Air Force Institute of Technology

AFIT Scholar

Theses and Dissertations | Student Graduate Works

3-22-2019

Testing the Fault Tolerance of a Wide Area Backup Protection System using SPIN

Kenneth James

Follow this and additional works at: https://scholar.afit.edu/etd

 Part of the Databases and Information Systems Commons, and the Information Security Commons

Recommended Citation

James, Kenneth, "Testing the Fault Tolerance of a Wide Area Backup Protection System using SPIN" (2019). *Theses and Dissertations*. 2264.
https://scholar.afit.edu/etd/2264

This Thesis is brought to you for free and open access by the Student Graduate Works at AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact richard.mansfield@afit.edu.

TESTING THE FAULT TOLERANCE OF A WIDE AREA BACKUP PROTECTION SYSTEM USING SPIN

THESIS

Kenneth James, 2d Lt, USAF

AFIT-ENG-MS-19-M-034

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENG-MS-19-M-034

TESTING THE FAULT TOLERANCE OF A WIDE AREA BACKUP PROTECTION SYSTEM USING SPIN

THESIS

Presented to the Faculty

Department of Electrical and Computer Engineering

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

in Partial Fulfillment of the Requirements for the

Degree of Master of Science in Cyber Operations

Kenneth James, BS

2d Lt, USAF

March 2019

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AFIT-ENG-MS-19-M-034

TESTING THE FAULT TOLERANCE OF A WIDE AREA BACKUP PROTECTION SYSTEM USING SPIN

THESIS

Kenneth James, BS
2d Lt, USAF

Committee Membership:

Dr. Kenneth M. Hopkinson
Chair

Dr. Scott R. Graham
Member

Dr. Douglas D. Hodson
Member

AFIT-ENG-MS-19-M-034

Abstract

Cyber-physical systems are increasingly prevalent in daily life. Smart grids in particular are becoming more interconnected and autonomously operated. Despite the advantages, new challenges arise in the form of defending these assets. Recent studies reveal that small-scale, coordinated cyber-attacks on only a few substations across the U.S. could result in cascading failures affecting the entire nation. In support of defending critical infrastructure, this thesis tests the fault tolerance of a backup protection system. Each transmission line in the system incorporates autonomous agents which monitor the status of the line and make decisions regarding the safety of the grid. Various malfunctions that could occur from real-life attacks are simulated in the grid with the intent of determining its ability to successfully respond to faults despite adversity. The SPIN model checker is used to execute all combinations of fault location and malfunctions to determine which types can occur, and how many, before the system is unable to properly clear a fault. With results analyzed, the decision making process of the model is revised to increase its fault tolerance.

Table of Contents

Page

List of Figures

List of Tables

TESTING THE FAULT TOLERANCE OF A WIDE AREA BACKUP PROTECTION SYSTEM USING SPIN

I. Introduction

1.1 Motivation

The instantaneous access to electricity with the flick of a switch is a technological marvel. Electricity is delivered to the user at the proper voltage and frequency without so much as an afterthought. Yet, this is only possible through the extensive design and installation of generators, transformers, sensors, switches, and transmission lines which all operate cohesively [3]. Due to our dependency on a reliable energy source for day-to-day activities, severe consequences may result when equipment within this complex network experiences failure. For example, the Northeast Blackout of 2003 caused a few billion dollars in infrastructural damage, leaving an estimated 50 million people without power. There are also social costs to consider that can't be measured with a dollar amount, such as massive traffic congestion, businesses closing, and communication channels being rendered useless [4]. This catastrophe was caused by the domino effect of a few small transmission line failures, which reinforces the conclusions made by a 2014 Federal Energy Regulatory Commission (FERC) study regarding the vulnerability of the US power grid system [5].

While we may not be aware of it, our electrical infrastructure faces constant adversity. This may come in the form of operator error, equipment malfunction, or extreme weather. An additional aspect to consider that is becoming increasingly relevant is the cyber security of our critical infrastructure [6]. New technologies

have been incorporated into power protection systems, allowing autonomous, wide-area networking. However, this comes with associated risks. For example, using digital relays instead of mechanical relays opens the possibility of cyber-based attacks. Compared to threats due to natural causes that at least require physical proximity, cyber intrusions are more uncertain since they can be executed from any range [7].

To combat cases of primary protection failure and cascading effects, various layers of backup protection are interwoven into grids. "Protection" in the context of power engineering does not refer to the ability to prevent abnormal activity from happening. Rather, it refers to the ability to minimize damage and other factors once abnormal conditions have been detected. Equipment is designed such that power systems are resilient in the face of challenge. This is known as protective relaying [8].

Despite the integration of backup protection, power grid systems still have flaws. Researchers and engineers have investigated backup protection systems to improve their resiliency. One way to measure the resiliency of these systems is through adversarial thinking. Adversarial thinking, the strategic approach of putting oneself in the shoes of the attacker, has been heavily studied with regards to cyber security and game theory [9, 10]. While power grid systems don't always face threats from a human attacker, this method of thinking still applies. By predicting what types of danger these systems can face, then evaluating all possible outcomes, resiliency can be assessed and improved.

1.2 Problem Statement

Backup protection systems employing autonomous protective relays have varying degrees of tolerance. Testing a system's fault tolerance with an exhaustive search of failure scenarios reveals weaknesses. When taking into consideration the inter-dependencies between autonomous agents, this task is combinatorically massive. In

this thesis, the SPIN model checker is used to explore all possible failure scenarios of a backup protection system. With patterns analyzed, recommendations are made to improve the resiliency of the system.

1.3 Research Questions

In this thesis, the following three questions are investigated.

- RQ1: *How can the model checking approach be applied to test the fault tolerance of a backup protection system?*
- RQ2: *What specific vulnerabilities lie within the backup protection system?*
- RQ3: *Using the model checking approach, what improvements can be made to the model to increase its resiliency?*

1.4 Organization

This thesis is sectioned as follows.

Chapter II provides a background on the various types of protection relays found within power grid systems and how they function, the specific backup protection system being examined, and the application of SPIN. It also discusses related work regarding other uses of model checkers.

Chapter III examines the methodology of how the fault tolerance of the grid is tested. Specifically, what variables are used in the experiment, the steps that SPIN takes when executed, limitations of the model, and assumptions made.

Chapter IV discusses results of this experiment. The results are analyzed to determine what combinations of malfunctions are required to cause the backup protection system to fail. With failure conditions examined, a revision is made to the decision making process of the model to improve its fault tolerance.

Chapter V provides concluding remarks on the research accomplished, states potential avenues for future work, and discusses the research questions presented in 1.3.

II. Literature Review

This chapter presents a thorough background on the topics relevant to this thesis; namely, backup protection and the model checking approach. It covers the use of protection relays, the specific backup protection system used for this research, the application of model checking, and SPIN.

2.1 Protection Relays

Protection relays are embedded within electric distribution circuits that are designed to detect power system conditions of an abnormal or dangerous nature, and report the activity [11]. They are typically located in substations and protect the transmission lines connecting those substations. Protective relays are configured to detect specific types of dangerous activity, such as a short circuit or voltage dip. Intelligent Electronic Devices (IEDs) are autonomous agents which compile information about current, voltage, or frequency from these relays, and make decisions regarding the safety of the power grid [12, 13]. This information is communicated across agents to ensure wide-area protection [14].

When IEDs sense any abnormal state on a transmission line, they send a message to the circuit breaker. The circuit breaker interrupts the flow of current through the line. With no current, the affected area is isolated, preventing any further equipment damage [15]. This is known as "tripping" a line. There are two types of protective relays that are relevant to this research. These are distance relays, and directional relays.

2.2 Distance Relays

Distance relays function by observing the impedance of a line. Impedance is a measure of the opposition a line provides to the current flowing through it when voltage is applied. Since impedance and current are inversely related, a drop in impedance results in a rise in current. If a line's impedance falls below a certain threshold, a distance relay determines that a fault has occurred. Within the context of this experiment, a fault is assumed to be a short circuit.

