

Cloudy Computing: Leveraging Weather Forecasts in Energy Harvesting Sensor Systems

Navin Sharma, Jeremy Gummeson, David Irwin, and Prashant Shenoy

University of Massachusetts, Amherst

{nksharma,gummeson,irwin,shenoy}@cs.umass.edu

Abstract—To sustain perpetual operation, 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 and its hardware components are not energy-proportional, since it cannot precisely scale its usage to match its supply. Instead, the system must choose when to satisfy its demands based on its current energy reserves and predictions of its future energy supply. In this paper, we explore the use of weather forecasts to improve a system’s ability to satisfy demand by improving its predictions.

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 evaluate our model for both energy sources in the context of two different energy harvesting sensor systems with inelastic demands: a testbed that leases sensors to external users and a lexicographically fair sensor network that maintains steady node sensing rates. We show that using weather forecasts for predictions in both solar- and wind-powered sensor systems increases each system’s ability to satisfy its demands compared with existing strategies.

I. INTRODUCTION

Energy harvesting systems collect and store environmental energy to sustain continuous operation without access to external power sources.¹ Energy harvesting has been well-studied in the context of sensor networks, because they are often deployed in remote locations with no access to the power grid. Energy-neutral systems always consume less than or equal to the energy they harvest [10]. 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 hardware components are *energy-proportional*, such that their energy consumption scales linearly with their workload’s intensity [3]. 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 [6], [9], [10], [12], [13], [19], [20], [21]. 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.

Instead, 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 sensing requests from system users, or internal objectives, such as maintaining a stable workload for a long period of time. The former is relevant to sensor testbeds [7], [18], [22] that operate off harvested energy, since external users dictate the workload. 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 [8]. As others have noted, workload scheduling algorithms in energy harvesting systems with inelastic demands are highly sensitive to energy harvesting predictions [15].

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 [8], [11], [15], [16]. 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 sensor 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. Our hypothesis is that energy harvesting predictions derived from weather forecasts for large regions improve nearby systems’ ability to satisfy their demand over the time-scales of hours to days, when compared against predictions derived from the immediate past.

¹This work is supported in part by the National Science Foundation under grants #CNS-0916577, #CNS-0855128, and #CNS-0834243.

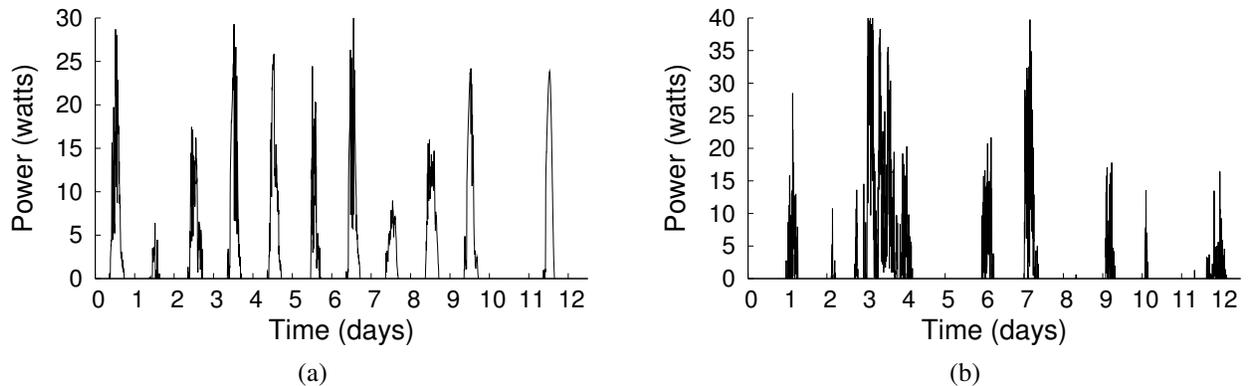


Fig. 1. Power generated during a 12 day period in October, 2009 from our solar panel (a) and wind turbine (b).

In evaluating our hypothesis this paper makes the following contributions.

