



Contents lists available at ScienceDirect

# Sustainable Computing: Informatics and Systems

journal homepage: [www.elsevier.com/locate/suscom](http://www.elsevier.com/locate/suscom)

## Leveraging weather forecasts in renewable energy systems

Navin Sharma<sup>a,\*</sup>, Jeremy Gummesson<sup>b</sup>, David Irwin<sup>b</sup>, Ting Zhu<sup>c</sup>, Prashant Shenoy<sup>a</sup><sup>a</sup> CS Department, University of Massachusetts Amherst, MA 01003, United States<sup>b</sup> ECE Department, University of Massachusetts, Amherst, MA 01003, United States<sup>c</sup> CSEE Department, University of Maryland, Baltimore County, MD 21250, United States

### ARTICLE INFO

#### Article history:

Received 14 November 2012

Received in revised form 8 June 2014

Accepted 4 July 2014

#### Keywords:

Energy harvesting

Weather forecast

Energy prediction

Green computing

### ABSTRACT

Systems that harvest environmental energy must carefully regulate their usage to satisfy their demand. Regulating energy usage is challenging if a system's demands are not elastic, since it cannot precisely scale its usage to match its supply. Instead, the system must choose how to satisfy its demands based on its current energy reserves and predictions of its future energy supply. In this paper, we show that prediction strategies that use weather forecasts are more accurate than prediction strategies based on the past, and are capable of improving the performance of a variety of systems. We analyze weather forecast, observational, and energy harvesting data to formulate a model that translates a weather forecast to a solar or wind energy harvesting prediction, and quantify its accuracy. We then compare the performance of three types of energy harvesting systems—a lexicographically fair sensor network, an off-the-grid sensor testbed, and a solar-powered smart home—using prediction models based on both forecasts and the past. In each case, forecast-based predictions significantly improve system performance.

© 2014 Elsevier Inc. All rights reserved.

### 1. Introduction

Energy harvesting systems collect and store environmental energy to either sustain continuous operation without external power sources or reduce energy consumption from burning “dirty” fossil fuels.<sup>1</sup> Harvesting environmental energy is useful for a diverse range of cyber-physical systems. For instance, past research focuses on energy harvesting sensor networks, since they are often deployed in remote locations without access to the power grid [27,10]. As another example, cloud data centers are integrating renewable energy to offset the growing monetary costs and carbon emissions from rising electricity demands [21,11,3]. Finally, net metering combined with time-of-use (TOU) pricing models provide strong financial incentives for home owners to augment grid power with on-site renewables [31].

Energy-neutral systems always consume less than or equal to the energy they harvest [13]. An underlying goal of most energy harvesting systems is to operate as close to energy-neutral as possible to prevent downtime from battery depletions. The strategy a system uses to achieve energy-neutral operation depends on

the specific characteristics of its energy source, battery, hardware components, and workload. Achieving energy-neutral operation is simple if an energy source produces power faster than a system can consume it. Unfortunately, environmental energy sources, such as solar and wind, are intermittent and vary significantly over time due to weather conditions. As a result, these energy sources typically do not produce enough power to continuously operate a system's hardware components.

Instead, systems must adapt their energy usage over time to ensure they do not consume more energy than they are able to harvest and store. Ideal systems are *energy-proportional*, such that their energy consumption scales linearly with their workload's intensity [5]. Thus, a system with elastic workload demands achieves energy-neutral operation by changing the intensity of its workload, and hence its energy usage, at fine time-scales to match the energy it harvests. Prior work on energy harvesting primarily focuses on systems with energy-proportional components that have elastic workload demands [8,12,13,15,16,26–28]. Maintaining energy-neutral operation in a system with *inelastic workload demands* using components that are not energy-proportional poses new challenges, since the system is unable to precisely change the intensity of its workload and energy usage to match the energy it harvests.

The system must choose how to satisfy its workload's demands based on its current and expected energy supply. Inelastic demands derive from either external requests, such as satisfying requests from system users, or internal objectives, such as maintaining a stable workload for a long period of time. The former

\* Corresponding author. Tel.: +1 413 545 4753.

E-mail addresses: [nksharma@cs.umass.edu](mailto:nksharma@cs.umass.edu) (N. Sharma), [jgummesson@ecs.umass.edu](mailto:jgummesson@ecs.umass.edu) (J. Gummesson), [irwin@ecs.umass.edu](mailto:irwin@ecs.umass.edu) (D. Irwin), [zt@umbc.edu](mailto:zt@umbc.edu) (T. Zhu), [shenoy@cs.umass.edu](mailto:shenoy@cs.umass.edu) (P. Shenoy).<sup>1</sup> This paper is an extension of a previous conference paper [22].

is relevant to both off-the-grid sensor testbeds [9,23,29] and energy harvesting smart homes, since external testbed users or home occupants, respectively, dictate the workload's energy demands. The latter is relevant to lexicographically fair energy harvesting sensor systems, since the primary goal is to maintain steady and fair node sensing rates for a target time period [10]. As others have noted, workload scheduling algorithms in energy harvesting systems with inelastic demands are highly sensitive to energy harvesting predictions [18].

While past work recognizes the need for accurate energy harvesting predictions, prior prediction methods derive from the underlying idea that the past is an accurate predictor of the future [10,14,18,19]. While the past is accurate for both sufficiently short, i.e., seconds to minutes, and sufficiently long, i.e., months to years, time-scales, we show in Section 2 that predictions derived from weather forecasts are more accurate at the medium-length time-scales, i.e., hours to days, relevant to a large class of energy harvesting systems. Our empirical findings match the same intuition that causes people to tune into a nightly weather forecast, rather than step outside, to find out the expected weather for the next few days. In this paper, we design a method for leveraging weather forecasts to improve the performance of energy harvesting systems. In particular, as discussed below, we (i) analyze historical weather data to identify the time-scales when forecasts are most accurate, (ii) develop a model that maps a specific forecast to an energy harvesting prediction, and (iii) quantify the performance improvements from using our model in three real-world case studies.

**Analyze Historical Weather Data.** We analyze extensive traces of past forecast and observational data from the National Weather Service (NWS), as well as fine-grain solar and wind energy harvesting and observational data from our own deployment. We use these traces to quantify how well both weather forecasts and multiple variants of predictions using the immediate past predict the weather phenomena—sky condition and wind speed—that most impact solar and wind energy harvesting at time-scales ranging from 1 h to 72 h in the future.

To demonstrate the broad applicability of our approach, we analyze data from five locations within the United States with five distinct climate profiles. These locations include Chicopee Falls, Massachusetts, Daytona Beach, Florida, Phoenix, Arizona, Norfolk, Nebraska, and Seattle, Washington. We find that in all cases NWS forecasts in these regions are a better predictor of the future than existing prediction strategies based on the immediate past over NWS forecast time-scales for both sky condition and wind speed.

**Formulate Forecast→Energy Model.** We use our observational data to correlate (i) weather forecasts for our region with our own local weather observations and (ii) our own local weather observations with the energy harvested by our solar panel and wind turbine deployment. We use both data sets to formulate a simple model that predicts how much energy the solar panel and wind turbine will harvest in the future given weather forecasts every hour from 1 h to 72 h in the future. To evaluate the forecast accuracy of remote regions, we assume that the model we develop using our own local data collection, as well as the NWS weather observations, are accurate.

**Case Studies.** We quantify the benefits of using energy harvesting predictions based on weather forecasts in the context of three different types of energy harvesting cyber-physical systems with inelastic demand. The first system is a deployed off-the-grid testbed [23] we have built as part of the NSF GENI prototype [20]. The second system is inspired by recent work on lexicographically fair energy harvesting sensor systems [10]. The third system is a solar-powered smart home that minimizes electricity costs for TOU pricing by deciding when to draw power from the grid versus an on-site battery.

In each case, we compare the performance of a forecast-based approach with one or more prediction models that use the immediate past to predict the future. For each system, we find that our forecast-based approach is significantly better than any prediction scheme based on the immediate past for the system's relevant performance metrics, e.g., requests satisfied, combined length of power outages, and monetary cost.