Impedance is not constant. It varies due to many factors, such as length and temperature of the line. As a result of constant variation, different zones of protection are employed in distance relays. There are primary and secondary zones, which are referred to as Zones 1 and 2.

Traditionally, a Zone 1 distance relay is set to cover the first 70-90% of the line length [16, 17]. When a fault is detected, the relay reacts instantly and sends a command to trip the line. The distance covered intentionally falls short of the entire length of the line. Since impedance varies, a Zone 1 relay set to cover exactly 100% of the line length may over reach and detect a fault on the next line. In this scenario, the relay would mistakenly trip the local line, where a fault hasn't occurred. Setting the reach to 70-90% accounts for this variability.

To ensure protection of the ends of each line, Zone 2 relays are used. These are typically set to cover 100% of the local line, plus 20% - 50% of the shortest adjacent line. Zone 2 relays operate with a short time delay, which allows Zone 1 relays enough time to clear the fault before reacting [18].

For the purposes of this experiment, Zone 1 relays have been set to cover 80% of the line length. Zone 2 relays have been set to cover 100% of the line length, as well as 20% of an adjacent line. These relays are configured to be uni-directional. Since IEDs on either end of a line in the grid face inward, the middle 60% of each line is

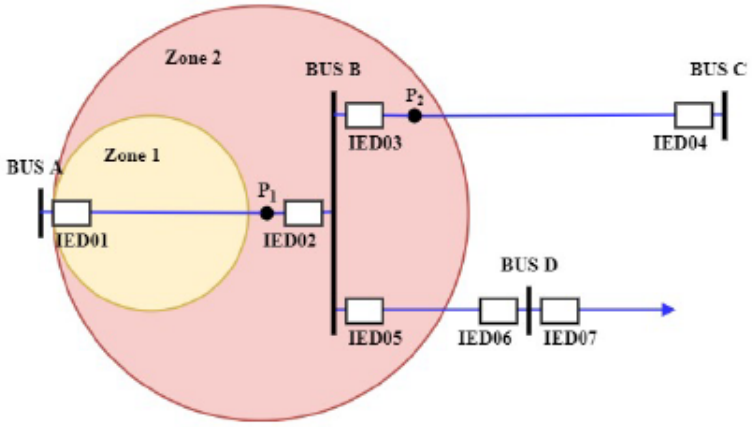

Figure 1. Zone Distance Relays [1]

a region protected by overlapping Zone 1 distance relays. The first and last 20% of each line is not protected by overlapping Zone 1 distance relays, but does fall within the range of Zone 2 distance relays from adjacent lines.

Figure 1 is an illustration of distance relay protection. If a fault were to occur at P1, it would fall outside of the radius of IED01's Zone 1 distance, but within the Zone 2 distance. P2 is an example of a fault within a Zone 2 distance, but on an adjacent line.

2.3 Directional Relays

Current can flow in either direction in a power grid system. When a fault occurs in a grid, the current will always flow towards the location of the fault. Since the role of the IEDs is to detect these faults, directional relays are used to indicate which way the current is flowing. The direction of the current is referred to from the perspective of an IED, either down the line or towards the bus. For example, suppose a single fault occurred in Figure 1, at P1. IEDs 01, 02, 04, and 06 would detect the current as flowing down the line. IEDs 03, 05, and 07 would detect it flowing towards the bus.

2.4 Backup Protection

To minimize the risk of cascading failure described in 1.1, power systems often employ backup protection [11, 19]. This typically comes in two forms: local and remote backup protection.

With local backup protection, multiple relays of the same function operate in parallel in the same station. A station could have multiple Zone 1 relays, as an example. In doing so, all that is required for the correct operation to occur is for one of the devices to function properly. The proximity of redundant devices minimizes delays in the case of malfunction. However, this proximity has drawbacks as well. While some malfunctions that systems face will affect devices individually, others will have a greater scale. For example, extreme weather could disable an entire station, making the use of local redundant devices ineffective.

Remote backup protection presents a balance to the strengths and weaknesses of local protection. As the name implies, this form of backup protection involves devices located at remote stations. An application of this has already been discussed, in the form of Zone 2 relays. Having overlapping zones of responsibility originating from different transmission lines increases the chance of detecting faults despite malfunction. By design, remote protection incorporates a time delay to allow local relays to react first. This could be a disadvantage, because a fault can cause equipment damage in the time it takes for the remote relay to react, if the primary protection fails.

Local and remote backup protection are used together, which gives a power grid flexibility in protecting its assets. These redundant processes are not implemented with the intent of preventing malfunction in a power system, but rather to increase tolerance. In this context, tolerance refers to the ability of a backup protection system to sustain varying levels of malfunction before the system can no longer properly respond to a fault.

2.5 The Wide-Area Backup Protection System

As described in previous research, the cyber-physical system examined in this experiment is a "regional decentralized peer-to-peer negotiating WABP multi-agent system" that incorporates "local and adjacent line, first and second zone, distance protection and directional protection systems as well as fault states from additional lines" [2]. The topographical structure of the WABPS was designed in resemblance of the IEEE 14-Bus System, as seen in Figure 2. However, certain features of the 14-Bus system have been abstracted away.

For example, the 14-Bus system is composed of 14 buses and 17 transmission lines. It also includes the use of generators and transformers. The model used here, displayed in Figure 3, condenses some of the buses and lines. The dotted squares are reduced to a single substation each. Additionally, the generators and transformers have been removed. These abstractions have been made to minimize the effect of the state-space explosion problem, which the model checking approach is vulnerable to. Designing the most simplistic model as possible to verify for correctness is imperative when using a model checker [20, 21, 22]. The WABPS modification of the IEEE 14-Bus system contains 10 buses (referred to as substations in the figure) and 15 lines.

Each of the 15 lines in the system has an IED positioned on either end, resulting in a total of 30 IEDs. Each IED has a corresponding Line Decision Agent (LDA), and each bus has a Regional Decision Agent (RDA). As the name suggests, these agents make decisions on how to react when a fault in the grid occurs. Each LDA collects information from the IEDs on its local and adjacent lines, through coordination with the RDA, and performs calculations to determine whether or not the local line is faulted. Adjacent in this context means that two lines share a bus. For example, lines 1 and 3 are adjacent.

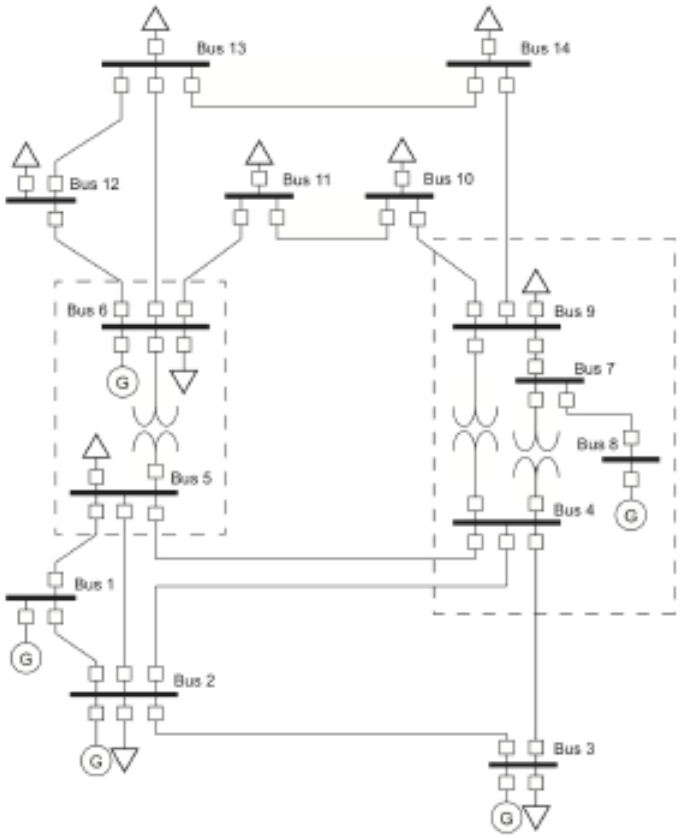

Figure 2. IEEE 14-Bus System [2]

