Report from Dagstuhl Seminar 14292

Network Attack Detection and Defense: Securing Industrial Control Systems for Critical Infrastructures

Edited by

Marc Dacier[1], Frank Kargl[2], Hartmut König[3], and Alfonso Valdes[4]

1 **Qatar Computing Research Institute (QCRI), Qatar**
2 **Universität Ulm, DE,** frank.kargl@uni-ulm.de
3 **BTU Cottbus-Senftenberg, DE,** koenig@informatik.tu-cottbus.de
4 **University of Illinois – Urbana, US,** avaldes@illinois.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 14292 "Network Attack Detection and Defense: Securing Industrial Control Systems for Critical Infrastructures".

The main objective of the seminar was to discuss new approaches and ideas for securing industrial control systems. It is the sequel of several previous Dagstuhl seminars: (1) the series "Network Attack Detection and Defense" held in 2008 and 2012, and (2) the Dagstuhl seminar "Securing Critical Infrastructures from Targeted Attacks", held in 2012. At the seminar, which brought together members from academia an industry, appropriate methods for detecting attacks on industrial control systems (ICSs) and for limiting the impact on the physical components were considered. A central question was whether and how reactive security mechanisms can be made more ICS- and process-aware. To some extent it seems possible to adopt existing security approaches from other areas (e. g., conventional networks, embedded systems, or sensor networks). The main question is whether adopting these approaches is sufficient to reach the desired level of security for ICSs. Detecting attacks to the physical components and appropriate reactions to attacks are new aspects that need to be considered as well. The main result of the seminar is a list of recommendations for future directions in ICS security that is presented in this report.

Seminar July 13–16, 2014 – http://www.dagstuhl.de/14292
1998 ACM Subject Classification K.6.5 Security and Protection, C.2.0 General, J.7 Computers in Other Systems
Keywords and phrases Security, Intrusion Detection, Critical Infrastructures, Industrial Control Systems, SCADA, Vulnerability Analysis, Malware Assessment, Attack Response and Countermeasures
Digital Object Identifier 10.4230/DagRep.4.7.62
Edited in cooperation with Rens van der Heijden

Network Attack Detection and Defense: Securing Industrial Control Systems for Critical Infrastructures, *Dagstuhl Reports*, Vol. 4, Issue 7, pp. 62–79
Editors: Marc Dacier, Frank Kargl, Hartmut König, and Alfonso Valdes

Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Executive Summary

Marc Dacier
Rens van der Heijden
Frank Kargl
Hartmut König
Alfonso Valdes

From July 13–16, 2014, more than 30 researchers from the domain of critical infrastructure security met at Schloss Dagstuhl to discuss the current state of security in industrial control systems.

Recent years have highlighted the fact that security precautions of information and communication technology (ICT) in many critical infrastructures are clearly insufficient, especially if considering targeted attacks carried out by resourceful and motivated individuals or organizations. This is especially true for many industrial control systems (ICS) that control vital processes in many areas of industry that are relying to an ever-larger extent on ICT for monitoring and control in a semi or fully automated way. Causing ICT systems in industrial control systems to malfunction can cause huge economic damages or even endanger human lives. The Stuxnet malware that actually damaged around 1000 Uranium enrichment centrifuges in the Iranian enrichment facility in Natanz is the most well-known reported example of an ICT attack impacting ICS.

This situation led to increased efforts in research which also resulted in a number of Dagstuhl seminars related to this topic of which this seminar is a follow-up event, namely two Dagstuhl seminars on "Network Attack Detection and Defense" in 2008 and 2012 and one on "Securing Critical Infrastructures from Targeted Attacks" held in 2012. The main objective of our this latest seminar was to discuss new approaches and ideas on how to detect attacks on industrial control systems and how to limit the impact on the physical components. This is closely coupled to the question of whether and how reactive security mechanisms like Intrusion Detection Systems (IDS) can be made more ICS- and process-aware. To some extent it seems possible to adopt existing security approaches from other areas (e. g., conventional networks, embedded systems, or sensor networks) and one of the questions is whether adopting these approaches is enough to reach the desired security level in the specific domain of industrial control systems, or if approaches specifically tailored for ICS or even single installations provide additional benefits.