## 2. The case for using forecasts

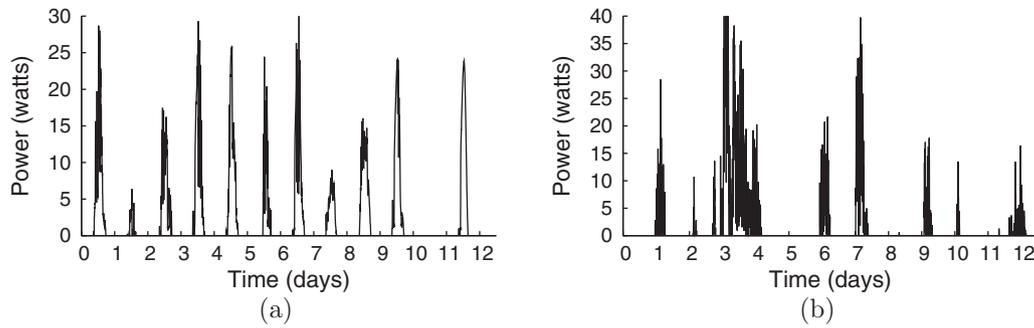
To motivate the use of weather forecasts for prediction, we analyze both forecast and observational data from the year 2008 to compare the accuracy, at different time-scales, of predictions based on NWS forecasts with predictions based on the past. Others have noted that over appropriate time-scales and under ideal conditions the past predicts the future for both solar [4,6,13] and wind [14] power. However, our analysis leads to four observations that motivate the use of forecasts, instead of the past, for predictions over time-scales of hours to days. We use data from an extended deployment of a weather station, wind turbine, and solar panel on the roof of the Computer Science Building at the University of Massachusetts Amherst, as well as data from NWS observations, the National Digital Forecast Database, and the National Solar Radiation Data Base. Our observational traces are available upon request from <http://traces.cs.umass.edu> and the NWS traces are available upon request from <http://www.nws.noaa.gov/ndfd/>.

Our weather station reports wind speed and solar radiation every 5 min, while the NWS reports an observation every hour and an archival forecast every 3 h for each region of the country for the last few years. Each forecast includes predictions every 3 h from 3 h to 72 h in the future. Additionally, real-time forecasts are available every hour from 1 h to 72 h in the future. Unless otherwise noted, we use our own weather station's observations for Amherst, Massachusetts, and NWS observations for other regions. While our weather station and the NWS report a variety of weather metrics, we focus on the two metrics with the most direct relationship to the energy our solar panel and wind turbine harvest: sky condition, as a percentage of cloud cover between 0% and 100%, and wind speed, in miles per hour. We show how these metrics impact solar and wind energy harvesting in Section 3.

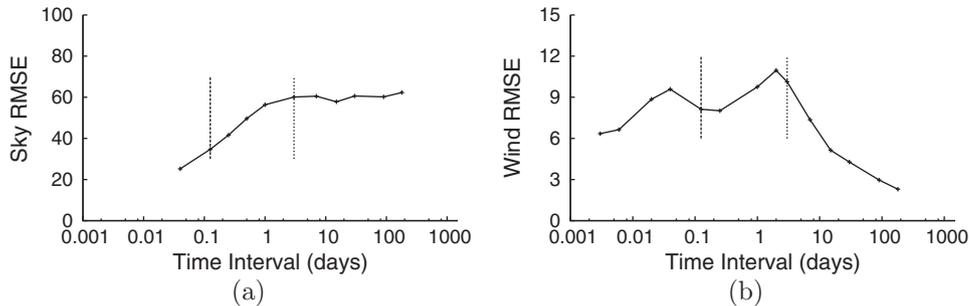
To compare with forecast-based predictions, we first define a basic prediction strategy, which we term *past predicts the future* or *PPF*. The basic PPF strategy predicts that a weather metric's value in the next  $N$  time units will exactly match the observations of that metric from the last  $N$  time units. For solar energy harvesting, there are multiple variants of this basic approach in prior work. We discuss these variants in Section 5, which adjust the basic PPF strategy to adapt to seasonal variations in sunlight [13,14,26] or sudden changes in cloud cover [19]. In this section, we focus on the basic PPF model only to motivate the limitations of solely using the past to predict the future. For wind energy harvesting, we have found no variants of the basic PPF approach in prior work.

The accuracy of the PPF model or any of its variants is dependent on the climate at a specific location. For example, a PPF model for solar power may be more accurate in areas with consistent sunlight and little variation in weather patterns, such as the desert in Australia [6], while a PPF model for wind power may be more accurate in areas likely to be in the path of a jet stream. Regardless of the area, though, prediction strategies without the aid of detailed weather forecasts must inherently rely on the past. Both our intuition and our empirical measurements lead to our first observation: there are many areas, including Amherst, Massachusetts, that do not have consistent weather patterns.

**Observation #1:** *Sky condition and wind speed show significant inter-day and intra-day variations, as a result of changing weather in*



**Fig. 1.** Power generated during a 12 day period in October, 2009 from our solar panel (a) and wind turbine (b).



**Fig. 2.** The error in sky condition (a) and wind speed (b) when using the past to predict the future for different time intervals in 2008 at 1 h and 5 min granularities, respectively, for Amherst, Massachusetts.

Amherst, Massachusetts, as well as other regions, including Arizona, Florida, Washington, and Nebraska.

While we expect wind to be intermittent, the data for the regions we examine also shows significant variations in the sky condition observed by the NWS both within each day and between days. As an example from our deployment, Fig. 1(a) and (b) shows the solar and wind power we harvest, respectively, during a 12 day period in October. As expected, wind is highly variable, with the wind turbine harvesting the most energy on days 3, 4, and 7, while harvesting lesser amounts on days 1, 6, 9, 10, and 12. The turbine harvests nearly zero energy on days 2, 5, 8, and 11. Surprisingly, despite its diurnal nature, solar power shows significant variations as well due to cloud cover, with the solar panel harvesting less than half its maximum possible energy on days 2, 3, 7, 8, and 11, with significant variations within each day. Our solar panel actually harvests no energy on day 11.

Even when the solar panel or wind turbine harvest the same amount of energy on two different days, the profile of power generation within each day is variable. For example, on both day 3 and 4 our solar panel harvests similar amounts of energy, but the power profile for day 4 is more consistent and less variable than day 3. Overall, the solar panel and wind turbine harvest less than 1/2 their rated daily maximum on 40% and 75% of the days, respectively. While we chose a 12 day period to enhance the readability of the graph, we have witnessed a similar degree of day-to-day variation throughout our solar panel and wind turbine deployment.

**Observation #2:** Using PPF to predict the future is least accurate at medium-length time-scales from 3 h to 1 week.

To evaluate the accuracy of the PPF model we focus on Amherst, Massachusetts, and calculate the root mean squared error (RMSE) between the average value of both sky condition and wind speed over an interval from  $t=0$  to  $t=N$  and from  $t=N$  to  $t=2N$  for all possible intervals of length  $2N$  in the year 2008, given that our observational data has a granularity of 5 min. RMSE is a standard statistical measure of the accuracy of values predicted by a model with respect to the values observed. Intuitively, the RMSE's value

quantifies the PPF model's accuracy at different time-scales. For instance, an RMSE of zero for an interval of length  $N$  indicates that for all possible intervals of length  $N$  during the year the metric's average in the previous interval exactly predicts the metric's average in the next interval. The closer the RMSE is to zero for a particular interval duration, the more accurate the past predicts the future for that interval.

Fig. 2(a) and (b) shows the RMSE for sky condition and wind speed, respectively, as a function of time interval duration  $N$  ranging from 5 min to 6 months. Notice that we plot both graphs on a log scale. The analysis shows that predictions based on the past are most accurate at both short (<2 min) and long time-scales (>10 days), and are least accurate in between. Moreover, wind speed predictions based on the past tend to get better over long-term, whereas sky condition predictions remain almost same after 10 days. For both sky condition and wind speed, the maximum inaccuracy occurs between 3 h and one week, as indicated by each graph's vertical lines.

**Observation #3:** Over NWS forecast time-scales of 3 h to 3 days, sky condition and wind speed forecasts are better predictors of the future than the PPF model.

We next show that NWS forecasts for the medium-length time-scales of hours to days are more accurate than the PPF model. To quantify the relative accuracy of weather forecasts, we use NWS forecast data from three months in different seasons—January, April, and September 2008—at all five of our locations. Fig. 3(a) shows the RMSE between the observational sky condition and the sky condition from the NWS forecasts, as a function of the forecast time horizon<sup>2</sup> for multiple regions. Similarly, Fig. 3(b) shows the RMSE between the observational sky condition and the sky condition using PPF for the same regions. As expected, the accuracy of the sky condition forecast decreases as the time horizon increases.

<sup>2</sup> We use "time horizon" and "time interval" interchangeably throughout the paper.

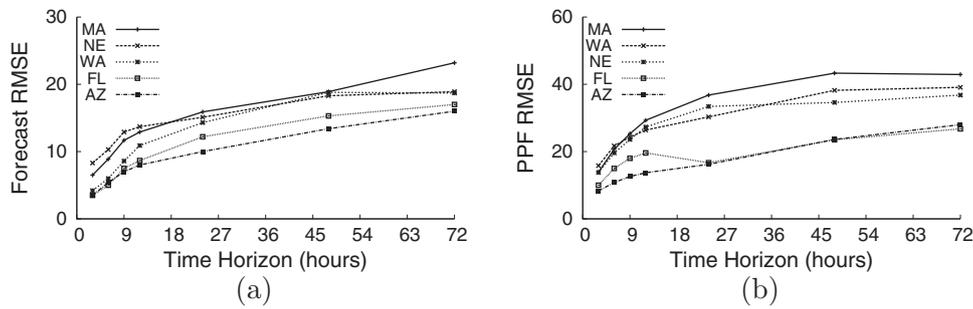


Fig. 3. RMSE between the observed sky condition and the sky condition predicted by NWS forecasts (a) and the PPF model (b) for different time intervals in 2008 at 1 h granularities.