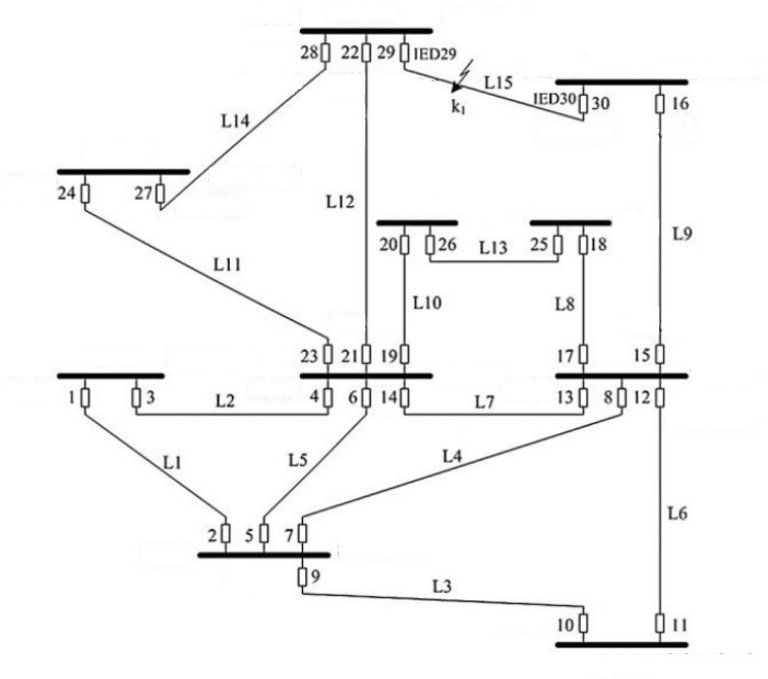

Figure 3. WABPS Modified IEEE 14-Bus System [2]

More information about the role of LDAs and RDAs in the WABPS can be found in [2]. For the purposes of this research, certain aspects of LDAs and RDAs have again been abstracted away. For the WABPS model used here, IEDs utilize the three protection relays discussed: the Zone 1 distance relay, Zone 2 distance relay, and directional relay. In the scenario of a fault, the two IEDs on a line work together as an LDA. The use of RDAs can be obfuscated for the intent of this research. Figure 4 displays how IEDs and LDAs are implemented in this model of the WABPS.

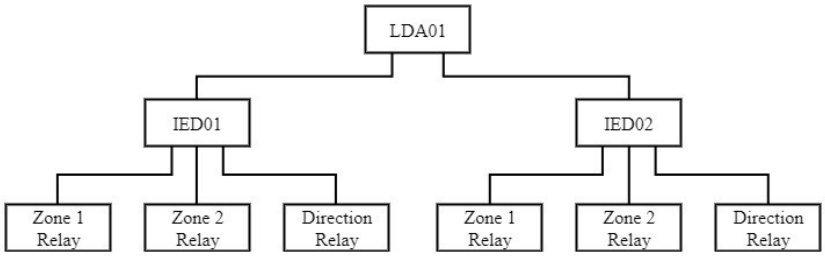

Figure 4. Relationship between IEDs and LDAs [1]

2.6 Model Checking

Model checking is often used to support hardware and software development. For complex technologies, especially those that are interdependent with other systems, it is infeasible for a human to test all possible scenarios. The use of a model checker allows engineers to examine how their systems operate under given sets of specifications. This can expose troublesome situations that the engineer wasn't otherwise aware of, such as deadlock or race conditions. While software errors may be seen as an annoyance, there certainly are scenarios in which undetected bugs led to catastrophe. For example, the Ariane-5 missile crashed less than minute after launch in 1996, due to an unintentional integer overflow in the control software. Similarly, Intel suffered from a bug in the calculation of a floating point unit in the early 90s. This mistake cost about $475 million in processor replacements, as well as their reputation [21].

As a result of the damage caused by undiscovered software bugs, extensive research has been conducted regarding the application of model checking techniques. This idea of conducting thorough verification of complex software and hardware systems is commonly referred to as formal methods. Formal methods are backed by both the Federal Aviation Authority (FAA) and National Aeronautics and Space Administration (NASA). In an investigation conducted by FAA and NASA, the parties state that "Formal methods should be a part of the education of every computer scientist and software engineer, just as the appropriate branch of applied maths is a necessary part of the education of all other engineers."

Many companies presently follow that recommendation and incorporate formal methods into their business. Amazon, for example, applies the model checker known as TLA+ to Amazon Web Services [23]. SPIN specifically has been used in many real world applications. For example, NASA has applied SPIN to verify critical software in Mars Exploration Rovers and Deep Impact, the former of which is an ongoing mission.

The software has also been used in an investigation regarding a malfunction of the 2005 Toyota Camry, as well as in verifying medical device transmission protocols [24].

2.7 SPIN

There are several types of IED malfunctions, many IEDs which could be affected, and many locations in the grid where faults could occur. Evaluating how the WABPS reacts to each possible scenario of malfunctions is a combinatorically massive feat, for which a model checking approach is appropriate.

SPIN, a term which was initially created as an abbreviation of Simple PROMELA Interpreter, is one such model checker. SPIN is typically not used to convey the meaning of the acronym; rather, it is a stand-alone term for the name of the model checker. PROMELA, similarly an abbreviation of Process MetaLanguage, is used to model nondeterministic behavior in distributed systems. The creator of SPIN, Gerard Holzmann, claims that "PROMELA is not a programming language. It is a language for building verification models [20]."

Unlike conventional programming languages, PROMELA does not execute in a linear fashion. For example, in C or Java, the first true condition under an 'if' statement will be executed. In PROMELA, these conditions are known as guard statements. Any true guard statement can be executed, regardless of the order in which they are written. Guard statements are often used to modify the same variables. Due to shared memory and the nondeterministic manner in which these statements are selected, different executions of PROMELA code will lead to different results. There are two different modes in which SPIN can be used: simulation and verification.

In simulation mode, SPIN randomly selects a certain set of conditions and tests them against a predefined verdict. If the system fails to meet the verdict under the given set of conditions, a trail file is created that specifies exactly which steps were

taken that led to the failure. The user can also manually specify the set of conditions met to see the results of that scenario. For example, suppose that in a PROMELA file, there are two variables, x and y. Both x and y are integers that can take the values between 1 and 100. The verdict set for the program is that both variables are even numbers. In simulation mode, SPIN could randomly select $x = 43$ and $y = 10$. SPIN would then compare these values to the verdict, determine that the scenario failed, then leave a trail file mentioning which values for each variable were selected. This is one simulation run. This, of course, is a very simple example of SPIN might be used. In practice, SPIN would be used to evaluate much more comprehensive models.

In verification mode, SPIN conducts a simulation run for each possible combination that the model can attain. As the name implies, it is used to verify the correctness of a model. Keeping with the current example, SPIN would evaluate all combinations of $1 \leq x \leq 100$ and $1 \leq y \leq 100$.

While SPIN (and the model checking approach in general) provides great utility in verifying models for correctness, there are limitations that one should take into consideration. One limitation, as mentioned previously, is the state-space explosion problem. System verification can be extremely resource intensive, as investigated in [25]. In the simple example with x and y, there are 10,000 possible simulation runs to evaluate. In a model with more variables, there could very easily be several million possible states. Beyond the issue of simply waiting a long period of time for verification mode to finish, the computer could run out of memory before completion. With this in mind, one should take caution when using SPIN on a complex model. Possible solutions would be to verify small pieces of the model individually, or to introduce abstractions. This presents an issue of balance, as a model should be detailed enough to convey the properties of the real system, yet not so detailed that it is too large to be evaluated thoroughly.

This balance between detail and abstraction in a model is a segue to the next limitation, which is that this approach of system verification only verifies the system model, not the actual system. For this reason, one should be careful when designing the model to ensure that it accurately represents the system. Limiting factors and abstractions made in the model should be thoroughly explained. As the authors of [21] state, "any verification using model-based techniques is only as good as the model of the system." Additionally, when analyzing and making conclusions from the results of model checking, one should avoid extrapolation. Patterns observed from using this approach could be predictive of real world behavior, but are not necessarily indicative.

