The Effects of Variation in Load Demand and Line Failure on the Electric Power System During a Pandemic

Audrey Ahlenius, Geneve Lauby, and Brooks Lewis

ECE 4320

December 9, 2022

Table of Contents

Executive Summary

Purpose of Report

This report is to serve as an operating procedure for Georgia's power grid should Georgia face a pandemic, such as COVID-19, again. Risk management and planning for "worst-case" scenarios prevents disorganization during a time of change in normalcy and allows for structured procedure based on data driven results. This report details operation, maintenance, and repair recommendations by analyzing the effect of parameters pertaining to load demand, line failure, and crew allocation at various levels of severity.

Therefore, the purpose of this report is to:

- Perform a parameter sensitivity analysis for various parameters and plot the trend between value of the parameter and total cost of operation at that value
- Provide recommendations for operation of the grid based on analysis and correlation between the parameters and total cost of operation

Methods and Findings

After initial modification to decrease randomness of crew allocation, a loop was used to change the values of the parameter for each simulation and summarize the resulting scores in a matrix. The data collected was put into a spreadsheet and graphs were created to display trends between parameters. From our findings we concluded that the following variables had a correlation between total cost of operation and increasing severity:

- Last Turn
- Pmu
- Minimum Failure Probability
- Fail Probability
- Failure Rate Increase
- Crew Quarantine Penalty
- Proximity Sickness Likelihood

These parameters were deemed critical to operation and were assessed for operational recommendations.

Future Recommendations

Through analysis of the simulation data collected and insight from personal experience living through a pandemic, observations were used to form recommendations for a cost-saving and reliable electric grid. In a real life grid operation setting, it would be impossible to predict the transmissibility and severity of the disease, but by combining data with personal experience, a cautious approach is proposed and the most pertinent recommendations are as follows:

- Prioritize line maintenance at the beginning of the pandemic to prevent continued failure later in the pandemic when crews become more limited
- Monitor failure rates of newly repaired lines as well as monitor the failure rates of currently aging lines
- Maintain accurate data on failure probabilities to create an accurate expectation for pandemic operations and choose suppliers for components that ensure the lowest failure rate.
- Monitor the lifespan of line components, and collect data on how their failure rate increases over time.
- Limit the number of crews to critical power outages and limited maintenance in the later stages of the pandemic to decrease the risk of exposure and maintain crew morale.

Introduction

The modern electric power grid is one of the most complex engineering feats still in use today in part due to its reliance on societal patterns to dictate supply and demand of electricity. With a system so influenced by those that use it, when there is a deviation in normalcy, such as a pandemic, the effects on the power grid are difficult to predict. The COVID-19 pandemic has taught the world much about the impact of pandemics on the electric power grid and has allowed for change and adaptation in a post-pandemic world.

One of the largest impacts the COVID-19 pandemic had on the power grid was the decrease in demand for electricity due to stay-at-home orders leading to a temporary halt in industrial and commercial activities. With the exception of hospitals, many large commercial and industrial entities were forced to shut down, leading to higher grid voltage without the large voltage drops from commercial and industrial structures. As a result, hospitals were seeing voltage higher than the range allowed for their equipment, leading to automatic shutoff, a potentially deadly consequence as hospitals neared capacity. Similarly, as people began working from home, there became a shift in peak demand at residential homes from before 9am and after 5pm to consistent demand all day. These factors combined resulted in a 4-5% decrease in peak hour demand in Texas according to a study by the Electric Reliability Council of Texas (ERCOT). This change in demand made it difficult for both utilities and long-term system planners and traders to predict the estimated reduction in generation so that demand is still met.

Though COVID-19 led to a reduction in demand for electricity, electric power grid reliability became more important than ever. As the shift to work-from-home increased, access to reliable internet became a necessity. Additionally, hospitals required consistent power for 24/7 operation. To maintain these needs, line crew allocation to restore power or do routine maintenance had to be weighed against potential exposure risk. If the number of workers per crew were limited, this could pose safety hazards and increase restoration time. Conversely, if a large team of workers were sent out and became exposed, then there may be a lack of skilled labor to resolve issues, again, increasing restoration time. This balance between being cautious in allocating the available workforce and being able to resolve issues in a timely manner was another consideration that those working on the grid during a pandemic had to contemplate.

Summary of Pandemic Model

The model provided as part of the project is intended to simulate the power grid in Georgia. Although not an exact representation of the real power grid, the behavior is typical of the real life counterpart. The simulation is intended to provide researchers with a test model in order to accurately estimate the effects of future pandemics on the power grid.

The simulation uses DC power flow equations to model the grid, so no reactive power is considered. The simulation uses 20 buses, 6 generators, and 30 transmission lines. Load demands vary during different times of day, as well as during different weeks of the year. The system considers generator operation costs, but neglects generators not in use as well as any ramp up costs. The generator cost coefficients are the same for each time period.

