Operating a Power System During a Pandemic

Ethan Taylor, Zachary Remick
School of Electrical and Computer Engineering
Georgia Institute of Technology
ethan.taylor@gatech.edu
zremick3@gatech.edu

***Abstract*– This is a recommendation report intended to provide operating procedures for maintaining and repairing a power system during a range of pandemic scenarios. These recommendations are based on a power system simulation related to varying load demands, line failures, and crew quarantine probabilities. This simulation provides insights on how to handle future situations similar to what we are currently experiencing with the COVID-19 pandemic.**

I. Introduction

With the following information provided in this report, we hope that power systems across the country will be given insight on how to handle decisions regarding different pandemic scenarios.

Pandemics have a broad impact on electric power grids, as well as almost every other industry across the nation. System reliability is one of the largest issues that pandemics force upon power grids [5]. This unreliability comes in the form of a harsh impact on the workforce, making it more difficult for crews to perform necessary tasks [5]. Along with this, utilities face varying and unpredicted load demands, as well as a decrease in overall revenue [5]. Some utilities were more prepared than others for the impacts of COVID-19. Hopefully the insights from this report will encourage and promote future planning.

The findings in this report will focus on three different pandemic scenarios. These scenarios consist of negligible risk of crew exposure, moderate risk of crew exposure, and high risk of crew exposure. Any crew exposed is required to quarantine and thus not available to be utilized for the remainder of the simulation. The previously mentioned crew exposure probabilities are 0.1%, 5%, and 35% respectively.

Given these scenarios, the power system will be altered in order to achieve the lowest cost of operation, while also keeping the number of quarantine crews as low as possible. In order to achieve this, security constraints, line flow limits, and crew allocations will be altered depending on the pandemic scenario. Recommendations for these alterations will be outlined in the following report.

The data used in this report will come from performing a simulation of a power system modeled over the state of Georgia, consisting of 30 transmission lines and 20 buses. The simulation steps through 30 time periods.

II. Base Case

This section of the report looks into establishing a baseline for the results of the simulation. For this part of the experiment, the probability of crew exposure was set to 0.1%. This probability should resemble what a crew might experience while not in the midst of a pandemic.

Also, crew allocations were consistent across all of the simulations for the base case. Crews were randomly allocated to repair lines, weighted by how close they were to being repaired, while attempting to have enough crews per outaged line to fully repair the line in one time period. The remaining crews were then allocated to the remaining lines to perform maintenance, weighted by failure probabilities.

A. Security Constraints

Altering the line contingencies resulted in different outcomes of the simulation. For the base case, security constraints of *N-0*, *N-1*, and *N-2* were tested. For this part of the report, line flow limits and crew allocations did not change.

All results in the following table are calculated by averaging the results of 10 different simulation runs, each containing 30 turns. The final score is determined by summing the generation cost and load shedding cost and then multiplying that sum by a penalty based on the number of crews quarantined at the end of the simulation. This is how all data will be reported for the rest of the report.

TABLE 1
AVERAGE RESULTS OF VARYING SECURITY CONSTRAINTS FOR LOW RISK OF CREW EXPOSURE

Security	Final Score ($)	Total Load Shedding (MW-turn)	Average Line Outages (per turn)	Number of Remaining Crews
N-0	3.11E+7	4,287.9	2.426	7.303
N-1	1.11E+7	1,014.6	0.1166	9.9
N-2	5.54E+8	50,216.2	0.0833	9

The results from Table 1 show almost all crews remained after each simulation, which was due to the very small probability of exposure. We can see that *N-1* security performed the best for our power system. The load shedding for *N-1* security was kept relatively low compared to *N-0* and *N-2* security. The simulated power system is set up to handle *N-1* contingencies, so it makes sense that this setting gave us the best result.

B. Line Flow Limits

For this simulation, the probability of a line failing increases as the line flow surpasses what the line is rated for. This is the case until the rated flow is surpassed by 40%, which causes the line to instantly fail.