2.8 Related Work

The model checking approach is commonly used in various areas of research in present day. Petri Nets are one such example, in which mathematical modeling is used to verify system correctness. While not quite as powerful as other forms of model checkers, they have been used in research regarding power system analysis [26, 27, 28]. Bayesian networks, also used for the research of fault diagnosis for power systems in [29], are another verification tool. Other researchers employ Markovian Models [30, 31], or PRISM [32] for similar purposes.

SPIN has been widely recognized and applied since its introduction in 1991. It was awarded the ACM System Software Award in 2002 [24]. The most recent version was released in December 2018. It's been used by researchers in recent years for deadlock detection in the scheduling of autonomous vehicles [33], as well as to verify the status protocol of network nodes [34].

III. Methodology

This chapter covers how the testing of the WABPS is conducted. In this experiment, different types of component and IED failures are simulated in the grid to determine whether or not the WABPS is able to properly locate a fault for a number of scenarios. The component and IED failures resemble errors that could occur from real-life factors. Limitations of the model are stated, as well as assumptions made to reduce the problem space.

3.1 Variables

There are two independent variables in this experiment: component failures and total IED failures. The dependent variable is how many times the WABPS incorrectly diagnoses a fault for each combination of component/IED failures.

Component Failures.

Component failures are representative of real life malfunctions that could occur due to transmission error, faulty equipment, malicious activity, or any other reason. For example, an IED fails to detect a fault where one occurs, an IED detects a fault where one does NOT occur, or an IED doesn't collect any data at all [13]. A component failure is simulated in the grid by changing the initialized state of a single relay for an IED to an incorrect state. The possible states for protection relays are listed in Tables 1 and 2. An example of a simulated component failure would be changing an IED's Zone 1 distance relay state from "FAULT" to "NO_FAULT."

Directional Relay States

LINE_FAULT	Line side fault detected
BUS_FAULT	Bus side fault detected
NO_FAULT	No fault detected
NO_DATA	No data received

Table 1. Directional Relay States

Zone Distance Relay States

FAULT	Fault detected
NO_FAULT	No fault detected
NO_DATA	No data received

Table 2. Zone Distance Relay States

Total IED Failures.

Total IED failures are representative of real life scenarios where an IED is rendered completely inoperative. These situations could occur as a result of extreme weather, cyber-based attacks, or vandalism [35, 36]. Since a total IED failure resembles a situation in which the device cannot contribute any information, it is simulated in the grid by changing the states of all three relays for an IED to "NO_DATA."

3.2 Configurations

This thesis directly follows the research of [13] and [1]. In the former, SPIN was applied to see how the model checking approach could be used to assess the reliability of the WABPS. Only Line 15 was tested, with the assumption that it was representative of any line in the grid. In the latter, the model was improved by incorporating more potential relay states to examine resiliency. Additionally, all 15 lines in the grid were tested. In both of these previous works, one of the limiting assumptions was that faults occurred only within the overlapping Zone 1 relays on each line. This research specifically focuses on the more vulnerable regions outside

of the overlapping Zone 1 relays. Since there are 15 lines, and two regions on each line to test (the first and last 20% of each line), there are 30 total line segments in which faults are simulated to occur. Each location is notated by the line number and either the letter "a" or "b." "a" refers to the end of the line with the lower numbered IED, while "b" refers to the end with the higher numbered IED. For example, Line 1a refers to the region on Line 1 on the end with IED 1. A portion of the WABPS, specifically Line 1, is shown again in Figure 5 for the reader's convenience.

For each of those 30 locations, SPIN is run in verification mode, as described below, for all combinations of 0 to 3 component failures and 0 to 3 total IED failures. This is to examine how resilient the WABPS is in face of different levels of malfunction. The range of 0 to 3 is used because it gives a good spread of results. As one might expect, the WABPS fails in almost every trial if it experiences too many malfunctions. Since a failure rate of nearly 100% is resource intensive and not useful to the researcher, the combination of failures is kept to a low to moderate rate.

The number of scenarios in which the WABPS misdiagnoses a fault for each region and each configuration of component and IED failures is recorded. It should be noted that a scenario is recorded when any line in the grid is misdiagnosed, not necessarily the faulted line. For example, suppose a fault is simulated in the region of Line 6a. If the WABPS correctly identifies that fault, but mistakenly identifies a fault somewhere else on the grid, it is counted as a scenario failure.

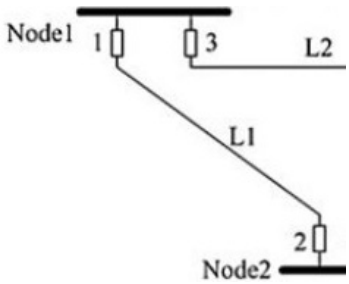

Figure 5. Line 1

3.3 Simulation Run

This experiment is conducted by running SPIN in verification mode, which iterates through all possible simulation runs. Each run is performed in the following manner.

1. A fault is assumed to occur at a specified location in the WABPS (Line 1a-15b). The protection relays for all 30 IEDs are initialized in PROMELA to their proper values in response to the fault. For example, suppose the fault is simulated within the Line 1a region. IEDs 1 and 2 would both have the directional relay set to "LINE_FAULT." IED 1 would have both distance relays set to "FAULT." IED 2 would have the Zone 1 relay set to "NO_FAULT," and the Zone 2 relay set to "FAULT." The other 28 IEDs are set according to their position relative to the fault.

2. The conditions to verify for correctness are set. Each of the 15 LDAs has a verdict indicating whether its corresponding line contained the fault or not.

3. A specified number of protection relays experience a component failure. A random IED is selected. A random relay is selected for that IED. Then, that relay is changed to a random value. This is repeated for the specified number of component failures.

4. A specified number of IEDs experience a total IED failure. A random IED is selected. All three relays for that IED are changed to "NO_DATA." This is repeated for the specified number of total IED failures.

5. PROMELA calculates several factors using the now modified state of all IEDs. Each LDA uses these factors to make a conclusive statement on whether or not its respective line was faulted. The logic behind this algorithm is discussed in detail in 4.1.

6. The decision that each LDA made is then compared to the verdict defined earlier. If all 15 LDAs were correct, the WABPS succeeded in reacting to the fault, and the simulation stops here. If any LDA was incorrect, the scenario is counted as a failure.

7. If the simulation was not a success, SPIN leaves a trail file indicating which configurations led to this result, as well as exactly what steps each LDA took in the decision making process. For example, Line 1a was tested for one component failure and one total IED failure. A code snippet from one of the trail files left by SPIN is displayed in Figure 6. These files range between 200 to 300 lines in length, thus only the first few lines are included. This file indicates that the combination of IED 3's Zone 2 relay changing to "FAULT" and IED 1 turning off was sufficient enough to cause to WABPS to come to an incorrect decision on locating the fault.

This is one execution for a specified region, number of component failures, and number of total IED failures. SPIN continues until ALL possible combinations of relay state and IED failures for the given specification have been selected.

3.4 Limitations

With the SPIN settings used in this experiment, a simulation run is finished once either all predetermined verdicts have been validated, or until one fails. A simulation run will not continue once it is determined that the conditions for a single verdict

```
1 proc 0 = :init:
2 using statement merging
3       L01 State = FAULTED
4       Corrupted Component IED3, sec = FAULT
5       Turned Off IED1
```

Figure 6. Trail File Snippet

have not been met. In the context of this research, if a particular scenario would result in three LDAs coming to the wrong conclusion for their lines, the simulation stops after only one. Not only will SPIN determine the incorrect reading of just one line, but it will do so for whichever line happens to be calculated first.