Analyze Historical Weather Data. We analyze extensive traces of past forecast and observational weather data from the National Weather Service (NWS), as well as fine-grain solar and wind energy harvesting and observational weather data from our own deployment. We use these traces to quantify how well both weather forecasts and 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 3 hours to 72 hours. We find that NWS forecasts in the regions we examine are a better predictor of the future than the immediate past at these 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 entire region with our own local weather observations and (ii) our own local weather observations with the energy harvested by our deployed solar panel and wind turbine. We use both data sets to formulate a simple model that predicts how much energy our solar panel and wind turbine will harvest in the future given weather forecasts every 3 hours from 3 hours to 72 hours in the future.

Case Studies. We quantify the benefits of using energy harvesting predictions based on weather forecasts in the context of two different types of energy harvesting sensor systems with inelastic demand. The first system is a deployed sensor testbed [18] we have built as part of the NSF GENI prototype [17], while the second is inspired by recent work on lexicographically fair energy harvesting sensor systems [8]. For both systems, we find that our forecast-based approach is at least 25% better compared with predictions based on the immediate past for each system’s relevant performance metrics, e.g., requests satisfied and/or combined length of power outages.

In Section 2, we motivate the use of NWS weather forecasts by comparing them with predictions based on the immediate past. Section 3 then formulates a model that translates a weather forecast to an energy harvesting prediction for our solar panel and wind turbine, while Section 4 presents two case studies that demonstrate the benefits of deriving energy

harvesting predictions from weather forecasts. Finally, Section 5 presents related work and Section 6 concludes.

II. 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 [2], [4], [10] and wind [11] 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 at Amherst, as well as data from NWS observations and the National Digital Forecast Database. Our observational traces are available 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 at 5 minute granularities, while the NWS reports an observation every hour and a forecast every 3 hours for every region of the country for the last 4 years. Each NWS forecast includes predictions every 3 hours from 3 hours to 72 hours 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 units of miles per hour. We show how these metrics impact solar and wind energy harvesting in Section III.

While we have found multiple approaches in prior work that use the immediate past to predict the future, the basic approach, which we term *past predicts the future* or *PPF*, we compare against 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. We discuss variants of

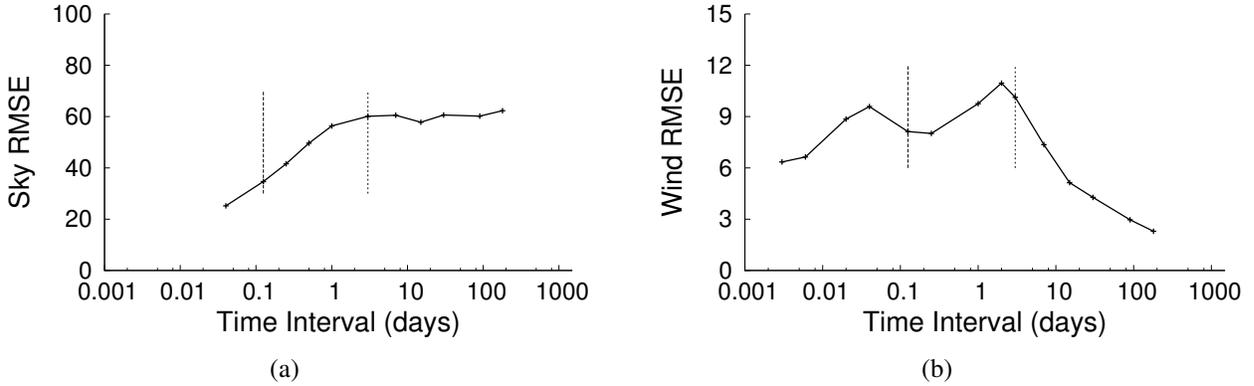


Fig. 2. The error in sky condition (a) and wind speed (b) when using the past to predict the future for different time intervals in 2008 at 1 hour and 5 minute granularities, respectively, for Amherst, Massachusetts.

this basic approach for solar power prediction in Section V that adapt to seasonal variations in sunlight [10], [11], [19] or sudden changes in cloud cover [16]. We have found no prior work that focuses on variants of the PPF model for wind speed predictions.

The accuracy of the PPF model 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 [4], 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: Both sky condition and wind speed show significant inter-day and intra-day variations, as a result of changing weather conditions in Amherst, Massachusetts, as well as other regions we examine, including Arizona, Florida, Washington, and Nebraska.

While we expect wind speeds 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 own deployment, Figures 1(a) and 1(b) show the solar and wind power we harvest, respectively, during a 12 day period in October, 2009. As expected, wind power 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 aggregate 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 $\frac{1}{2}$ 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 since the beginning of our solar panel and wind turbine deployment 4 months ago.