Line failures are determined somewhat randomly by the simulation. Lines with a status of less than 1 are unavailable to the user until repaired. At the start of each simulation, each line is randomly assigned an exogenous failure probability to simulate severe weather and external events. During each turn, the simulation randomly determines which lines fail based on the failure probability. If a line fails, it may cause a cascading failure by overloading surrounding lines. The user can perform maintenance or repair on lines specified each turn as well. If a line has not failed, maintenance crews decrease the failure probability by a multiplicative factor (not to fall below the minimum failure probability). If a line has failed, a maintenance crew will increase a line's status by a specified amount. Each crew sent in the field can be infected by the disease and then become unavailable for the remainder of the simulation.

The simulation is scored based on load shedding costs, generation costs, and a penalty for the number of quarantined crews. The generation and load shedding costs are summed over the time period then multiplied by the quarantine penalty. A lower score means a more effective strategy for operating the grid.

Throughout the simulation, the user is given access to DC power flow solvers and DC security constrained unit commitment problems. The user can run the simulation with a graphical interface. The user can then click buttons to change turns, and use the command line to specify values for power generation and crew allocations. The user can then view the outcomes of their decisions, as well as view downed lines in the map window. The user can choose to run the simulation in the command line if testing multiple simulations very quickly.

Description of Experimental Set Up

The goal of this report is to perform a parameter sensitivity analysis. However, there is randomness worked into simulation that affects the results. To perform an effective parameter sensitivity analysis, actions were taken to limit the randomness. First, the crew allocation was derandomized and second, the simulation was run at least five times to account for randomness.

The allocation of the maintenance crews were based on the failure probability of each line. This allocation was done by finding the failure probability for each line, sorting them in descending order, and then only assigning crews to lines that have a failure probability above the threshold of 0.01. The decision of 0.01 as the threshold probability was chosen after doing tests to see the range of common failure probabilities. The chosen threshold in most scenarios allows for more than one line to have maintenance or repair available but avoids having all maintenance/repair crews deployed. This balance allows some lines to be repaired but avoids a large number of crews potentially being exposed to disease. The code for assigning crews is in Appendix A.

In addition to reducing the variability of crew allocation, each parameter was tested at least five times to increase the data points for analysis. The original goal was to test each parameter ten times, but after doing a test run, it took 40 minutes for each parameter. Since that was an unrealistic goal based on our team's availability, the five turns were chosen instead.

For each parameter tested, a range of six to eleven values were tested. The code loop used fprintf for the file setParameters.m to change the value of the parameter. When changing the parameter tested, the code must be manually changed to insert the variable in the correct part of the parameters file. The code for changing parameters and running the simulation is in Appendix B. When comparing how the system reacts to different parameter values, the scores were compared. The next section will go through the results of each parameter. The analysis was run for each parameter, except for TimesOfDay and StartWeek, due to errors that occurred when changing those values. For the full analysis of the parameters, only parameters that showed a trend between parameter values and total score were evaluated. The simulation was run with N-1 contingency in every test for consistency. The goal of the experiment set up was not to have the most effective system, but to have the least random system.

When evaluating a single parameter, all other parameter values were held constant. Table 1 shows the constant values used in all simulation runs. Table 1 also shows the range of values used when evaluating an individual parameter.

Table 1. List of parameters with constant values and variation ranges

Parameter	Constant Value	Variation Range
lastTurn	30	25-35
nCrew	10	5-15

Psigma	0.2	0.1-0.3
Pmu	1.5	1-2
TimesOfDay	[1 9 19]	Not tested
StartWeek	40	Not tested
minFailureProbability	0.0008	0.003-0.012
failProb	0.02	0.01-0.028
outlierFailProb	0.08	0.03-0.12
outlierFraction	0.05	0.025-0.07
failureRateIncrease	1.01	1.01-1.1
instaFail	1.10	0.8-1.7
cascadeFailureStatus	0.67	0.27-0.87
Ptol	0.1	0.05-0.3
crewQuarantinePenalty	0.1	0.1-0.7
proximitySicknessLikelihood	0.05	0.03-0.09
repairEffectiveness	0.34	0.14-0.74
maintenanceEffectiveness	0.4	0.2-0.9

Before running the full parameter sensitivity analysis, a randomness test was performed to see the extent of the randomness after all changes were made. The result of that test, where no parameters were changed, resulted in a standard deviation of 3.18E+8 and an average value of 3.65E+8. The standard deviation is fairly large given the mean of the data set, so while the code was removed of some randomness, there are still more random components impacting the results. This randomness confirms that it was a good choice to run the simulation multiple times for each parameter.

The analysis of parameters is performed to determine which parameters are most affected by the score. If changing a parameter increases the score, the simulation set up is less effective and more money is spent or more crews had to quarantine. If the score decreases, the parameter change positively affects the grid.

Parameter Sensitivity Analysis

Parameters with No Correlation

For each parameter, a record of the parameter values tested and the subsequent scores after running the simulation were taken. A graph was then made to compare each parameter value and average score. A trendline was found for analysis of the existence of a correlation or lack thereof. If there was an expected correlation, then the parameter was deemed relevant. After testing 16 parameters, the parameters with no correlation were found to be the following:

- Number of Crews
- Psigma
- Cascade failure status
- Ptol
- Repair Effectiveness
- Maintenance Effectiveness

Some of these parameters were expected to have more of an impact on the final score, but may have shown no correlation due to the long runtime of each parameter's simulation and the time constraint for the project. Due to this, there is not enough data to prove a correlation. The following parameters were shown to have either an expected or unexpected correlation and therefore will be evaluated for full analysis.

