CHEMICAL ENGINEERING TRANSACTIONS
VOL. 77, 2019
A publication of AIDIC
The Italian Association of Chemical Engineering
Online at www.cetjournal.it
Guest Editors: Genserik Reniers, Bruno Fabiano

ISBN 978-88-95608-74-7; **ISSN** 2283-9216
DOI: 10.3303/CET1977149

Investigating Cyber Threats in a Nuclear Power Plant

Susan S. Adams*, Nicole Murchison, Robert J. Bruneau

Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM 87185
smsteve@sandia.gov

Malicious cyber-attacks are becoming increasingly prominent due to the advance of technology and attack methods over the last decade. These attacks have the potential to bring down critical infrastructures, such as nuclear power plants (NPP's), which are so vital to the country that their incapacitation would have debilitating effects on national security, public health, or safety. Despite the devastating effects a cyber-attack could have on NPP's, it is unclear how control room operations would be affected in such a situation. In this project, the authors are collaborating with NPP operators to discern the impact of cyber-attacks on control room operations and lay out a framework to better understand the control room operators' tasks and decision points. A cyber emulation of a digital control system was developed and coupled with a generic pressurized water reactor (GPWR) training simulator at Idaho National Laboratories. Licensed operators were asked to complete a series of scenarios on the simulator in which some of the scenarios were purposely obfuscated; that is, in which indicators were purposely displaying inaccurate information. Of interest is how this obfuscation impacts the ability to keep the plant safe and how it affects operators' perceptions of workload and performance. Results, conclusions and lessons learned from this pilot experiment will be discussed. This research sheds light onto about how cyber events impact plant operations.

1. Introduction

Nuclear Power Plant (NPP) systems are becoming increasingly digital including plant monitoring sensors, displays available to operators and control of devices within those systems. In the early implementation, digital control systems were put in place to supplement or replace existing analog systems for improved data collection, reconfiguration and remote access capability. Digital control systems have historically been standalone with no need to interface with outside systems, as their sole purpose was to automatically regulate a physical process. As such, these early systems were designed with no cybersecurity measures in place that would protect them from modern day cyber threats due to an implicit assumption that they would be isolated from the Internet and other outside network influences (Watts, 2003). However, digital isolation is becoming increasingly difficult in the age of wireless networking (e.g., wi-fi, Bluetooth, cellular data networks) and the ubiquitous Internet of Things (IoT), where virtually any technological device has some wireless connectivity.
The Online Trust Alliance estimates that nearly 160,000 cyber incidents occurred in 2017 (Online Trust Alliance, 2018). Electrical power has been targeted by cyber-attacks from foreign nations in order to threaten the safety of civilians and create chaos (Park, Summers & Walstrom, 2017). Nuclear Power Plants are attractive targets for foreign adversaries since, compared to nonnuclear plants, NPPs have much broader safety consequences because of the immediate biological consequences and overall public fear of after-effects of nuclear-materials release. One way to circumvent this problem is to layer protections around the systems determined to require security. Adding layers of security can have a positive impact in terms of increasing the difficulty of attacks against the system. Some NNP systems have cyber-security controls built in (e.g., password protection). Additionally, Industrial Control Systems use boundary protection mechanisms to secure these systems (analogous to layered physical security). However, those controls generally provide minimal protection and there still may be weaknesses (see McLeod, 2016). Foreign adversaries are determined, innovative, and persistent. And they have been successful.
In 2010, a cyber-attack, using a worm dubbed Stuxnet, successfully infiltrated an Iranian nuclear facility and took control of the nuclear centrifuges. In addition, to prevent the operators from noticing the attack, Stuxnet

Paper Received: 20 December 2018; Revised: 27 April 2019; Accepted: 3 July 2019

Please cite this article as: Adams S., Murchison N., Bruneau R., 2019, Investigating Cyber Threats in a Nuclear Power Plant, Chemical Engineering Transactions, 77, 889-894 DOI:10.3303/CET1977149

took control of the displays in the control room and mimicked normal operation (Karnouskos, 2011). This obfuscation of the displays was crucial. For a cyber-attack to be successful, the human operators must be duped into believing that there is no need to intervene. Successful implementation of this type of exploit would degrade or remove corrective actions that would normally be performed by the operator to mitigate abnormal, dangerous or degrading conditions. Fortunately, human operators are not solely reliant on the digital displays and make use of redundant displays, which include analog displays in the control room and local displays in the plant. The questions are, 'How difficult is it for operators to recognize contradictory evidence, determine which information is correct, and aptly respond to the actual plant conditions? How will this influence their perceived workload and perceived performance?' The study described in this paper begins to explore the answers to these questions.