Observation #2: Using PPF to predict the future is least accurate at medium-length time-scales ranging from 3 hours 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 minutes. RMSE is a standard statistical measure of the accuracy of values predicted by a model with respect to the values observed. Intuitively, the value of the RMSE quantifies the accuracy of the PPF model 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 average of the metric in the previous interval exactly predicts the average of the metric 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.

Figures 2(a) and 2(b) show the RMSE for sky condition and wind speed, respectively, as a function of time interval duration N ranging from 5 minutes 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 minutes) and long time-scales (> 10 days), and are least accurate in between. For both sky condition and wind speed,

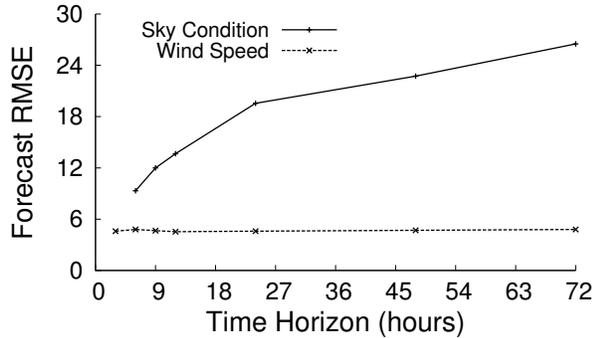


Fig. 3. RMSE between the observed sky condition and wind speed and those predicted by NWS forecasts from 3 hours to 72 hours in the future.

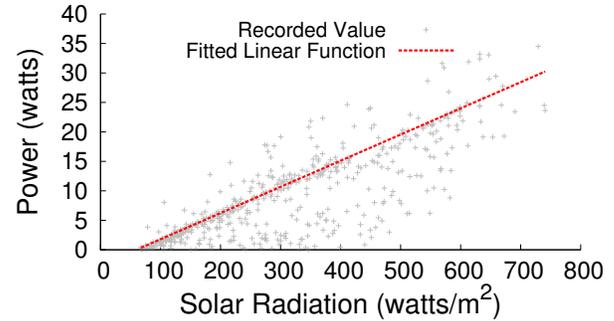


Fig. 4. Relationship between the solar radiation our weather station observes and the power generated by our solar panel.

the maximum inaccuracy occurs between 3 hours and one week, as indicated by each graph’s vertical lines.

Observation #3: Over the forecast time-scales from 3 hours to 3 days provided by the NWS, 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—for Chicopee Falls, Massachusetts. Chicopee Falls, at 20 miles away, is the closest NWS site to Amherst. We first compare the accuracy of a forecast for sky condition with the accuracy of the PPF model from Figure 2(a). Figure 3 shows the RMSE between the observational sky condition and the sky condition from the NWS forecasts, as a function of the forecast time horizon. Since our weather station does not report sky condition, we use the hourly NWS observations of sky condition at Chicopee Falls, Massachusetts. As expected, the accuracy of the sky condition forecast decreases as the time horizon increases. Since the RMSE of the sky condition forecast (<30) is less than the RMSE of the PPF model from Figure 2(a) between 3 hours and 3 days (~ 60) we conclude that the forecast is a better predictor than the past for sky condition in Amherst, Massachusetts

We next compare the accuracy of the NWS forecast for wind speed with the accuracy of the PPF model from Figure 2(b). Figure 3 shows the RMSE between the observational wind speed and the wind speed from the NWS forecast, as a function of the forecast time horizon. As the figure shows, the accuracy of the wind speed forecast does not vary significantly for any future time horizon. The cumulative distribution of errors in the wind speed forecast echoes the point, since the error for each time horizon is roughly equivalent, with 80% of the errors being less than 7mph. Since our wind turbine only generates power at wind speeds greater than 7 mph, we also examined the wind speed forecast accuracy after filtering out lower wind speeds. Our analysis shows that the accuracy of the wind speed

forecast only increases as we raise the filtering threshold. Since the RMSE of the NWS wind speed forecast (<6) is less than the RMSE of the PPF model from Figure 2(b) between 3 hours and 3 days (~ 11), 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 be able to improve the performance of energy harvesting systems with inelastic demands that make workload decisions over 3 hour to 3 day time horizons.