Since the RMSE of the sky condition forecast (<20) is less than the RMSE of the PPF model from Figs. 2(a) and 3(b) between 3 h and 3 days (~40–60) we conclude that the forecast is a better predictor than the past for sky condition in every region we study, including Amherst, Massachusetts.

We next compare the accuracy of the NWS forecast for wind speed with the accuracy of the PPF model. Fig. 4(a) shows the RMSE between the observational wind speed and the wind speed from the NWS forecast, as a function of the forecast time horizon for multiple regions. Similarly, Fig. 4(b) shows the RMSE between the observational wind speed and the predicted wind speed using the PPF model for the same regions. As the figures show, the accuracy of the wind speed forecast does not vary significantly for any future time horizon. Since the RMSE of the NWS wind speed forecast (<6) is less than the RMSE of the PPF model from Fig. 2(b) and 4(b) between 3 h and 3 days, we conclude that the NWS forecast is a better predictor than the past for wind speed in Amherst, Massachusetts, which leads to our final observation.

**Observation #4:** We conclude that using weather forecasts as a basis for prediction should improve the performance of energy harvesting systems with inelastic demands that make workload scheduling decisions over 3 h to 3 day time horizons.

### 3. Forecast → energy model

To leverage our observations from the previous section, we formulate models that predict the energy our solar panel and wind turbine will harvest given a NWS weather forecast. Note that our models are based on our specific solar panel and wind turbine, as well as the weather forecasts at our location. In addition, we assume an unobstructed solar panel and wind turbine not affected by shade from trees or buildings. Since we derive our model parameters empirically, they depend on the specific characteristics of our deployment, and are not directly useful for other deployments. While the methods we use for building our models are applicable to other deployments, the accuracy we report is dependent on the specific characteristics of our location’s climate. Further, since

we deploy our harvesting equipment in an open area, we do not evaluate the effect of local conditions, such as shade from foliage or wind shear from surrounding buildings, on our model.

Before discussing our model, we briefly describe our energy harvesting deployment, which consists of a battery, solar panel, and wind turbine. Air-X manufactures our wind turbine, and rates its maximum power output as 400 watts in 28 mile per hour winds. The turbine uses an internal regulator to govern the power delivered to the battery to prevent overcharging when the battery voltage increases beyond a threshold of 14.1 volts. Kyocera manufactures our solar panel, and rates its maximum power output as 65 watts at 17.4 volts under full sunlight. We connect the solar panel to a deep-cycle battery through a TriStar T-60 charge controller, which protects the battery from overcharging. Our battery has an ideal capacity of 1260 watt-hours.

The purpose of our deployment is to measure power harvested over time. As a result, we must ensure that the battery is never full, since a full battery cannot harvest and store energy. To prevent our system’s battery from becoming fully charged, we use an additional T-60 load controller in conjunction with a 60 watt automotive bulb to bleed the battery’s energy. The controller connects the load to the battery at 13.6 volts and disconnects at 12.1 volts to ensure the battery stays charged to 55% of its capacity. The final component of our measurement system is a HOBO U30 wireless data logger. The logger measures battery voltage, using a built-in analog-to-digital converter, and electrical current, using an external current transducer for each energy source. The logger measures each quantity every 30 s and stores a 5 min average locally. Each hour, the logger uploads its log file to a server hosted by HOBO, where data is publicly available for viewing through the HOBO web interface.

#### 3.1. Sky condition → solar power model

We base our model for solar energy on a simple premise: if the sky condition reports a cloud cover of  $N\%$  then the observed solar radiation, as well as our solar panel’s power production, will be  $(100 - N)\%$  of the maximum possible under ideal cloudless skies. For

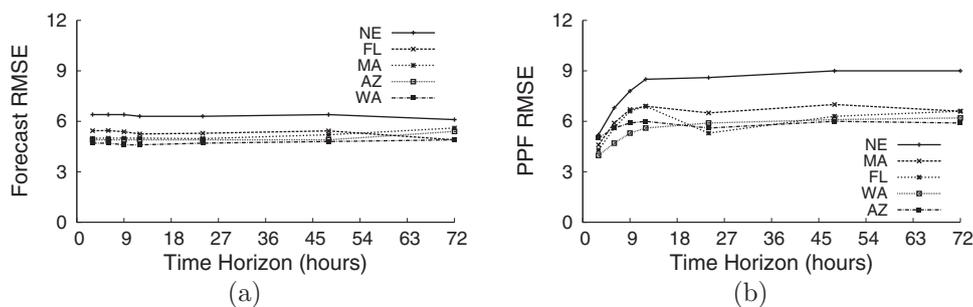


Fig. 4. RMSE between the observed wind speed and those predicted by NWS forecasts (a) and using the past to predict the future (b) for different time intervals in 2008.

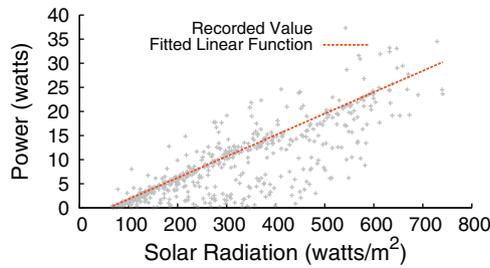


Fig. 5. Relationship between the solar radiation our weather station observes and the power generated by our solar panel.

example, if the 3 h forecast predicts a sky condition with 50% cloud cover, and the maximum possible solar power production is 60 watts over that 3 h interval, then the solar power prediction for that 3 h interval will be  $60 * 0.5 = 30$  watts. Given our simple premise, to formulate our model we must first estimate the maximum possible solar power production at any time of the day and year, given the tilt of the earth’s axis and the sun’s diurnal nature. Since our solar panel deployment has not been active for an entire year, we use our weather station’s traces of solar radiation to construct our model.

3.1.1. Computing solar power from solar radiation

We first derive the relationship between the solar radiation our weather station observes and the power our solar panel produces using our trace data, as shown in Fig. 5. The relationship should be linear, since our solar panel produces energy in proportion to the solar radiation with a constant factor loss due to inefficiency. As expected, the relationship we observe is close to linear. We use the least-squares approach to fit the following regression line to the data, which we use to convert the solar radiation our weather station observes to the solar power our panel produces, where power is in units of watts and solar radiation is in units of watt/m<sup>2</sup>.

$$SolarPower = 0.0444 * Radiation - 2.65 \tag{1}$$

3.1.2. Computing the maximum possible solar power

We next derive an estimate for the maximum solar power possible at a given time of the day and year. The value is dependent on

multiple factors, including the time of the day, day of the month, month of the year, and geographic location. While highly accurate models that take into account all of these factors are possible, we use a simple approximation that assumes the change in position of the sun relative to a specific location does not vary significantly within any single month. Thus, we use a profile for a single sunny day in each month of the year as the baseline for computing the ideal maximum power on any day of that month. We select a single sunny day with no cloud cover for each month from the year 2008 using our weather station data and observational data from National Solar Radiation Database.

Fig. 6(a) shows the profile of solar power our panel would harvest on three perfectly clear and sunny days in January 2008, May 2008, and September 2008. Similarly, Figs. 6(b)–(d) show the profile of solar power our solar panel would harvest in the other locations we study, including Arizona, Nebraska, and Washington, on three perfectly clear and sunny days in January 2008, May 2008, and September 2008. For the graph, we convert the solar radiation observed by our weather station and the observational radiation data from National Solar Radiation Database on these days to the expected solar power harvested by our solar panel using Eq. (1) from above. We find that power is quadratically related to the time of day. Since daylight hours change throughout the year, the power profile for a sunny day also changes. Of the three months in the figure, May has the maximum possible potential for power generation since it is nearest to the summer solstice, while January has the least possible potential for power generation since it is nearest to the winter solstice. For each month, we fit the quadratic function below, where *a*, *b*, and *c* are the parameters of the quadratic function, and *Time* is in hours after 12 am. The parameters *a*, *b*, and *c* for each month in Amherst, Massachusetts are given in Table 1.

$$MaxPower = a * (Time + b)^2 + c \tag{2}$$

3.1.3. Solar model

To complete our model, we compute the power our solar panel generates using the equation below, where *MaxPower* is in units of watts from Eq. (2) and *SkyCondition* is the percentage cloud cover from the NWS. Fig. 7(a) compares the observed solar power generated by our panel with the solar power predicted by our