This has pros and cons. It is beneficial, because it greatly speeds up the process of verification. If a simulation fails, SPIN immediately moves on to the next one, similar to a 'continue' condition in a for/while loop in programming. The downside is that it presents each simulation run with a binary result. Either a run failed, or it did not. This prevents one from investigating the extent of failure to which a simulation run achieves. The WABPS is said to fail the strong correctness test if any line is misdiagnosed, while it fails the weak correctness test only if the faulted line is misdiagnosed. With this distinction, one can differentiate severe failures from those that are merely annoyances. The application of SPIN in this experiment does not separate the two types of correctness failures, so the results may make the WABPS appear to be more vulnerable than it actually is. Future work could entail redefining the SPIN settings and investigating strong vs. weak correctness failures.

This experiment does not incorporate the use of Zone 3 relays, nor does it account for the back reach of distance relays. Both abstractions are made to reduce the problem space. These aspects of relay protection could potentially be added to the model in future work, though it would amplify the currently existing vulnerability of this experiment to the state-space explosion problem.

3.5 Assumptions

All faults are assumed to occur on the ends of each line, outside of the overlapping Zone 1 relays. This assumption is not made to reduce the complexity of the problem, but rather because these regions have not been tested in the previous research. Inves-

tigating how the WABPS reacts to faults in these areas is of interest because it now activates Zone 2 relays from neighboring lines, which previously was not the case.

All transmission lines are assumed to be the same length. This assumption is specifically made to normalize the reach of Zone 2 relays. In practice, Zone 2 relays protect 100% of the local line length plus 20% of the shortest adjacent line, not necessarily 20% of each specific adjacent line. Thus, a short transmission line will have more of its area covered by neighboring Zone 2 relays than a longer transmission line. Given that the IEEE 14-Bus system is a model topography, not a diagram of a physical implementation, the length of each transmission line is undefined. For the simplicity of having each Zone 2 relay cover the same distance, transmission lines in the grid are left as an undefined, yet equal, length in this experiment.

Lastly, directional relays are assumed to be able to detect faults in the WABPS with no regard to proximity. A directional relay implemented in a real world application will naturally have a prescribed range of effectiveness. The signal received from faults occurring outside of that operating range will grow increasingly weak as the distance grows, to the point that a directional relay will be unable to detect that a fault has occurred at all. Incorporating an effective range for directional relays would run into the same issue as before; the transmission lines do not have a defined length. Thus, directional relays in this experiment are intended to detect any fault within the grid despite location.

IV. Results & Analysis

		Line 13b			
		Component			
		0	**1**	**2**	**3**
Total IED	**0**	0	35	9,938	69,226
	1	5	2,353	99,149	108,067
	2	138	46,874	109,597	128,622
	3	1,825	116,325	129,087	160,365

Table 3. Line 13b Results

Table 3 displays results for Line 13b. While the experiment was completed for all thirty regions, only one is included for brevity. Beyond the aspect of brevity, there is no need for further data compilation. Analysis can be conducted with minimal data, considering that one line in the grid is representative of any line. The x-axis represents the number of component failures specified for a verification check in SPIN, while the y-axis represents the number of total IED failures. The numbers in the table are the amount of simulation runs in which the WABPS was unable to properly diagnose all lines in the grid for the given configuration. For example, SPIN used verification mode with the specification of one component failure and one total IED failure and found 2,353 simulation runs in which the WABPS was unsuccessful in reacting to a fault in the Line 13b region.

The process in which PROMELA determines whether or not the WABPS can properly clear the fault is entirely deterministic. Thus, one should expect to receive the exact same results upon repeated verification runs. However, there may be variations in the number of failures reported across simulations. This is primarily due to the mathematical complexity of this research problem. Certain configurations of malfunctions will have hundreds of thousands of potential scenarios. In cases like these, the computer could very well likely run out of memory and terminate before

all scenarios have been evaluated. Despite imperfect testing conditions, results should be similar enough between verification runs to draw conclusions.

4.1 Black Box Approach to Analysis

A few observations can be made from inspecting these results. First, the combination of zero component/IED failures led to zero scenarios in which the WABPS was unsuccessful. This should be expected, since the system is supposed to function perfectly when there aren't any malfunctions. While it seems trivial in this case, this detail can be useful early in the model checking process, when ensuring that the model designed satisfies the properties for correctness.

The table above is useful for giving a quantifiable value of how resilient the WABPS is with varying levels of stress. However, equally significant are the trail files left by SPIN, due to their utility for the researchers. For example, consider the effect that each of the two variables have on the outcome. When comparing the effect of component failures vs. total IED failures, it's clear that component failures result in significantly more simulations where the WABPS is unable to function properly. This initially seems counter-intuitive, as a single total IED failure affects three protective relays, while a component failure only affects one.

The reason for this stems from how each variable affects the system. A total IED failure changes all three relays to "NO_DATA," while a component failure can change the relay to any other value. This indicates that LDAs are much more likely to come to the incorrect decision as a result of receiving false information than receiving no information at all.

An examination of the trail files shines some insight as to why this occurs. One behavior observed is that each LDA incorporates the least amount of information possible when making its decision. If an LDA is confident on the status of its line

after compiling the information from its two IEDs, then the decision making process stops here. An LDA only consults neighboring lines if it did not gather enough information from its two IEDs to confidently decide.

In application, this isn't necessarily a bad thing. If a fault occurs on a line, the circuit breaker should be notified as quickly as possible to trip the line, so that damage to the equipment is minimized. Within the scope of this experiment, this tendency to decide as quickly as possible is a noticeable issue. It leads to simulations in which an LDA receives false data and immediately decides, whereas it could have come to the correct decision had it gathered information from adjacent lines. This is a strong conclusion that can now be used to improve the resiliency of the model.

There remains many conclusions to be made from analyzing the trail files. For example, which types of component failures are more likely to cause a failed simulation run, or comparing the number of false positives to false negatives. However, the black box approach to analyzing results was only used in this section to show that the model checking approach makes it possible to do so. Given that the internal workings of the WABPS are known to the researcher, it is much more practical to take a white box approach to analysis, which is discussed next.

4.2 Decision Making Process

Up until this point, the research conducted has been in setting up an appropriate model such that the WABPS can be evaluated with the model checking approach. Next, the results of this initial phase of the experiment are used to improve the model to make the system more fault tolerant. To do so, a further examination of the logic behind each LDA is required. This section covers Step 5 of the simulation run in 3.3. For clarity, it is broken up into three steps.

Finding the Fault Value.

First, each IED calculates an Action Factor (AF). The AF is an aggregate of the information that each IED collects from its three relays. A higher AF indicates a higher threat of a fault on the IED's local line.

- If both the directional and Zone 1 relays of an IED detect a fault on the local line, then $AF = 2$, regardless of the status of the Zone 2 relay.
- If both the directional and Zone 2 relays of an IED detect a fault on the local line, but the Zone 1 relay doesn't, then $AF = 1$.
- If the directional relay does not indicate a fault on the local line, and neither of the distance relays detect a fault, then $AF = -2$.
- If there is any other combination of the three protection relays, then $AF = 0$.

AF	**Directional Relay**	**Zone 1**	**Zone 2**
2	LINE_FAULT	FAULT	ANY
1	LINE_FAULT	NO_FAULT	FAULT
-2	BUS_FAULT or NO_FAULT	NO_FAULT	NO_FAULT
0	Any other combination		

Table 4. Action Factor Calculation

The calculation of an IED's Action Factor is summarized in Table 4. Each LDA takes the sum of the two AFs determined on its line. This new factor is the line's Fault Value, $\mathrm{F}_{\mathrm{out}}(i)$. From here, $\mathrm{F}_{\mathrm{out}}(i)$ is used to determine the state of the line. There are four possible states - *fault*, *suspect*, *special*, and *normal*. *Fault* and *normal* are end states, while *suspect* and *special* are transitional states.

- If $F_{\mathrm{out}}(i) > 2$, the line is in the *fault* state.
- If $F_{\mathrm{out}}(i) = 2$, the line is in the *suspect* state.

- If $0 \leq F_{\text{out}}(i) < 2$, the line is in the *special* state.
- If $F_{\text{out}}(i) < 0$, the line is in the *normal* state.