III. FORECAST \rightarrow ENERGY MODEL

To leverage our observations from the previous section, we now formulate models that predict the power 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. 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 may be 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

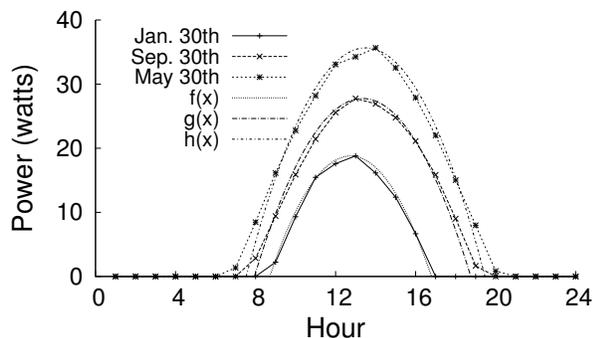


Fig. 5. Profile for solar power harvested on clear and sunny days in January, May, and September, and the quadratic functions $f(x)$, $g(x)$, and $h(x)$ we fit to each profile, respectively.

battery from overcharging. Our battery has an ideal capacity of 1260 watt-hours.

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 seconds and stores a 5 minute 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.

A. Sky Condition \rightarrow 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 example, if the 3 hour forecast predicts a sky condition with 50% cloud cover, and the maximum possible solar power production is 60 watts over that 3 hour interval, then the solar power prediction for that 3 hour 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.

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 Figure 4. The relationship should be linear, since our solar panel produces energy in proportion to the solar radiation with a constant factor loss due to

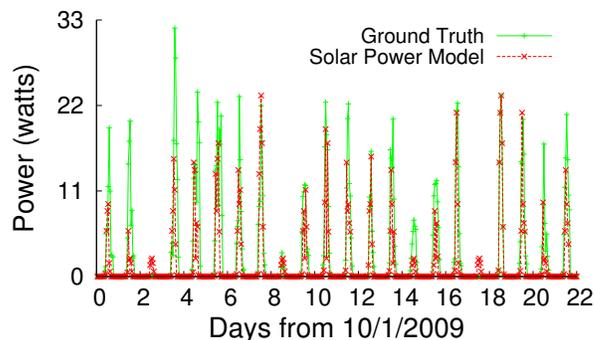


Fig. 6. Power output from our solar panel and the power output predicted by our solar model over the first 3 weeks of October, 2009. The dark foreground line is our solar power model, while the lighter background line is the ground truth.

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^2 .

$$\text{SolarPower} = 0.0444 * \text{Radiation} - 2.65 \quad (1)$$

2) *Computing the Maximum Possible Solar Power:* We next derive our 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.

Figure 5 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. For the graph, we convert the solar radiation observed by our weather station on these days to the expected solar power harvested by our solar panel using equation (1) from above. We find that power appears 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, 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 for each month are given in Table I, and $Time$ is in hours after 12am.

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

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

TABLE I

VALUES FOR A, B, AND C IN OUR QUADRATIC SOLAR POWER MODEL, WHICH IS A FUNCTION OF THE TIME OF DAY FOR EACH MONTH OF THE YEAR.

3) *Solar Model*: To complete our model, we compute the solar power our panel generates using the equation below, where $MaxPower$ is in units of watts from equation (2) and $SkyCondition$ is the percentage cloud cover from the NWS. Figure 6 compares the observed solar power generated by our panel with the solar power predicted by our model. 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.

$$Power = MaxPower * (1 - SkyCondition) \quad (3)$$

B. Wind Speed \rightarrow 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. Figure 7 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 [1].

$$Power = 0.01787485 * (WindSpeed)^3 - 3.4013 \quad (4)$$

We fit the cubic power curve in equation (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. As future work, we plan to investigate the use of wind gust readings to indirectly derive a measure of turbulence and provide a more accurate prediction. Note that the wind turbine stops producing power near 28 miles per hour, so our function ramps down to 0 at that point.