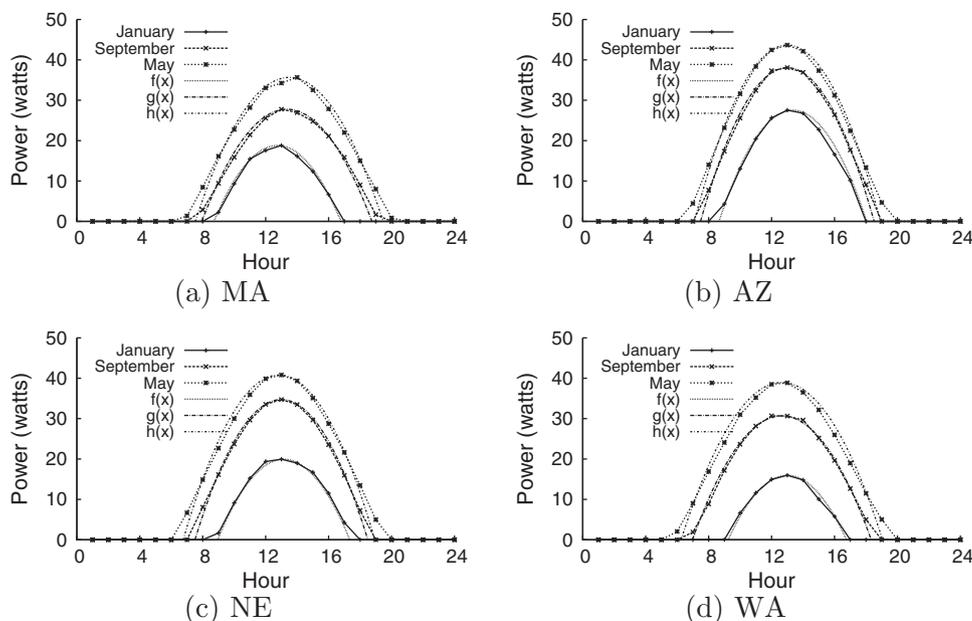


Fig. 6. Profile for solar power harvested on clear and sunny days in January, May, and September, and the quadratic functions *f(x)*, *g(x)*, and *h(x)* we fit to each profile, respectively.

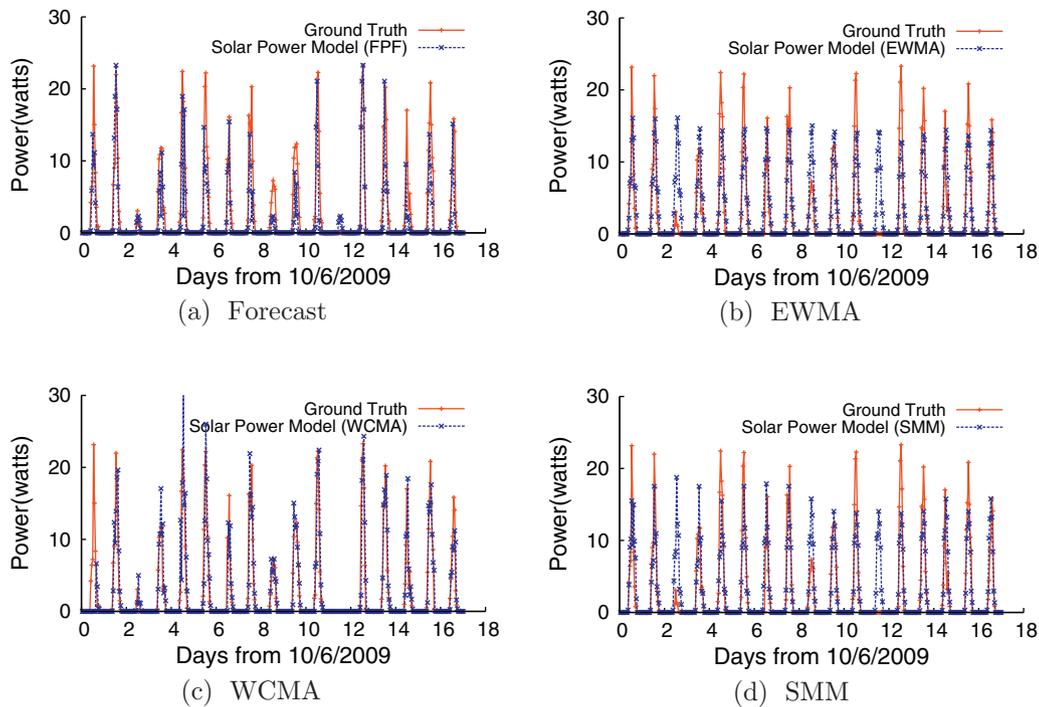


Fig. 7. Power output from our solar panel and the power output predicted by different prediction models.

model for Amherst, Massachusetts (three hours into the future). The graph demonstrates that the daily average difference between each observed and predicted value is small. Further, the model tends to be conservative when incorrect: the predictions are generally less than the observations, which reduces battery depletions from incorrect predictions.

$$Power = MaxPower * (1 - SkyCondition) \tag{3}$$

3.2. Wind speed → wind power model

Our wind power model is simpler than our solar model, because, as opposed to sky condition, both our weather station and the NWS forecast report wind speed. Fig. 8 shows the recorded power output of the wind turbine for different recorded wind speeds, as well as curves showing the power ratings for the turbine in both turbulent and steady winds. Wind power production is known to be a cubic function of the wind speed [2].

$$Power = 0.01787485 * (WindSpeed)^3 - 3.4013 \tag{4}$$

We fit the cubic power curve in Eq. (4) to the observed data using the least-squares method to generate our wind power model, where Power is in units of watts and WindSpeed is in units of miles

per hour. Our cubic function is nearly half-way in between the rated power curves for turbulent and steady winds. Fig. 9(a) compares the observed wind power generated by our wind turbine with the wind power predicted by our forecast model for Amherst, Massachusetts (three hours into the future). The graph demonstrates that the daily average difference between each observed and predicted value is small. Further, the model tends to be conservative when incorrect: the predictions are generally less than the observations, which reduces battery depletions from incorrect predictions. Note that the wind turbine stops producing power near 28 miles per hour, so our function ramps down to 0 at that point.

3.3. Compensating for forecast errors

Our solar and wind power models convert an observed sky condition and wind speed to the expected solar and wind power generated by our deployment. To convert a forecast for sky condition and wind speed to a prediction for solar and wind power we multiply the output of both models with an error constant  $\alpha$ . We base our  $\alpha$  constant for each forecast time horizon on the RMSE for sky condition and wind speed forecasts in the previous section. Thus, the greater the expected error in the forecast at a particular future time, the smaller the value of  $\alpha$  in our model. We use  $\alpha = 0.8$

Table 1 Values for a, b, and c in our quadratic solar power model.

Month	a	b	c
January	-1.15	-12.75	21.45
February	-1.15	-12.75	29.13
March	-1.15	-12.75	35.97
April	-1.25	-13.5	43.72
May	-1.1	-13.5	43.5
June	-1.1	-13.5	43.4
July	-1	-13.5	40.35
August	-1.15	-13.5	40
September	-1.15	-13.5	36.32
October	-1	-13.35	27
November	-1.45	-12	22.66
December	-1.15	-12.5	16.79

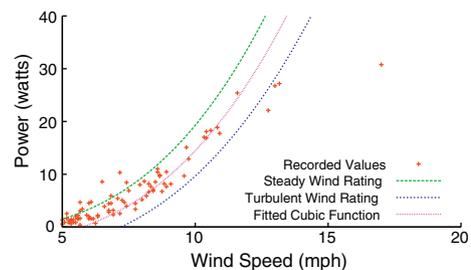


Fig. 8. Power output from our wind turbine and the power output predicted by our wind power model.

in our model, which we derive from the accuracy of forecasts in our region.

### 3.4. PPF variants

In addition to the basic PPF model, we evaluate three PPF variants from prior work: Exponentially Weighted Moving Average (EWMA), Simple Moving Median (SMM), and Weather Conditioned Moving Average (WCMA). We briefly describe each approach.