Last Turn

The last turn parameter determines how many times the simulation reevaluates. The simulation was run with values of 25 to 35 to see how the change in length of the simulation changes the score. The output of the simulation runs is shown in Figure 1. The results show that there is an increase in score with an increase of turns. This matches intuition because the longer the pandemic is, the more damage can be done. The more often crews have to go repair lines and risk exposure, the more likely they are to have to quarantine.

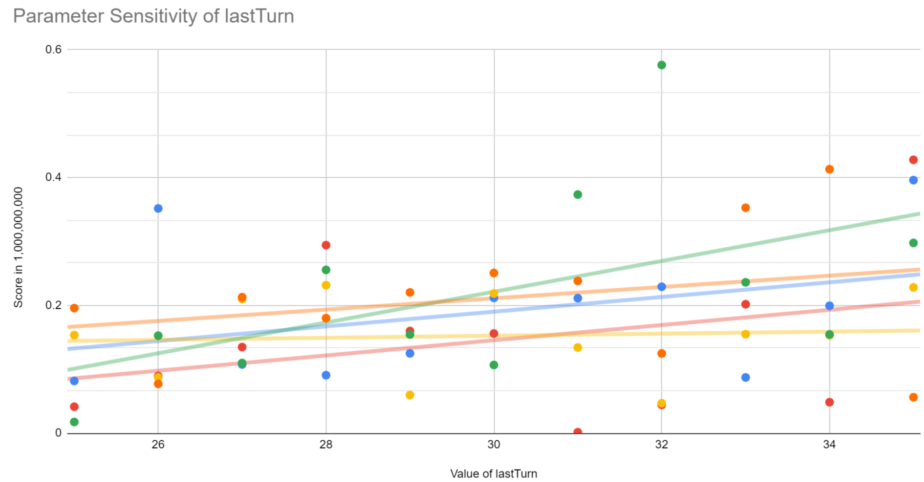

Figure 1. Parameter sensitivity of lastTurn showing increasing scores with increasing turns

However, this is only the case when the other values are left constant. In real life, as a pandemic goes on, there is a period where the likelihood of getting sick increases, making it even riskier to send out repair crews. This change would amplify the problem.

The results show that with more time, there will be more risk of quarantine and higher costs. Based on this parameter, the recommendation for operation is to limit the number of crews exposed to the pandemic if the time scale is longer. While the experimental setup set the probability of failure threshold at 0.01, a higher value would prevent crews from unnecessary exposure. It could also be beneficial to increase the threshold as time goes on and the risk increases. A strategy of prioritizing maintenance of lines in the early days of the pandemic to prevent continued failure later in the pandemic with crews becoming more limited could be successful as well.

Pmu

The value Pmu is the multiplicative offset of the random term driving load demand variation. In simpler terms, with increasing Pmu, there is increasing load demand variation. The values for Pmu for each run and the resulting scores are shown in Figure 2. There is a general positive trend of values, meaning the simulation becomes less effective as the value increases. However, there is also a lot of variation. Based on this initial graph, there is a use to see the variation of scores for each value of Pmu. The standard deviation of scores for each Pmu is shown in Figure 3.