When a line is in one of the end states, the LDA is confident enough to make a decision regarding the status of its line, and PROMELA moves on to Step 6 in the simulation run. Either a fault has occurred on the line (*fault*) or one has not occurred on the line (*normal*). If the line is in one of the transitional states, the LDA was unable to gather sufficient information from its local IEDs, so it consults neighboring lines before coming to a decision.

Resolving the *Special* state.

If a line is in the *special* state, each IED on that line calculates a Certification Factor (CF), displayed in 1. The CF is a measure of how directional relays on neighboring lines report the flow of the current. If most of the directional relays from neighboring lines indicate current flowing towards the local line, there is more reason to believe that is where the fault is located.

$$CF = \frac{i}{n} = \begin{cases} \geq 0.5 & \textit{convert to suspect} \\ < 0.5 & \textit{convert to normal} \end{cases} \tag{1}$$

The total number of IEDs on neighboring lines is represented by the variable n. The number of those IEDs whose directional relay points towards the local line is represented by the variable i. If both IEDs on a line in the *special* state have CFs that are greater than or equal to 0.5, the line converts to the *suspect* state. If at least one IED on the line has a CF less than 0.5, the line is set to the *normal* state.

Resolving the *Suspect* State.

If a line is in the *suspect* state, the WABPS waits to make a decision until the states of all neighboring lines have been resolved. A process of elimination is then used to determine where the fault occurred. There are three cases that could occur:

- One of the neighboring lines was determined to be in the *fault* state.
- All of the neighboring lines were determined to be in the *normal* state.
- One of the neighboring lines was also determined to be in the *suspect* state.

In the first case, the local line is deemed to be in the *normal* state. In the second case, the local line is deemed to be in the *fault* state. In the third case, the two *suspect* lines need to be compared to each other. The first comparison is the fault value. The line with the larger F_{out} is determined to be the faulted line, while the other is clear. If the two lines have an equal F_{out}, the distance relays of each line's IEDs are compared. The line which has more relays set to "FAULT" converts to the *fault* state, while the other moves into the *normal* state. It is not possible for two neighboring lines with same number of distance relays activated to have made it this far into the decision making process, thus all possible scenarios have been accounted for.

4.3 Types of Failures

Now that the algorithm the WABPS uses has been discussed, the types of system failures can be further examined. Scenarios in which the WABPS incorrectly identified an unfaulted line (false positives) are referred to as Type I failures. Scenarios in which a faulted line was not identified (false negatives) are referred to as Type II failures. This section is divided into the three possible initialization states that a line can have in the experiment.

Faulted Lines.

Consider the malfunctions required for any given faulted line to be incorrectly diagnosed as safe. The only possible scenario in which this could occur is if

$$F_{\text{out}}(i) < 0$$

Suppose that a faulted line does not fulfill this condition, and ends up in the *suspect* state. It cannot be incorrectly diagnosed as *normal* at this point, because that would require a neighboring line to be in the *fault* state. If so, the simulation run would have already been concluded due to a false verdict. Thus, we can narrow down the combinations leading to a false positive. Table 5 shows the starting values for IEDs on a faulted line.

	Directional Relay	**Zone 1**	**Zone 2**	**AF**
IED X	LINE_FAULT	FAULT	FAULT	2
IED Y	LINE_FAULT	NO_FAULT	FAULT	1

Table 5. Initialization of IEDs on the Faulted Line

For $F_{\text{out}}(i) < 0$ to be true, a minimum of three malfunctions must occur. Initially, IED X has $AF = 2$, while IED Y has $AF = 1$. If IED X has its directional relay changed to either "BUS_FAULT" or "NO_FAULT," and both of its distance relays changed to "FAULT," then $AF = -2$. $F_{\text{out}}(i) = -1$, which is sufficient for the line to transition to the *normal* state. This requires three component failures. Alternatively, IED X could suffer a total IED failure, while IED Y changes its values to those required for $AF = -2$. This requires 1 total IED failure and 2 component failures.

The discovery that a Type II failure can only occur with at least 3 malfunctions, and only specific sets of at least 3 malfunctions, leads to the conclusion that the WABPS is more fault tolerant than upon initial inspection. While results in the

range of hundreds of thousands may initially seem disparaging, it's clear that a vast majority of those scenarios are false positives. System failures are equivalent from the perspective of SPIN, but they are not equal in practical applications. Tripping a transmission line that was working properly could be costly, but the consequences are much lower compared to the equipment damage caused by a faulted line left unchecked.

Proximal Lines.

In this context, a proximal line is any non-faulted neighboring line to the faulted line which has one of the Zone 2 relays activated. Consider the malfunctions required for a proximal line to be misdiagnosed as one that is faulted. Table 6 shows the starting values of the IEDs on a line of this type.

	Directional Relay	**Zone 1**	**Zone 2**	**AF**
IED X	LINE_FAULT	NO_FAULT	FAULT	1
IED Y	BUS_FAULT	NO_FAULT	NO_FAULT	-2

Table 6. Initialization of IEDs on an Proximal Line

A proximal line, with $F_{\text{out}}(i) = -1$, can suffer from a Type I failure in two scenarios.

$$F_{\text{out}}(i) > 2$$

$$0 \leq F_{\text{out}}(i) \leq 2$$

Meeting the first condition guarantees that a Type I failure will occur. This condition can be met with a minimum of only two component failures. If IED Y's directional relay changes to "LINE_FAULT," and its Zone 1 relay changes to "FAULT," the action factor will raise from -2 to 2, thus $F_{\text{out}}(i) > 2$. These two malfunctions

together are sufficient for the line to be placed in the *fault* state. If either IED on the line faces a total IED failure, this condition cannot be met.

Meeting the second condition makes a Type I failure likely, but does not guarantee it. This condition can be met with a fewer number of malfunctions. $0 \leq F_{\text{out}}(i) \leq 2$ can be fulfilled by a single total IED failure or a single component failure. There are eight possible single malfunctions that would raise the $\text{F}_{\text{out}}(i)$ to the specified range. For example, a total IED failure for IED Y, or changing the Zone 1 relay for IED X to "FAULT." The line either goes to the *suspect* state, or is diagnosed as *normal*, in which case it is not a Type I failure. In the *suspect* state, success is only possible if the line is adjacent to the faulted line. If so, the proximal line in question will be properly diagnosed as *normal*. If the proximal line is not adjacent to the faulted line, two things can occur. Either all of its neighbors were diagnosed as *normal*, thus misdiagnosing it as faulted, or at least one of its neighbors was also put in the *suspect* state. In the second case, either the local line or the neighboring *suspect* line will be misdiagnosed as faulted.

Distal Lines.

In this context, distal refers to any line either not adjacent to the faulted line, or an adjacent line with neither of the Zone 2 relays activated. Consider the malfunctions required for a distal line to be misdiagnosed as one that is faulted. Table 7 shows the starting values of the IEDs on a line of this type.

	Directional Relay	**Zone 1**	**Zone 2**	**AF**
IED X	LINE_FAULT	NO_FAULT	NO_FAULT	0
IED Y	BUS_FAULT	NO_FAULT	NO_FAULT	-2

Table 7. Initialization of IEDs on a Distal Line

As with proximal lines, there are two conditions that could cause a Type I failure

for a distal line.

$$F_{\text{out}}(i) > 2$$

$$0 \leq F_{\text{out}}(i) \leq 2$$

Since the initialization between a proximal line and a distal line only differs by the value of a single distance relay, the circumstances in which they are misdiagnosed are very similar. The minimum number of malfunctions required to meet the first condition is three, instead of two. The number of single malfunctions that would meet the second condition is still eight.

Due to how easily non-faulted lines can be misdiagnosed, it's clear that the majority of failures reported in the results are Type I, not Type II. With this in mind, the proposed revision to the decision making process, discussed next, specifically addresses some of the vulnerable scenarios non-faulted lines face.