The seminar brought together junior and senior experts from both industry and academia, covering different scenarios including electrical grids, but also many other control systems like chemical plants and dike or train control systems. Apart from the detection and prevention of attacks by both security and safety mechanisms, there was an extensive discussion on whether or not such systems should be coupled more strongly from a security perspective. It was also argued that there exists a very diverse space of application domains, many of which have not yet been subject to much study by security researchers, for various reasons. Many of these discussions were triggered by plenary or short talks, covering topics from the state of the art in ICS security, forensics in ICS, security assessments, and the new application domain of flood management.

Apart from talks and subsequent discussions, a number of working groups were organized during the seminar, intended to address specific issues in the field. In total, there were four

working groups, each of which provided a summary of their results included in this report. The first was on forensics, discussing how attacks can be detected and analyzed after the fact. A second working group addressed the issue of security and risk management, analyzing why existing IT security approaches do not work for ICS and discussing potential improvements. Industry 4.0 and the wide range of new and non-classical ICS use cases was the topic of a third working group, which discussed the new security challenges arising from these emerging research topics. Finally, there was a working group on the detection of cyber-physical attacks; a core question here were advantages and disadvantages of process-aware intrusion detection mechanisms. The group also discussed the interaction between intrusion detection, intrusion response, and security management.

Based on the talks, discussions and working groups, the Dagstuhl seminar was closed with a final plenary discussion which summarized again the results from the working groups and led to a compilation of a list of open issues that participants consider necessary to be addressed. Those issues partly overlap with the list of open issues identified in the seminar proposal but also uncovered many new challenges that may become highly relevant research topics and may lead to a new agenda for future research. Those issues are discussed at the end of this report.

2 Table of Contents

3 Plenary Talks

3.1 Examples of Cyber-attacks on SCADA Systems for the Electrical Grid and their Consequences

Gunnar Björkman (ABB – Mannheim, DE)

Main reference EU FP7 Project VIKING
URL http://www.kth.se/ees/omskolan/organisation/avdelningar/ics/research/cc/v

The presentation gave a number of examples of possible cyber-attacks on SCADA system used for the supervision and control of the electrical grid. It should be noted that these were examples of possible cyber-attacks, not a description of real incidents although some the attacks were inspired from real events like Stuxnet. The included scenarios have been verified by major transmission grid operators and found to be realistic in the meaning that they would be possible to conduct and that the reported consequences are realistic. The attack scenarios are a subset of the story boards developed in the FP7 project VIKING which ended in 2011. The scenarios described in the presentation had been selected to represent attacks of different categories like Denial of Service, Social Engineering, lack of Security Training, etc. and to demonstrate both major and minor consequences from the triggered black-outs in the society. For each scenario the probability for attack success calculated applying the CySeMoL method on the given ICS configuration was reported. The Cyber Security Modeling Language (CySeMoL) is also a result from the VIKING project. In addition to the scenario description and the attack success, two types of societal consequences were described in the presentation. One was a monetary cost which illustrates the lost Gross Domestic Product caused by the blackout and the other a logarithmic value calculated from impacted number of people and black-out duration that indicates non-economic consequences in the society. The latter measure closely resembles the Richter scale and gives an intuitive feeling for the seriousness of an attack on the power grid.

3.2 Towards Resilient Control of Critical Infrastructures

Alvaro Cárdenas Mora (University of Texas at Dallas, US)

The protection of industrial control systems is usually achieved via a series of safety, fault-tolerant, and robust control mechanisms. These solutions were design under the assumption of a non-strategic adversary and targeted mostly random failures.

In this talk we discussed possible new directions on how to extend these previous solutions and adapt them to be resilient to strategic adversaries that will try to evade anomaly detection schemes while at the same time maximize their negative impacts. The talk also described a generic approach to identify the vulnerability of systems at the physical and control layer via controllability and observability concepts, and then proposed resilient control algorithms that can survive their mission critical objectives even when faced with successful attacks partially compromising the system. The speaker showed use-cases in industrial process control of anaerobic and chemical reactors, and in frequency control as well as demand-response control in the power grid.

