Proceedings of the ASME 2015 International Design Engineering Technical Conferences &
Computers and Information in Engineering Conference
IDETC/CIE 2015
August 2-5, Boston, Massachusetts, USA

SAND2015-0353C

DETC2015-46107

***DRAFT***

# AUTONOMOUS MICROGRID DESIGN USING CLASSIFIER-GUIDED SAMPLING

**Peter B. Backlund**
Senior Member of Technical Staff
Sandia National Laboratories
Albuquerque, NM, USA

**John P. Eddy**
Principal Member of Technical Staff
Sandia National Laboratories
Albuquerque, NM, USA

## ABSTRACT

Identifying high-performance, system-level microgrid designs is a significant challenge due to the overwhelming array of possible configurations. Uncertainty relating to loads, utility outages, renewable generation, and fossil generator reliability further complicates this design problem. In this paper, the performance of a candidate microgrid design is assessed by running a discrete event simulation that includes extended, unplanned utility outages during which microgrid performance statistics are computed. Uncertainty is addressed by simulating long operating times and computing average performance over many stochastic outage scenarios. Classifier-guided sampling, a Bayesian classifier-based optimization algorithm for computationally expensive design problems, is used to search and identify configurations that result in reduced average load not served while not exceeding a predetermined microgrid construction cost. The city of Hoboken, NJ, which sustained a severe outage following Hurricane Sandy in October, 2012, is used as an example of a location in which a well-designed microgrid could be of great benefit during an extended, unplanned utility outage. The optimization results illuminate design trends and provide insights into the traits of high-performance configurations.

## 1. INTRODUCTION

Technological advancements, environmental pressures, and economic incentives are causing a shift from large, centralized power generation plants to smaller, distributed energy resources (DERs) [1]. However, haphazard placement of grid-integrated DERs, such as fuel-based generators, renewable sources, and storage components, creates many problems regarding maintenance, safety, and central dispatch control. Therefore, it may be beneficial to organize a collection of DERs into a microgrid (MG), thus enabling them operate as a single self-controlled entity [2]. An MG is defined as a collection of interconnected DERs and loads that may operate in one of two modes: grid-tied mode or islanded (autonomous) mode [3]. In grid-tied mode, MGs can provide congestion relief, postponement of increased generation and transmission capacity, and fast response to load changes [2].

In autonomous mode, the islanded MG can continue to supply power during planned or unplanned grid outages due to scheduled maintenance or faults, respectively [4]. Planned outages can be scheduled to occur during "off-peak" hours when loads are expected to be relatively small. However, unplanned outages can, by definition, occur at any time, and the cost of installing the capacity to fully accommodate worst-case peak-loads may be prohibitive. The impact on consumers becomes very significant when the outage lasts for an extended period of time, such as one that might follow a natural disaster or severe weather event.

The focus of this work is on identification of system-level MG designs that perform well during unplanned, extended periods of disconnect from the main utility grid. The variables that comprise the design problem are the sizes and locations of DERs and the topology of the microgrid. Designs are sought that reduce average loads not served (LNS) in islanded mode and do not exceed a pre-defined installation cost.

There are many sources of uncertainty in microgrid design and operation. For example, the magnitudes of loads, durations of outages, and reliability of generators are all uncertain. In addition, no generation or distribution components in the system are 100% reliable. For example, natural gas generators may fail during operation, or overhead lines may fault due to falling tree branches. Lastly, outages can last anywhere from a few seconds to weeks (in rare cases), making identification of a design for the most likely scenario difficult. In this work, these uncertainties are addressed by running a discrete event




Sandia National Laboratories is a multi-mission laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000

simulation of MG operation over a very long time horizon. The frequencies and durations of all failures, including generation components, transmission components, and the main grid, are generated by sampling from representative probability distributions. Using this evaluation model, a large number of outages are simulated and average performance of a given MG design can be computed and used as an objective function in an optimization algorithm.

A drawback of this approach is that each evaluation of the objective function requires many simulations of a grid outage. This computational expense makes simulation-based optimization and design space exploration difficult, because common metaheuristic optimization methods such as genetic algorithms may require a large number of function evaluations to identify high-performance solutions. Furthermore, the simulation is black-box in that it cannot be readily modeled using analytical equations. It is not known if the objective function is multi-modal or convex. This property rules out the possibility of using traditional optimization techniques.

Metamodels, also known as surrogate models, can be used in place of computationally expensive simulations to increase computational efficiency in support of engineering design optimization [5]. Traditionally, metamodels are developed by fitting a surface to a set of training points that are generated from an expensive base model. However, this approach relies on the assumption of continuous design variables and a smooth objective to approximate [6]. Surrogate modeling techniques for mixed discrete/continuous functions remains a challenge, and existing methods have only been shown to be viable for problems with small numbers of categorical variables that have few discrete choices [7].

To address this challenge, classifier-guided sampling (CGS) [8, 9], an optimization technique appropriate for computationally expensive problems with discrete variables and discontinuous responses, is used to identify MG configurations that reduce average LNS while not exceeding a specified cost constraint. The classifier is a probabilistic model that serves as an inexpensive approximation of the expensive simulation. This method is similar to direct-search metamodel-based design optimization techniques [5], but is better suited to solve problems that have combinatorial variables and discontinuous objective functions (such as the MG design problem studied here).

As a test case, the city of Hoboken, NJ, which experienced a week-long outage in the aftermath of Hurricane Sandy in October 2012, is modeled and simulated. The optimization results enable the identification of design trends among top-performing solutions and generate useful insights into the types of MG configurations that reduce average LNS during extended grid outages.

In the following section, the MG simulation tool that is used to assess MG performance is described. In Section 3, the CGS algorithm is discussed in detail. In Section 4, the microgrid simulation model and CGS are used together to identify high-performance MG designs for Hoboken, NJ that reduce average LNS during an extended main grid outage. Concluding remarks are provided in Section 5.

## 2. MICROGRID SIMULATION MODEL

The model that is used to simulate grid outages and assess the performance of candidate MG designs is referred to here as the Performance and Reliability Model (PRM). The PRM is a simulation code written in C++ that is used to statistically quantify the performance and reliability of an MG operating in autonomous (islanded) mode. The PRM allows the performance of an MG to be quantified in terms of fuel usage, renewables penetration, renewables spillage, and other operational characteristics. MG reliability can be quantified in terms of frequency and magnitude of load lost on a tier-by-tier, load-by-load, or bus-by-bus basis, or on an aggregated basis over the MG. The PRM also supports calculation and reporting of individual equipment reliability statistics. The PRM simulation relies on a representation of an unreliable power utility, specifically through the definition of failure and repair modes. The PRM typically simulates thousands of such utility outages to ensure that the calculated statistics are stable.

The PRM models systems behavior as a discrete sequence of events in time. At each event, appropriate logic is executed which may result in the scheduling of more events. The simulation proceeds until all events have been executed or until some explicit stopping criterion is satisfied. This type of simulation is minimalistic in the sense that logic is only executed in response to a known event. This approach contrasts to some other simulation methods whereby regular time steps are taken and at each, tests are run to see if anything should be done and only then is appropriate logic executed.

The primary input to the PRM is an MG topology specification. A topology specification includes descriptions of components such as electrical lines, busses, switches, transformers, generation assets and fuel sources, batteries, and inverters, in addition to the details of how these components are interconnected. Most components in the PRM can be modeled as potentially unreliable. To specify an unreliable component, failure modes must be defined where each mode represents a specific type or mode of component failure. Each failure mode specifies a failure time generator and a repair time generator. In this context, a "generator" refers to a sampling scheme, typically associated with a parametric statistical distribution.

