Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

SELECTION OF A PLANNING HORIZON FOR A HYBRID MICROGRID USING SIMULATED WIND FORECASTS

Mumtaz Karatas

Turkish Naval Academy
Tuzla, Istanbul, 34942, TURKEY

Emily M. Craparo
Dashi I. Singham

Naval Postgraduate School
1411 Cunningham Road
Monterey, CA, 93943 USA

ABSTRACT

Hybrid microgrids containing renewable energy sources represent a promising option for organizations wishing to reduce costs while increasing energy security and islanding time. A prime example of such an organization is the U.S. military, which often operates in isolated areas and whose reliance on a fragile commercial electric grid is seen as a security risk. However, incorporating renewable sources into a microgrid is difficult due to their typically intermittent and unpredictable nature. We use simulation techniques to investigate the performance of a hypothetical hybrid microgrid containing both wind turbines and fossil fuel based power sources. Our simulation model produces realistic weather forecast scenarios, which we use to exercise our optimization model and predict optimal grid performance. We perform a sensitivity analysis and find that for day-ahead planning, longer planning horizons are superior to shorter planning horizons, but this improvement diminishes as the length of the planning horizon increases.

1 INTRODUCTION

Significant increases in energy requirements due to the exponential growth of the world's population (Demirbas 2007) and depletion of fossil fuel resources have led to an increased interest in both energy savings and adoption of renewable energy sources (RESs) (Gadelovits et al. 2014). While both approaches address the issue of energy scarcity, RESs are favored as a means of reducing dependency on fossil fuels (Li et al. 2009). Today, RESs supply around 14% of the total world energy demand (Panwar et al. 2011, Goldemberg et al. 2000) and this proportion is expected to increase to 30-80% by 2100 (Panwar et al. 2011, Manzano-Agugliaro et al. 2013).

Among the various types of RESs, wind power is one of the fastest-growing technologies with a recent annual growth rate of 34% (Lenzen 2010). In 2012, wind power accounted for about 39% of the renewable power capacity added globally; solar and hydroelectric power each accounted for approximately 26% (REN21 Steering Committee 2013).

A hybrid microgrid supplies power from power generation devices, often utilizing renewable sources, and storage systems. A microgrid operator can realize significant fuel savings by maximizing the runtime efficiency of generators and supplementing this production with renewable power and intelligent drawdowns from a storage device. In addition to potential cost savings, a microgrid can enhance an organization's energy security and islanding time. For this reason, hybrid microgrids containing RESs are an attractive prospect for the U.S. military, which often operates in isolated areas and relies on a fragile commercial electric grid. However, the potential benefits of including RESs in a microgrid are often difficult to realize due to their intermittent nature. One remedy to the intermittence issue invovles the use of an energy storage device, such as a pumped hydroelectric or compressed air energy storage system. Such a system enables energy produced during times of abundant renewable energy availability to be stored and utilized at a later time. Another option is to connect the microgrid to a commercial grid, thereby providing an alternative energy source and also allowing any excess production to be sold to the grid.

A key component of a hybrid microgrid is the Energy Management Center (EMC). The EMC acts as a mediator between a facility's load side, generation side, and energy storage system, and it can partially overcome the issue of intermittence by properly planning grid operations. Ideally, an operating plan should anticipate any fluctuation in the renewable power output and respond accordingly. By forecasting the future wind speed with reasonable accuracy, one can predict the future wind power production and use this prediction to plan generator operations, drawdowns from a storage system, and purchases from a commercial grid.

This paper uses simulation techniques to investigate the performance of optimal day-ahead operating plans for a hypothetical hybrid microgrid containing wind turbines, dispatchable fuel-based generators, an energy storage system, and a connection to a commercial grid. To predict grid performance we utilize an optimization model that takes as input an ensemble weather forecast. An ensemble weather forecast consists of a finite sample of equally-likely weather scenarios that captures both the predicted outcome and the uncertainty inherent in this prediction. The optimization model then generates an operating plan for the microgrid that is designed to perform well in expectation across all of these weather scenarios. Because such an optimization model must, in practice, be run in an iterative manner in order to incorporate updated forecast information, we perform a sensitivity analysis to determine the effect of the planning horizon length on the microgrid's performance. The historical weather forecast data used in this study was produced by the Global Ensemble Forecasting System (GEFS) developed by the National Centers for Environmental Prediction (NCEP) (Hamill et al. 2013).

Optimization models are often tested using simulated data in order to determine their sensitivity to their inputs. Simulated data provides a means for experimenting with a model to determine which parts of the optimal solution are sensitive to the inputs, and which parts are robust to incorrect inputs. For our model, the primary input of interest is a time series of wind forecasts from an 11-member ensemble forecast. Simulating complex and dependent time series is a challenging problem. Fitting parametric models often requires large amounts of data and limiting assumptions. We rely on an algorithm that is able to generate similar time series from just one set of data while maintaining a similar dependence structure to the original series. This enables sensitivity analysis of the model's response to the input series by generating simulation replications of multidimensional time series. The model can then be tested on these time series which have a similar dependence structure, but are significantly different from the original data.

The rest of the paper is structured as follows: Section 2 reviews the simulation methodology used to generate wind forecasts. In Section 3 we develop a mathematical model of the hybrid microgrid performance. Results are provided in Section 4. Section 5 presents concluding remarks.