The purpose of the EWMA variant is to adapt to seasonal variations in output. Thus, EWMA divides a day into slots and predicts the energy for a particular slot as the weighted sum of the energy harvested in the same slot on  $N$  previous days [13]. EWMA assigns an exponentially decaying weighting factor to each previous day, since the recent past tends to provide more accurate predictions. We choose each slot's length to be 60 min, since the environmental variation within each hour is typically small and to provide a fair comparison with our forecast-based approach. For EWMA, the predicted energy at a time slot  $t$  on  $i$ th day is given as:

$$E_{predict}^t(i) = \alpha E_{observe}^t(i-1) + (1-\alpha)E_{predict}^t(i-1) \quad (5)$$

We empirically find that a weighting factor  $\alpha = 0.1$  minimizes the RMSE between the observed and predicted energy for our deployment based on historical data. SMM is another variant of PPF, which predicts the energy for a particular slot as the median of energy harvested in the same slots on  $N$  previous days. SMM is more robust than EWMA to high fluctuations or other anomalies in the time series data. Finally, WCMA is a variant of EWMA, which uses the current day's, as well as previous days', observational data to make predictions [24,25]. In contrast to EWMA, WCMA considers the weather conditions of previous slots in the current day. Thus, it performs better than EWMA during inconsistent or fluctuating weather conditions. WCMA predicts energy for any time slot  $t$  on  $i$ th day as:

$$E_{predict}^t(i) = \alpha E_{observe}^{t-1}(i) + GAP_k(1-\alpha)M_D^t(i) \quad (6)$$

Where  $\alpha$  is a weighting factor similar to EWMA, and  $M_D^t(i)$  is the mean of the observed values in time slot  $t$  over the past  $D$  days.  $GAP_k$  is a factor that depends on past  $k$  slots and measures

the present weather conditions compared to the same conditions over the previous days. Recas et al. [24] provide a detailed description of calculating  $GAP_k$ . For WCMA, we find the optimal values of three parameters— $\alpha$ ,  $D$ ,  $k$ —that minimize the RMSE between the observational energy and the predicted energy as  $\alpha = 0.4$ ,  $D = 6$ , and  $k = 15$  for our solar panel, and  $\alpha = 0.9$ ,  $D = 7$ , and  $k = 14$  for our wind turbine. Similar to EWMA, we assume a slot duration of 60 min. We use the optimal values for the EWMA and WCMA parameters in our evaluation.

Fig. 7 compares the observed solar power generated by our panel with the solar power predicted by all four prediction models: (a) Forecast Predicts Future (FPF), (b) Exponentially Weighted Moving Average (EWMA), (c) Weather Conditioned Moving Average (WCMA), and (d) Simple Moving Median (SMM). Since the figure plots predictions only three hours into the future, it represents a best case scenario for the prediction models based on the past. The figure demonstrates that the forecast-based approach and WCMA perform significantly better than the EWMA or SMM model, especially when environmental conditions change. Although WCMA's accuracy is similar on average to our forecast-based approach, it over-predicts on most days, which results in frequent battery depletions in energy harvesting systems.

Similarly, Fig. 9 compares the observed wind power generated by our turbine with the wind power predicted by all four prediction models. Again, WCMA and our forecast-based approach provide similar prediction accuracy for wind power, while EWMA and SMM are much less accurate due to the wind's intermittent nature. WCMA also suffers from over prediction with wind energy, while our forecast-based approach tends to under predict when it is inaccurate. As we show in our case studies, over prediction causes unexpected battery depletions that decrease system performance. Our case studies also indicate that WCMA has worse performance at longer time-scales, e.g., day-ahead predictions.

### 3.5. Forecast-WCMA hybrid

As the previous section indicates, the WCMA variant of PPF performs well when weather does not vary significantly, while a purely forecast-based approach performs well for dramatic weather

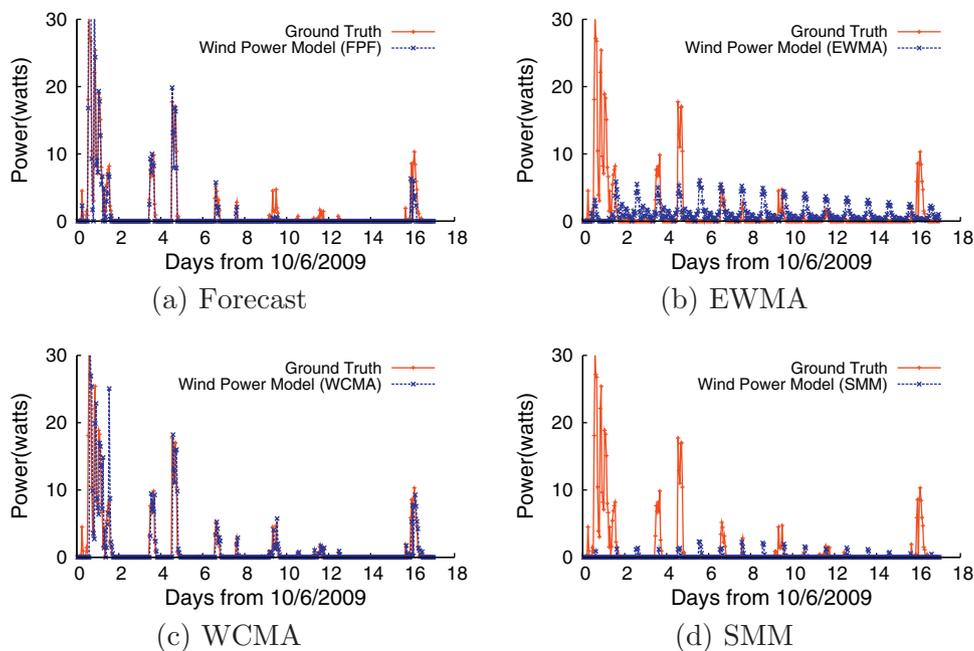
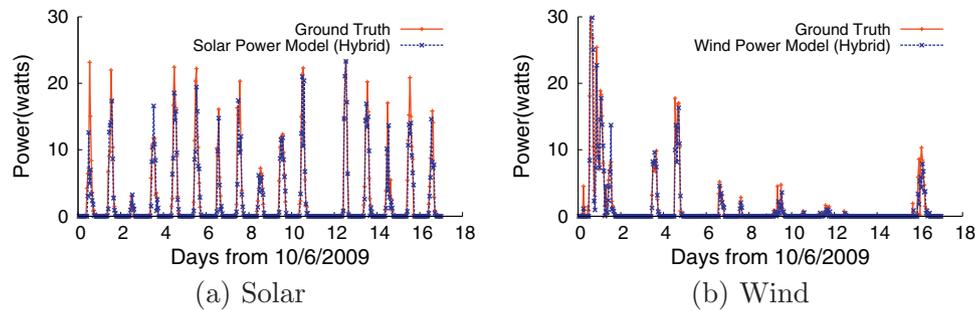


Fig. 9. Power output from our wind turbine and the power output predicted by different prediction models over the first 3 weeks of October, 2009.



**Fig. 10.** Power output from our solar panel (a) and wind turbine (b) and the power output predicted by the hybrid prediction model over the first 3 weeks of October, 2009.

variations. To gain the benefits of both approaches, we introduce a Forecast-WCMA hybrid approach. The approach assigns weights to each technique that vary dynamically based on prior prediction accuracy. More formally, the model predicts energy for a time slot  $t$  as:

$$E^t_{Hybrid} = \beta E^t_{Forecast} + (1 - \beta) E^t_{WCMA} \quad (7)$$

$$\beta = \frac{e^{t-1}_{WCMA}}{e^{t-1}_{Forecast} + e^{t-1}_{WCMA}} \quad (8)$$

where  $E^t_{Hybrid}$ ,  $E^t_{Forecast}$ , and  $E^t_{WCMA}$  represent the energy prediction for time slot  $t$ , using the hybrid-, forecast-, and WCMA-based approach, respectively, and  $e^{t-1}_{Forecast}$  and  $e^{t-1}_{WCMA}$  represent the absolute value of the prediction error for the previous slot for each approach. Fig. 10 demonstrates that the hybrid approach provides better prediction accuracy than either the forecast-based approach or WCMA. For the experiment, the RMSE for the hybrid approach (2.32 for solar, and 2.84 for wind) is lower than the RMSE for the forecast-based (FPF) approach (3.12 for solar, and 2.95 for wind) or for WCMA (2.96 for solar, and 3.05 for wind).

#### 4. Case studies

We evaluate our models from the previous section in the context of three types of energy harvesting systems with inelastic demand: an off-the-grid testbed that we have built as part of the NSF GENI prototype [23] that leases virtualized resources to users, a lexicographically fair sensor network inspired by recent work [10], and a smart home with on-site renewables. For each system, we quantify how much the use of forecast-based predictions increases the system's relevant performance metrics when compared with both the PPF model and a conservative approach that does not use predictions and only makes decisions based on the current battery level. For each case study, we predict hour-by-hour energy harvesting one day in advance. Thus, each day's predictions are a combination of a single 1-, 2-, 3-... 23-, 24-h prediction. Finally, we quantify the impact of battery capacity on performance. We refer to our forecast-based model as PPF in our graphs. Note that our objective is not to optimize any specific objective function but to demonstrate that better predictions lead to better performance for real-world energy harvesting sensor systems.

##### 4.1. ViSE testbed

ViSE, which stands for Virtualized Sensing Environment, is a 4-node high-power sensor testbed we built as part of GENI [20]. A ViSE sensor node consists of a conventional x86-processor connected to multiple high-power sensor nodes, including a radar, video camera, and weather station. Our ViSE radar is based on a CASA radars, which while smaller than Doppler radars, require mounting on large towers above any obstructions. As with other GENI testbeds, ViSE leases external users access to a slice of its nodes upon request.