2. Pilot study protocol

2.1 Participants

Two NPP operators familiar with the generic pressurized water reactor (GPWR) simulator and its functionality were recruited to participate. One participants served as a supervisor reactor operator (SRO), who selected the appropriate procedure(s) given the plant conditions and called out instructions, and the other participant served as a reactor operator (RO), who checked and responded to plant conditions under the SRO's instructions. The participants worked together as a team/crew.

2.2 Apparatus and stimuli

The experiment was conducted using the GPWR simulator, which provides a full-scope model of a nuclear plant simulation, along with procedures analogous to those used in the plant. The simulator was located at the Idaho National Laboratory Human System Simulation Laboratory (INL HSSL; see Figure 1). Control boards were represented on top screen displays and 15 panels were linked together to make up a full-scale representation of the control room.

Figure 1. Idaho National Laboratory Human System Simulation Laboratory

The abnormal and emergency condition procedures were pre-printed and bound into procedure books. Additional procedures were available on a laptop, as needed.

2.3 Scenarios

Two fault scenario types were designed for use in this study, one involving an Interconnected System Loss of Coolant Accident Residual Heat Removal (IS LOCA RHR) system and one involving the Interconnected System Loss of Coolant Accident Pressure Operated Relief Valve (IS LOCA PORV). These scenarios were selected because they are consistent with analyses of severe accidents, which are characterized by conditions that can result in a release of radionuclides to the public.

The IS LOCA RHR scenarios were initiated by the abnormal opening of the valves RHR1 and RHR2. The two normally closed valves were gradually opened over a period of 20 seconds, and then locked in the open position for the remainder of the scenario. This resulted in primary coolant flowing through the RHR system, which is not rated for operational pressure, leading to a LOCA. The IS LOCA PORV scenarios involved the PORV being stuck in the fully-open position while at full reactor coolant system (RCS) pressure, leading to a LOCA. The experimental conditions were:

- Obfuscation: True indications of plant conditions (non-obfuscated) vs. false-normal (obfuscated)
- System failure type: IS LOCA RHR vs. IS LOCA PORV

There were three types of trials: Non-obfuscated trials (baseline), obfuscated trials (experimental), and non-obfuscated distractor trials. The non-obfuscated and obfuscated trials used the RHR and PORV scenario types. The distractor trials used other malfunction scenarios (available in the GPWR simulator.) Distractor

trials were added to preclude anticipation and progressive learning of the RHR and PORV conditions over the course of the experiment.The baseline and experimental trials consisted of the following:

- RHR-1 (SNP194) with true indications of the failure (no obfuscation)
- RHR-2 (SNP193) with false-normal indications (obfuscation)
- RHR-3 (SNP192) with false-normal indications (obfuscation)
- PORV-1 (SNP191) with true indications of the failure (no obfuscation)
- PORV-2 (SNP188) with false-normal indications (obfuscation)
- PORV-3 (SNP189) with false-normal indications (obfuscation)

2.4 Hypotheses

1. The ability to keep the plant in a safe condition (by both operator actions and automatic plant safety system actions) will be degraded in the obfuscated scenarios compared to baseline scenarios.
2. Self-Rated Task Performance will be rated lower in obfuscated scenarios compared to baseline scenarios.
3. NASA TLX ratings will reflect higher workload in obfuscated scenarios compared to baseline scenarios.

2.5 Procedure

The participants completed three days of experimentation. On day one, the participants read through and signed the informed consent, filled out a demographic questionnaire and completed a practice scenario, in which they were re-familiarized with the simulated control room panels. The participants were told to treat the indicators as if they were all digitally controlled. The participants were instructed to respond to malfunctions and faults in each scenario as they would in a real NPP control room. After this re-familiarization task, the participants completed the NASA Task Load Index (NASA TLX; Hart, 1988) and Self-Rated Task Performance (SRTP; Dumas, 2015) questionnaires. The participants were encouraged to ask questions to ensure that they understood the questionnaire probes and rating scales. Next, the participants completed the non-obfuscated scenarios and a distractor scenario on the first day of the experimental trials. The non-obfuscated scenarios were performed prior to the obfuscated scenarios so as not to impugn the legitimacy of the baseline (Lammers & Badia, 2004). The remainder of the week was dedicated to experimental and distractor scenarios. The run order of the scenarios was randomized. Audio/video were captured for both participants. At the end of each scenario, the participants filled out the NASA-TLX and SRTP questionnaires. After completion of the scenarios, a verbal walkthrough using a modified version of Applied Cognitive Task Analysis (ACTA; Militello & Hutton, 1998) was completed, modeled after previous research (Stevens-Adams, et al., 2015). Finally, at the very end of the study, the participants were debriefed and asked for their input and feedback about the scenarios to inform future studies.