In addition to the physical layout and reliability characteristics of an MG, the PRM requires three configured controllers to define how the grid operates during different phases. A controller implements and executes the logic applied during the three phases of operation: grid-tied operation, MG startup operation, and autonomous MG operation.

Grid-tied operation includes any period of normal operations occurring in the time leading up to the first simulated utility outage and the time period between utility outages. A utility outage is a complete failure of the bulk power grid leaving the microgrid to operate in autonomous mode. Microgrid startup operation occurs in the period immediately following the onset of a utility outage. This controller implements the logic that manages the formation of the microgrid (starting of required MG assets, connecting busses together, connecting or reconnecting renewables, etc.). Microgrid operation is the period immediately following the microgrid startup phase. This controller handles operation of the fully formed autonomous MG.

The process of handling a utility outage begins with operating the Grid-Tied Controller (GTC) for the time leading up to the start of the utility outage. The GTC performs a relatively small amount of work, as its only job is to manage the failures and repairs of those pieces of equipment that function during normal operations. This includes components such as lines, transformers, and renewable generators, but excludes components such as backup diesel generators as these do not operate during normal operations. The GTC serves to provide a reasonable estimate of the state of these pieces of equipment at the onset of the next utility outage.

At the beginning of the next sampled utility outage, the startup controller (SC) is entered. The SC does no generator dispatch beyond the simplistic strategy of starting them all immediately. The SC also manages the interconnection of all busses and assets. This is necessary because the MG is not technically formed at the onset of the outage. Generators may or may not start according to their startup probabilities, and restart attempts may be made if configured to do so. Each generator may have a delay associated with it to represent the spin-up time that must pass before it can be connected to its bus. After spin-up, a generator can be immediately connected to its bus if the bus is not yet energized. If the bus is energized (i.e., another generator is already connected to it), then a synchronization delay is employed.

Once startup is complete, control passes to the sustained microgrid controller (MC). The MC reacts to all events during autonomous operation and rebalances the grid at each one. Events include failures and repairs, load changes, renewable output changes, fuel outages, refueling events, and so on. Most importantly, the MC must determine what generation assets should be operational, and then it must make use of these operational assets to serve loads.

The logic that determines what generation assets should be operating is handled by a dispatcher. The dispatcher is executed when the current grid load conditions are such that it is below a user-defined dispatch threshold or above the reserve power requirement. Once the operational assets are known, the MC determines how much power is coming from each; the running fossil generators are all assumed to run in a droop control mode and thus will run at the same utilization rate.

The first action the dispatcher takes when determining which generation assets should be operational is to account all currently running dispatchable generation units. A dispatchable unit in this context is a fossil unit such as a diesel or natural gas generator. Renewables are not considered dispatchable by the PRM. The dispatcher minimizes the changes from the currently running set in an attempt to prevent frequent start and stop actions on generators. The controller must account for currently producing generators as well as those that are in their startup phase but have not yet connected or synchronized on the grid. Any predicted power deficit (PPD) or excess (PPE) must account for these generators as though they were producing.

Using this information, the dispatcher attempts to determine whether or not it would like to start generators, stop generators, or leave things the way they are. The generator dispatch rules are summarized in Table 1. Once it has made this determination, it attempts to execute. In the case where it decides to leave things the way they are, it exits without action. If it believes that generators should be started, it attempts to do so according to the logic that is summarized in Table 2. If the dispatcher has determined that some generators should be stopped, then it enters the generator stop logic, which is summarized in Table 3.

**Table 1: Generator dispatch rules**

| |
|---|
| Generator START attempts are made if ANY of the following are true:<br>• There are fewer that the than the minimum allowable number running<br>• The maximum predicted load is greater than the maximum possible generation of the currently running generators<br>• The maximum predicted load is greater than the reserve power requirement and the average predicted load is closer to the maximum than the minimum |
| Generator STOP attempts are made if NONE of the start clauses are true and ANY of the following are true:<br>• The minimum predicted load is less than the minimum desirable generation of the currently running set. (Each generator has a minimum desirable running rate. In the case of a diesel generator, this is usually a rate below which wet stacking begins to occur and efficiency drops sharply)<br>• The average predicted load is closer to the minimum than the maximum |
| Lastly, no START or STOP attempts are made if ANY of the following are true:<br>• The dispatcher would like to stop some but to do so would violate the minimum allowable number of running generators<br>• The maximum predicted load is greater than the reserve power requirement but the average load is closer to the minimum than the maximum |

**Table 2: Generator START logic**

| |
|---|
| The START logic begins by gathering a list of all "startable" generators. A startable generator is one that:<br>• Is not running<br>• Is functional<br>• Has at least enough fuel to run for the minimum runtime specified as a parameter to the dispatcher<br>• Is microgrid enabled if dealing with a non-isolated bus |
| Once the list is determined, it is sorted by appropriateness of the generators to meet the load. A generator (A) is favored over another generator (B) if:<br>• It is large enough to cover the PPD and (B) is not<br>• Both are large enough to cover the PPD and (A) is smaller than (B)<br>• Both are too small to cover the PPD and (A) is larger than (B) |
| Once the list is sorted, the winning generator is at the top. It is scheduled for start. The PPD is reduced by that generator's recommended maximum running rate and the remaining list is resorted using the new PPD. This process is repeated until the required minimum running generators are running and the PPD is covered or there are no more generators that can be started. |

**Table 3: Generator STOP logic**

| |
|---|
| The STOP logic begins by gathering a list of all "stoppable" generators. A stoppable generator is one that:<br>• Is running<br>• Is not so large that to shut it down would result in shedding more generation than desired |
| Once this list is determined, it is sorted by appropriateness of each generator. A generator (A) is favored for shutdown over another generator (B) if:<br>• It is larger (remember that both are small enough that stopping them will not result in a PPD) |
| Once the list is sorted, the winning generator is at the top. It is scheduled for shutdown. The PPE is reduced by that generators recommended maximum running rate and the remaining list is resorted using the new PPE. This process is repeated until the required minimum number of running generators has been reached, the PPE is gone or there are no more generators that can be shut down without resulting in a PPD. |

If there is excess power and the dispatcher has opted to keep it online, then the excess power may be used to serve previously dropped loads or to charge UPSs and batteries. If after all sources are accounted for in full there is not enough generation to meet load, then loads are dropped if possible. A load can only be dropped if there is a switch in the path between it and the bus. The logic for dropping loads is similar to the logic for starting and stopping generators in that a list of droppable loads is created and sorted by priority and size. Loads are dropped one at a time until enough has been shed.

The MC's job is complete when the utility outage ends. There is a small amount of logic that must be carried out for reconnection to the main grid. All microgrid paths that were closed during startup must be opened and all backup generators must be shut down. These tasks are all performed instantaneously (i.e., no simulation time passes). Once reconnection is complete, control is passed back to the GTC until the next utility outage occurs.

The primary outputs of the PRM are computed quantities and statistics. All unreliable entities gather statistics detailing the outages they suffer, their durations, and the amount of downtime attributable to each failure mode. All levels of the grid for which it makes sense to express load service statistics track load served and not served.

All generator types have statistics computed for them in addition to unreliability statistics. They include run time statistics, utilization rate statistics, energy production statistics, statistics describing the time for which the generator was not in use, and statistics for efficiency.

Any of the quantities and statistics that are computed by the PRM may be used for optimization purposes. In the context of emergency autonomous operation, which is the focus of this paper, the impact of MG topology (generators, lines, busses, and how these components are interconnected) on average LNS is the quantity of interest.