3.3 Security of Train Control Systems

Stefan Katzenbeisser (TU Darmstadt, DE)

In the talk Stefan Katzenbeisser gave an overview of the technical systems providing safe train operations in Germany. In particular, he discussed safety and security problems raised by these deployed systems. Furthermore, he also reported on an ongoing work that attempts to provide IT security recommendations for the railway industry, ranging from risk analysis over security aware design up to security management aspects. This set of recommendations is currently proposed as a standard in the DIN German Institute for Standardization. Finally, open research problems in this domain were surveyed.

4 Short Talks

4.1 Performing Forensic Investigations of Industrial Control Systems

Heiko Patzlaff (Siemens – München, DE)

Performing computer forensic investigations in industrial control systems (ICS) presents various challenges. Starting with Stuxnet, the general lack of tools and procedures for analyzing security incidents in industrial settings has become apparent. The talk gave an introduction to ICS. It looked at the challenges that arise when one needs to analyse security incidents in industrial products. In particular, what data is available in these systems? How do you acquire this data? How do you transfer and analyse it? And what conclusions can you draw from it? The talk provided some examples of real world cases. And it presented results from a research project aimed at developing tools and approaches for performing computer forensic investigations in ICS that led to the development of a new forensic platform for the investigation of such incidents.

4.2 Automatic Analysis of Unknown Network Protocols

Konrad Rieck (Universität Göttingen, DE)

Joint work of Rieck, Konrad; Krueger, Tammo; Gascon, Hugo; Krämer, Nicole
Main reference T. Krueger, H. Gascon, N. Krämer, K. Rieck, "Learning stateful models for network honeypots," in Proc. of the 5th ACM Workshop on Security and Artificial Intelligence (AISec'12), pp. 37–48, ACM, 2012.
URL http://dx.doi.org/10.1145/2381896.2381904

Proprietary protocols in ICS are a major hurdle for traffic analysis and intrusion detection. Without detailed knowledge of message formats and protocol states, there is little chance that attacks can be spotted in the traffic of these protocols. This talk provided an overview of techniques for automatic protocol reverse engineering. Different approaches for inferring message formats and protocol state machines from network traffic are presented and possible applications for securing ICS were outlined.

4.3 Are you threatening my Hazards?

Marina Krotofil (Hamburg University of Technology, DE)

Joint work of Krotofil, Marina; Larsen, Jason
Main reference M. Krotofil, J. Larsen, "Are You Threatening my Hazards?," in Proc. of the 9th Int'l Workshop on Security (IWSEC'14), LNCS, Vol. 8339, pp. 17–32, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-319-09843-2_2

Cyber-physical systems (CPS) are characterized by an IT infrastructure controlling effects in the physical world. On one hand, embedded computers enable governing of physical applications to achieve desired outcomes. On the other hand, physical systems can be instructed to perform actions that are not intended in the same way. Cyber-attacks on physical systems are correspondingly called cyber-physical attacks. The implications of this class of cyber-attacks (the ability to inflict physical damage) is the main difference between cyber-physical and conventional cyber-attacks.

In the context of CPS, safety systems have the critical function of detecting dangerous or hazardous conditions and taking actions to prevent catastrophic consequences on the users and the environment. Relationship between safety and security is usually considered in the context of dependable computing with a focus on IT or system-design. The limitation of this approach is lacking the view of what is happening in the physical world and whether the system remains safe. In our approach we take physical processes into account and propose a simple mental model of the relationship between security assisting why the understanding the hazards within the systems is crucial to designing its security architecture.