C. Compensating for Forecast Errors

Our solar and wind power models convert an observed sky condition and wind speed to the expected solar and wind

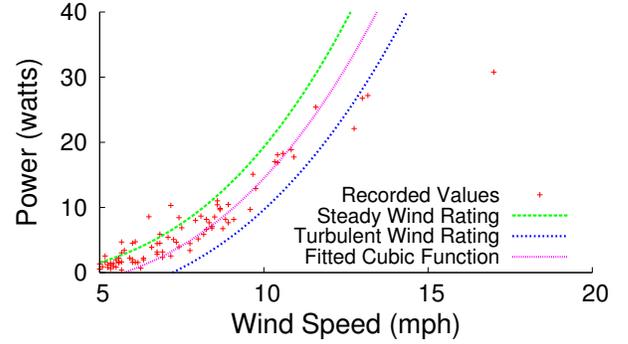


Fig. 7. Power output from our wind turbine and the power output predicted by our wind power model over the first 3 weeks of October. The graph shows the rated power curves from the wind turbine's manual [1] for steady and turbulent wind, as well as our fitted curve.

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 α . We base our α 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 α in our model.

IV. CASE STUDIES

We evaluate our models from the previous section in the context of two types of energy harvesting sensor systems with inelastic demand: a testbed that we have built as part of the NSF GENI prototype [18] that leases sensors to users, and a lexicographically fair sensor network inspired by recent work [8]. For both systems, we quantify how much the use of forecast-based predictions increases each 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.

A. ViSE Testbed

ViSE, which stands for Virtualized Sensing Environment, is a 4-node high-power sensor testbed we have built as part of GENI [17]. 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. 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 a specific 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 a multi-testbed experiment, and thereby waste any resources previously 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

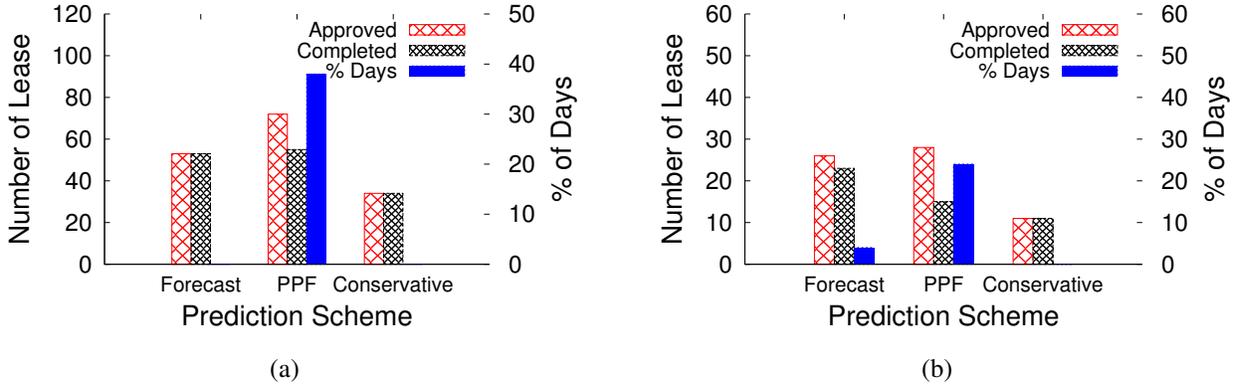


Fig. 8. Since forecast-based predictions are more accurate than PPF, using them to make workload decisions results in improved performance for both solar (a) and wind (b) power. In both cases, when using PPF, ViSE tends to approve more requests but service less of them, while when using forecast-based predictions ViSE tends to only approve requests it is able to service. A conservative approach completes every request it approves without running out of energy, as indicated by the right-most Y-axis which measures the percentage of days each approach depletes its battery, but completes almost half as many requests.

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 throughout its duration. Note that our prediction-based approach not only ensures that there is enough energy at the end of the lease, but also at all times throughout the lease’s duration based on the energy harvested while the lease is active.

To evaluate the benefits of our models from Section III relative to both the conservative approach and 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 [5]. 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 that starts at 50% capacity and is capable of running our node at full utilization for 2 hours.

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 hours. We discuss ViSE’s approach

to virtualizing sensors in recent work [18]. 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 lease requests that it believes it can satisfy based on a conservative approach, the NWS forecast-based model, or the PPF model, where we assume each virtual sensor 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 any node 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 both solar and wind energy harvesting. Figures 8(a) and 8(b) show the number of lease requests ViSE approves, as well as the number of leases ViSE completes for the solar panels and wind turbines, respectively. The experiments show that using forecast-based predictions results in better performance than either the conservative approach or the PPF model. With solar power, the PPF model approves more leases than our forecast-based approach, but completes only three-fourths of the leases it approves without any 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. Figure 8(b) shows better results for wind power, since PPF is less useful for predicting wind speed.