## 3. CLASSIFIER-GUIDED SAMPLING OPTIMIZATION

Classifier-guided sampling (CGS) [8, 9] is a stochastic optimization technique suitable for discrete optimization problems with computationally expensive objective functions. CGS achieves efficient global optimization by using a Bayesian classifier to provide categorical predictions of the performance of candidate solutions prior to expensive evaluation. The classifier enables CGS to focus on promising solutions and avoid wasteful evaluations on poor-performing ones. A brief overview of Bayesian classifiers is provided next, followed by a detailed explanation of the CGS algorithm.

### Bayesian Classifiers

In machine learning, a classifier is used to predict categorical class labels to test points that have known feature attributes but unknown class labels [10]. A classifier is trained using a set of feature vector / class label pairs that are generally obtained experimentally. In the context of design optimization, a set of design point / objective function value pairs can be used as training points.

A Bayesian classifier uses a factorization of probability distributions to predict the categorical performance of a candidate configuration based on all previously evaluated points. Consider a $K$ category classifier. If $c$ is the discrete class variable, let $c_k$ be the class $k$ (i.e., a specific instance of $c$) for $k \in \{1, 2, ..., K\}$. The classification is performed in a $D$-dimensional design space, and $\mathbf{x} = [x_1, x_2,\ldots, x_D]$ is a vector of design variables. If $\hat{\mathbf{x}}$ is a specific design instance of $\mathbf{x}$, Bayes' formula can be used to estimate the posterior probability of the class $c_k$ given $\hat{\mathbf{x}}$, $P(c_k| \hat{\mathbf{x}})$, according to:

$$P(c_k \mid \hat{\mathbf{x}}) = \frac{P(\hat{\mathbf{x}} \mid c_k) P(c_k)}{P(\hat{\mathbf{x}})} = \frac{P(\hat{\mathbf{x}} \mid c_k) P(c_k)}{\sum_{k=1}^{K} P(\hat{\mathbf{x}} \mid c_k) P(c_k)} \tag{1}$$

where $P(c_k)$ is the prior probability of any randomly selected point belonging to class $c_k$, and $P(\mathbf{x}|c_k)$ is the class conditional probability of a design instance given the class label. Design $\hat{\mathbf{x}}$ is classified as a member of class $c_k$ that has the highest $P(c_k| \hat{\mathbf{x}})$ when compared to all other classes.

In this work, the prior probabilities, $P(c_k)$, are set according to a constant discrete uniform distribution such that:

$$P(c_k) = \frac{1}{K}, \quad \forall k \in \{1, 2, ..., K\} \tag{2}$$

where $K$ is the number of performance categories.

$P(\mathbf{x}|c)$ is a $D$-dimensional joint distribution that must be estimated from a training set of design vector / class label pairs. Due to the large number of training points required to achieve an accurate predictor, it is advantageous to make conditional independence assumptions about the design variables and refactor $P(\mathbf{x}|c)$ into a product of univariate distributions. Specifically, by assuming that all design variables are independent of each other, $P(\mathbf{x}|c)$ reduces to:

$$P(\mathbf{x} \mid c) = P(x_1 \mid c) P(x_2 \mid c) \ldots P(x_D \mid c) \tag{3}$$

Using the factorization in Eq. (3) in Eq. (1) is a special case of Bayesian classifier known as the naïve Bayes classifier [11]. The task of estimating the $D$ distribution parameters is achieved by assigning class labels to all design points that have been evaluated with the expensive simulation. For example, assume a design variable $x$ has domain {0, 1} and we want to estimate $P(x|c_{\text{good}})$ where $c_{\text{good}}$ is a class label assigned to 'good' design points. This distribution can be modeled with a probability mass function with two parameters: $\theta_{x=0,c=\text{good}}$ and $\theta_{x=1,c=good}$. The distribution parameter $\theta_{x=0,c=\text{good}}$ is estimated according to:

$$\theta_{x=0,c=good} = \frac{\beta + \#(x=0, c=good)}{\alpha + \#(x=1, c=good) + \beta + \#(x=0, c=good)} \quad (4)$$

where #($x$ = 0, $c$ = good) and #($x$ = 1, $c$ = good) are the number of times the training points ($x$ = 0, 'good') and ($x$ = 1, 'good') appear in the training set, respectively. The parameters $\alpha$ and $\beta$ represent initial counts and can be used to initialize the class conditional probability distributions when prior knowledge of the distribution probabilities is available. The distribution parameter $\theta_{x=1,c=\text{good}}$ is estimated similarly.

The process is generalized to variables with larger domain sizes as follows. If the discrete design variable $x_i$ has cardinality (domain size) $C$, then $x_i^j$ is the $j^{th}$ level in the domain of $x_i$, where $j \in \{1, 2, ..., C\}$. Furthermore, if the initial counts are represented by the vector $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \ldots, \alpha_C]$, then the discrete distribution parameters for $x_i^j$ given class $c_k$ are estimated by Eq. (5):

$$\theta_{x_i = x_i^j, c = c_k} = \frac{\alpha_j + \#(x_i = x_i^j, c = c_k)}{\sum_{j=1}^{C} \left[ \alpha_j + \#(x_i = x_i^j, c = c_k) \right]} \quad (5)$$

The simplest setting for the initial counts is to set them all to unity [12]. Doing so sets all of the class conditional probability distributions to uniform before any training points are added to the classifier. This approach generally makes sense unless there is reason to believe, possibly through prior experimentation, that some variable values have higher likelihood than others given the class label.

## Global Optimization with Bayesian Classifiers

CGS uses a Bayesian classifier to achieve efficient design space exploration and optimization. Each newly evaluated point is assigned a class label (e.g., 'good' / 'bad') depending on its objective function value, and the point is added to the classifier training set. The updated classifier is then used to screen each new candidate solution based on the posterior probability of the design's class prior to expensive evaluation. Furthermore, new candidate points are generated by sampling the class conditional probability distributions that comprise the classifier. By sampling the distributions that are trained with high-performance solutions, new points are generated that are likely to improve the objective function. In general, the first sample will be drawn from a uniform random distribution over the input domain. With each new point that is evaluated, assigned a class label, and added to the classifier training set, the classifier improves its ability to generate high-performance solutions and filter out low-performance solutions.

Figure 1 shows a flow chart of the CGS method. It begins by instantiating the class-conditional probability distributions as uniform discrete distributions. The next step is to sample the class conditional probabilities of the high-performance class to generate a candidate solution. In Step 2, the candidate's posterior probability of being 'good' or not is determined by the classifier. In Step 3, the candidate solution is accepted or rejected for evaluation based on two criteria. First, the solution is checked against all previously evaluated solutions to avoid repeat evaluations of the same solution. Second, if the candidate solution's posterior probability of being 'good' is below a threshold, it is rejected. The threshold value is determined for each candidate solution by sampling from a uniform distribution ranging from zero to one. If the candidate solution is accepted, the new design point is evaluated with the expensive simulation in Step 4. Otherwise, the method returns to Step 1 to generate a new point.

In Step 5, a class label is assigned to the newly evaluated point based on the fitness of the design, and the classifier training set is updated. At a minimum, two classes are required, but more may be used if desired. For this paper, each evaluated design is given a class label of either 'good' or 'bad'. The labels are determined by assigning the top $N$ solutions a label of 'good' and all others a label of 'bad', where $N$ is a user-defined constant. Therefore, new solutions that outperform those in the set of top $N$ will replace those that are inferior (i.e., their class label will be reassigned from 'good' to 'bad').