3. Preliminary findings

It should be noted that there were significant differences between traditional NPP control rooms and the experimental setup that may have influenced results. These differences included:

- The experimental control room only used two operator positions, one reactor operator and one supervisor reactor operator. A shift crew would normally include one SRO and two RO’s at a minimum.
- The directionality of the alarm sounds in the simulated control room were different from what they were accustomed to; namely, in an actual control room, the alarms would be localized to the board at which their attention needs to be focused. Additionally, panels with alarm indicators from other areas of the plant (such as the Reactor Auxiliary Building) were missing from the simulated environment.
- Operators normally have access to all procedures at the beginning of their shift; in the experimental control room, the experimenters had to locate a few procedures in the middle of a scenario.

Researchers analyzed audio/video files and questionnaire data for the baseline and obfuscated scenarios. In addition, the simulation output of parameters, such as temperature and pressure, that were identified as critical and relevant to the scenario were captured. The operators were responsible for monitoring and adjusting these parameters to ensure that enough cooling was reaching the reactor core to prevent a reactor core melt condition. In addition, several plant safety system actions were automatically enacted during the scenario. The data from the distractor scenarios was not analyzed.

3.1 IS LOCA RHR

For these analyses, the obfuscated scenarios were combined and/or averaged. For all the scenarios, the plant was kept in a safe state such that the core always had enough cooling. This is contrary to the hypothesis that

the ability to keep the plant in a safe state would be degraded in the obfuscated conditions.The NASA TLX scores were analyzed for both participants (see Figure 2a). Note that there was confusion regarding the 'performance' construct (the participants realized that the scale was reversed halfway through the experiment) so that data was not analyzed. There was no statistically significant difference in ratings between baseline and the experimental obfuscated scenarios, contrary to the hypothesis that workload would be perceived as higher in the obfuscated conditions. In fact, the participants had higher ratings for the baseline scenario, indicating higher perceived workload, which is opposite of the hypothesis. However, that is likely because the baseline condition was the first scenario and the participants indicated that they were having difficulty locating the correct panels and indicators. The Self-Rated Task Performance questionnaire data was also consolidated across participants and analyzed for the baseline and obfuscated scenarios (see Figure 2b). There were no statistically significant differences between the baseline and experimental conditions, contrary to the hypothesis that performance would be perceived as lower in the obfuscated conditions. In fact, the participants had lower ratings for the baseline scenario, indicating lower perceived performance, which is opposite of the hypothesis. However, again, that is likely because the baseline condition was the first scenario and the participants indicated that they were having difficulty locating the correct panels and indicators and thus felt that their performance suffered.

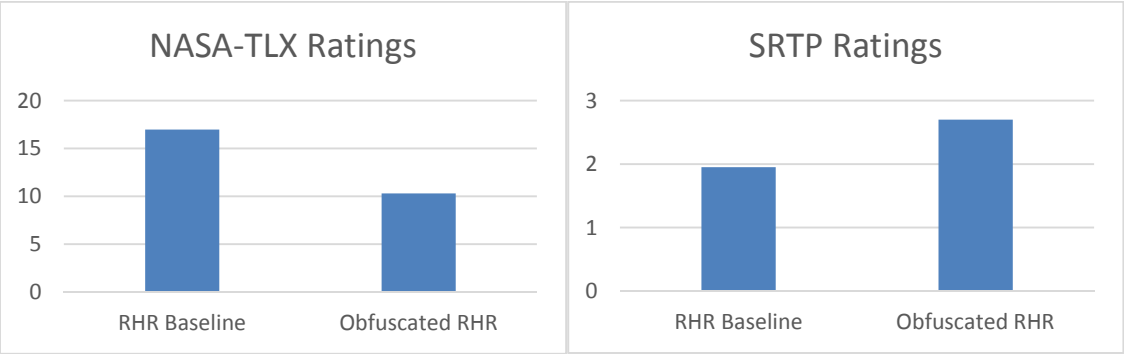

Figure 2a and 2b. NASA TLX (2a) and Self-rated task performance (2b) scores for IS LOCA RHR scenarios

3.2 IS LOCA PORV

As above, the obfuscated scenarios were combined and/or averaged. For all the scenarios, the plant was kept in a safe state such that the core always had enough cooling. This is contrary to the hypothesis that the ability to keep the plant in a safe state would be degraded in the obfuscated conditions.The NASA TLX scores were analyzed for both participants (see Figure 3a). Again, note that there was confusion regarding the 'performance' construct (the participants realized that the scale was reversed halfway through the experiment) so that data was not analyzed. There was no statistically significant difference in ratings between baseline and the experimental obfuscated scenarios, contrary to the hypotheses that workload would be higher in obfuscated conditions. However, while not significant, the ratings were in the expected direction.The Self-Rated Task Performance questionnaire data was consolidated across participants and analyzed for the baseline and obfuscated scenarios (see Figure 3). There were no statistically significant differences between the baseline and experimental conditions, contrary to the hypothesis that perceived performance would be rated lower in the obfuscated conditions. However, while not significant, the ratings were in the expected direction.