4.4 Revision to the Algorithm

Now that the causes of failure have been analyzed, the next step is to make improvements to the model to increase its resiliency. One suggestion is to make a small adjustment to how the Action Factor of an IED is calculated. Currently, any IED with a directional relay reporting a line side fault and an activated Zone 1 relay will have $AF = 2$. This is independent of the Zone 2 relay reading. Assigning the highest possible AF while only considering two out of the three relays is rash and leads to many of the Type II failures previously examined. Table 8 shows proposed values for calculation of the Action Factor.

A table with proposed values for calculation of the Action Factor is shown below.

This proposed table restricts the conditions necessary for an IED to be assigned

AF	Directional Relay	Zone 1	Zone 2
2	LINE_FAULT	FAULT	**FAULT**
1	**LINE_FAULT**	**FAULT**	**NO_FAULT**
1	LINE_FAULT	NO_FAULT	FAULT
-2	BUS_FAULT or NO_FAULT	NO_FAULT	NO_FAULT
0	Any other combination		

Table 8. Proposed Action Factor Calculation

an Action Factor of 2. In turn, this will lead to noticeably fewer situations in which component failures cause lines to be incorrectly diagnosed as faulted. It should be noted that the second row in the table (an addition) can't normally be achieved by an IED. If a Zone 1 relay has been activated, then the Zone 2 relay by default has also been activated since it covers a greater area. Thus, this possibility can only be achieved as a result of component failure. This ties back to the approach of adversarial thinking. By expecting lines to mistakenly fulfill this condition, and reducing the ramifications of the condition, less failures should occur. With this in mind, the added row should only help the system in preventing false positives, and should not deter it from correctly identifying positives. To see how significant this proposed change is in the system's fault tolerance, the experiment will be conducted again using the new Action Factor calculation.

4.5 Revised Results

Line 13b

		Component			
		0	**1**	**2**	**3**
Total IED	**0**	0 (-0)	30 (-5)	8,559 (-1,379)	61,609 (-7,617)
	1	5 (-0)	2,230 (-123)	92,535 (-6,614)	101,158 (-6,909)
	2	138 (-0)	45,645 (-1,229)	106,449 (-3,148)	118,498 (-10,124)
	3	1,825 (-0)	113,742 (-2,583)	127,070 (-2,017)	127,199 (-33,166)

Table 9. Line 13b Revised Results

Table 9 displays the results of testing the WABPS again with the proposed revisions. The first number in each cell is the number of failed simulation runs for each configuration, as in Table 3. The number in parentheses is the difference between the new results and the old. For example, the verification run with two component failures led to 8,559 failed simulation runs, which is 1,379 less than previously reported. As expected, the number of failure scenarios is slightly lower than in the previous iteration.

The change in the Action Factor calculation made one specific contribution leading to the increased fault tolerance of the system. As discussed in 4.3, the minimum number of component failures previously necessary for $F_{\text{out}}(i) > 2$ for Proximal lines was two. This was possible if IED Y's directional relay changed to indicate a line side fault and if the Zone 1 relay activated. These two malfunctions give the IEDs on the line Action Factors of 1 and 2, whose sum is high enough for a Type I failure. However, this same combination of malfunctions under the new calculation results in Action Factors of 1 and 1, whose sum is NOT high enough for a Type I failure. Thus, the minimum number of component failures required for this condition is raised from two to three, a clear indication of increased resiliency.

It is apparent from the results that the revision led to less system failures in scenarios where the model only had one component error. This is due to the set of circumstances discussed in 4.3 that could lead Proximal/Distal lines to meet the second condition, $0 \leq F_{\text{out}}(i) \leq 2$. The revision reduces the number of single component failures leading to the second condition from eight to seven.

The contribution to the WABPS from this thesis slightly increased its fault tolerance. Further improvements are certainly possible, yet would likely require a different approach to problem solving. The Action Factor calculation was designed to be adaptive, allowing the system to react to various combinations of malfunctions. As it is

already quite efficient in detecting faults, further improvements would be in the form of expanding the model, rather than revising it. For example, incorporating the use of Zone 3 relays in the Action Factor calculation. While a third distance relay would no doubt improve the WABPS' fault tolerance, it would also significantly increase the complexity of the problem space. Other such additions to the model would be just as computationally intensive.

V. Conclusion

5.1 Contributions

In this thesis, a backup protection system employing autonomous agents to achieve wide area communication is stress-tested. Various types and levels of real-life malfunction are represented as component and IED failures. The backup protection system's response to a fault is then modeled, with the malfunctions incorporated. Each of the 15 Line Decision Agents in the system performs a decentralized, logical decision making process to determine whether or not it was able to properly diagnose the line for a fault. Due to the countless number of combinations of malfunctions and fault locations, the SPIN model checker was used to iterate through all possible scenarios. With results compiled, the logic behind the algorithm was broken down to examine exactly what combinations of malfunctions led to failed simulation runs. Specific vulnerabilities of the backup protection system were then uncovered. The logic behind the algorithm was slightly modified, which increased the number of malfunctions required for a specific failure scenario from two to three. Allowing the backup protection system to face a higher level of adversity while still being able to operate correctly is a clear indication of improved fault tolerance.

5.2 Review of Research Questions

- RQ1: *How can the model checking approach be applied to test the fault tolerance of the WABPS?*

  The WABPS was tested through the use of the PROMELA model which determined theoretical responses to various fault scenarios. SPIN, used in verification mode, executed all possible scenarios and recorded simulation runs in which the system was unable to properly respond to the fault.

- RQ2: *What specific vulnerabilities lie within the WABPS?*

  Specific vulnerabilities were analysed in 4.3. Categorizing lines in the system into Faulted, Proximal, and Distal lines allowed for further examination of vulnerability. The initialization for each type of line was discussed, as well as which Fault Values resulted in incorrect verdicts. With the conditions listed for incorrect verdicts, the number and type of malfunctions required for each condition were studied.

- RQ3: *Using the model checking approach, what improvements can be made to the model to increase its resiliency?*

  In RQ2, specific combinations of malfunctions leading to simulation failures were uncovered. With these in consideration, a revision was made to the logic behind the PROMELA model to increase the system's resiliency. SPIN was used again to show that the WABPS did indeed become more fault tolerant.

5.3 Future Work

As mentioned in Chapter 3, many assumptions and limitations were necessary to focus on the problem space of this research. Future work could entail tackling one of these areas to improve the model.

- Incorporate further sources of backup protection. Section 2.4 detailed local and remote backup protection. The current model does not embody the idea of local backup protection at all. This could be added in by having a Zone 1 relay, and a backup Zone 1 relay. An example Action Factor calculation table for this is proposed in [1]. Remote backup protection can be expanded on through the use of Zone 3 distance relays. Tertiary distance protection is commonly implemented in real life, but is not present in the current PROMELA model.

Admittedly, this avenue of future work could prove to be troublesome. Adding more protection relays to each IED would broaden an already complex model making it more resource-intensive to perform verification runs on.

- Find a way to differentiate between strong and weak correctness failures. As covered in 3.4, SPIN does not allow for simulation runs to continue once a verdict is violated. Weak correctness failures overlap with strong failures. Thus, a simulation run can pass the weak correctness test, but fail the strong test. It is currently not possible to tell how many strong correctness failures would have passed the weakness test, which could be useful information when evaluating the extent of failure of the WABPS.

- Define the length of transmission lines in the system. 3.5 discloses that lines in the model are of an arbitrary, but equal, length. This abstraction is made so that the operating range of Zone 2 relays and directional relays are simplified. This line of future work would make the model much more realistic.

Bibliography

1. K. Elliott, "Evaluation of Resiliency in a Wide-Area Backup Protection System Via Model Checking," Ph.D. dissertation, Air Force Institute of Technology, 2018.

2. X. Tong, X. Wang, R. Wang, F. Huang, X. Dong, K. M. Hopkinson, and G. Song, "The study of a regional decentralized peer-to-peer negotiation-based wide-area backup protection multi-agent system," *IEEE Transactions on Smart Grid*, 2013.

3. S. Borlase, *Smart Grids: Infrastructure, Technology, and Solutions*. CRC Press, 2013.

4. P. Hines, J. Apt, and S. Talukdar, "Trends in the history of large blackouts in the United States," *IEEE Power and Energy Society 2008 General Meeting: Conversion and Delivery of Electrical Energy in the 21st Century, PES*, vol. 15213, pp. 1–8, 2008.