The talk covered such aspects as trustworthiness or veracity or process measurements, interactions of the cyber/physical systems ("unexpected physics"), hidden impact data ("unseen" influence of components on each other), hazard-aware security zoning. The latter one allows to harmonize security and safety lifecycles. The system remains secure if updated often (e. g. patched); the system remains safe if untouched. A granular architecture can be created by tracing specific hazards back through a cyber-physical system matching specific devices and specific data flows caring data related to a specific hazard. Components associated with severed hazards must be protected more vigorously.

4.4 Secure Our Safety: Building Cyber Security for Flood Management

Dina Hadziosmanovic (TU Delft, NL)

In this short talk Dina Hadziosmanovic presented the key idea of an upcoming project. This project aims at improving the cyber security of critical infrastructures by bridging the gap between safety and security risk management and monitoring. The project uses the context of flood management to provide integrated decision support for incident response related to cyber threats, based on both safety and security science. Firstly, it works on enriching network security monitoring with safety context information. Here, the context consists of static information about the underlying physical process, as well as dynamic information about safety threats (i. e., extreme hydrometeorological conditions). Secondly, the project addresses how to update safety incident response by procedures that include information from security monitoring in assessing the expected effectiveness of responses. The integration of the two innovations will enable adequate responses to flood defence security threats.

4.5 Security Assessment and Intrusion Detection for Industrial Control Systems

René Rietz and Andreas Paul (BTU Cottbus, DE)

A targeted-oriented improvement of the IT security of Industrial Control Systems (ICS) requires an initial evaluation of the current level of security. Although relevant industry standards in the ICS security domain (e. g., ISO/IEC 27000 series) have already made reference to the need for a comprehensive security analysis, there is no appropriate approach available, yet. This presentation presented a structured, multi-step process for quantitative security assessment of ICS networks. The process yields resilient values, called security indicators, considering the present network topology as well as the technical configuration of the involved industrial network devices. Security indicators can be used to point out potential security vulnerabilities. They also aim to compare the security level of different systems or infrastructures and simulate the application of various measures for improvement. As a novel measure to enhance the security of ICS, an anomaly-detection-based IDS approach was further presented in this talk.

4.6 ICS Security – Challenges, State of the Art and Requirements

Ulrich Flegel (Infineon Technologies – München, DE)

Security for industrial control system faces a radically different landscape than security for modern information systems. Although these old acquaintances again raise their ugly faces. The talk characterized ICS and summarized the main challenges that security practitioners face when securing such systems. Following it described the current approaches and pointed out their challenges. Ulrich Flegel then summarized the state of the art of available solutions as a baseline to improve upon and finally provided a set of requirements – or rather limitations – that new solutions need to meet.

5 Working Groups

5.1 Security Consequences and Quantitative Risk Analysis

Moderation: Stefan Katzenbeisser (TU Darmstadt, DE)
Minutes: Rens van der Heijden (Universität Ulm, DE)

5.1.1 Problem Statement

The goal of this working group was to discuss how to analyze the risk posed by cyber attacks on industrial control systems, with a focus on the potential consequences of these attacks. The quantification of these consequences should help to improve risk assessment

methodologies, and it will allow the design of better detection and preventive measures. This also allows the inclusion of attacks that rely on the attack prevention or safety systems to cause damage, i. e., by triggering a response from these systems the attacker causes financial damage.

5.1.2 Discussion Topics

The discussion covered several topics with opposing viewpoints, specifically regarding how risk analysis should be performed and how much of the physical processes should be considered. We discussed a separation of an ICS into a cyber and a physical component, with limited interaction between them. It was proposed that the physical component mainly needs to be analyzed for safety issues, which can be analyzed using fault trees and analyses like Failure Mode and Effects Analysis (FMEA). Similarly, in the cyber component, attack trees can be employed, as well as analyses like CORAS.

Another question that was discussed is how the cost of different identified security consequences should be considered in risk assessment. It was found that the cost for the compromise of different assets is significantly different in systems with multiple stakeholders, such as a smart metering scenario. This led to a discussion of whether risk assessments should include a focus on societal cost and how quantitative assessments can be adapted for these purposes.

The working group noted that the interaction between the cyber and the physical component is especially interesting for detecting attacks, and therefore should also be considered during risk analysis. Several related important challenges in risk assessment for ICS can be partially addressed in this way.