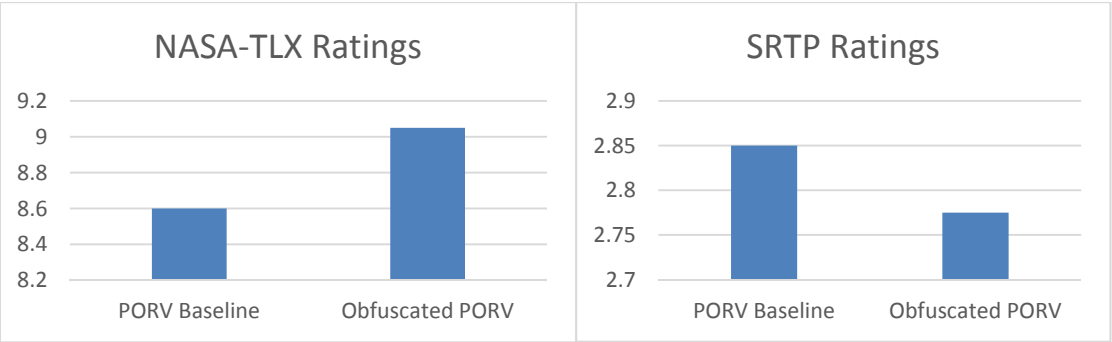

Figure 3a and 3b. NASA TLX (3a) and Self-rated task performance (3b) scores for IS LOCA PORV scenarios

3.3 Verbal walkthrough

The participants completed the verbal walkthroughs as a team for each scenario. The participants were not told about the obfuscation nor the purpose of the experiment during these walkthroughs. They first outlined the major events in each of the scenarios and, for each major event, discussed what cues were relevant and how they came to their decision to move forward. For all the scenarios, the participants stated that the alarms and other key indicators were crucial to helping them identify what the issue was. In addition, they relied heavily on the procedures to determine the correct steps to take to mitigate any issues. In fact, the participants stated that their job is to follow the procedures in response to any issue and do not generally engage in any kind of diagnosis of the plant until after the plant is stabilized. For a few of the scenarios, the participants did mention in the verbal walkthrough that there was something awry with the indicators and for one scenario joked that 'it could be a cyber-attack'.

3.4 Knowledge elicitation and debriefing

After the end of the verbal walkthrough, participants were debriefed and queried about how the experiment could have been improved. The participants were told about the purpose of the study and that some of the scenarios had included obfuscation, in which the indicators had been manipulated to purposely display incorrect or inconsistent information. The participants were then asked for their input and feedback regarding the scenarios and whether changes could be made so that the scenarios were more realistic for future studies. The participants explained that their main task is to keep the plant safe, which they rely on the procedures to help them do. Even so, the participants admitted that they spent the first experimental day trying to diagnose what was happening in the plant but then modified their strategy in later experimental sessions to be more realistic in their approach in which they simply followed the procedures. The participants also noted that, to make the scenarios more realistic, it would have been helpful to have a larger crew (typically a crew consists of 5 operators) because it was very difficult for them to perform the experiment with just the two of them. They also suggested that an experimenter should play the role of an outside operator who can provide some valuable information about what is happening in the field or other parts of the plant. The participants did note that they wondered if some sort of obfuscation was happening during a few of the scenarios, and said that one in particular brought to mind Three Mile Island (the incident which the IS LOCA PORV scenario was based).

4. Discussion

This pilot study set the foundation for assessing the influence of cyber-attacks in nuclear power plant operations. The operators did react to the erroneous indications and verbally acknowledged that some of the indications showed responses inconsistent, when considering their control actions. These same indicators were those that were mentioned as important to understanding the scenario progression and choosing the correct path forward. In addition, simulator logs revealed that the plant safety protection systems enacted some automatic actions which suggests that the operators, in combination with that plant safety systems, were successful in keeping the plant safe. Operators also reported their priority, and the crux of their job is to follow procedures and keep the plant safe. Diagnosis happens with an incident response team after the plant is stabilized to safe conditions. In a more realistic operational environment with additional operators available, the plant operators may have an ability to focus more on the discrepant indications to better diagnose the actual plant condition.There were many limitations to this study which may be the reason why the hypotheses were not supported. For one, this study involved only a single, two-person team and the participants' change in strategy half-way through the experiment might have accounted for the results. This topic warrants further research and lessons learned from this study will help to inform future studies.