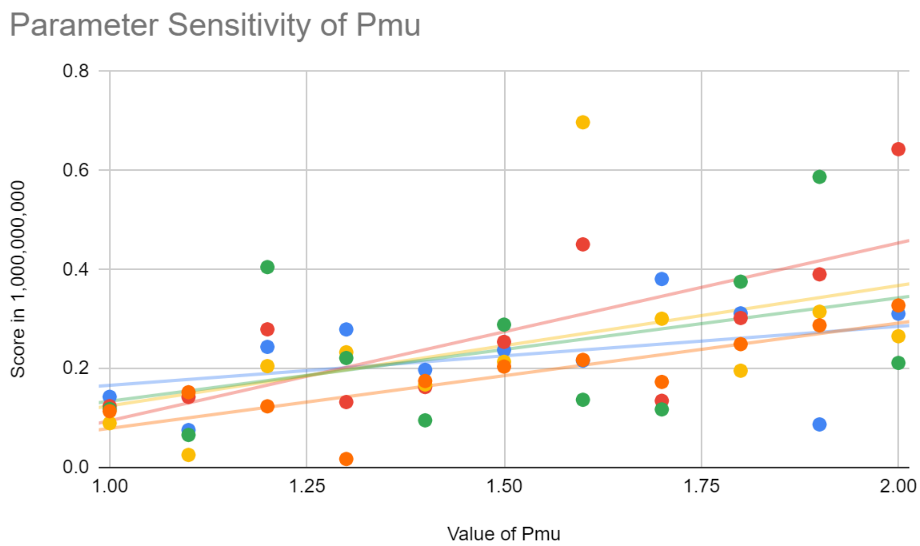

Figure 2. Parameter sensitivity of Pmu showing a positive trend

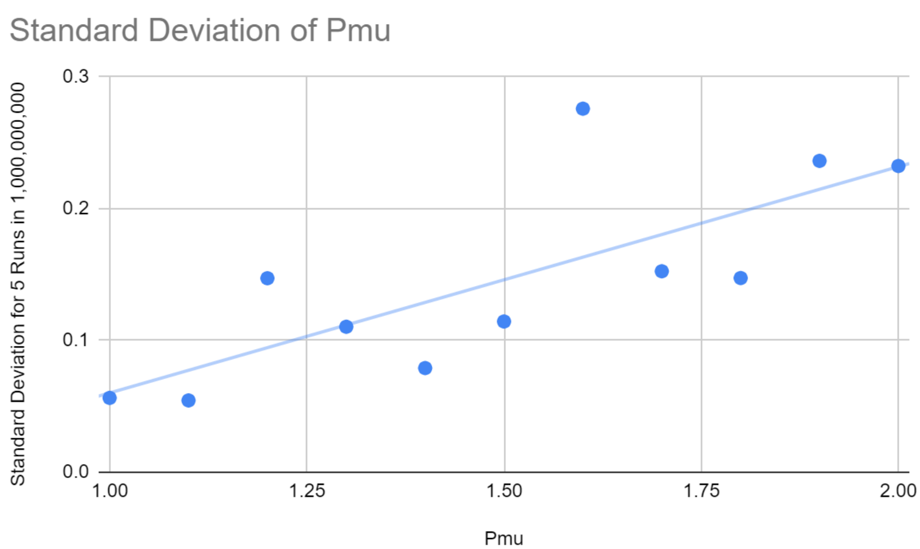

Figure 3. Standard deviation of Pmu showing a positive trend

The results from Figure 3 show that there is in fact a general increase in variation of final scores. While there is a general trend, there are not enough data points to draw a definitive conclusion. Since Pmu also impacts the load variation, evaluating a value other than score may show a more clear trend.

The recommendation for operation based on this is to make efforts to limit randomness with increasing Pmu. In a real life grid operation setting, it is impossible to eliminate all randomness, but it is possible to decrease randomness in crew allocation and other controllable factors. Decreasing randomness does not automatically award a better score, but it does decrease the amount of surprises that grid operators face in a pandemic setting.

Minimum Failure Probability

The minimum failure probability determines the floor of the possible failure probabilities. If a line is newly repaired or maintained, it will start with this minimum chance of failure. When the simulation begins, lines are randomly assigned failure rates between this value and the failure probability value. Based on the initial hypothesis, increasing this value should degrade the performance, and decreasing it should allow for better scores in the simulation. The results are shown in Figure 4.

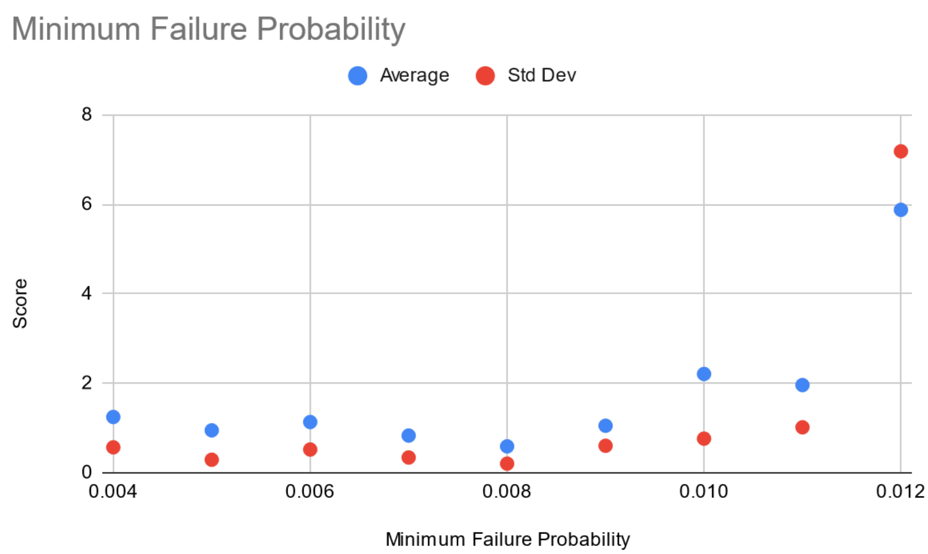

Figure 4. Average and standard deviation values for Minimum Failure Probability

This data does not present a very clear indicator that the minimum failure probability directly affects the outcome of the simulation. One conclusion that can be gathered from this simulation is that increasing the minimum failure probability close to the value of failure probability creates more randomness and has the potential for a higher score, although not explicitly guaranteed.

The recommendation based on this data is for the system operator to closely monitor the failure rates of newly repaired lines based on simulation, as well as monitor the failure rates of currently aging lines. If newly repaired lines are failing anywhere close to the average failure rates, then the system operator should be sure to be very cautious making repair and maintenance decisions as there is an increased risk of severe failures.

Fail Probability

The failure probability determines the maximum initial failure probabilities for lines. When the simulation begins, lines are randomly assigned failure rates between the minimum failure probability and this value. Based on the initial hypothesis, increasing this value should allow for worse scores, and increase the standard deviation of all scores. This value was tested with a wider range than minimum failure probability, ranging from .01 to .028 while minimum failure probability remained the same each time. The results are shown in Figure 5.