TABLE 2
AVERAGE RESULTS OF VARYING LINE FLOW LIMITS FOR N-1 SECURITY

Percent Over Initial Limits	Final Score ($)	Total Load Shedding (MW-turn)	Average Line Outages (per turn)	Number of Remaining Crews
10%	6.738e+6	592.98	0.12667	10
20%	1.300e+7	1,331.4	0.2	9.8
30%	1.2988e+7	1,515.5	0.19667	9.8

In order to determine the impact line flow limits have on the system operation, all of the line flow limits were relaxed to allow greater flow than what they are rated for. To test the effects this had on the system, the flow limits of all lines were relaxed to 110%, 120%, and 130% of the initial flow limits. To keep things simple, the system was kept at *N-1* security and followed the crew allocation specifications as stated before.

The results in *Table 2* show us that relaxing the line flow limits can improve or impair the results of the simulation. If we relax the limits by 10%, we see a significant improvement in the final score and load shedding, decreasing roughly by a factor of two compared to Table 1. However, after relaxing the limits past 10%, it seems that the system levels out at a consistent final score but increases in load shedding.

Given this information, we can assume that it would be beneficial to relax the line flow limits of the system by a small amount - around 1%-10%. This would be fine to do for the base case considering the other constraints. Since the probability of crew exposure is so low, we can focus on allocating as many crews that we have available to repair and maintain lines, meaning that we can relax the flow limits up to a certain point with little or no consequences.

This could have benefits seen across the system. Perhaps if a line goes down, the surrounding lines will be able to pick up some of that slack. Also, this slightly relaxed flow limit might better suit the system during a short period of consumption spikes, perhaps during the mornings. *Figure 1* shows what these spikes look like.

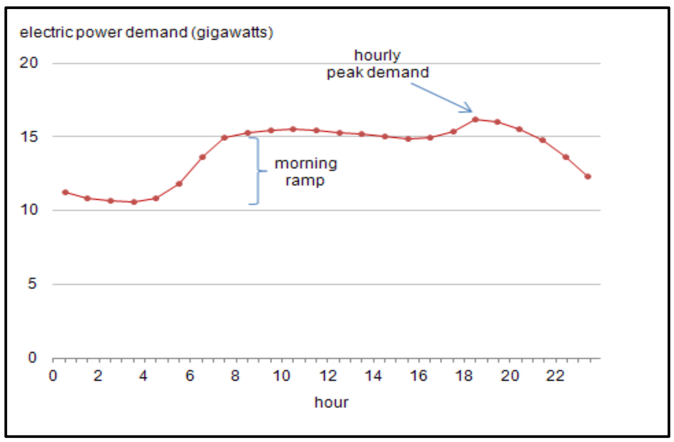

Figure 1. Plot showing the hourly electric load curve on 10/22/2010 [3].

The above base case was examined to see how our fictitious power system would operate under normal scenarios. With the increasing severity of the pandemic, we are likely to see differing results. In the following sections of the report, we will examine how to best allocate crews, what kind of security constraints to run, and how far we should relax the line flow limits.

III. Moderate Probability of Exposure

For this section of the report, a case where the crews have a moderate risk of exposure will be examined. This probability of exposure will be 5%. This means we are expecting to see more crews exposed, thus having fewer crews remaining at the end of the simulation. This result could change depending on how crews are allocated. For this case, different security constraints, line flow limits, and crew allocations will be tested.

A. Security Constraints

In order to see how the increased probability of exposure alters the results of the simulation, we can keep the line flow limits set to their initial values and keep the crew allocations the same as Section II. However, we will see how operating the system at *N-0*, *N-1*, and *N-2* security affects the simulation for moderate risk of crew exposure.

TABLE 3
AVERAGE RESULTS OF VARYING SECURITY CONSTRAINTS FOR MODERATE RISK OF CREW EXPOSURE

Security	Final Score ($)	Total Load Shedding (MW-turn)	Average Line Outages (per turn)	Number of Remaining Crews
N-0	3.814e+7	2,460.9	0.61667	2
N-1	2.795e+7	1,401.5	0.29667	1.9
N-2	9.907e+8	52,106	0.32222	1.5

The results from this test show us how an increased crew exposure rate impacts the system operation. For all security scenarios, we see how hard the crews are impacted while using a normal technique to allocate them to various lines for maintenance and repairs. With a moderate risk of exposure, we only have roughly 20% of crews available after the simulation ends. This is not a sustainable method to operate the power system, which tells us changes have to be made in how crews are allocated.