One of these challenges is that the amount of possible attacks and the lack of information regarding actual attacks mean that quantifying the risk is a difficult challenge compared to normal risk analysis common for security. This can be compensated by the fact that the interaction points between the cyber and the physical system are limited. In addition, this allows a tighter integration of safety risk analysis and security risk analysis.

Another challenge is that the typical computation of risk (defined as cost times probability of occurrence) is not a very good metric for security, since the probabilities are unknown, and often hard to estimate. This leads to a very low return-on-investment, which makes the improvement of security a problematic issue. However, although many different attacks can have low probability, the attacker needs to exploit only one of them. Risk assessment should take into account the physical components and the safety risks associated with them.

The participants identified several classes of security consequences:

- Blackout or shutdown of production process
- Damage to the environment/equipment
- Non-optimal operation (higher cost/lower product quality) within production parameters
- Forcing the use of backup processes that are lower quality

Several concrete ideas to improve risk assessment were proposed in the working group.

First, risk analysis should include the cost for the attacker, and identify measures that increase this cost significantly. This especially includes the analysis of potential feedback channels for the attacker; it should be difficult for the attacker to determine the state of the system, so that she cannot immediately determine the success of different phases of the attack. This requires the attacker to build a testbed (as done with Stuxnet).

A second proposal was to reason from the physical system and determine the potential points of attack. This allows integration with safety analysis on the one hand, and development of more resilient physical systems on the other.

Standardization to provide better generic risk assessments and security in general. Current standards (e. g., Modbus) allow custom systems on top of them, and this makes interoperability hard. Interoperability tests do exist, but differing implementations of standards make both security and risk assessment harder and more costly.

One point of disagreement in the discussions was how useful techniques like threat analysis and attack trees are. These approaches are commonly employed for risk analysis in IT-security, but they may not transfer to ICS as well as expected, due to the highly targeted attacks that should be considered. On the other hand, these methodologies allow the use of standardized risk assessments.

5.1.3 Conclusion and Open Challenges

The working group discussion showed that risk analysis for security in ICS requires specific attention beyond employing standardized techniques from IT-security. Several challenges for risk assessment were identified. First, the fact that different stakeholders have different inputs, which leads to widely different risks for different assets, was identified as an open question. Second, the disconnect between risk assessments for safety and for security, respectively for the physical and the cyber component, is an area requiring improvement. Furthermore, it appears that it is especially hard to model risks that do not lead to a conclusive damage, but rather to non-optimal operation. Finally, there was strong criticism expressed regarding the reliance on attack trees, especially considering the fact that there are very few incidents that can be used as a basis for estimating attack probabilities.

5.2 Industry 4.0

Moderation: Nils Aschenbruck (Universität Osnabrück, DE)
Minutes: Alfonso Valdes (University of Illinois – Urbana, US)

This working group addressed the rise of "Industry 4.0", the associated diversification of industrial control systems (ICSs) and cyber-physical production systems (CPPSs). ICSs are characterized in this context as cyber monitoring and control of processes that interface with the physical world. They are targeted specifically at the industry market, which excludes for example medical devices. Some claim that compared to SCADA, ICSs refer to a more distributed management architecture, although this was controversial. Industry 4.0 refers to a broader and more research-oriented field, although it is still targeted at industry. The term refers to a next generation of factories and farms, which are characterized by highly automated and highly customized networked machines. Customization may even include a end-user specification of products, with checks that the final product meets constraints. Lastly, cyber-physical production systems (CPPSs) are also production- and processing-oriented; some argued they are the same as ICSs. Overall, the exact terminology that should be used is not universally agreed upon and changes over time – therefore, the working group put a focus on use cases that are appropriate for the study of these systems.

5.2.1 New Security Challenges

One of the core issues for each type of system discussed in the working group is the rise of new security challenges by the increased degree of networking introduced by all these ideas. As discussed in the plenary sessions, the increasing connectivity of ICSs in the past has already given rise to a multitude of security issues for current systems. It is important for research to preemptively seek out new security challenges that may arise from future developments.