2 SIMULATION OF WIND FORECASTS

This section describes how we simulated wind forecasts for use in our optimization model. To test our model's robustness, we take wind forecasts generated from real weather models and generate simulated forecasts that are not identical to the real forecasts, but maintain similar patterns and dependence structures. We want to maintain the dependence within a given forecasts series generated by a model, and across the different times the forecasts were collected. This means that the simulated forecasts will have similar qualitative aspects to the real forecasts. If the forecasts on two adjacent days are similar to each other, the simulated forecasts, while different, will approximately maintain this correlation. The advantage of this algorithm is that is allows us to perform sensitivity analysis on a single set of real forecasts without making assumptions or fitting a parametric model to the data.

We have 11 forecast predictions from each day of 2012, each of which predicts the wind speed at an altitude of 80 meters (the approximate height of a wind turbine) for up to 72 hours into the future. These real forecasts can be used in the optimization model to determine the energy policy. Our intention is to use time series simulation methods to generate replicated forecasts that incorporate the statistical properties of the real forecasts. Two parameters are key to our algorithm implementation. The parameter λ is the degree of "correlation" the simulated data has with the real data. A value of $\lambda = 1$ means that the simulated data follows the the real data exactly, $\lambda = 0$ means the simulated time series has a random trajectory, and $\lambda = -1$ mean the data goes in the opposite direction of the forecasts (this could lead to negative windspeed forecasts which would not be relevant in this case). We choose a value of $\lambda = 0.5$, so that the simulated data somewhat follows the trajectory of the real data, but not too closely, as we want to test the model's sensitivity to different forecasts.

The parameter σ is the variation associated with individual simulated forecasts. If $\sigma = 0$ the algorithm would return the same simulated forecast for a given value of λ. The perturbations to the forecasts are normally distributed with variance σ^2 to allow for multiple replications to be generated. We choose a value of $\sigma = 0.5$ for our experiments. The values of $\lambda = 0.5$ and $\sigma = 0.5$ provide forecasts that have a qualitatively and quantitatively similar structure, while differing enough from the original forecasts to provide a meaningful test for model robustness. As an example of the types of forecasts that are simulated, consider the example in Figure 1 that shows a real ensemble of forecasts, along with simulated forecasts. Note that we do not want the simulated forecasts to be too close to the real forecasts or to be too close to one another.

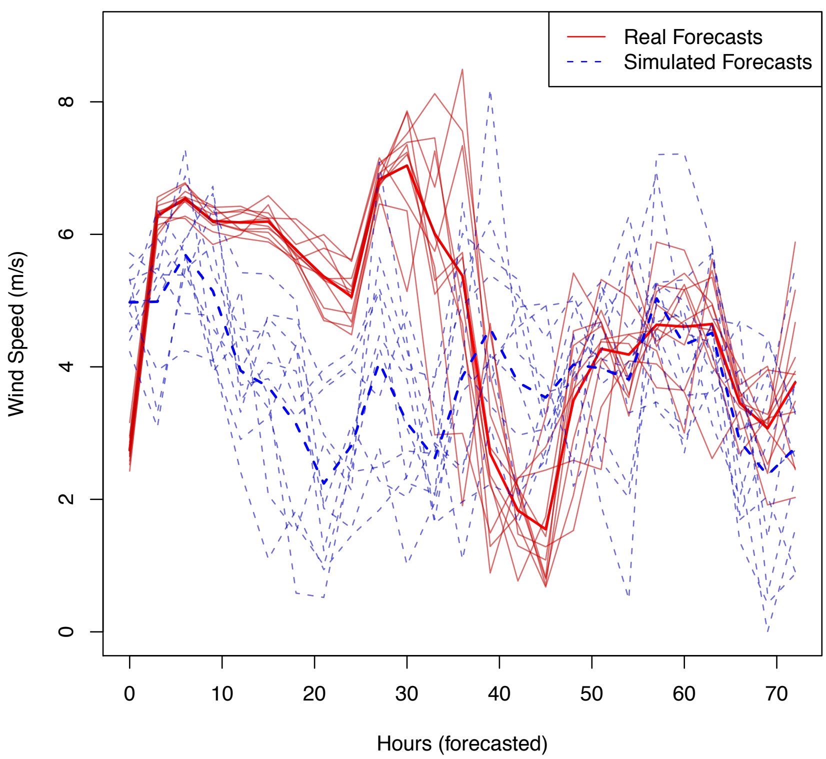

Figure 1: Sample plot of real forecast data, along with simulated forecasts. Each line corresponds to a particular member of the forecast ensemble. The darkened lines correspond to the mean forecasts across the ensemble.

We use an algorithm developed in Schruben and Singham (2014) to generate the simulated forecasts. We briefly describe the algorithm here, and refer the reader to that paper for more details. The algorithm works by mapping each dimension of the data (in this case, each day forecasts were collected) into one dimension of a hypercube. For the 2012 data, this is a 366-dimensional hyperrectangle. There are twenty five points in the time series at times $t = 0, 3, \ldots, 72$ corresponding to the forecast times at each three hour interval in the future for three days. Let this set of "real" data at a given time t be $\mathbf{x}_t$, which is a vector of length $n = 25$ representing the forecast length. We wish to generate simulated forecasts $\mathbf{y}_t$. The algorithm is recursive, so we initialize $\mathbf{y}_t$ as $\mathcal{N}(\mathbf{x}_t, \sigma^2\mathbf{I}_n)$ where $\mathbf{I}_n$ is the identity matrix with dimension n.