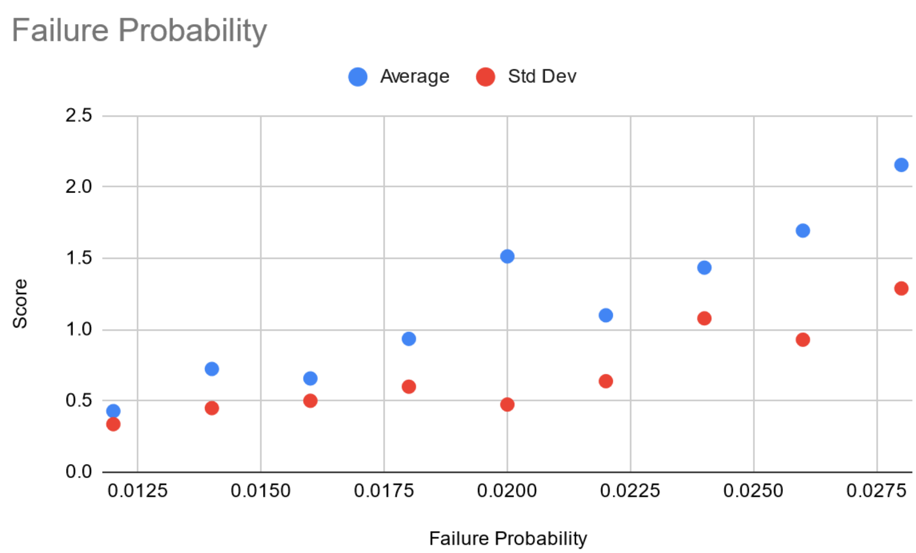

Figure 5. Average and standard deviation values for fail probability

This data does agree with the initial hypothesis. There is a clear trend that increasing this parameter allows for worse scores as well as higher standard deviations. The increase in standard deviations indicates that this parameter affects the randomness of the score and does not simply increase the score (although a score increase is observed). This has a similar effect to minimum failure probability, and is a very straightforward correlation.

The recommendation to system operators is to maintain accurate data on failure probabilities to create an accurate expectation for pandemic operations. Operators should choose suppliers for components that ensure the lowest failure rate. Operators should also monitor the failure rates of lines during a pandemic scenario, and if the maximum failure rate is increasing randomness should be expected.

Failure Rate Increase

The failure rate increase parameter determines how much the line failure probability increases each turn. The initial failure probabilities are assigned at the beginning of the simulation, and this constant multiplicative factor increases it every additional turn. If the parameter is set to a high value, lines will degrade quicker and be more likely to fail. This parameter should show an overall increase in score as the value is increased. This test was conducted over a range of 1 percent to 10 percent increases, since it is not realistically possible for the factor to be less than 1. The results are shown in Figure 6.

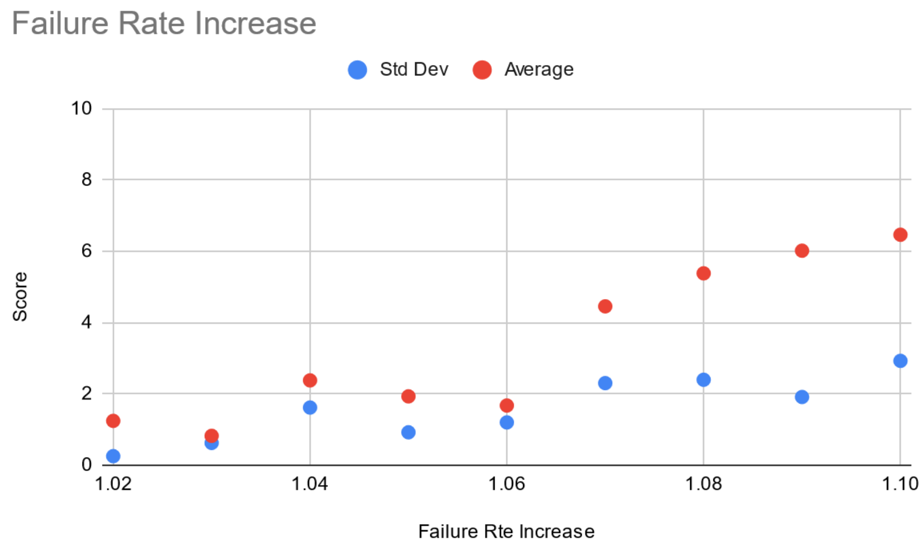

Figure 6: Average and standard deviation values for failure rate increase

This test overall shows that a higher failure rate increase leads to higher scores. The standard deviations mostly remained the same, so this indicates the parameter does not cause as much randomness as other factors discussed in this report. This experiment indicates a direct correlation between failure rate increase and score increase, and indicates that system operators should be concerned if line degradation rates increase.

The recommendation to system operators would be to monitor the lifespan of line components, and collect data on how their failure rate increases over time. System operators can expect to have almost guaranteed increases in costs if they use components that degrade quicker, and there is very little increase in randomness. A 10 percent increase in line degradation could triple the cost of operating the system, thus this parameter is very important to keep track of.