A lease ensures users access to an isolated partition of the testbed's resources for some duration. Since GENI coordinates access to multiple testbeds, it is important that each testbed satisfy its leases to enable multi-testbed experiments, since not satisfying a lease may stall an experiment, and thereby waste any resources reserved on other testbeds.

When using harvested energy, ViSE must approve or reject lease requests from external users based on its available energy supply. The workload is an example of inelastic demand, since ViSE cannot change its decision to accept or reject a lease request based on new conditions after the initial decision is made. Further, ViSE must make each decision without complete knowledge of its future energy supply. A conservative approach is to reject all requests with durations greater than each node's expected operating time based on its current reserve of stored energy. However, a conservative approach may reject more requests than necessary if some knowledge of the future energy supply is known. An alternative approach uses predictions of the future to determine whether or not to approve each request, either using our model or a PPF-inspired model. When using predictions, ViSE approves lease requests if both the existing energy in the battery and the predicted energy harvested while the lease is active is sufficient to satisfy the lease over its duration. Note that our prediction-based approach not only ensures that there is enough energy at the end of the lease, but also throughout the lease based on the energy harvested while the lease is active.

To evaluate the benefits of our models from Section 3 relative to both the conservative approach and the variants of the PPF model, we ran simulations based on our ViSE node's power characteristics using our traces of solar and wind power. Each ViSE node consumes 115 watts at full utilization when the radar is transmitting, with the radar consuming 50 watts by itself [7]. Neither the radar nor the compute node are energy-proportional. The radar consumes either 0 watts when off or 50 watts when transmitting, while the main node's power consumption scales linearly from 45 watts at idle to 65 watts at 100% utilization. Since a single wind turbine or solar panel from our deployment is only sufficient to run our node for a few hours each week, for our experiments we assume the use of 5 identical solar panels or 5 identical wind turbines. We assume a battery capacity capable of running our node at full utilization for 2 h.

For our experiments, ViSE makes decisions to accept or reject lease requests at the beginning of each day, where each lease reserves a virtual sensor, i.e., an isolated sliver of the sensor's resources, for 24 h. We discuss ViSE's approach to virtualizing sensors in recent work [23]. We assume that the workload includes enough lease requests each day to completely consume the maximum possible energy the solar panels or wind turbines can produce. At each decision point, ViSE only accepts leases that it believes it can satisfy based on a conservative approach, the NWS forecast-based model, or the variants of the PPF model (basic PPF, EWMA, WCMA, SMM, and Hybrid), where we assume each virtual sensor

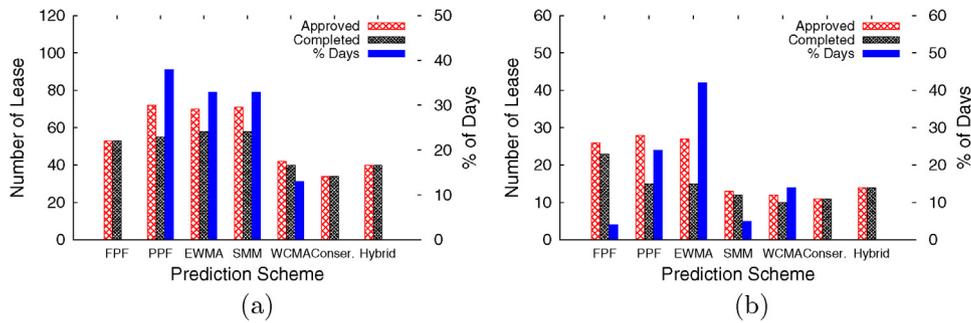


Fig. 11. Number of approved and completed leases, and % of days the node runs out of energy for different prediction schemes for both solar (a) and wind (b) power.

will operate at 1/24th of full utilization for the duration of the lease. The performance metrics we use to evaluate the different approaches are (i) the number of leases ViSE approves and (ii) the percentage of approved leases ViSE satisfies without nodes running out of energy. Ideally, ViSE should never approve a lease that it does not have the energy to complete.

We evaluate ViSE separately for solar and wind energy. Fig. 11(a) and (b) shows the number of lease requests ViSE approves, and number of leases ViSE completes for the panels and turbines, respectively. The experiments show that using forecast-based predictions results in better performance than either the conservative approach or the variants of the PPF model. With solar power, PPF and its variants approve more leases than our forecast-based approach, but complete only three-fourths of the leases it approves without a node running out of energy. While the conservative approach completes all of the leases it approves without ever depleting any node's battery, it completes only half of the leases of PPF. In contrast, our forecast-based approach combines the best characteristics of both: it completes nearly as many leases as the PPF model without depleting any batteries. Fig. 11(b) shows better results for wind power, since PPF is less useful for predicting wind speed.

Note that an approach, such as PPF, that over predicts the available energy tends to approve more jobs than it could complete, and, as a result, runs out of the energy more often. In contrast, an approach, such as FPF or conservative, that under predicts the available energy might approve less requests, but is able to complete almost all of them, while rarely running out of energy.

#### 4.2. Lexicographically fair sensor systems

Computing lexicographically fair sensing rates in energy harvesting sensor systems has been studied recently for both static and mobile networks [10]. Put simply, an assignment of sensing rates to nodes is lexicographically fair if it is impossible to increase the sensing rate of any node without decreasing the sensing rate of another node due to either bandwidth or energy constraints.

For energy harvesting systems, the primary constraint is that each node must maintain energy-neutral operation. An assignment of lexicographically fair rates will not be valid if the energy harvesting behavior changes due to weather conditions. However, recomputing sensing rates for all nodes is time-consuming and network-intensive, since it requires gathering the most recent energy harvesting data from each node, computing new rates, and distributing the new rates to all nodes in the network.

Thus, prior work sets a fixed coarse-grain time period, e.g., 1 day [10], to recompute the fair rate for each node based on the most recent energy harvesting information. The duration of the time period balances the expense of resetting rates globally with the risk of any node running out of energy due to stale or imprecise energy harvesting information. As a result, maintaining steady network-wide sensing rates for a fixed time period represents an instance of inelastic demand, since the system is not able to continuously vary the sensing rates for all nodes, which ultimately determine the energy demand, to precisely match the energy supply. The recomputation may use either a conservative approach that only takes into account current battery reserves or a prediction model that accounts for expected future energy [10].

We evaluate the use of our forecast-based predictions in a lexicographically fair sensor network in simulation. In this case, the performance metrics we use are the (i) maximum rate allocated in a 24 h interval and (ii) the number of 24 h intervals where one or more nodes run out of energy. We examine a networked setting based on a deployment of five conventional TelosB motes, with the same power characteristics as the nodes in [28] in a simple tree topology using the distributed algorithm developed by Fan et al. [10] to compute the network-wide rates every 24 h. We view our use of only five nodes as conservative: increasing the number of nodes in the network also increases the benefits of better predictions, since, as the number increases, more nodes are capable of depleting their battery. For the TelosB simulation, we scale down the power output to 1% of the power produced by both our 60 watt solar panel and our 400 watt wind turbine to better match the characteristics of the TelosB's low energy demand.

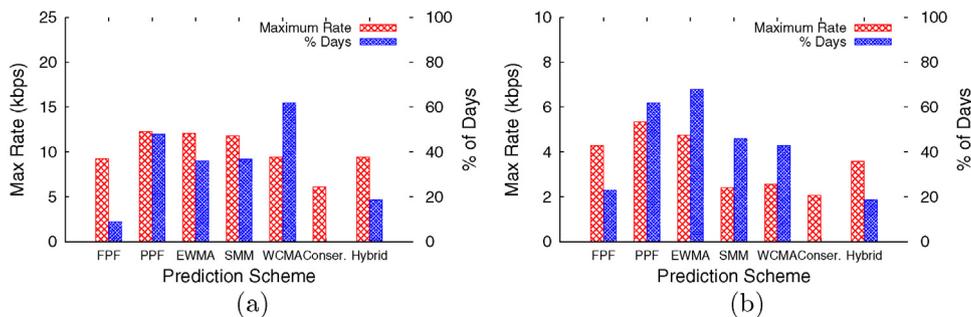


Fig. 12. Maximum sensing rate and % of days one or more nodes run out of energy for different prediction schemes for both solar (a) and wind (b) power.

Fig. 12(a) shows that solar power predictions based on PPF overestimate the proper maximum rate, causing at least one node to deplete its battery on nearly 50% of the days. In contrast, setting the rates using our forecast-based approach results in at least one node depleting its battery on less than 5% of all days, while maintaining 80% of the rate set using the PPF approach. As expected, a conservative approach never depletes any batteries but sets a rate near 45% of the rate using PPF. Fig. 12(b) shows similar results for wind. However, since wind is more intermittent than solar, both the forecast-based and PPF approach have at least one node that runs out of power on more days. In both cases, the standard deviation of sensing rates each day in the PPF model ( $\sim 7.4$ ) are more than the standard deviation of the sensing rates using our forecast-based approach ( $\sim 6.1$ ). Thus, our forecast-based approach maintains more consistent rates between each 24 h period than the PPF model, which is an important goal for steady rate allocation.