The process repeats if a convergence criterion does not end the cycle. Stopping criteria can include reaching a predetermined number of design evaluations, achieving a desired objective function value, or failing to improve the best design by some percentage after a predetermined number of new designs are evaluated.

In light of the discussion above, the choice of a Bayesian classifier over other techniques becomes clear. A Bayesian classifier is ideal for CGS because it provides a probability that a test point belongs to a class, while many other classification methods only provide the class label. Furthermore, sampling the distributions of 'good' solutions effectively generates new candidate solutions that have a high likelihood of improving the objective function value.

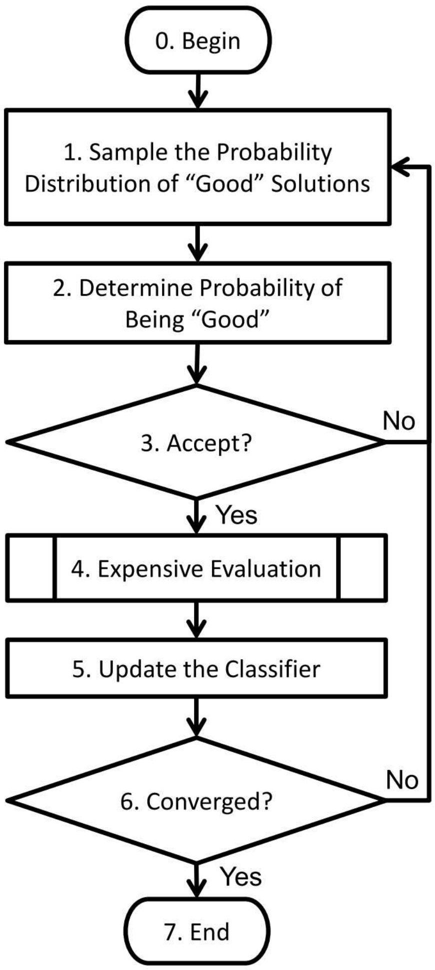


**Figure 1: Classifier-guided sampling algorithm**

## 4. HURRICANE SANDY CASE STUDY

In October of 2012, Hurricane Sandy brought intense wind, rain, storm surges and claimed hundreds of lives. In the U.S., it is estimated that Sandy damaged or destroyed 650,000 houses, left 8.5 million customers without power, and caused 72 deaths [13]. It also left millions of homes and business without electric power. The storm was particularly devastating to the small city of Hoboken, NJ, which sits on the Hudson River across from Manhattan. Sandy’s storm surge was over 17 feet above mean sea level, and as a result, Hoboken was flooded for several days following the storm. The city was disconnected from the main utility grid during this time. A well-designed MG may have been able to provide autonomous support to critical infrastructure during this outage. However, the decisions about what DERs to install, where to install them, and how they should be interconnected have significant implications on their cost and ability to serve loads during autonomous MG operation. In this section, this problem is formulated as a single-objective optimization problem, and CGS is used to identify configurations that reduce LNS subject to an upper limit on installation cost. This work is an extension of a previous study in which a thorough energy surety analysis was performed for the City of Hoboken by Sandia National Laboratories [14].

### Microgrid Design Problem Parameters

There are 55 buildings in Hoboken that are designated as critical and should remain operational during an emergency. The buildings include emergency services, pump stations (for flood control), affordable and senior housing units, grocery stores, gas stations, and government buildings. A complete listing of the buildings, their locations, and their estimated continuous loads is given in Appendix A, Table A.1.

The problem addressed here is to design a microgrid that minimizes average LNS, in kWh per hour of outage, of the 55 buildings without exceeding a limit on installation cost. There are two types of design variables that define this problem. First, the sizes and locations of natural gas generators are considered. Natural gas generators are chosen as the generator type because the high-pressure natural gas pipeline that runs throughout Hoboken is assumed to remain operational during an emergency, as it did during Hurricane Sandy. If installed, the generators are to be co-located with the 55 critical buildings. Generator installation costs are based on their size. The second type of design variable defines the topology of the microgrid. K-means clustering was used to group the 55 potential generator sites into 13 sub-grids based on geographic location. All buildings are assumed to be connected within each sub-grid, and each sub-grid can only be connected to adjacent sub-grids. The cost to connect each sub-grid is based on the cost per foot of underground cable. The building sub-grid assignments are tabulated in Table A.1 and shown on a city map in Figure A.1.

In total, there are 153 design variables. Each variable is binary in the sense that the choice for each is to either do nothing or install something (either a generator or MG connection). There are 131 generator installation design variables, and 22 MG connection design variables. All 153 design variables and their installation costs are tabulated in Appendix B, Table B.1. An upper limit on installation cost is set at \$8M. Because of the emergency nature of this problem, fuel cost is not considered because fuel use is a short-term need that will be a relatively low cost in the long run.

Load not served is evaluated by running a PRM simulation on candidate microgrid designs that are proposed by the optimizer. Time-dependent hourly load data for each building is estimated based on the estimated peak loads presented in Table A.1. Utility outage frequencies and durations are sampled from probability distributions such that they have an average duration of 1 week and occur an average of every 100 years. The PRM is used to simulate grid-tied and autonomous operation continuously over a period of 100,000 years. Therefore, the expected number of outages that are simulated with each run is 1,000. Recall from Section 2 that a large number of simulations are required because PRM considers all components that comprise the MG to be unreliable (i.e., they fail and are subsequently repaired according to user-defined probability distributions). Therefore, a simulation of a single outage does not provide a meaningful estimate of "average" microgrid performance. The expected outage count was selected by increasing the total simulation time until stable results were observed between runs of identical configurations.

## Results and Discussion

In this section, CGS is used to identify Hoboken MG configurations that reduce LNS and do not violate the cost constraint. To validate the effectiveness of CGS, a genetic algorithm (GA) was also used to solve the optimization problem. The GA implementation used is the JEGA single-objective solver available in the Dakota toolkit [15].

For CGS, the parameter $N$ (number of solutions with the 'good' label) is set to 50. To promote broad search early in the solution process, the first 100 candidates are sampled randomly from uniform distributions. After these initial 100 samples, the classifier is used to guide the search for the remainder of the optimization. For the GA, the population size is 50, and the crossover and mutation rates are 0.8 and 0.1, respectively.

Both optimization methods handle the cost constraint by always preferring feasible solutions to infeasible ones. When comparing two infeasible solutions, the one that exceeds the cost limit by a lesser amount is preferred. This approach is similar to that which is proposed by Deb [16].

Performance comparison of CGS and the GA is achieved by executing a set of rate of convergence tests in which the current best known solution versus the number of objective function evaluations is recorded. The two methods are executed five times each with a fixed upper limit of 10,000 objective function evaluations, and the average results of the 5 trials are computed. Repeat evaluations of previously assessed solutions are not performed and are therefore not included in the rate of convergence results. This test provides a visual measure of how quickly each method identifies high-performance configurations. In Figure 2, the results of a rate of convergence test are shown.