We define the angle between $\mathbf{y}_{t-1}$ and $\mathbf{x}_t$ as $\boldsymbol{\theta}_t$, and the distance between these points as R_t. Next, we define the angle between $\mathbf{x}_{t-1}$ and $\mathbf{x}_t$ as $\boldsymbol{\theta}'_t$, and the distance between these points as R'_t. Hence, $R_t\boldsymbol{\theta}_t = \mathbf{x}_t - \mathbf{y}_{t-1}$ and $R'_t\boldsymbol{\theta}'_t = \mathbf{x}_t - \mathbf{x}_{t-1}$. This means that $R_t\boldsymbol{\theta}_t$ is the vector between the simulated data in the previous time lag and the real data at time t, while $R'_t\boldsymbol{\theta}'_t$ is the vector between the real data at times $t-1$ and t. The algorithm generates the path of the simulated data as a linear combination of the direction of data following the same direction as the real data (from $\mathbf{x}_{t-1}$ to $\mathbf{x}_t$), and the direction of moving towards the real data ($\mathbf{y}_{t-1}$ to $\mathbf{x}_t$).

Lastly, we let $\boldsymbol{\varepsilon}_t$ be distributed as $\mathcal{N}(\mathbf{0}_n, \sigma^2\mathbf{I}_n)$ where $\mathbf{0}_n$ is a vector of zeros with length n. We then simulate windspeed forecasts using the following equation:

$$\mathbf{y}_t = \mathbf{y}_{t-1} + \lambda R_t\boldsymbol{\theta}_t + (1-\lambda)R'_t\boldsymbol{\theta}'_t + \boldsymbol{\varepsilon}_t.$$

The values of $\mathbf{y}_t$ follow the path of $\mathbf{x}_t$ with some variation in the hyperrectangle. We note that by mapping each day's forecasts into a path in a hyperrectangle, we are able to capture the dependence between these forecasts. For example, today's forecast for windspeed two days from now is likely to be correlated with tomorrow's forecast for two days from now. Additionally, there is correlation in the windspeed within a given day's forecast. Rather than constructing individual models for all the possible dependencies, we map all the days forecasts into one path that contains these dependencies, and simulate forecasts that follow the real forecasts. We note that we do not model dependence between the different ensemble members, and so each ensemble member is simulated separately.

This algorithm was used to simulate arrivals to an emergency department in Schruben and Singham (2014), for movement of soldiers in a military border crossing scenario in Singham, Thompson, and Schruben (2011), and is generally applicable to many complex multidimensional time series. Adjustments to λ and σ can be made on an ad hoc basis.

3 MICROGRID PERFORMANCE MODEL

We evaluate the performance of a hypothetical microgrid using a variation on a discrete-time stochastic mixed-integer linear optimization model first described by Bouaicha (2013). This model prescribes optimal operating schedules of fuel-based generators given an ensemble of forecasts for wind power production. The objective of the model is to minimize the expected monetary cost incurred over the planning horizon, where the expectation is taken with respect to the members of the forecast ensemble.

3.1 Microgrid Components

We consider a microgrid consisting of a number of fuel-based generators with different operating characteristics, a wind farm, an energy storage device, and a connection to a commercial power grid. The commercial grid can provide power in case of shortages and can purchase excess power produced by the microgrid.

An important limitation of fuel-based generators is their inability to rapidly transition from an idle state to an operational state. After a generator is powered on, it must undergo a "warm-up" period during which it is running and consuming fuel, but is not connected to the grid. During this time the generator reaches a safe operating temperature and its power output is stabilized. To account for this limitation, the optimization model ensures that after being powered on, each generator must run for a prescribed number of time periods before contributing to the total power output. Thus, it is important to carefully plan generator operations in order to minimize excess fuel consumption. For simplicity, we assume a constant (deterministic) hourly demand and only consider uncertainty introduced by wind power production.

3.2 Optimization Model

We now highlight a few key characteristics of the optimization model we use to evaluate our hypothetical microgrid. Full details of this model, including a mathematical formulation, are available in Bouaicha (2013). To account for the uncertainty inherent in planning based on weather forecasts, the model utilizes an ensemble weather prediction representing a finite sample of the uncountably infinite set of possible weather outcomes. The model then prescribes an optimal operating schedule for the components of the microgrid, including:

- when to turn each generator off or on, and at what speed to operate it,
- when to schedule a charge or discharge of the energy storage device, and at what rate,
- when and how much energy to buy from the commercial grid, and
- when and how much energy to sell to the commercial grid.

Decisions regarding generators are constant across all weather forecast scenarios; all other decisions are allowed to vary by scenario in order to reflect short-term operational decisions that can be made "on the fly." Our implementation differs slightly from that in Bouaicha (2013) in that we allow the energy purchased from or sold to the commercial grid to vary by scenario, whereas Bouaicha ensures that the amount of energy purchased or sold in each time step is constant across all scenarios.

The model's objective is to minimize the expected total cost incurred over the planning horizon. Costs are incurred from fuel consumption during generator warm-up and production periods, usage of the energy storage device, and purchases from the commercial power grid. Additionally, revenue may be obtained if excess power is sold to the commercial power grid.

The optimization model's constraints are designed to model the physical functionality of the microgrid, as well as various operational limitations. Constraints fall into three main categories reflecting limitations on power production, generator operations, and energy storage.