Crew Quarantine Penalty

The crew quarantine penalty is the penalty factor percentage for quarantined crews due to the poor impact quarantine has on the quality of life of line workers. The product of the crew quarantine penalty and the number of crews quarantined at the end of the simulation is added to the final score. This means that if there is a high penalty percentage, there will be a larger addition to the final score compared to if there were a lower penalty percentage, given the same amount of quarantined crews. Therefore, if this parameter is relevant to the score, the trend should show that as crew quarantine penalty increases, so does the score. The range tested for the cascade failure status was 0.1 to 0.7 at a 0.1 interval. The base value was 0.1. Each value was run for 5 turns. Each individual run is shown in Figure 7 and the average is shown in Figure 8.

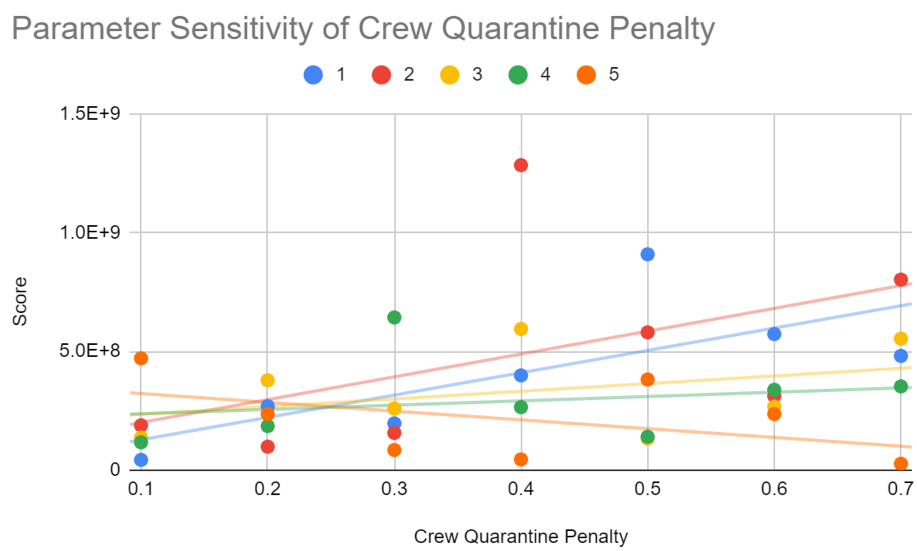

Figure 7. Parameter sensitivity of crew quarantine penalty showing positive correlation

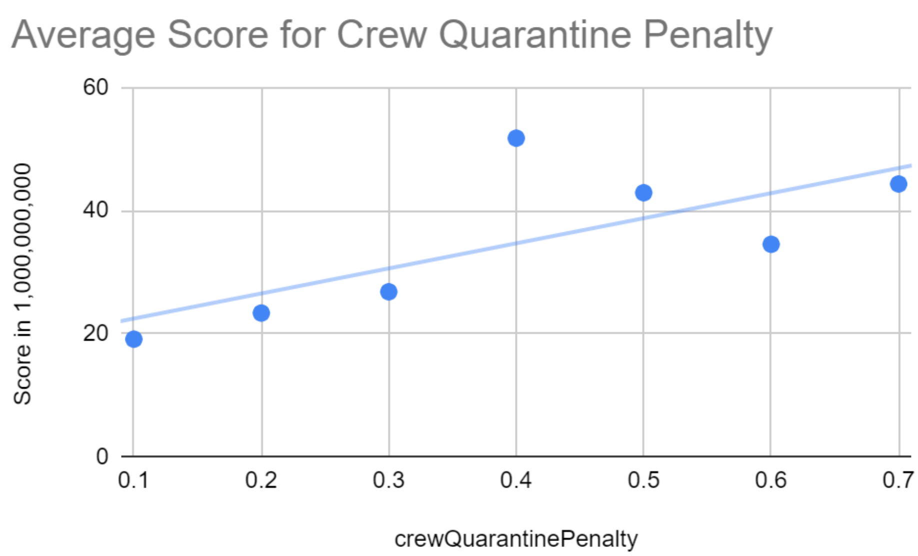

Figure 8. Average score for crew quarantine penalty

The results of the simulation display that the average scores increase as the crew quarantine penalty increases, confirming the hypothesis. It can therefore be concluded that there is a correlation between crew quarantine penalty and final score. With this parameter so influential in the overall score, variables affecting the crew quarantine penalty should also be considered. Therefore, minimizing the number of crews exposed to the disease should be prioritized as this value can lead to a large penalty when multiplied with the crew quarantine penalty.

The recommendation for operation based on this is to consider limiting the number of crews to critical power outages and limited maintenance to decrease the risk of exposure and maintain crew morale. With the transmissibility of the disease unknown, a possible high crew quarantine penalty due to the poor impact on the crew's quality of life, such as unfavorable symptoms and low morale, could have a substantial effect on the total cost of generation and load shedding. In a real life grid operation setting, it is impossible to predict where or when exposure will take place, but by allocating crews in a way that reduces the risk of exposure, a lower chance of crew absenteeism is possible making maintenance and repair faster in the long term.