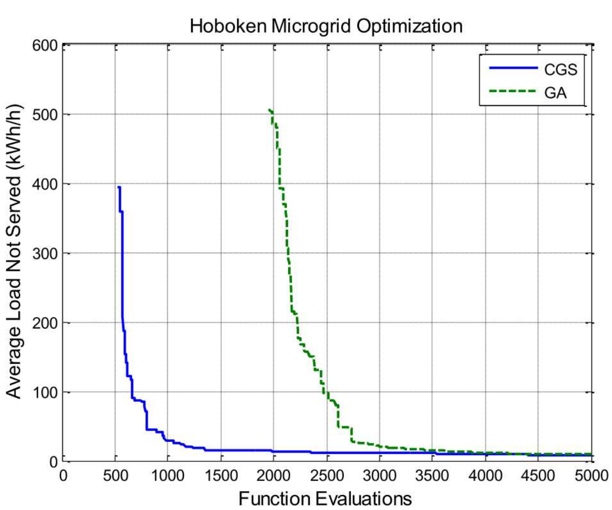


**Figure 2: Rate of convergence results**

On average, CGS identifies high-performance (low LNS) solutions with significantly fewer objective function evaluations than the GA. Neither method immediately identifies solutions that satisfy the cost constraint in all five trials. Therefore, the curves in Figure 2 begin at the lowest number of evaluations needed for all five trials to identify feasible solutions. Beyond the first 5,000 function evaluations, the curves level off at almost identical levels of average LNS until the algorithms were terminated at 10,000 evaluations.

The overall "best" solution found by CGS is indicated in Table B.1. This configuration has an estimated installation cost of \$7.91M, and the PRM simulation estimates that it results in average LNS of 5.06 kWh/h (kWh per hour of outage). Considering that the estimated peak building loads in Table A.1 sum to approximately 9.9 MW, this level of LNS is relatively low.

This configuration implements 34 natural gas generators with a total installed generating capacity of 7.5MW. This capacity is significantly less than the sum of the estimated static loads in Table A.1. However, the loads used in the PRM simulation are time-dependent (hourly), and were estimated using a combination of metered data and the quantities in Table A.1. Since the peak loads for each building occur at a different times of the day, the overall time-dependent peak load for all building combined is significantly less than the total static load.

The best known solution forms 8 microgrids by connecting some, but not all, of the clusters (MG1-MG8, MG2-MG9-MG3, MG7-MG4-MG12, MG11, MG5, MG6, MG1UB, and MG10). Connecting some of these grids enables generators in one cluster to serve loads in another. Therefore, the connections give the search algorithm more and better choices about how to optimally match installed capacity to estimated loads.

Another noteworthy feature of this configuration is that Building 7 (B7), the sewage treatment plant, does not have any additional generation capacity installed and is not connected to any other sub grids. The plant currently has an existing diesel backup generator that is large enough to cover its expected peak load. However, an additional NG generator or a connection to another bus could reduce LNS by providing redundant power sources (the single existing diesel generator could fail). This feature sheds light on the difficulty that the tight cost constraint places on the decision making process. For B7, the cheapest generation option is nearly \$1M, and the least inexpensive sub-grid connection option is roughly \$672k (MG8-MG11). Either of these options would consume a significant percentage of the \$8M cost limit, and the CGS result indicates that spending the money on other assets does more to reduce LNS.

## 5. CONCLUSIONS

Design of an autonomous microgrid is a significant challenge due to a large number of design choices and uncertainty related to loads and component reliability. In this paper, a discrete event simulation was used to model and assess the average performance of a candidate microgrid design by simulating a large number of stochastic utility outages over a long time horizon. To mitigate the computational expense that is inherent to this approach, classifier-guided sampling (CGS), a Bayesian classifier-based optimization algorithm, was used to identify high-performance configurations that do not exceed a pre-determined installation cost. As a test case, a microgrid is designed for the city of Hoboken, NJ, which suffered an extended outage in the wake of Hurricane Sandy in 2012. Results show that CGS is able to identify high-performance solutions with fewer objective function evaluations than the genetic algorithm implementation.

## ACKNOWLEDGMENTS

This work was supported by Sandia National Laboratories under the auspices of the Laboratory Directed Research and Development program. The authors thank Jason Stamp, Michael Baca, Ross Guttromson, Jordan Henry, Richard Jensen, Karina Muñoz-Ramos, Ben Schenkman, and Mark Smith for their contributions to the City of Hoboken Energy Surety Analysis effort. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND TBD.

## REFERENCES

1. Lasseter, R.H. and P. Paigi. *Microgrids: A Conceptual Solution*. in *35th Annual IEEE Power Electronics Specialists Conference*. 2004. Aachen, Germany.

2. Lasseter, R., et al., *Integration of Distributed Energy Resources: The CERTS MicroGrid Concept*. 2002, Lawrence Berkeley National Laboratory (whitepaper).

3. Katiraei, F. and M.R. Iravani, *Power Management Strategies for a Microgrid With Multiplle Distributed Generation Units.* IEEE Transactions on Power Systems, 2006. **21**(4): p. 1821-1831.

4. Pecas Lopes, J.A., C.L. Moreira, and A.G. Madureira, *Defining Control Strategies for MicroGrids Islanded Operations.* IEEE Transactions on Power Systems, 2006. **21**(2): p. 916-924.

5. Wang, G.G. and S. Shan, *Review of Metamodeling Techniques in Support of Engineering Design Optimization.* Transactions of the ASME, 2007. **129**: p. 370-380.

6. Forrester, A.I.J. and A.J. Keane, *Recent Advances in Surrogate-Based Optimization.* Progress in Aerospace Sciences, 2009. **45**: p. 50-79.

7. Swiler, L.P., et al., *Surrogate models for mixed discrete-continuous variables*. 2012, Sandia National Laboratories: Sandia Report SAND2012-0491.

8. Shahan, D.W., P.B. Backlund, and C.C. Seepersad, *Classifier-Guided Sampling for Discrete Variable, Discontinuous Design Space Exploration*, in *ASME International Design Engineering Technical Conferences*. 2013: Portland, OR.

9. Backlund, P.B., D.W. Shahan, and C.C. Seepersad, *Classifier-Guided Sampling for Discrete Variable, Discontinuous Design Space Exploration: Convergence and Computational Performance.* Engineering Optimization, 2014. **Available Online April, 2014 (www.tandfonline.com)**.

10. Kotsiantis, S.B., *Supervised Machine Learning: A Review of Classification Techniques.* Informatica, 2007. **31**: p. 249-268.

11. Zhang, H. *The Optimality of Naive Bayes*. in *Proceeding of the Seventeenth International Florida Artificial Intelligence Research Society Conference*. 2004. AIAA Press, Maimi Beach.

12. Barber, D., *Bayesian Reasoning and Machine Learning*. 2012: Cambridge University Press.

13. Blake, E.S., et al., *Tropical Cyclone Report: Hurricane Sandy*. 2013, National Hurricane Center, 12 February 2013.

14. Stamp, J., et al., *City of Hoboken Energy Surety Analysis: Preliminary Design Summary*. 2014, Sandia National Laboratories: Sandia Report, SAND2014-17842.

15. Bohnhoff, W.J., et al., *Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.1 User's Manual*. 2014, Sandia National Laboratories: SAND2014-4633.

16. Deb, K., *An efficient constraint handling method for genetic algorithms.* Computer Methods in Applied Mechanics and Engineering, 2000. **186**: p. 311-338.