3.2.1 Power production and demand satisfaction

The model ensures that during each time step, the total energy produced is at least as great as the total load. Energy is produced by generators, wind turbines, purchases from the commercial grid, and discharge of the energy storage device; it is consumed by demand satisfaction, sales to the commercial grid, and charging of the energy storage device.

3.2.2 Generator operation

As described in Section 3.1, fuel-based generators must undergo a warm-up period following activation before they can contribute to the power grid. The optimization model ensures that an appropriate warm-up period is enforced each time a generator is turned on. Additionally, the model enforces minimum and maximum operating speeds for each generator, and it limits the total number of changes to each generator's operating speed over the planning horizon.

3.2.3 Energy storage

The model calculates the total energy stored in the energy storage device at each time step and ensures that this quantity is nonnegative and no greater than the maximum capacity of the energy storage device. Additionally, it enforces a minimum and maximum charge rate each time the energy storage device is charged or discharged.

3.3 Rolling Horizon Optimization

Weather forecasts are provided for a short time frame, typically 24 to 120 hours, and their accuracy is not constant over this time frame. In particular, the uncertainty associated with a forecast grows with the lead time. Figure 3.3 shows an example 96-hour forecast with 11 ensemble members. Note that although the wind speed predicted by the 11 members is fairly consistent for short lead times, the ensemble members begin to diverge at a lead time of around 60 hours and are substantially different by 96 hours for this particular example. Although the exact time at which ensemble members begin to significantly diverge varies from forecast to forecast, this qualitative behavior is consistent in weather forecasting.

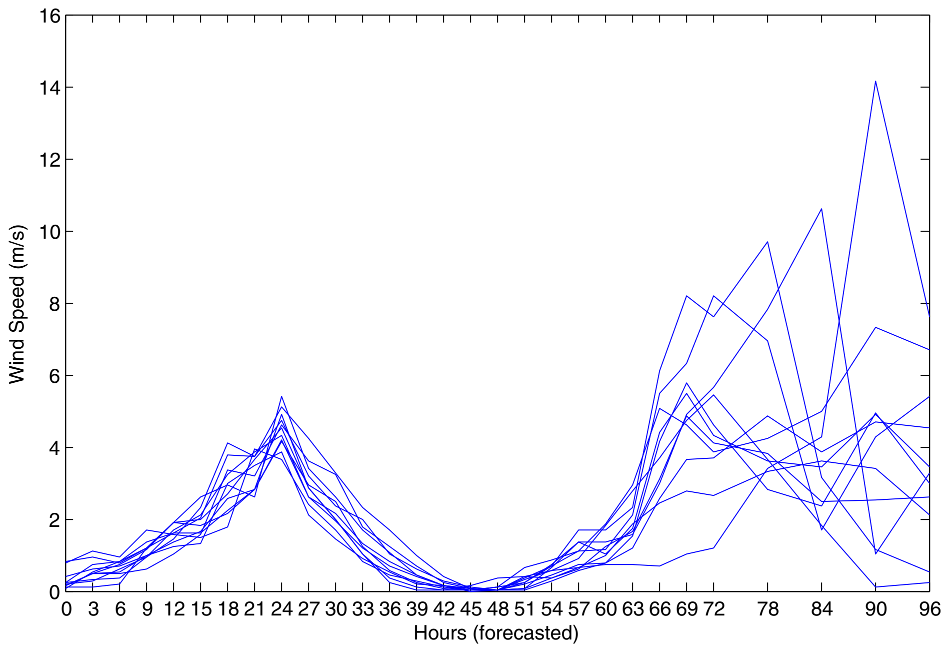

Figure 2: Example wind speed predictions from an 11-member ensemble forecast. Note that although the ensemble members are largely consistent in their predictions for short lead times, they differ substantially at long lead times.

The difficulty of making accurate weather predictions far in advance necessitates that any planning model that relies on these predictions be run in an iterative fashion. In rolling horizon optimization (sometimes called receding horizon optimization), an optimization model is solved over a time horizon known as the *planning horizon*. The resulting plan is then executed over a shorter time frame known as the *execution horizon*. After the execution horizon has elapsed, the model is solved again over the planning horizon using updated input data, and the process repeats.

Rolling horizon optimization has a number of advantages. A key advantage in our application is that rolling horizon optimization does not rely on perfect information about the future, but instead incorporates improved information as it becomes available. Moreover, even if perfect information is available, rolling horizon optimization is less computationally costly than optimizing over the full time frame of interest and is sometimes used for this reason alone.

Two parameters dictate how a sequence of rolling horizon optimization runs proceeds: the length of the planning horizon and the length of the execution horizon. When planning microgrid operations, it is advantageous to utilize updated weather forecasts as they become available. Thus, a natural choice for the execution horizon is time that passes between forecast updates. In this paper we utilize forecasts that are generated every 24 hours, so our execution horizon is 24 hours. Choosing a planning horizon, however, is less straightforward. Long planning horizons are favored in many applications because they reduce the likelihood that highly suboptimal "myopic" decisions will be executed. However, because the quality of weather forecasts deteriorates substantially with the lead time, there is reason to believe that a very long planning horizon may not be ideal for planning microgrid operations. In particular, with a long planning horizon, one runs the risk of subordinating near-term decisions, which will be executed during periods of low uncertainty, to account for possibilities expressed in highly uncertain data pertaining to the future. Thus, we perform a computational study to determine an appropriate planning horizon for planning microgrid operations based on weather forecast data.