Proximity Sickness Likelihood

The proximity sickness likelihood is the probability that an allocated crew will be exposed to disease and have to quarantine for the rest of the simulation. The higher the probability, the more transmissible the disease is. If multiple crews are working on the same line and one crew gets exposed, all crews will be sent to quarantine. This means that if there is a high proximity sickness likelihood, then the final score should be higher as less crews will be available to restore power and there will be a larger crew quarantine penalty. Therefore, if this parameter is relevant to the score, the trend should show that as proximity sickness likelihood increases, so does the score. The range tested for the cascade failure status was 0.03 to 0.09 at a 0.01 interval. The base value was 0.05. Each value was run for 5 turns. The results are shown in Figures 9 and 10.

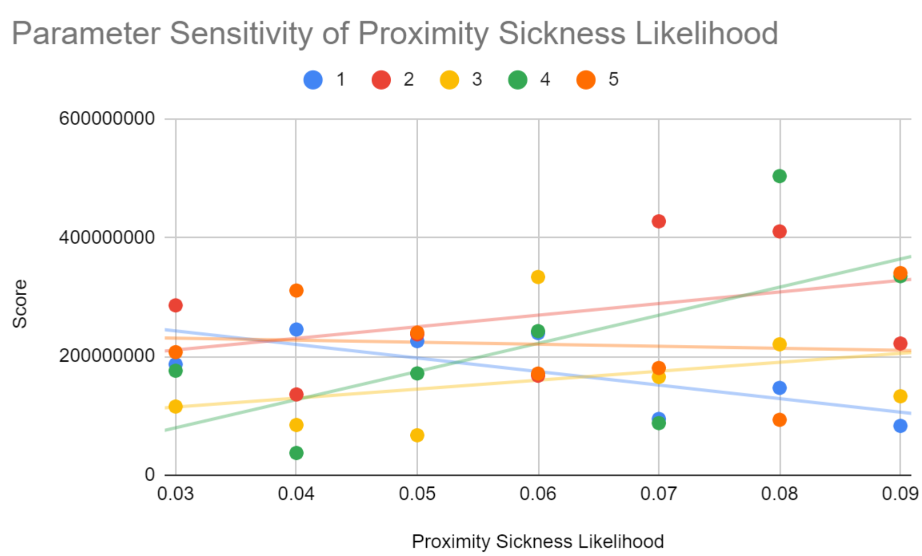

Figure 9. Parameter sensitivity of proximity sickness likelihood

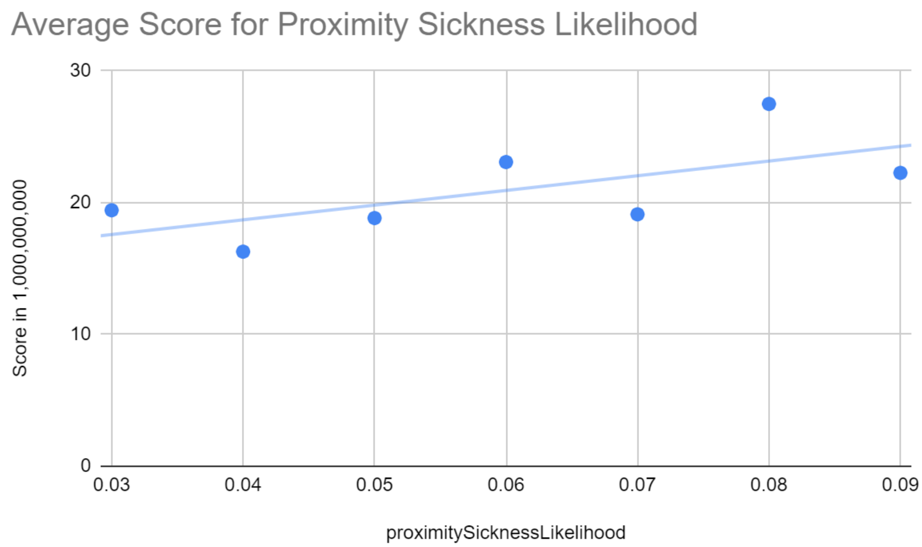

Figure 10. Average for proximity sickness likelihood

The results of the simulation display that the average scores increase as the proximity sickness likelihood increases, which confirms the hypothesis. We can therefore conclude that there is a correlation between proximity sickness likelihood and final score. This means that if the disease is highly transmissible, there is a greater chance of an increased cost of generation and load shedding as more crews will have to quarantine due to exposure.

The recommendation for operation based on this is to consider limiting the number of crews to critical power outages and limited maintenance to decrease the risk of exposure to COVID, similar to the crew quarantine penalty. Additionally, allocating only one crew to an area would limit the number of people infected if exposed as if one person contracts the disease, everyone working on that line has to quarantine. In a real life grid operation setting, it is impossible to predict how transmissible the disease is, so though these suggestions would lead to a longer time to restore power, the long term benefits of having an available crew for future turns is more valuable to the total cost of operation.

Limitations to Analysis

While this analysis gives grid operators a good starting point for operation procedures during a pandemic, there are limitations due to the assumptions made and the limited resources and time.