Due to the number of crews decreasing rapidly through the simulation, we also see a slight jump in the *Final Score* and *Total Load Shedding.* This is due to not having a sufficient number of crews to repair lines during one turn or enough crews that are available to perform maintenance to decrease the probability of failure for certain lines.

Once again, we see that operating at *N-1* security results in the most ideal outcome for the system. Since this is the case, we will operate the system in this way for moderate risk of crew exposure.

B. Crew Allocations

To begin correcting the negative results of the increased probability of exposure, crews must be allocated to repair and maintain lines in a different manner. In order to determine a proper solution, line flows were held at their initial limits and the system was operated at *N-1* security.

By altering the way we allocate crews, we can see a large jump in the number of crews remaining at the end of the simulation. By simply stopping the general maintenance of lines, we were able to double or even triple the number of crews remaining at the end of the simulation, which is up from an average of 1.9. The crews were still randomly allocated to repair lines, weighted by how close they were to being repaired, while attempting to have enough crews per outaged line to fully repair the line in one time period.

Although we had more success with keeping crews un-exposed with this method, we saw an increase in the average load shedding by two to four times. This is likely because we stopped performing preventative maintenance and thus suffered from an increase in line failures. The average numerical results of operating the simulation in this manner are displayed in the following table.

TABLE 4
AVERAGE RESULTS OF NEW CREW ALLOCATION METHOD

Final Score ($)	Total Load Shedding (MW-turn)	Average Line Outages (per turn)	Number of Remaining Crews
6.2955e+7	4,105	0.70667	5.1

We can also see from *Table 4* that the *Final Score* did not increase as dramatically as the *Total Load Shedding* did. This is because we had many more crews remaining at the end of the simulation, which does not impact the final score as dramatically.

C. *Line Flow Limits*

Just as in the base case, relaxing the line flow limits impacts the results of the simulation. It seems that there is a sweet spot for this specific system, which was found to be line flow limits set to 10% above their initial values. This resulted in the best outcome for both this case and the base case.

TABLE 5
AVERAGE RESULTS OF NEW CREW ALLOCATION METHOD

Percent Over Initial Limits	Final Score ($)	Total Load Shedding (MW-turn)	Average Line Outages (per turn)	Number of Remaining Crews
10%	6.3769e+7	4,374.1	0.82667	6.6

The results from *Table 5* are promising for the system. We see a very slight increase in the average *Total Load Shedding* and *Final Score*. We see a significant increase in the average number of remaining crews, up almost 30% compared to no relaxation of the line flow limits. This is probably due to the ability to push the lines slightly harder before failure, which means that we could hold off longer before having to send crews out to repair the lines.

Running the power system this way, we are able to achieve a much more sustainable operating point for a moderate risk of crew exposure. Although we see a rise in the average final score and load shedding, we are able to sustain the number of crews we have at the end of the simulation.

IV. HIGH PROBABILITY OF EXPOSURE

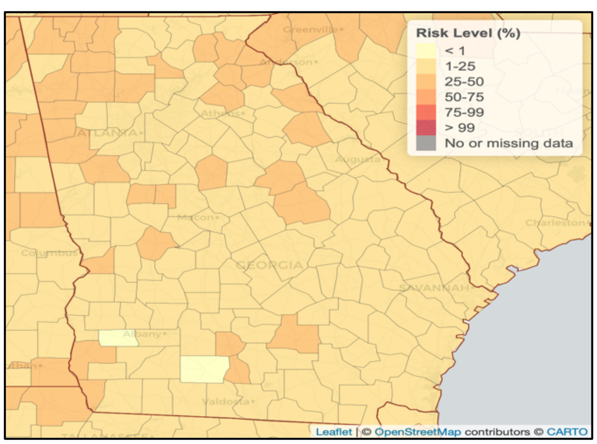

Figure 2. Map showing the risk level of gatherings of 15 or less people in the state of Georgia for COVID-19 [2].

For this section of the report, a case where the crews have a high risk of exposure will be examined. This probability of exposure will be 35%. For this case, the crews face a much larger risk while performing maintenance or repairs on outaged lines. As shown in *Figure 2*, this case is very relevant to the current situation we are facing with COVID-19 [2].