4 RESULTS

We now study the impact of the planning horizon length on the quality of the schedules produced for the hybrid microgrid described in Section 3. We implement the optimization model in the General Algebraic Modeling System (GAMS) environment and solve it using CPLEX 12.2.0.2. Using the simulation method described in Section 2, we generate thirty 72-hour ensemble weather forecasts for each day in a 60-day period. We then solve our optimization model in a rolling horizon manner using an execution horizon of 24 hours and planning horizons of 24, 27, 30, 33, 36, 48, 60 and 72 hours. We evaluate the performance of the plans resulting from each planning horizon in terms of actual cost, i.e., the cost that would be incurred if the microgrid were operated in accordance with the resulting plans and under observed weather conditions. In the absence of observed wind speeds at an 80-m altitude in the historical database (and thus in our simulated forecasts), we estimate actual costs using a "most accurate" member of each forecast ensemble. We designate the ensemble member whose prediction at a 24-hour lead time most closely matches the ensemble mean in the first time step of the next forecast as the "most accurate" and treat its prediction as ground truth. Because the initial time step in a weather forecast represents the forecasting model's best estimate of the current conditions based on assimilated data from all available observations, this is a reasonable proxy for observed data.

In addition to comparing the actual costs resulting from various planning horizons, we also compute a "best-case" plan and its resulting cost for each simulated forecast set. The best-case cost is the lowest total cost that could be achieved if the wind speed were known in advance with perfect accuracy for the entire 60-day period of interest and the optimization model were run with a 60-day planning horizon. Although one could not hope to achieve the best-case cost in practice, it provides a benchmark by which to compare the costs resulting from the various planning horizons.

Figure 3 gives an example of power production (left) and load (right) schedules resulting from a best-case plan (top) and a plan generated with a 24-hour planning horizon (bottom). In this instance, demand is constant at 2000 kW, and the purchase and selling prices of energy are also constant. Production costs for the three generators vary, but are always less than the cost of purchasing from the commercial grid (excluding the generator's warm-up period). Note that in the best-case production schedule, Generator 1 is shut down at around time step 55 due to the large influx in wind power that is upcoming. In the actual production schedule, Generator 1 continues to run during this time. Conversely, the actual schedule discountinues usage of Generator 3 around time step 85, resulting in a significant amount of energy purchased from the grid. The best-case schedule, on the other hand, continues to run Generator 3 during this time period. Overall, the best-case schedule makes heavier use of the storage device and purchases less energy from the commercial grid.

To investigate the impact of the planning horizon on the total cost incurred, we examine three different microgrid configurations. These configurations are summarized in Table 1. Configuration 1 represents a baseline configuration. In Configuration 2, the purchase and selling prices for the commercial grid are modified to make the commercial grid a more unattractive option. In Configuration 3, both demand and the size of the wind farm are doubled in order to reflect a larger installation with a more significant renewable component. Generator attributes are constant across all configurations: Generator 1 produces energy at a cost of \$0.10 per kWh (after the warm-up period), Generator 2 produces energy at a cost of \$0.09 per kWh,