The first assumption was that the simulation models Georgia Power's grid. The simulation provides great insight into how different parameters affect the grid, the actual Georgia grid is much more complex. This parameter analysis could not be applied to other systems without further evaluation.

Another assumption was made that the score would be a fair value by which to judge the parameters. While the results of the scores provided insight into many parameters, the analysis could benefit by digging deeper into what changed about the results. The score is composed of both the quarantine penalty and the cost. However, when only analyzing scores, it is unclear in many cases which of those two values were more impacted.

The results are also limited by the team's ability to run simulations. For example, doing 10 runs would be too time consuming for the team to complete as each parameter would take 40 minutes to complete. Instead, only five runs for each parameter were done. Even then, many parameters required multiple runs as changing the value too much would make the simulation crash. In addition, there were a maximum of 11 values tested per parameter. An analysis of more values could give different solutions.

Finally, better code and data analysis could be performed. Due to the limited knowledge of the three team members and the limited time to complete the project, sacrifices were made in the code. Updating the entire setParameters file was not the most efficient way to change variables and most likely contributed to increased run times. While the data was sufficient for this project, if decisions were made based on the analysis, much more data would be needed to make a confident recommendation. In addition, the team had limited knowledge of data analysis. Averages and standard deviations were used to evaluate the data, but a more sophisticated analysis could lead to different results and recommendations. In a different setting, subject matter experts would be brought in to assist with the analysis.

Conclusion

In conclusion, to prevent power grid failures during a pandemic scenario, a multi-faceted approach should be used. Throughout our analysis, we found that multiple parameters affected the sensitivity of the system. In a general sense, utilities should maintain data on line failure rates and lifespans, as well as run simulations to estimate the effects of different failures. In the early stages, system operators should run simulations and acquire testing data relating to both the grid and pandemic. Utilities should cooperate with local health officials and share data on both disease transmission rates as well as weaknesses in the power grid. Accurate data on line failure rates, average component lifespan, and other relevant factors should be collected. Utilities should limit their maintenance to critical failures and lines most at risk of failure or monetary loss. Repair crews should be allocated scarcely, but can be fine tuned based on disease transmission rates and quarantine lengths. Although our analysis provides a window into the ideal pandemic operating procedure, more testing should be performed, potentially with more data science, power engineering, and healthcare opinions considered. The simulation should be adjusted to provide more accurate metrics to assess what specifically utilities need to address, as well as more simulations performed on higher powered computers to increase efficiency.

References

Aziz, A., & Than, A. (n.d.). *Covid-19 implications on Electric Grid Operation*. IEEE Smart Grid. Retrieved December 8, 2022, from https://smartgrid.ieee.org/bulletins/may-2020/covid-19-implications-on-electric-grid-operation

Wormuth, B., Wang, S., Dehghanian, P., Barati, M., Estebsari, A., Filomena, T. P., Kapourchali, M. H., & Lejeune, M. A. (2020, November 30). *Electric power grids under high-absenteeism pandemics: History, context, response, and opportunities*. IEEE access : practical innovations, open solutions. Retrieved December 8, 2022, from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8545261/

Appendix

Appendix A - Crew Allocation Code

```
%% Allocate the crews to do maintenance/repair
crewAllocation = zeros(ncrew,1);
% Get the indices of the lines in order of decreasing failure probability
failure_probability = mpcs{currentTurn}.failure_probability;
[~,failure_sort] = sort(failure_probability,'descend');
threshold = 0.01;
counter = 1;
for i=1:length(failure_probability)
   if (failure_probability(i)>threshold)
       if (counter<=ncrew)
           crewAllocation(counter) = i;
           counter = counter + 1;
       end

   end
end
```

Appendix B - Change Parameters and Run the Simulation

```
display_gui = false;
scores = zeros(5,10);
for i = 1:5
   n = 1;
   for val = 30:40
       fid = fopen('setParameters.m','w+');
```

```
        fprintf(fid,"function [opts, lastTurn, ncrew] =
setParameters"+newline+"lastTurn = 30;"+newline+"ncrew =
10;"+newline+"opts.Psigma = 0.2;"+newline+"opts.Pmu =
0.2;"+newline+"opts.TimesOfDay = [1; 9; 19];"+newline+"opts.StartWeek =
"+string(val)+";"+newline+"opts.minFailureProbability =
0.0008;"+newline+"opts.failProb = 0.02;"+newline+"opts.outlierFailProb =
0.08;"+newline+"opts.outlierFraction = 0.05;"+newline+"opts.FailureRateIncrease
= 1.01;"+newline+"opts.instaFail = 1.20;"+newline+"opts.cascadeFailureStatus =
0.67;"+newline+"opts.Ptol = 0.1;"+newline+"opts.crewQuarantinePenalty =
0.1;"+newline+"opts.proximitySicknessLikelihood =
0.05;"+newline+"opts.repairEffectiveness =
0.34;"+newline+"opts.maintenanceEffectiveness = 0.4;"+newline+"end");
        fclose(fid);
        [mpcs,Pd,mpcs_specified,Pd_specified,crewAllocation,ncrew,score] ...
            = power_system_pandemic('benchmark_approach',display_gui);
        scores(i,n) = score;
        disp(i)
        disp(score)
        n = n+1;
    end
end
```