For this scenario, utilities will have to make larger changes in the way they operate their systems. Crews will be at a high risk of exposure, meaning that utilities could see a steep decline in their workforce very quickly. This could be very impactful on utilities that have smaller workforces and lower numbers of available crews to begin with.

A. *Line Analysis*

In order to begin with the simulation of this scenario, we need to perform some upfront analysis of our system. To do this, we are examining the line flow limits (in MW) that each line is set with.

TABLE 6
ORIGINAL LINE FLOW LIMITS BY LINE

Line Number (1-15)	Flow Limit (MW)	Line Number (16-30)	Flow Limit (MW)
1	150	**16**	420
2	230	**17**	150
3	720	**18**	340
4	650	**19**	240
5	400	**20**	430
6	1020	**21**	430
7	570	**22**	160
8	690	**23**	360
9	470	**24**	570
10	710	**25**	240
11	270	**26**	230
12	420	**27**	430
13	190	**28**	200
14	180	**29**	470
15	830	**30**	410

Since we know what the lines are rated for, we make assumptions on what lines are the most important depending on high they are rated. We can also determine which lines are not important depending on how low they are rated.

We can see from *Figure 3* that the highest rated lines are those connected to Atlanta: Line 6, Line, Line 10, and Line 15.

This makes sense because these would be the most densely populated areas in the state of Georgia. The reverse of this is true for what would be the less densely populated regions across the state, which would be the plains region.

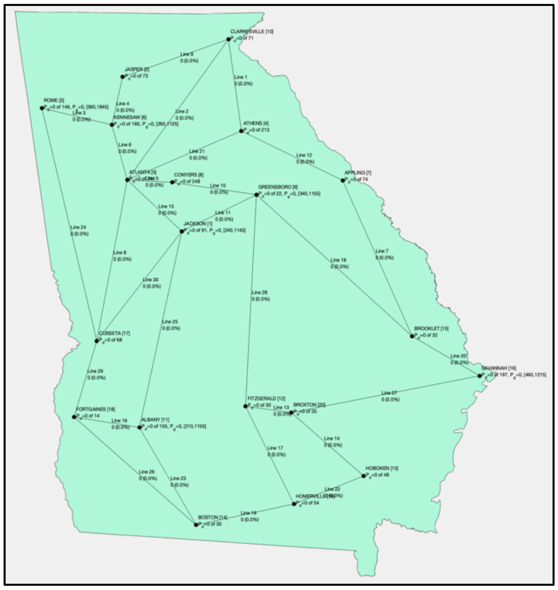

Figure 3. Map of the fictitious power system model of the state of Georgia.

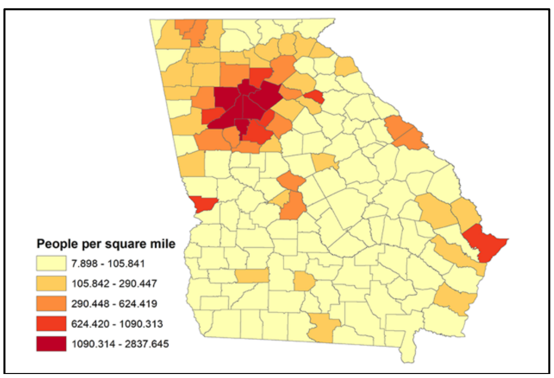

Figure 4. Map showing the population density of the state of Georgia [4].

If we overlay *Figure 3* and *Figure 4*, we can see that the lines with the higher limits are on located in the same geographical areas as the higher populated regions. This information will help us determine which lines should have high and low priorities.

B. *Crew Allocations*

For the high probability of crew exposure scenario, we are still going to halt preventative maintenance on all lines. This will help keep our crews unexposed by not performing tasks that are not crucial to the operation of the system. We are also going to utilize the information about which lines are important and not important, which we discovered in the previous section.

In order to allocate crews for this scenario, we are going to place the highest priority on the lines that are rated for the highest power flow. These are the lines that we will dispatch crews out to immediately upon failure. We will do this because failure of the more important lines could result in greater load shedding and a higher likelihood of a cascading failure across the system. The lines rated with the lowest power flow will be put on the backburner of the system. These lines will not be top priority to be repaired. The failure of these lines should be less catastrophic to the system.