B. Lexicographically Fair Sensor Systems

Computing lexicographically fair sensing rates in energy harvesting sensor systems has been studied recently for both static and mobile networks [8]. 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

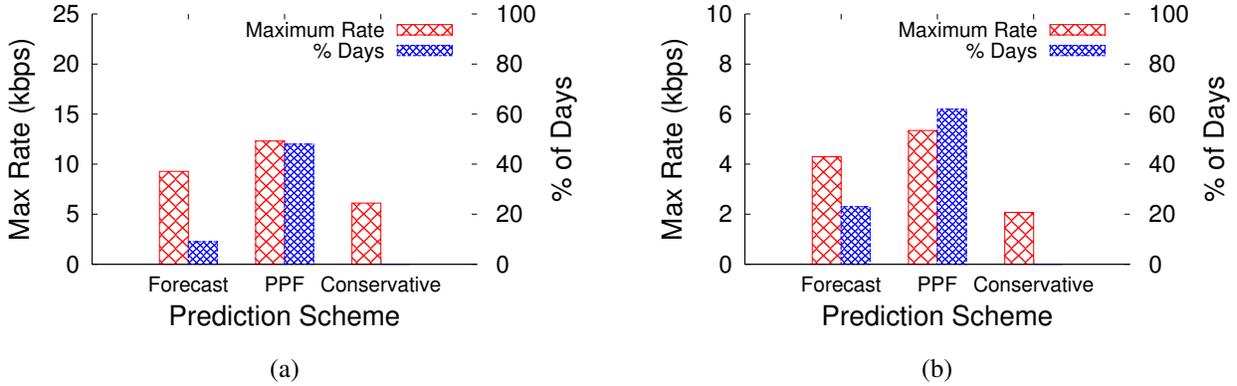


Fig. 9. Using PPF, a lexicographic sensor network sets a higher sensing rate than other approaches each day, but has at least one node that runs out of energy on nearly 50% of all days. In contrast, when using our NWS forecast-based predictions, the system has at least one node run out of energy on less than 5% of all days while maintaining a steady sensing rate only 20% less than the PPF approach. While the conservative approach never has a node run out of energy, as indicated by the right-most Y-axis which measures the percentage of days each approach depletes at least one battery, it sets the rate each day at almost half the rate of the forecast-based approach.

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 [8], 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 the current battery reserves or a prediction model that accounts for the expected future energy supply [8].

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 hour interval and (ii) the number of 24 hour 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 [21] in a balanced tree topology using the distributed algorithm developed by Fan et al. [8] to compute the network-wide rates every 24 hours. 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.

Figure 9(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. Figure 9(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 (~ 7.4) are more than the standard deviation of the sensing rates using our forecast-based approach (~ 6.1). As a result, our forecast-based approach also maintains more consistent rates between each 24 hour time period than the PPF model, which is an important goal for steady rate allocation.

V. 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 [10], [19], [21]. However, while Moser et al. [15] 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 [4], [23]. 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 [10], [14]. 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. [10] 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. [16] use a historical model for solar radiation 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. We plan to study how these variants of PPF compare with our forecast-based approach as part of future work.

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.

VI. 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 two case studies of sensor systems with inelastic workload demands, and show that for both solar and wind power our models combine the best characteristics of the PPF model and a conservative approach. While our work here is based primarily on extensive data analysis and simulation, as part of future work, we plan to integrate our ViSE testbed with the NWS forecast web service that allows programmatic querying of live forecast data.

As part of our ongoing work, we are studying the combination of workload and hardware characteristics where forecasts provide the most benefit. For example, the capacity of the battery affects how well a forecast-based approach does relative to a conservative approach. With an infinite capacity we expect the conservative and forecast-based approaches to be nearly equivalent for the performance metrics we cite, since

the system will always eventually use its store of energy. We are also studying different ways to leverage predictions in sensor systems. For instance, in our case studies we use a fixed 24 hour time interval to recompute the rates in a lexicographically fair sensor network, but we could also use our predictions to instead calculate the appropriate time interval before recomputing new rates. Finally, we plan to investigate if using other parameters in the NWS forecast, in addition to sky condition and wind speed, enhance our predictions, as well as study the use of hybrid solar panel and wind turbine systems, since our traces indicate that wind power generation often increases at the same time that solar power generation decreases, e.g., during thunderstorms or in the evening.