#### 4.3. Smart home

Our final case study examines how to use predictions to efficiently manage renewable energy in smart homes. Many homes today that use renewable solar energy leverage net metering to reduce costs. Net metering allows homes to sell excess energy back to the grid when the home does not need it. The price reductions from net metering are critical in providing additional financial incentives to reduce the cost of on-site renewables in homes. Unfortunately, net metering is not a scalable solution. The electric grid must always balance supply and demand by dispatching generators to match demand as it varies. At large scales across thousands of homes, electricity demand is highly predictable and changes slowly throughout the day, which permits well-planned generator dispatch schedules. However, incorporating significant amounts of volatile renewable energy sources into the grid disrupts dispatch schedules and may destabilize grid operations. As a result, state laws often place caps on both the total number of participating customers and/or the total amount of energy contributed per customer [1]. After exceeding these caps, utilities are not obligated to purchase excess power. For example, Washington state caps the total number of participating customers at 0.25% of all customers.

In related work [31], we propose combining a small on-site battery with time-of-use electricity pricing to lower utility costs, and provide a financial incentive similar to net metering. The goal is to determine when to operate off the battery, e.g., when prices are high, versus the grid and when to charge the battery from the grid, e.g., when prices are low. However, the approach requires accurate predictions of future solar generation to make effective charging and discharging decisions. Here, we apply and evaluate the performance of each of the solar generation prediction methods. We experiment with each methods in the context of simple residential TOU pricing model used by the Ontario Electric Board (OEB). The OEB divides rates into three categories: on-, mid-, and off-peak. The on-peak rate is 10.7 ¢/kWh from 7 am to 11 am and from 5 pm to 9 pm, the mid-peak rate is 8.9 ¢/kWh from 11 am to 5 pm, and the off-peak rate is 5.9 ¢/kWh from 9 pm to 7 am. The OEB sets a different fixed ratio for on-, mid-, and off-peak rates in the summer (May 1st–October 31st) and winter (November 1st–April 30th), and on weekends and holidays. However, the exact rates change on a monthly basis according to generation costs and demand. Residential TOU pricing is still a nascent concept; the rates above still do not accurately reflect the price of energy, which is much more volatile in wholesale energy markets. In these markets, spot prices vary as little as every 5 min and may differ by orders of magnitude each day.

Given electricity rates, the charging algorithm determines one day in advance a schedule for when the home should use the battery versus the grid for power, and when to charge the battery from

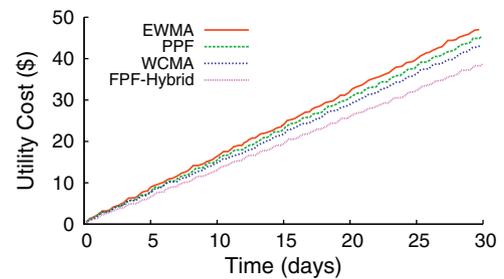


Fig. 13. The cost of electricity from the utility for different prediction models. The forecast-based hybrid model results in a 12–23% cost decrease compared to others.

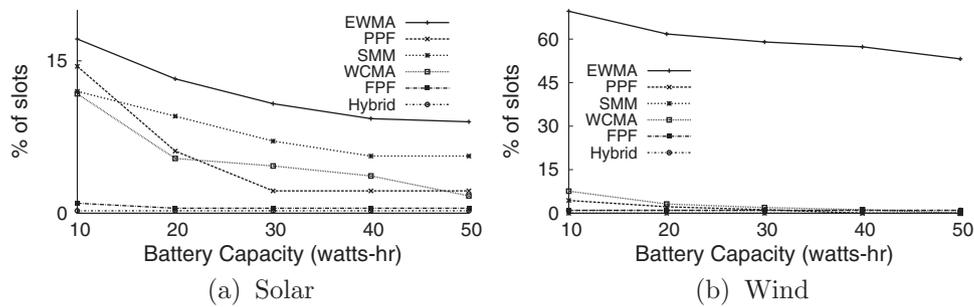
the grid. The inputs to the algorithm are predictions of how much solar energy the home will harvest the next day, how much energy the home will consume, the TOU rate plan as described above, and the battery's capacity and current energy level. We assume a 12 kWh battery array, which is similar to the capacity of battery in an electric car. We then use second-by-second traces of home power consumption from a real home, and quantify the cost savings from each prediction model. We scale up one month of traces from our solar panel deployment by 17 to align with the aggregate energy consumption of the home. The final result is in Fig. 13, which shows that using OEB rates, the forecast-based Hybrid prediction model again results in the lowest cost with a bill of \$38.60. In contrast, the WCMA model increases costs by 12% (\$43.27), the PPF model increases costs by 18% (\$45.43), and finally the EWMA model increases costs by 23% (\$47.55).

We view our results as conservative, since they are based on today's simple TOU pricing plans. More volatile TOU pricing plans that better reflect the current price of energy would improve the results, since there would be a greater penalty for misprediction. That said, a 12–23% reduction in electricity bills provides additional incentives for incorporating sophisticated prediction strategies into smart homes.

#### 4.4. Battery capacity

Battery capacity affects the performance of an energy harvesting system by storing the surplus energy when energy demand is less than the energy supply, and by providing extra energy when energy demand is greater than the supply. For example, our energy harvesting systems predict a fixed amount of energy available at one hour granularities. If the system harvests more energy than predicted, it is able to store the extra energy in its battery. Similarly, if the system harvests less energy than expected, it may use the reserve energy in its battery energy to mitigate the impact of the misprediction. Since battery cannot store more energy than its capacity, harvested energy may be wasted when the surplus energy exceeds the battery capacity. Sharma et al. [21] shows the energy vs voltage graph for a lead-acid battery, similar to the battery used in this paper. A larger capacity battery is able to store more surplus energy and, thus, prevent the system from depleting its energy reserves when its demand exceeds the supply.

To study the impact of battery capacity in energy harvesting systems, we divide a day into 24 1-h slots. At the end of a slot, our energy harvesting system predicts the incoming energy for the next slot and sends this information to the application. Please note that the application's energy demand for any slot is limited to the energy predicted by the energy harvesting system for that slot. For this experiment, we do not care about the nature or type of the application; we assume that its demand is equal to the energy predicted by the system. Fig. 14 plots the percentage of slots the battery ran out of energy for different prediction models used to predict the energy for next slot by our energy harvesting system.



**Fig. 14.** % of slots our energy harvesting system runs out of energy for different battery capacities. Each slot is of 1-h duration and the experiment is run over the first 3 weeks of October, 2009

For solar energy, our prediction model (based on the weather forecast) outperforms the PPF model and its variants, whereas for wind energy, our model and two variants of the PPF model—WCMA and SMM—perform equally well. Even for wind energy, our model performs better than WCMA and SMM for low capacity batteries (10 watts-h). Though EWMA performs equally well, compared to other variants of PPF for solar energy, it performs poorly for wind energy with the battery depleting its energy reserves on almost half of the days. Thus, EWMA is suitable for consistent weather patterns, but not for inconsistent or intermittent weather patterns. WCMA performs better than all other PPF models—PPF, EWMA, SMM—for both solar and wind energy.

## 5. Prior work

We know of no prior work that evaluates the use of forecast-based predictions in energy harvesting systems. Much of the prior work on energy harvesting sensor systems assumes elastic workload demands that do not require predictions, since the system continually adapts its workload's intensity and energy usage to match its energy supply [13,26,28]. However, while Moser et al. [18] assume perfect future knowledge of an energy source and do not investigate prediction strategies, they do note that scheduling algorithms for workloads with inelastic demands are highly sensitive to the accuracy of predictions. While our observation about the inter- and intra-day variations in solar radiation hold for Amherst, Massachusetts, prior work on solar harvesting assumes diurnal behavior that is more consistent than we observe [6,30]. In these areas, the NWS forecast-based approach may be less effective.

Most prior work focuses on simple prediction schemes, such as the PPF model, based on the immediate past [13,17]. As we show, the simple PPF approach is not as accurate as a NWS forecast-based approach for either solar or wind power at time-scales of hours to days. Kansal et al. [13] maintain an exponentially weighted moving average (EWMA) for solar power to achieve energy-neutral operation in a system with elastic workload demands. The EWMA approach is a variant of PPF that adapts to seasonal variations in solar radiation. However, EWMA does not account for drastic changes in weather that the NWS forecast predicts. Noh et al. [19] use a historical model for solar radiation, akin to WCMA, that maintains an expectation for each time slot in a day based on the previous day's solar radiation reading, but down-scales all future time-slots in a day by  $N\%$  if it records a solar radiation reading  $N\%$  less than expected.

