoy. 2017. SunShade: Enabling Software-defined Solar-powered Systems. In *ICCPS*.
[14] R. Tonkoski and L. Lopes. 2011. Impact of Active Power Curtailment on Overvoltage Prevention and Energy Production of PV Inverters Connected to Low Voltage Residential Feeders. *Renewable Energy* 36 (June 2011).
[15] R. Tonkoski, L. Lopes, and T. El-Fouly. 2010. Droop-based Active Power Curtailment for Overvoltage Prevention in Grid Connected PV Inverters. In *IEEE ISIE*.
[16] R. Tonkoski, L. Lopes, and T. El-Fouly. 2011. Coordinated Active Power Curtailment of Grid Connected PV Inverters for Overvoltage Prevention. *IEEE Transactions on Sustainable Energy* 2, 2 (April 2011).
[17] R. Tonkoski, L. Lopes, and D. Turcotte. 2009. Active Power Curtailment of PV Inverters in Diesel Hybrid Mini-grids. In *IEEE Electrical Power and Energy Conference*.