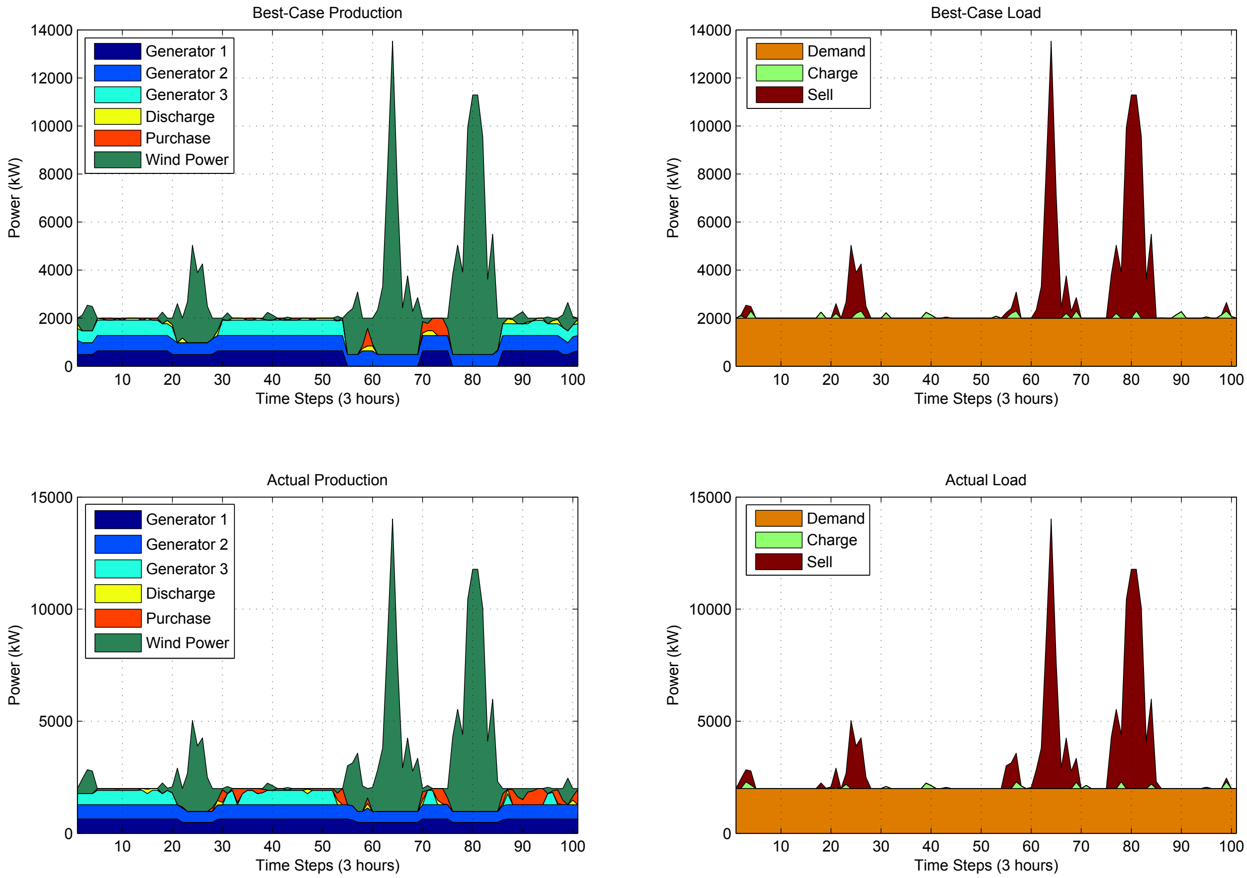

Figure 3: An example of the optimal power production (left) and load (right) schedules resulting from a best-case plan (top) and a plan generated with a 24-hour planning horizon (bottom). Power production sources include three generators, discharge of the storage device, purchases from the commercial grid, and wind power. Power consumers include demand, charge of the storage device, and sales to the commercial grid.

and Generator 3 produces energy at a cost of \$0.14 per kWh. Each generator has a minimum production capacity of 490 kW while running, and a maximum production capacity of 640 kW.

Figure 4 summarizes the results of our computational experiments, where the optimization model was terminated when the objective value was proven to be within 1.5% of the optimal objective value. The left side of Figure 4 displays cost data. Blue dots represent actual values for the individual simulated scenarios for each planning horizon, while red dotted lines represent best-case values for each simulated scenario. Solid lines represent average quantities over the 30 simulated scenarios. Note that although a planning horizon of 24 hours results in highly suboptimal performance, increasing the planning horizon to 33 hours yields a large improvement in solution quality for all configurations considered. Increasing the planning horizon beyond 33 hours does not yield substantial additional benefit for the architectures considered.

The right side of Figure 4 displays data concerning purchases from the commercial grid. Although purchases from the commercial grid are not explicitly minimized in the optimization model, the costs are structured so as to make the commercial grid an unattractive option, and thus large purchases from the grid reflect inefficiencies in the planning process. Similar to the cost results, increasing the planning horizon results in a decrease in the amount of energy purchased from the commercial grid up to a planning horizon of 33 hours, at which point the marginal benefit of increasing the planning horizon is negligible.

Table 1: Microgrid Configurations.

	Configuration 1	Configuration 2	Configuration 3
Demand (kW)	1000	1000	2000
Purchase cost ($ per kWh)	0.12	0.18	0.18
Selling price ($ per kWh)	0.08	0.06	0.06
Size of wind farm	baseline	baseline	2x baseline

A comparison of the results for Configurations 1 and 2 reveals that when the costs pertaining to the commercial grid are made less attractive, the total cost incurred increases, as would be expected. However, this effect is substantially more pronounced for a 24-hour planning horizon than for any other planning horizon. With a sufficiently long planning horizon, the optimal plan avoids making purchases from the commercial grid and shifts to other means of production; this shift is clearly reflected in the right side of Figure 4. (Note that the exponent differs on the vertical axes for Configuration 1 and Configuration 2.)

Comparing the results for Configurations 2 and 3, we see that doubling both the demand and the wind production capacity results in a substantial increase in cost; in fact, the cost incurred more than doubles. This is due to the fact that although the wind production capacity has doubled, the generator production capacity remains constant across all configurations. Thus, in Configuration 3 the generators are able to satisfy a smaller fraction of the residual demand not satisfied by wind power. The only alternative is to purchase a larger fraction of energy from the commercial grid; note that this quantity approximately triples from Configuration 2 to Configuration 3.

5 CONCLUSIONS AND FUTURE WORK

This paper has described the use of simulation and optimization methods evaluate a hybrid microgrid containing both wind turbines and fuel-based generators. After creating realistic weather forecast scenarios using time series simulation methods, we use a rolling horizon optimization technique to create realistic grid operation schedules. We perform a sensitivity analysis to determine the impact of the planning horizon length on solution quality. Our experimental results indicate that although longer planning horizons are superior to shorter planning horizons, the marginal benefit of increasing the planning horizon decreases substantially as the planning horizon increases.

Note that we have considered only the uncertainty resulting from imperfect knowledge of future wind power production. Another important source of uncertainty is the demand, which we have assumed to be constant. In reality, demand is also uncertain and is often correlated with weather. Future work will extend our study to incorporate demand uncertainty. Such a study might involve using techniques similar to those described in this paper to produce demand forecasts, or it may involve a more detailed agent based simulation model to simulate consumers in the network.

Future work will also investigate the impact of energy storage capacity on microgrid robustness and efficiency and may provide operators with real-time guidance and control policies for microgrid operation.

ACKNOWLEDGMENTS

The authors thank Dr. Jonathan Moskaitis of the Naval Research Laboratory in Monterey, California, for his assistance in obtaining and interpreting the historical weather forecast data.