## APPENDIX A

## CRITICAL HOBOKEN BUILDINGS

**Table A.1: Buildings in Hoboken, NJ designated as critical during extended utility outage [14]**

| Bldg. # | Building Name | Type | Location | Estimated Load (kW) | Cluster # |
|---|---|---|---|---|---|
| 1 | Fire Engine Company 3 | Emergency | 1313 Washington Street | 150 | 1 |
| 2 | Fire Engine Company 4 | Emergency | 801 Clinton Street | 22.5 | 2 |
| 3 | Fire Headquarters | Emergency | 201 Jefferson Street | 37.5 | 3 |
| 4 | Fire Engine Company 1 | Emergency | 43 Madison Street | 45 | 6 |
| 5 | Police Headquarters | Emergency | 106 Hudson Street | 150 | 7 |
| 6 | University Medical Center | Emergency | 308 Willow Avenue | 1000 | 9 |
| 7 | Sewage Treatment Plant | Flood Control | Adams Street | 900 | 11 |
| 8 | Pump Station 5th Street | Flood Control | 500 River Road | 750 | 4 |
| 9 | Pump Station 11th Street | Flood Control | 83 11th Street | 15 | 1 |
| 10 | Pump Station H1 | Flood Control | 99 Observer Highway | 225 | 7 |
| 11 | Volunteer Ambulance Corps | Emergency | 707 Clinton Street | 15 | 2 |
| 12 | Hoboken City Hall | Operation | 94 Washington Street | 225 | 7 |
| 13 | Hoboken High School | Shelter | 800 Clinton Street | 150 | 2 |
| 14 | Wallace School | Shelter | 1100 Willow Avenue | 250 | 8 |
| 15 | Hoboken Homeless Shelter | Shelter | 300 Bloomfield | 45 | 12 |
| 16 | St. Matthew’s Church | Shelter | 57 8th Street | 15 | 5 |
| 17 | St. Peter and Paul Church | Shelter | 404 Hudson Street | 30 | 4 |
| 18 | A&P | Groceries | 614 Clinton Street | 45 | 9 |
| 19 | Kings 1 | Groceries | 325 River Street | 450 | 4 |
| 20 | Kings 2 | Groceries | 1212 Shipyard Lane | 450 | 1 |
| 21 | Sunoco | Gas Station | 1301 Willow Avenue | 15 | 8 |
| 22 | Multi-Service Center | Shelter | 124 Grand Street | 90 | 3 |
| 23 | Public Works Garage | Operation | 256 Observer Highway | 30 | 7 |
| 24 | Garage B | Parking Garage | 28 2nd Street | 90 | 7 |
| 25 | Garage D | Parking Garage | 215 Hudson Street | 225 | 7 |
| 26 | Garage G | Parking Garage | 315 Hudson Street | 150 | 4 |
| 27 | Midtown Garage | Parking Garage | 371 4th Street | 150 | 9 |
| 28 | Columbian Arms | Senior Housing | 514 Madison Street | 90 | 10 |
| 29 | Marion Towers | Senior Housing | 400 1st Street | 225 | 3 |
| 30 | Columbian Towers | Senior Housing | 76 Bloomfield Street | 150 | 7 |
| 31 | Housing Authority 1 | Affordable Housing | 655 6th Street | 450 | 1UB |
| 32 | Housing Authority 2 | Affordable Housing | 501 Marshall Drive | 90 | 1UB |
| 33 | Housing Authority 3 | Affordable Housing | 400 Marshall Drive | 45 | 1UB |
| 34 | Housing Authority 4 | Affordable Housing | 320 Marshall Drive | 67.5 | 1UB |
| 35 | Housing Authority 5 | Affordable Housing | 300 Marshall Drive | 90 | 1UB |
| 36 | Housing Authority 6 | Affordable Housing | 321 Harrison Street | 45 | 1UB |
| 37 | Housing Authority 7 | Affordable Housing | 311 Harrison Street | 45 | 1UB |
| 38 | Housing Authority 8 | Affordable Housing | 320 Jackson Street | 90 | 1UB |
| 39 | Housing Authority 9 | Affordable Housing | 310 Jackson Street | 90 | 1UB |
| 40 | Housing Authority 10 | Affordable Housing | 311 13th Street | 90 | 8 |
| 41 | Housing Authority 11 | Affordable Housing | 804 Willow Avenue | 90 | 2 |
| 42 | Fox Hill Housing | Senior Housing | 900 Clinton Street | 45 | 2 |
| 43 | 5 Church Towers | Affordable Housing | Grand Street | 45 | 9 |
| 44 | 10 Church Towers | Affordable Housing | Clinton Street | 150 | 9 |
| 45 | 15 Church Towers | Affordable Housing | Grand Street | 90 | 9 |
| 46 | Clock Towers | Affordable Housing | 300 Adams Street | 150 | 9 |
| 47 | Marineview 1 | Affordable Housing | 331 Hudson Street | 450 | 4 |
| 48 | Marineview 2 | Affordable Housing | 301 Hudson Street | 450 | 4 |
| 49 | Applied 1 | Affordable Housing | 111 Newark | 45 | 7 |
| 50 | Applied 2 | Affordable Housing | 1203-1209 Willow Avenue | 225 | 8 |
| 51 | YMCA (SROs) | Affordable Housing | 1301 Washington Street | 150 | 1 |
| 52 | Police Department Radio Repeater | Emergency | N/A | 450 | 7 |
| 53 | Fire Department Radio Repeater | Emergency | N/A | 15 | 5 |
| 54 | CVS | Pharmacy | 59 Washington Street | 150 | 7 |
| 55 | Walgreens | Pharmacy | 101 Washington Street | 90 | 7 |