C. *Line Flow Limits*

Given the extreme circumstances of this test case, the system could benefit by relaxing line flow limits by 10% - 20% greater than their initial values. Due to all buses in the system having multiple branches connected to them, if one connected line fails then we can rely on the other connected lines to pick up that load. For this to possibly happen, the greater relaxation of the line flow limits will be necessary.

Figure 5 shows the hourly demand curve before and after COVID-19. Although this simulation does not reflect what is shown in the figure, it is something that we can take note of. According to the figure, power consumption has decreased while experiencing COVID-19 [1]. This is a trend that will most likely be present during future pandemics as well.

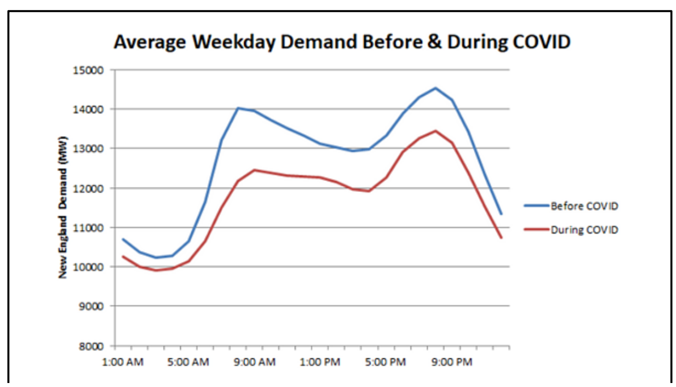

Figure 5. Plot showing the hourly demand curve for time periods before and during COVID-19 [1].

D. Security Constraints

In order to keep things consistent, the power system will continue to be operated at *N-1* security during the high probability of crew exposure scenario. With more test time and simulation results, we could perhaps determine a better contingency plan to operate this system at. However, for the previous two cases, *N-1* security has performed the best.

E. Results

In order to operate the system in the way that was previously described, we could no longer run the simulation with a function-based approach. The function-based approach allowed us to get more accurate results by averaging many different simulation runs. Instead of using this method, here we were required to step through each time period of the simulation and tailor how the system was operated.

With relaxed line flow limits set to 115% of their initial value and the system set to *N-1* security, we mainly altered how crews were allocated depending on the assumptions made in the *Line Analysis* section.

Using the ideas surrounding the difference of important and un-important lines, we were able to achieve a larger number of remaining crews but suffered from increased load shedding. If we allocated crews to repair important lines after failure, we ended up having around half of the crews remaining at the end of the simulation but suffered from load shedding between 10,000 MW - 20,000 MW. However, we did not see any cascading failures if we were able to repair the lines quickly, within one or two turns.

A slight spin on this technique was to not repair adjacent lines entering the same bus. This meant that we were sending crews out less often, which resulted in six to eight crews remaining at the end of the simulation. Although this was the case, we suffered from heavy load shedding and the system was more likely to become unstable via cascading failures.

With this knowledge, it is necessary to decide what is more important to your power system. Do you want to operate at a more cost-effective point by preventing load shedding? Or is the health of your crews more important? For every system there is going to be a solid middle ground that needs to be chosen, ensuring that most customers have continuous power while keeping the crews safe.

V. Recommendations and Results

This section of the report will examine the recommendations and results achieved from the findings from running the simulation.

A. Negligible Crew Exposure Risk

The results from our negligible risk analysis help determine a stable operating point when the risk of crew exposure is extremely low, or if there is no pandemic at all. The first step involves what kind of security to run on your power system. For the given simulation, *N-1* security provided the best results. However, this could vary across different systems. Also, contingency constraints could only be applied to specific parts of the system instead of the entire system, which could result in different outcomes.

Next, find the sweet spot to operate the line flow limits at. Some systems could benefit by larger increase and some could see benefit by small increase in the limits. The system in this simulation saw an improved operating point after increasing the line flow limits to 110% of their initial values. Increasing the line flow limits might take some strain off the system, allowing it to be more flexible.

Lastly, establish a good rule for crew allocations. The rule laid out in the base case seemed to work great for this system. Make sure that all lines currently down are being repaired with as much manpower as possible, while also sending out remaining crews to perform preventative maintenance on the system.