REFERENCES

Bouaicha, H. 2013. “Optimal Day-Ahead Scheduling of a Hybrid Electric Grid Using Weather Forecasts”. Master’s thesis, Monterey, California: Naval Postgraduate School.

Demirbas, A. 2007. “Energy Issues and Energy Priorities”. *Energy Sources, Part B: Economics, Planning, and Policy* 3 (1): 41–49.

Gadelovits, S., A. Kuperman, M. Sitbon, I. Aharon, and S. Singer. 2014. "Interfacing Renewable Energy Sources for Maximum Power Transfer—Part I: Statics". *Renewable and Sustainable Energy Reviews* 31:501–508.

Goldemberg, J. et al. 2000. *World Energy Assessment*. United Nations Development Programme.

Hamill, T. M., G. T. Bates, J. S. Whitaker, D. R. Murray, M. Fiorino, T. J. Galarneau, Y. Zhu, and W. Lapenta. 2013. "NOAA's Second-Generation Global Medium-Range Ensemble Reforecast Data Set". *Bulletin of the American Meteorological Society* 94:1553–1565.

Lenzen, M. 2010. "Current State of Development of Electricity-Generating Technologies: A Literature Review". *Energies* 3 (3): 462–591.

Li, H., H. C. Jenkins-Smith, C. L. Silva, R. P. Berrens, and K. G. Herron. 2009. "Public Support for Reducing US Reliance on Fossil Fuels: Investigating Household Willingness-to-Pay for Energy Research and Development". *Ecological Economics* 68 (3): 731–742.

Manzano-Agugliaro, F., A. Alcayde, F. Montoya, A. Zapata-Sierra, and C. Gil. 2013. "Scientific Production of Renewable Energies Worldwide: An Overview". *Renewable and Sustainable Energy Reviews* 18:134–143.

Panwar, N., S. Kaushik, and S. Kothari. 2011. "Role of Renewable Energy Sources in Environmental Protection: A Review". *Renewable and Sustainable Energy Reviews* 15 (3): 1513–1524.

REN21 Steering Committee 2013. *Renewables 2013: Global Status Report*. Paris: REN21 Secretariat.

Schruben, L., and D. Singham. 2014. "Data-driven Simulation of Complex Multidimensional Time Series". *ACM Transactions on Modeling and Computer Simulation* 24 (1): 5.

Singham, D., M. Thompson, and L. Schruben. 2011. "Applications of Flocking Algorithms To Input Modeling for Agent Movement". In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, 2443–2449. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

MUMTAZ KARATAS graduated from the Turkish Naval Academy in 2001. He received his MS degree in Industrial & Operations Engineering from University of Michigan and his Ph.D. degree in Industrial Engineering from Kocaeli University. He spent two years at the Naval Postgraduate School as a visiting researcher and postdoctoral fellow between 2011 and 2013. He is currently an Assistant Professor in the Industrial Engineering Department at the Turkish Naval Academy. His current research areas include combinatorial optimization and military operations research. His email address is mkaratas@dho.edu.tr.

EMILY M. CRAPARO is an Assistant Professor of Operations Research at the Naval Postgraduate School. She obtained her Ph.D., S.M., and S.B. in Aeronautics & Astronautics at the Massachusetts Institute of Technology. Prior to joining the Operations Research faculty at the Naval Postgraduate School, she completed a National Research Council Postdoctoral Fellowship in the same department. Her research areas include mathematical modeling, optimization, and military applications. Her email address is emcrapar@nps.edu.

DASHI I. SINGHAM is a Research Assistant Professor of Operations Research at the Naval Postgraduate School. She obtained her Ph.D. in Industrial Engineering & Operations Research at the University of California, Berkeley in 2010, and an M.A. in Statistics from Berkeley. Her B.S.E. is from Princeton University in Operations Research & Financial Engineering. Her research areas include simulation modeling and analysis and applied statistics. She is an Associate Editor at *IIE Transactions*, and a reviewer for *Operations Research, Informs Journal on Computing, ACM Transactions on Modeling and Simulation* and *Naval Research Logistics.* Her email address is dsingham@nps.edu.