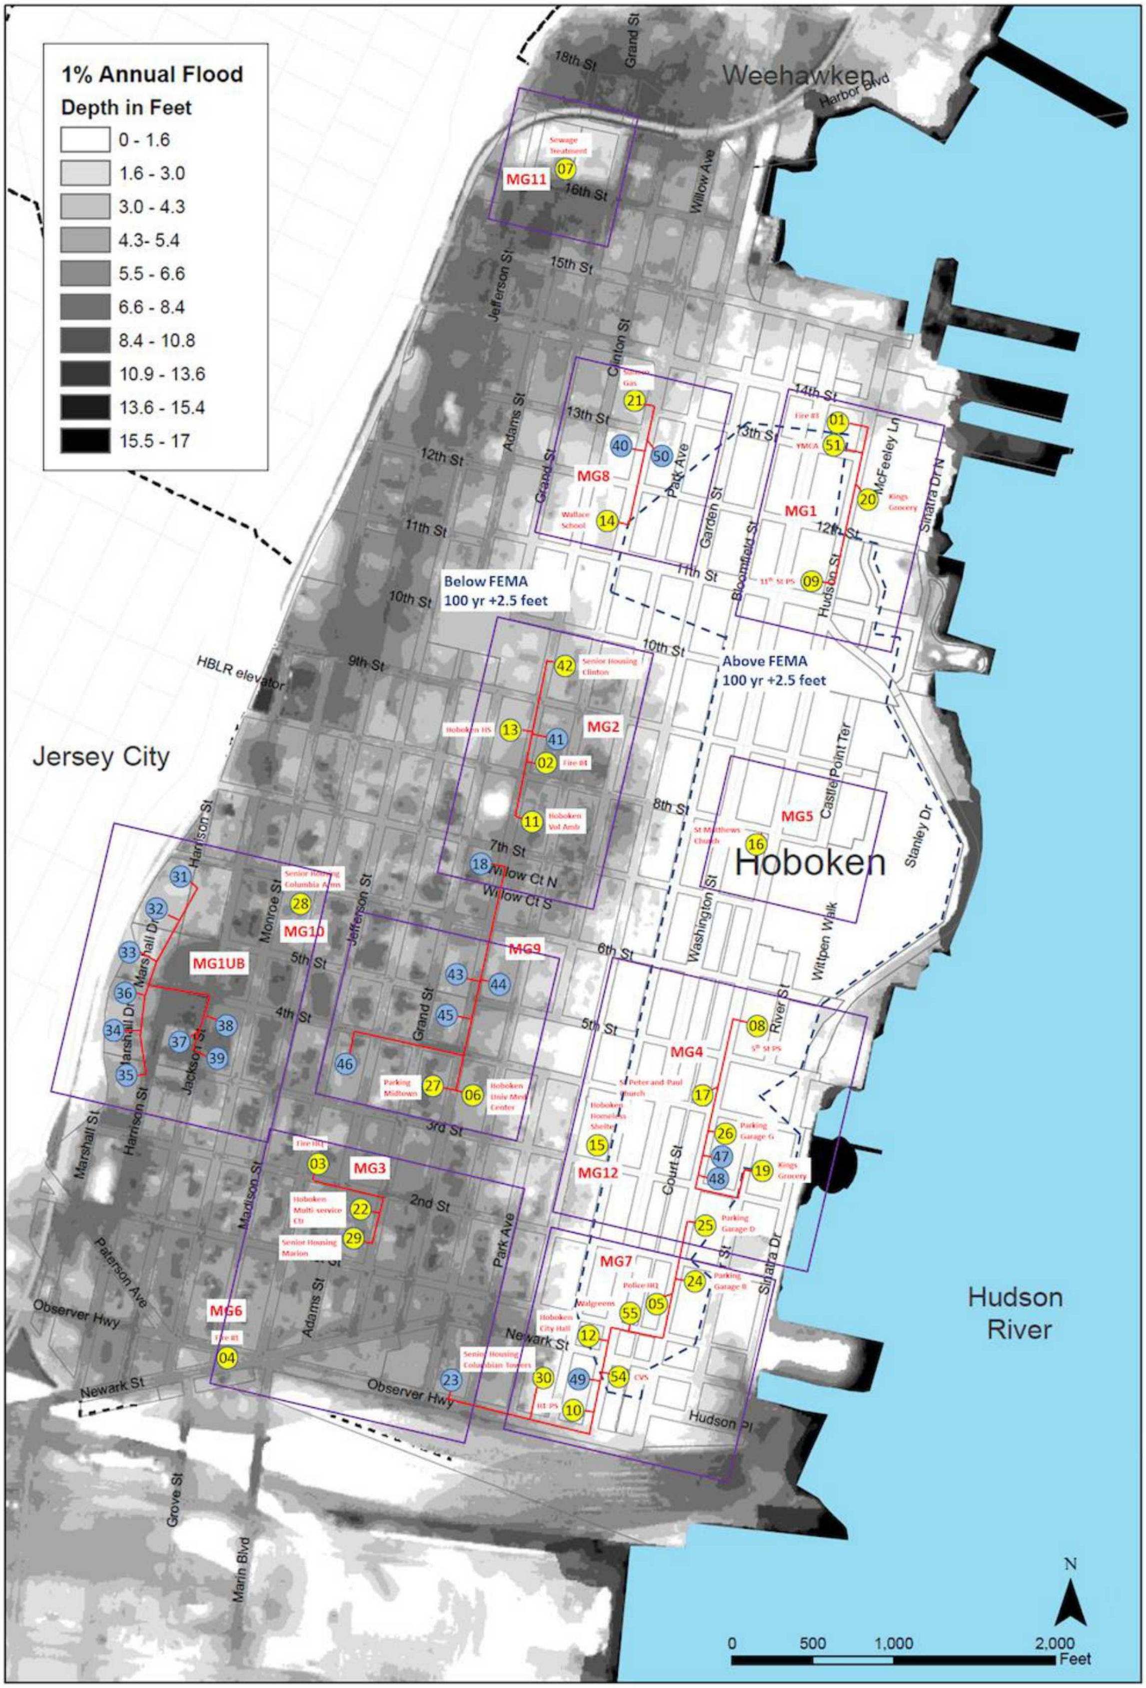


**Figure A.1: Critical Hoboken buildings shown on a flood map [14]**

## APPENDIX B

## HOBOKEN MICROGRID DESIGN VARIABLES

**Table B.1: Hoboken microgrid design variables and costs [14]**

| ID | Label | Option | Cost ($) |
|---|---|---|---|
| 1 | B12-NG1 | 200kW | $ 187,500 |
| 2 | B13-NG1 | 200kW | $ 187,500 |
| 3 | B13-NG2 | 200kW | $ 187,500 |
| 4 | B13-NG3 | 250kW | $ 217,500 |
| 5 | B13-NG4 | 250kW | $ 217,500 |
| 6 | B13-NG5 | 500kW | $ 396,000 |
| 7 | B14-NG1 | 30kW | $ 81,000 |
| 8 | B15-NG1 | 60kW | $ 107,100 |
| 9 | B15-NG2 | 60kW | $ 107,100 |
| 10 | B15-NG3 | 125kW | $ 147,600 |
| 11 | B15-NG4 | 150kW | $ 68,000 |
| 12 | B16-NG1 | 30kW | $ 81,000 |
| 13 | B16-NG2 | 30kW | $ 81,000 |
| 14 | B16-NG3 | 50kW | $ 97,800 |
| 15 | B16-NG4 | 75kW | $ 125,400 |
| 16 | B17-NG1 | 30kW | $ 81,000 |
| 17 | B17-NG2 | 50kW | $ 97,800 |
| 18 | B18-NG1 | 60kW | $ 107,100 |
| 19 | B19-NG1 | 300kW | $ 234,000 |
| 20 | B19-NG2 | 300kW | $ 234,000 |
| 21 | B19-NG3 | 600kW | $ 507,000 |
| 22 | B19-NG4 | 1500kW | $ 1,246,185 |
| 23 | B1-NG1 | 175kW | $ 183,000 |
| 24 | B1-NG2 | 400kW | $ 318,000 |
| 25 | B20-NG1 | 400kW | $ 318,000 |
| 26 | B20-NG2 | 600kW | $ 507,000 |
| 27 | B20-NG3 | 800kW | $ 648,000 |
| 28 | B20-NG4 | 800kW | $ 648,000 |
| 29 | B21-NG1 | 30kW | $ 81,000 |
| 30 | B22-NG1 | 100kW | $ 139,800 |
| 31 | B22-NG2 | 250kW | $ 217,500 |
| 32 | B23-NG1 | 40kW | $ 81,000 |
| 33 | B23-NG2 | 75kW | $ 125,400 |
| 34 | B24-NG1 | 125kW | $ 147,600 |
| 35 | B24-NG2 | 125kW | $ 147,600 |
| 36 | B25-NG1 | 300kW | $ 234,000 |
| 37 | B25-NG2 | 400kW | $ 318,000 |
| 38 | B26-NG1 | 200kW | $ 187,500 |
| 39 | B26-NG2 | 350kW | $ 264,000 |
| 40 | B27-NG1 | 200kW | $ 187,500 |
| 41 | B27-NG2 | 200kW | $ 187,500 |
| 42 | B27-NG3 | 275kW | $ 226,500 |
| 43 | B27-NG4 | 500kW | $ 396,000 |
| 44 | B28-NG1 | 125kW | $ 147,600 |
| 45 | B28-NG2 | 125kW | $ 147,600 |
| 46 | B28-NG3 | 150kW | $ 68,000 |
| 47 | B28-NG4 | 150kW | $ 68,000 |
| 48 | B28-NG5 | 175kW | $ 183,000 |
| 49 | B29-NG1 | 100kW | $ 139,800 |
| 50 | B29-NG2 | 250kW | $ 217,500 |
| 51 | B29-NG3 | 300kW | $ 234,000 |
| 52 | B29-NG4 | 300kW | $ 234,000 |
| 53 | B29-NG5 | 350kW | $ 264,000 |
| 54 | B29-NG6 | 350kW | $ 264,000 |
| 55 | B2-NG1 | 30kW | $ 81,000 |
| 56 | B30-NG1 | 200kW | $ 187,500 |
| 57 | B30-NG2 | 400kW | $ 318,000 |
| 58 | B31-NG1 | 125kW | $ 147,600 |
| 59 | B31-NG2 | 250kW | $ 217,500 |
| 60 | B31-NG3 | 175kW | $ 183,000 |
| 61 | B32-NG1 | 60kW | $ 107,100 |
| 62 | B33-NG1 | 100kW | $ 139,800 |
| 63 | B34-NG1 | 125kW | $ 147,600 |
| 64 | B34-NG2 | 275kW | $ 226,500 |
| 65 | B35-NG1 | 60kW | $ 107,100 |