REFERENCES

- [1] Southwest Windpower, Inc. AirX Land Owner's Manual. October, 2002.
- [2] J. Apt and A. Curtright. The Spectrum Of Power From Utility-scale Wind Farms And Solar Photovoltaic. Carnegie Mellon Electricity Industry Center, Technical Report CEIC-08-04, August 2004.
- [3] L. Barroso and U. Hözl. The Case For Energy-Proportional Computing. *Computer*, December 2007.
- [4] P. Corke, P. Valencia, P. Sikka, T. Wark, and L. Overs. Long-duration Solar-powered Wireless Sensor Networks. *EmNets*, June 2007.
- [5] B. Donovan, D. McLaughlin, M. Zink, and J. Kurose. OTGsim: Simulation Of An Off-the-Grid Radar Network With High Sensing Energy Cost. *SECON*, June 2008.
- [6] P. Dutta, et al. Trio: Enabling Sustainable And Scalable Outdoor Wireless Sensor Network Deployments. *IPSN*, April 2006.
- [7] E. Ertin, A. Arora, R. Ramnath, and M. Nesterenko. Kansei: A Testbed For Sensing At Scale. *IPSN*, April 2006.
- [8] K. Fan, Z. Zheng, and P. Sinha. Steady And Fair Rate Allocation For Rechargeable Sensors In Perpetual Sensor Networks. *SenSys*, November 2008.
- [9] X. Jiang, J. Polastre, and D. Culler. Perpetual Environmentally Powered Sensor Networks. *IPSN*, April 2005.
- [10] A. Kansal, J. Hsu, S. Zahedi, and M. Srivastava. Power Management In Energy Harvesting Sensor Networks. *Transactions on Embedded Computing Systems*, September 2007.
- [11] A. Kansal, D. Potter, and M. Srivastava. Performance Aware Tasking For Environmentally Powered Sensor Networks. *SIGMETRICS Performance Evaluation Review*, June 2004.
- [12] K. Lin, et al. Heliomote: Enabling Long-lived Sensor Networks Through Solar Energy Harvesting. *SenSys*, November 2005.
- [13] M. Minami, T. Morito, H. Morikawa, and T. Aoyama. Solar Biscuit: A Battery-less Wireless Sensor Network System For Environmental Monitoring Applications. *INSS*, June 2005.
- [14] Clemens Moser. Power Management In Energy Harvesting Embedded Systems. Ph.D. Thesis, ETH Zurich, March 2009.
- [15] C. Moser, D. Brunelli, L. Thiele, and L. Benini. Real-time Scheduling For Energy Harvesting Sensor Nodes. *RTSS*, December 2007.
- [16] D. Noh, L. Wang, Y. Yang, H. Le, and T. Abdelzaher. Minimum Variance Energy Allocation For A Solar-powered Sensor System. *DCSS*, June 2009.
- [17] The GENI Project Office. GENI System Overview. The GENI Project Office, September 2008.
- [18] N. Sharma, D. Irwin, and P. Shenoy. MultiSense: Fine-grained Multiplexing For Steerable Sensor Networks. *In Submission*, October 2009.
- [19] J. Taneja, J. Jeong, and D. Culler. Design, Modeling, And Capacity Planning For Micro-solar Power Sensor Networks. *IPSN*, April 2008.
- [20] C. Vigorito, D. Ganesan, and A. Barto. Adaptive Control For Duty-cycling In Energy Harvesting-based Wireless Sensor Networks. *SECON*, June 2007.
- [21] L. Wang, et al. AdaptSens: An Adaptive Data Collection And Storage Service For Solar-powered Sensor Networks. *RTSS*, December 2009.
- [22] G. Werner-Allen, P. Swieskowski, and M. Welsh. Motelab: A Wireless Sensor Network Testbed. *IPSN*, April 2005.
- [23] Y. Yang, L. Wang, D. Noh, H. Le, and T. Abdelzaher. SolarStore: Enhancing Data Reliability In Solar-powered Storage-centric Sensor Networks. *MobiSys*, June, 2009.