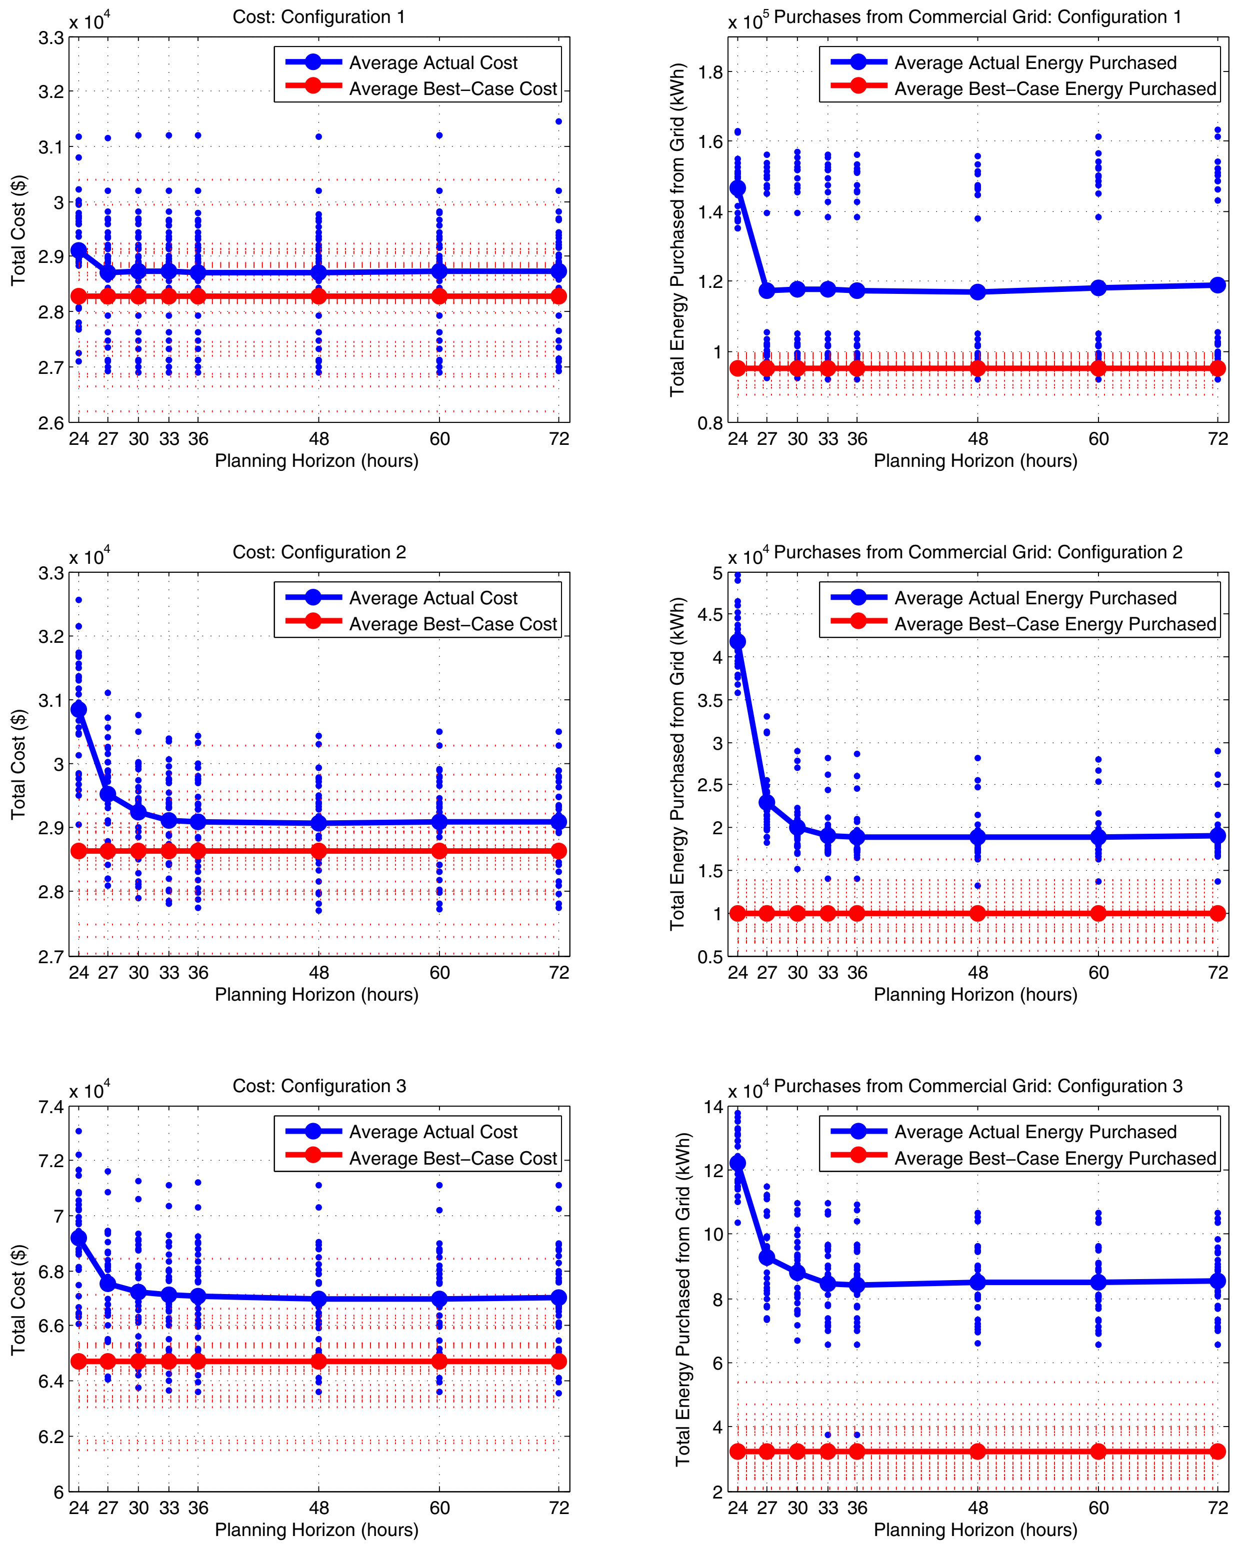

Figure 4: Actual and best-case costs (left) and energy purchases from the commercial grid (right) for various planning horizons. Blue dots represent actual values for individual simulated scenarios, while red dotted lines represent best-case values. Solid lines represent average quantities over the 30 simulated scenarios. All optimization runs were terminated when the objective value was proven to be within 1.5% of the optimal objective value.