| ID | Label | Option | Cost ($) |
|---|---|---|---|
| 66 | B36-NG1 | 100kW | $ 139,800 |
| 67 | B37-NG1 | 125kW | $ 147,600 |
| 68 | B37-NG2 | 275kW | $ 226,500 |
| 69 | B37-NG3 | 300kW | $ 234,000 |
| 70 | B38-NG1 | 125kW | $ 147,600 |
| 71 | B38-NG2 | 300kW | $ 234,000 |
| 72 | B39-NG1 | 125kW | $ 147,600 |
| 73 | B39-NG2 | 275kW | $ 226,500 |
| 74 | B39-NG3 | 300kW | $ 234,000 |
| 75 | B3-NG1 | 30kW | $ 81,000 |
| 76 | B40-NG1 | 300kW | $ 234,000 |
| 77 | B40-NG2 | 350kW | $ 264,000 |
| 78 | B40-NG3 | 400kW | $ 318,000 |
| 79 | B40-NG4 | 600kW | $ 507,000 |
| 80 | B40-NG5 | 600kW | $ 507,000 |
| 81 | B40-NG6 | 650kW | $ 607,500 |
| 82 | B40-NG7 | 650kW | $ 607,500 |
| 83 | B41-NG1 | 150kW | $ 68,000 |
| 84 | B41-NG2 | 175kW | $ 183,000 |
| 85 | B41-NG3 | 275kW | $ 226,500 |
| 86 | B42-NG1 | 60kW | $ 107,100 |
| 87 | B42-NG2 | 150kW | $ 68,000 |
| 88 | B43-NG1 | 125kW | $ 147,600 |
| 89 | B43-NG2 | 275kW | $ 226,500 |
| 90 | B43-NG3 | 300kW | $ 234,000 |
| 91 | B44-NG1 | 60kW | $ 107,100 |
| 92 | B45-NG1 | 200kW | $ 187,500 |
| 93 | B45-NG2 | 275kW | $ 226,500 |
| 94 | B46-NG1 | 200kW | $ 187,500 |
| 95 | B46-NG2 | 275kW | $ 226,500 |
| 96 | B46-NG3 | 500kW | $ 396,000 |
| 97 | B47-NG1 | 275kW | $ 226,500 |
| 98 | B47-NG2 | 350kW | $ 264,000 |
| 99 | B47-NG3 | 600kW | $ 507,000 |
| 100 | B47-NG4 | 1500kW | $ 1,246,185 |
| 101 | B48-NG1 | 300kW | $ 234,000 |
| 102 | B48-NG2 | 300kW | $ 234,000 |
| 103 | B48-NG3 | 600kW | $ 507,000 |
| 104 | B48-NG4 | 1000kW | $ 774,000 |
| 105 | B49-NG1 | 300kW | $ 234,000 |
| 106 | B49-NG2 | 750kW | $ 622,500 |
| 107 | B4-NG1 | 50kW | $ 97,800 |
| 108 | B4-NG2 | 50kW | $ 97,800 |
| 109 | B4-NG3 | 125kW | $ 147,600 |
| 110 | B4-NG4 | 150kW | $ 68,000 |
| 111 | B50-NG1 | 60kW | $ 107,100 |
| 112 | B51-NG1 | 200kW | $ 187,500 |
| 113 | B51-NG2 | 400kW | $ 318,000 |
| 114 | B52-NG1 | 400kW | $ 318,000 |
| 115 | B52-NG2 | 600kW | $ 507,000 |
| 116 | B52-NG3 | 1500kW | $ 1,246,185 |
| 117 | B52-NG4 | 1500kW | $ 1,246,185 |
| 118 | B53-NG1 | 30kW | $ 81,000 |
| 119 | B54-NG1 | 200kW | $ 187,500 |
| 120 | B54-NG2 | 400kW | $ 318,000 |
| 121 | B55-NG1 | 125kW | $ 147,600 |
| 122 | B5-NG1 | 100kW | $ 139,800 |
| 123 | B5-NG2 | 100kW | $ 139,800 |
| 124 | B6-NG1 | 30kW | $ 81,000 |
| 125 | B7-NG1 | 1200kW | $ 996,948 |
| 126 | B7-NG2 | 1500kW | $ 1,246,185 |
| 127 | B7-NG3 | 2000kW | $ 1,661,580 |
| 128 | B8-NG1 | 350kW | $ 264,000 |
| 129 | B8-NG2 | 350kW | $ 264,000 |
| 130 | B8-NG3 | 500kW | $ 396,000 |

| ID | Label | Option | Cost ($) |
|---|---|---|---|
| 131 | B8-NG4 | 600kW | $ 507,000 |
| 132 | MG 1-2 | Connect | $ 782,320 |
| 133 | MG 1-5 | Connect | $ 465,250 |
| 134 | MG 1-8 | Connect | $ 322,960 |
| 135 | MG 2-5 | Connect | $ 413,200 |
| 136 | MG 2-8 | Connect | $ 402,160 |
| 137 | MG 2-9 | Connect | $ 132,100 |
| 138 | MG 3-6 | Connect | $ 438,680 |
| 139 | MG 3-7 | Connect | $ 592,240 |
| 140 | MG 3-9 | Connect | $ 207,700 |
| 141 | MG 3-13 | Connect | $ 435,400 |
| 142 | MG 4-5 | Connect | $ 449,680 |
| 143 | MG 4-7 | Connect | $ 97,900 |
| 144 | MG 4-9 | Connect | $ 502,300 |
| 145 | MG 4-12 | Connect | $ 201,500 |
| 146 | MG 6-7 | Connect | $ 429,800 |
| 147 | MG 7-9 | Connect | $ 619,700 |
| 148 | MG 7-12 | Connect | $ 356,600 |
| 149 | MG 8-11 | Connect | $ 671,600 |
| 150 | MG 9-10 | Connect | $ 394,400 |
| 151 | MG 9-12 | Connect | $ 352,800 |
| 152 | MG 9-13 | Connect | $ 438,700 |
| 153 | MG 10-13 | Connect | $ 360,500 |

There are two types of design variables in this Hoboken MG design optimization problem: those that indicate installation of a new generator (IDs 1-131) and those that indicate connecting two building cluster microgrids (IDs 132-153).

Each variable has two discrete choices (install/don't install). Labels indicate the building and generator number or the two clustered microgrids to connect. E.g., the variable "B12-NG1" is a decision about whether to install a natural gas generator in Building 12, and MG1-2 is a decision about whether to connect microgrids 1 and 2.

The highlighted variables are those that were chosen by the overall best solution that was identified with CGS.