4.1 Lessons learned

The preliminary results do present a case for further studies to evaluate obfuscation in NPPs. To ensure the validity of the results in future studies, there were many lessons learned and recommendations from this development effort and pilot study that are outlined in Table 1.

4.2 Future work

The authors believe that further exploration into these areas is warranted. Future studies could test the impacts of cyber obfuscation in terms of the number and type of indicators that are impacted during an obfuscated attack and the sophistication of that attack. To address some of the Lessons Learned detailed above, the experiment could include an entire day of training on the simulator, and a day to establish baseline conditions. Additionally, having many days of experimentation with multiple operator teams would ensure statistical power.

Table 1. Lessons Learned from current study and Proposed Solutions moving forward

Lessons Learned	Proposed Solutions
Confusion was introduced by non-localized alarms	Provide a description of the simulator environment at the beginning of the experiment
Not enough time re-familiarizing operators with the simulator functionality and layout	Have a full day of refamiliarization trials
Operators reported a strategy change	In addition to added time to refamiliarize themselves with the simulated environment, ensure operators have no preconceived notions regarding the experimental purpose by updating the recruitment materials
Accessing procedures on a laptop was interruptive to operators	Ensure all relevant procedures are available, and validate the list with a SME
Operators reported diagnosis would typically not occur in a NPP until after the plant is stabilized to safe conditions	The qualitative data interview questions should be amended or changed such that insight into diagnosis is not elicited
Artificiality in scenarios due to erroneous indicators not realistically able to be infiltrated and inability to access data from elsewhere in the plant	The team should solicit SME input when creating scenarios for future studies and the scenario progression should be thoroughly scripted including an experimenter role that will provide information for indicators not located at the panels
Challenges having a two-operator team	Recruit a three-operator team; at a minimum there should be two ROs and one SRO

Acknowledgments

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
The authors would like to thank Idaho National Laboratory for the use of their lab and supporting in the execution of the study. In addition, the authors would like to thank the operators who participated in the study. Their insights and feedback was invaluable to the team.

References

Demas M W., Lau N., Elks C., 2015, Advancing human performance assessment capabilities for integrated system validation—A human-in-the-loop experiment, In: 9th American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation & Control and Human-Machine Interface Technologies (NPIC & HMIT).

Hart, S. G., Straveland, L.E., 1988, Development of the NASA-TLX (task load index): Results of Empirical and Theoretical Research. Chapter In: Advances in Psychology, Vol 52, North Holland-Amsterdam, 139-183.

Karnouskos, S., 2011, November. Stuxnet worm impact on industrial cyber-physical system security. In *IECON 2011-37th Annual Conference on IEEE Industrial Electronics Society* (pp. 4490-4494). IEEE.

Lammers W. J., Badia P., 2004, Experimental Designs: Single Subject Designs and Time-Series Designs, In: Fundamentals of Behavioral Research, Wadsworth, Belmont, CA, 14.1-14.28.

McLeod R., 2016, Issues in Assuring Human Controls in Layers-in-Defences Strategies, Chemical Engineering Transactions, 48, 925-930.

Militello L.G., Hutton R. J. B., 1998, Applied Cognitive Task Analysis: A Practitioner's Toolkit for Understanding Cognitive Task Demands, Ergonomics, 41(11), 1618-1641.

Online Trust Alliance, 2018, Cyber Incident & Breach Trends Report < https://otalliance.org/system/files/files/initiative/documents/ota_cyber_incident_trends_report_jan2018.pdf> accessed 24.09.2018

Park, D., Summers, J., Walstrom, M., 2017, Cyberattack on Critical Infrastructure: Russia and the Ukrainian Power Grid Attacks, University of Washington, <https://jsis.washington.edu/news/cyberattack-critical-infrastructure-russia-ukrainian-power-grid-attacks/>, accessed 09.24.2018.

Stevens-Adams S., Cole K., Haass, M., Warrender C., Jeffers R., Burnham L., Forsythe C., 2015, Situation Awareness and Automation in the Electric Grid Control Room, Procedia Manufacturing, 3, 5277-5284.

Watts D., 2003, Security & Vulnerability in Electric Power Systems, 35th North American Power Symposium, 2, 5.