5. R. Smith, "U.S. Risks National Blackout From Small-Scale Attack," *The Wall Street Journal*, 2014.

6. C. W. Ten, G. Manimaran, and C. C. Liu, "Cybersecurity for critical infrastructures: Attack and defense modeling," in *IEEE Transactions on Systems, Man, and Cybernetics Part A:Systems and Humans*, 2010.

7. G. Sorebo and M. Echols, *Smart Grid Security: An End-to-End View of Security in the New Electrical Grid*. Boca Raton: CRC Press, 2012.

8. S. H. Horowitz and A. G. Phadke, *Power System Relaying*, 4th ed. Wiley, 2014.

9. S. T. Hamman, K. M. Hopkinson, R. L. Markham, A. M. Chaplik, and G. E. Metzler, "Teaching Game Theory to Improve Adversarial Thinking in Cyberse-

curity Students," *IEEE Transactions on Education*, vol. 60, no. 3, pp. 205–211, 2017.

10. T. Scheponik, A. T. Sherman, D. DeLatte, D. Phatak, L. Oliva, J. Thompson, and G. L. Herman, "How students reason about Cybersecurity concepts," in *Proceedings - Frontiers in Education Conference, FIE*, 2016.

11. J. L. Blackburn and T. J. Domin, "Protective Relaying: Principles and Applications." CRC Press, 2007, ch. 1, pp. 1–36.

12. K. Hopkinson, X. Wang, R. Giovanini, J. Thorp, K. Birman, and D. Coury, "EPOCHS: A platform for agent-based electric power and communication simulation built from commercial off-the-shelf components," *IEEE Transactions on Power Systems*, 2006.

13. S. T. Hamman, K. M. Hopkinson, and J. E. Fadul, "A Model Checking Approach to Testing the Reliability of Smart Grid Protection Systems," *IEEE Transactions on Power Delivery*, 2016.

14. M. Begovic, "Trends in Power System Wide Are Protection," *IEEE PES Power Systems Conference and Exposition*, vol. 3, pp. 1612–1613, 2004.

15. R. Giovanini, K. Hopkinson, D. V. Coury, and J. S. Thorp, "A primary and backup cooperative protection system based on wide area agents," *IEEE Transactions on Power Delivery*, 2006.

16. A. Mahari and M. Sanaye-Pasand, "An Accelerated Single-Pole Trip Scheme for Zone-2 Faults of Distance Relays," *IEEE Transactions on Power Delivery*, 2017.

17. S. Vejdan, M. Sanaye-Pasand, and T. S. Sidhu, "Accelerated Zone II Operation of Distance Relay Using Impedance Change Directions," *IEEE Transactions on Power Delivery*, 2017.

18. S. H. Horowitz and A. G. Phadke, "Third zone revisited," *IEEE Transactions on Power Delivery*, vol. 21, no. 1, pp. 23–29, 2006.

19. J. F. Borowski, K. M. Hopkinson, J. W. Humphries, and B. J. Borghetti, "Reputation-based trust for a cooperative agent-based backup protection scheme," *IEEE Transactions on Smart Grid*, 2011.

20. G. J. Holzmann, *The SPIN Model Checker: Primer and Reference Manual*. Addison Wesley, 2004.

21. C. Baier and J.-P. Katoen, *Principles of Model Checking*. MIT Press, 2008.

22. A. Gmeiner, I. Konnov, U. Schmid, H. Veith, and J. Widder, "Tutorial on parameterized model checking of fault-tolerant distributed algorithms," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2014.

23. C. Newcombe, T. Rath, F. Zhang, B. Munteanu, M. Brooker, and M. Deardeuff, "How Amazon web services uses formal methods," *Communications of the ACM*, 2015.

24. "Official SPIN Website," 2019. [Online]. Available: http://spinroot.com/spin/success.html

25. T. C. Ruys, "Unit Testing for SPIN: runspin and parsepan," 2014. [Online]. Available: http://dx.doi.org/10.1145/2632362.2632382

26. G. Ramos, J. Sanchez, a. Torres, and M. Rios, "Power Systems Security Evaluation Using Petri Nets," *IEEE Transactions on Power Delivery*, 2010.

27. L. O. Matos and J. W. G. Sanchez, "Reconfiguration strategy for Fault Tolerance of power Distribution Systems using Petri net," in *2016 IEEE Ecuador Technical Chapters Meeting, ETCM 2016*, 2016.

28. Z. Jiang, Z. Li, N. Wu, and M. Zhou, "A Petri Net Approach to Fault Diagnosis and Restoration for Power Transmission Systems to Avoid the Output Interruption of Substations," 2017.

29. Z. Yongli, H. Limin, and L. Jinling, "Bayesian networks-Based approach for power systems fault diagnosis," *IEEE Transactions on Power Delivery*, 2006.

30. A. Khurram, H. Ali, A. Tariq, and O. Hasan, "Formal reliability analysis of protective relays in power distribution systems," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2013.

31. J. Wäfler and P. Heegaard, "A combined structural and dynamic modelling approach for dependability analysis in smart grid," in *Proceedings of the ACM Symposium on Applied Computing*, 2013.

32. A. Mahmood, O. Hasan, H. R. Gillani, Y. Saleem, and S. R. Hasan, "Formal reliability analysis of protective systems in smart grids," *Proceedings - 2016 IEEE Region 10 Symposium, TENSYMP 2016*, pp. 198–202, 2016.

33. M. Tsuji, K. Hasebe, and K. Kato, "Deadlock Detection in Scheduling of Last-Mile Transportation by Using Model Checking," pp. 1–6, 2017.

34. C. Bergenhem, "A status protocol for system-operation in a fault-tolerant system - Verification and testing with SPIN," *IEEE International Conference on Emerging Technologies and Factory Automation, ETFA*, 2012.

35. C.-W. Ten, K. Yamashita, Z. Yang, A. Vasilakos, and A. Ginter, "Impact Assessment of Hypothesized Cyberattacks on Interconnected Bulk Power Systems," *IEEE Transactions on Smart Grid*, 2017.

36. Z. Daria and K. Sroka, "The characteristics and main causes of power system failures basing on the analysis of previous blackouts in the world," *2018 International Interdisciplinary PhD Workshop (IIPhDW)*, pp. 257–262, 2018.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704–0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704–0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202–4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE *(DD–MM–YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From — To)*
21–03–2019	Master's Thesis	June 2017 — Mar 2019

4. TITLE AND SUBTITLE

Testing the Fault Tolerance of a Wide Area Backup Protection System using SPIN

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

James, Kenneth, 2d Lt, USAF

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way
Wright-Patterson AFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT-ENG-MS-19-M-034

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Intentionally Left Blank

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Distribution Statement A:
Approved for Public Release; Distribution Unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Cyber-physical systems are increasingly prevalent in daily life. Smart grids in particular are becoming more interconnected and autonomously operated. Despite the advantages, new challenges arise in the form of defending these assets. Recent studies reveal that small-scale, coordinated cyber-attacks on only a few substations across the U.S. could result in cascading failures affecting the entire nation. In support of defending critical infrastructure, this thesis tests the fault tolerance of a backup protection system. Each transmission line in the system incorporates autonomous agents which monitor the status of the line and make decisions regarding the safety of the grid. Various malfunctions that could occur from real-life attacks are simulated in the grid with the intent of determining its ability to successfully respond to faults despite adversity. The SPIN model checker is used to execute all combinations of fault location and malfunctions to determine which types can occur, and how many, before the system is unable to properly clear a fault. With results analyzed, the decision making process of the model is revised to increase its fault tolerance.

15. SUBJECT TERMS

model checking, SPIN, smart grid, backup protection, critical infrastructure

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Kenneth Hopkinson, AFIT/ENG
U	U	U	UU	54	19b. TELEPHONE NUMBER *(include area code)* (937) 255-3636, x4579; kenneth.hopkinson@afit.edu