Although the purpose of this report was not to give recommendations on how to operate a power system in a non-pandemic scenario, it provides great intuition that is useful in determining what actions to perform when a pandemic is present.

B. Moderate Crew Exposure Risk

While dealing with a moderate risk of crew exposure during a pandemic scenario, not much has to be altered from the base operation of the system. We continued to operate this system at *N-1* security with line flow limits set at 110% of their initial values. This operating point was chosen because it seemed to work the best, resulting in minimal load shedding and a lower number of crews becoming exposed. Of course, this is likely to vary depending on the system.

The most change took place in how crews were allocated to perform repairs and maintenance. For this scenario, preventative maintenance on all lines were halted in order to keep crews unexposed. A line repair model similar to the base case was used, which involved randomly allocating crews to repair lines, weighted by how close they were to being repaired, while attempting to have enough crews per outaged line to fully repair the line in one time period. This model could be modified by putting a capacity on how many crews could be sent to a single line. With this plan, lines couldn't be repaired as quickly, but we could prevent crews from infecting crews.

Our initial operation plan developed for the base case was able to handle the moderate risk of crew exposure scenario well with minor modifications. With an average of around seven crews remaining at the end of the simulation and relatively low load shedding, this would be a sustainable option for a power system, especially during shorter time periods.

C. *High Crew Exposure Risk*

The operation of the system has to be heavily modified when dealing with high probability of crew exposure. It is crucial for small-scale utilities to get this correct in order to preserve their limited number of crews.

Taking what we learned about the system in the previous two cases, we chose to operate this system at *N-1* security. Also, we knew we would want some more leeway with the line flow limits, so we chose to relax them to 115% of their initial value. This gave us extra wiggle room compared to the sweet spot of 110% we had been using in the other two scenarios. This might be good practice to carry over into other systems, as lines will eventually go unrepaired in this case.

As in the previous case, crews were no longer allocated to preform preventative maintenance on lines. In order to properly allocate crews to outaged lines, a line analysis was performed. We determined what lines might be important and what lines were likely less important. When using this technique on a real power system, there is most likely prior knowledge or intuition that helps distinguish the lines importance. With this information, we could pick and choose where we wanted to send our crews to repair lines.

Important things to consider while allocating crews for repair is how much power a line is carrying in the system. We chose to place a high priority on the highest rated lines, which helped prevent cascading failures across the system. We also chose to ignore lower rated lines that failed, which also were leading into the same bus as many other lines.

Operating the system in this way allowed us to maintain our crews, which is very helpful for a long-term pandemic scenario. However, the system did suffer from increased load shedding, raising the cost to operate the system dramatically. We could have lowered this load shedding by allocating more crews, which would likely result in more crew exposure. Determining the correct operating point during this scenario is going to be very system independent and require balancing between system cost and crew safety.

VI. Conclusion

Almost all utilities have encountered problems while trying to operate during COVID-19. Most modern utilities and people of this generation are walking on uncharted ground when it comes to being prepared for a large-scale pandemic. Hopefully through the understanding of this report, utilities will be provided with some insight on how they might want to conduct operations. We can take the knowledge established from simulations like this, combined with operating experience of those working through COVID-19, in order to form operating procedures that can be utilized during future pandemics.

References

[1] Braga, K., 2020. *Where's The Peak In The Era Of COVID?*. [online] Blog.greenenergyconsumers.org. Available at: <https://blog.greenenergyconsumers.org/blog/wheres-the-peak-in-the-era-of-covid>.

[2] Covid19risk.biosci.gatech.edu. 2020. *COVID-19 Event Risk Assessment Planning Tool*. [online] Available at: <https://covid19risk.biosci.gatech.edu>.

[3] Eia.gov. 2020. *Hourly Electric Load Curve*. [online] Available at: <https://www.eia.gov/todayinenergy/detail.php?id=43295>.

[4] Msm.edu. 2020. *Population Density Of Georgia*. [online] Available at: <https://www.msm.edu/RSSFeedArticles/2020/May/ga-counties-with-more-black-residents-higher-rates-covid-19.php>.

[5] Crsreports.congress.gov. 2020. *COVID-19: Potential Impacts On The Electric Power Sector*. [online] Available at: <https://crsreports.congress.gov/product/pdf/IN/IN11300>.