The techniques above do not apply to wind speed or wind power predictions, since the wind is more intermittent than solar radiation and not diurnal in nature. We know of no work that discusses prediction strategies for wind speed. The recent commoditization and emergence of micro-wind turbines, such as the 400 watt Air-X we use in our deployment, motivates further study of harnessing

wind power in sensor systems deployed at locations with ample wind but little sunlight, i.e., during the winter in the extreme north or south.

## 6. Conclusion

In this paper, we show how to leverage weather forecasts provided by the NWS to enhance the ability of energy harvesting sensor systems to satisfy their demand. We analyze observational weather data from our own weather station, energy harvesting data from our own solar panel and wind turbine, and NWS observational and forecast data. Our analysis shows that weather predictions based on NWS forecasts are more accurate than predictions based on the past in many regions of the United States, including Amherst, Massachusetts. To leverage NWS forecasts in sensor systems, we formulate a model for our solar panel and wind turbine that converts the forecast to an energy harvesting prediction. We then compare our models with other approaches in three case studies of systems with inelastic workload demands—an off-the-grid distributed testbed, a lexicographically fair sensor system, and a smart home—and show that for both solar and wind power our models our forecast-based approach improves system performance. Designing prediction models for obstructed energy harvesting is out of the scope of this paper. In future, we plan to extend our prediction models to consider obstructions due to trees and buildings. Finally, one advantage of our approach is that it does not require training data to learn a prediction model over time. Such models require a large amount of training data, e.g., many years, especially for solar, since each time of the day and each time of the year the sun has different solar capacity at a given location. Our model is more appropriate for new deployments that have not operated for many years, e.g., such as a homeowner that has newly installed solar panels.

## Acknowledgements

We would like to thank the anonymous reviewers for their insightful comments that improved this paper. This work was supported in part by NSF grants CNS-0916577, CNS-0855128, CNS-0834243, EEC-0313747, and CNS-1217791.

## References

- [1] Freeing the Grid: Best and Worst Practices in State Net Metering Policies and Interconnection Procedures. (2009). [www.newenergychoices.org/uploads/FreeingTheGrid2009.pdf](http://www.newenergychoices.org/uploads/FreeingTheGrid2009.pdf)
- [2] Southwest Windpower, Inc. AirX Land Owner's Manual, 2002.
- [3] B. Aksanli, J. Venkatesh, L. Zhang, T. Rosing, Utilizing green energy prediction to schedule mixed batch and service jobs in data centers, in: HotPower 2011 Proceedings of the 4th Workshop on Power-Aware Computing and Systems, 2011.

- [4] J. Apt, A. Curtright, The spectrum of power from utility-scale wind farms and solar photovoltaic, Carnegie Mellon Electricity Industry Center, 2004, Technical Report CEIC-08-04.
- [5] L. Barroso, U. Hölzle, The case for energy-proportional computing, *Computer* (2007).
- [6] P. Corke, P. Valencia, P. Sikka, T. Wark, L. Overs, Long-duration solar-powered wireless sensor networks, *EmNets* (2007).
- [7] B. Donovan, D. McLaughlin, M. Zink, J. Kurose, OTGsim: simulation of an off-the-grid radar network with high sensing energy cost, *SECON* (2008).
- [8] P. Dutta, et al., Trio: enabling sustainable and scalable outdoor wireless sensor network deployments, *IPSN* (2006).
- [9] E. Ertin, A. Arora, R. Ramnath, M. Nesterenko, Kansei: a testbed for sensing at scale, *IPSN* (2006).
- [10] K. Fan, Z. Zheng, P. Sinha, Steady and fair rate allocation for rechargeable sensors in perpetual sensor networks, *SenSys* (2008).
- [11] I. Goiri, K. Le, Md.E. Haque, R. Beauchea, T.D. Nguyen, J. Guitart, J. Torres, R. Bianchini, GreenSlot: Scheduling Energy Consumption in Green Datacenters SC, 2011.
- [12] X. Jiang, J. Polastre, D. Culler, Perpetual environmentally powered sensor networks, *IPSN* (2005).
- [13] A. Kansal, J. Hsu, S. Zahedi, M. Srivastava, Power management in energy harvesting sensor networks, *Trans. Embed. Comput. Syst.* (2007).
- [14] A. Kansal, D. Potter, M. Srivastava, Performance aware tasking for environmentally powered sensor networks, *SIGMETRICS Perform. Eval. Rev.* (2004).
- [15] K. Lin, et al., Heliomote: Enabling long-lived sensor networks through solar energy harvesting, *SenSys* (2005).
- [16] M. Minami, T. Morito, H. Morikawa, T. Aoyama, Solar biscuit: a battery-less wireless sensor network system for environmental monitoring applications, *INSS* (2005).
- [17] Clemens, Moser Power Management in Energy Harvesting Embedded Systems, *ETH Zurich*, 2009, Ph.D. Thesis.
- [18] C. Moser, D. Brunelli, L. Thiele, L. Benini, Real-time scheduling for energy harvesting sensor nodes, *RTSS* (2007).
- [19] D. Noh, L. Wang, Y. Yang, H. Le, T. Abdelzaher, Minimum variance energy allocation for a solar-powered sensor system, *DCSS* (2009).
- [20] The GENI Project Office GENI System Overview. The GENI Project Office, 2008.
- [21] N. Sharma, S. Barker, D. Irwin, P. Shenoy, Blink: managing server clusters on intermittent power, *ASPLOS* (2011).
- [22] N. Sharma, J. Gummeson, D. Irwin, P. Shenoy, Cloudy computing: leveraging weather forecasts in energy harvesting sensor systems, *SECON* (2010).
- [23] N. Sharma, D. Irwin, P. Shenoy, M. Zink, MultiSense: Fine-grained multiplexing for steerable camera sensor networks, *MMSys* (2011).
- [24] J. Recas, C. Bergonzini, T. Simunic Rosing, D. Atienza, Prediction and management in energy harvested wireless sensor nodes, *Wireless VITAE* (2009).
- [25] M. Ali, B. Al-Hashimi, J. Recas, D. Atienza, Evaluation and design exploration of solar harvested-energy prediction algorithm, *DATE* (2010).
- [26] J. Taneja, J. Jeong, D. Culler, Design, modeling, and capacity planning for micro-solar power sensor networks, *IPSN* (2008).
- [27] C. Vigorito, D. Ganesan, A. Barto, Adaptive control for duty-cycling in energy harvesting-based wireless sensor networks, *SECON* (2007).
- [28] L. Wang, et al., AdaptSens: An adaptive data collection and storage service for solar-powered sensor networks, *RTSS* (2009).
- [29] G. Werner-Allen, P. Swieskowski, M. Welsh, Motelab: a wireless sensor network testbed, *IPSN* (2005).
- [30] Y. Yang, L. Wang, D. Noh, H. Le, T. Abdelzaher, SolarStore: enhancing data reliability in solar-powered storage-centric sensor networks, *MobiSys* (2009).
- [31] T. Zhu, A. Mishra, D. Irwin, N. Sharma, D. Towsley, P. Shenoy, Efficient Renewable Energy Management for Smart Homes, University of Massachusetts Amherst, 2011, Technical Report.



**Navin Sharma** is currently a Postdoctoral Research Associate in the School of Computer Science at the University of Massachusetts Amherst. He received his M.S. and Ph.D. in Computer Science from the University of Massachusetts Amherst in 2010 and 2013, respectively. His research focuses on renewables powered green datacenters, distributed systems, green computing, smart homes and smart grids.



**Jeremy Gummeson** is currently a Research Scientist in the Networking and Mobility Lab at HP Labs in Palo Alto, California. He received his Ph.D in Computer Systems Engineering from the University of Massachusetts Amherst in 2014. His research interests are in mobile computing, wireless sensor networks, energy harvesting, and RFID.



**David Irwin** is an Assistant Professor in the Department of Electrical and Computer Engineering at the University of Massachusetts Amherst. He received his Ph.D. and M.S. in Computer Science from Duke University in 2007 and 2005, respectively, and his B.S. in Computer Science and Mathematics from Vanderbilt University in 2001. His research interests are broadly in experimental computing systems with a particular emphasis on sustainability.



**Ting Zhu** is an Assistant Professor in the Department of Computer Science and Electrical Engineering at the University of Maryland, Baltimore County. He received his Ph.D. from the Department of Computer Science and Engineering at the University of Minnesota in 2010. He received his B.S. and M.E. degree from the Department of Information Science and Electronic Engineering at the Zhejiang University in 2001 and 2004, respectively. His research interests include wireless networks, renewable energy, embedded systems, distributed systems, and security.



**Prashant Shenoy** received the B.Tech degree in Computer Science and Engineering from the Indian Institute of Technology, Bombay, in 1993, and the M.S and Ph.D. degrees in Computer Science from the University of Texas, Austin, in 1994 and 1998, respectively. He is currently a Professor of Computer Science at the University of Massachusetts. His current research focuses on cloud computing and green computing. He is a distinguished member of the ACM and a Fellow of the IEEE.