One important class of challenges relates to the physical consequences that attacks may have for the surroundings of a particular system. These consequences can often be considered as distinct from consequences for the controlled process, such as loss of production and loss of efficiency. A traditional example is the blackout scenario of a smart grid: although the loss of production is a serious problem for the energy provider, the much larger risk of a regional or national loss of the grid plays a big role. Similar issues exist in the different use cases that were discussed in this working group; for example, chemical manufacturing may lead to industrial spills, explosions may occur in many different manufacturing scenarios and pesticides may be released improperly in industrialized agriculture systems. Furthermore, critical infrastructures such as dyke control systems can fail to prevent or even amplify natural disasters such as floods. Traffic accidents are another risk that is associated with many of these use cases. These separate physical consequences are also referred to as societal harm. There is a fundamental distinction between the loss of production, such as manufacturing and agriculture systems, and the loss of control, which is a failure to prevent a disaster.

Privacy is another upcoming challenge for many systems, particularly those that interact closely with end-users. Examples include traffic monitoring systems, smart grids and building automation systems, all of which are typically operated by commercial or sometimes governmental entities. However, research has shown that the sensors in these types of systems can easily be used to deduce significant parts of peoples' lives. On the other hand, the collected data is often essential to the functionality offered by those systems. An open question is who owns this data, and what a fair trade-off is between user privacy and potential profit.

A related but distinct issue is that of confidentiality of data stored in Industry 4.0 systems. In addition to the customization data related to specific end-users, which relates closely to privacy, many ICSs transmit and store a huge amount of data about the monitored process. As the process is often central to the production of a product, this data is of significant value for industrial espionage. The data presents a new way through which company proprietary information may leak into the hands of an adversary. Protecting this information may be a key element in motivating strong security for Industry 4.0.

A final challenge that was discussed is the scalability of existing security mechanisms. Due to the extreme decentralization that is often associated with Industry 4.0, many existing security mechanisms may no longer be sufficient. For example, distributing and regularly exchanging key material and certificates for each and every sensor in a large factory is a difficult process to scale. In the case of industrialized agriculture, compromise of individual sensors or actuators is nearly unavoidable, due to the highly distributed nature of these components, both in space and in sheer number of networked devices. In addition to this challenge, data replication and synchronization will play a role, especially in scenarios where ad-hoc networks are used, as they may not be permanently connected.

To avoid a strong focus on the electrical grid use case, which is a common example, some additional usage scenarios covered by Industry 4.0 were discussed, as their requirements may differ significantly from the electrical grid. The discussed scenarios included future factories

and production systems, environmental monitoring and control, future farms and transport monitoring and control. These scenarios were contributed back to the plenary discussions and were used to avoid a strong focus on the electrical grid use case, whose control loop has very specific properties.

Future factories will feature enhanced customization for customers through increased intelligence in the individual components performing the production process. By performing small changes to the parameters provided to these components, a highly customizable process is created. Important applications of this include customized cars, planes and kitchens, which can be produced more effectively this way. Similarly, multiple companies may order customized components from a single manufacturer. In such a case, there is a significant risk of leaking proprietary information. The potential for collaboration between companies to increase the attack surface of a third company is also a risk that should be considered. A potential solution may be cooperative defense through specifically designed security policies.

5.3 (Real-time) Detection of CPS Attacks

Moderation: Alfonso Valdes (University of Illinois – Urbana, US)
Minutes: Rens van der Heijden (Universität Ulm, DE)

5.3.1 Problem Statement

This working group addressed several questions surrounding attacks on cyber-physical systems. The primary proposal for the working group was to address network attacks on industrial control systems, which was limited to real-time network-oriented detection to prevent overlap with the working group on (mainly host-based) forensics.

5.3.2 Discussion Topics

The discussion was started with a follow-up from the results of a proceeding working group on Industry 4.0 and diversity of ICS, to avoid an extensive discussion on which systems we should discussed. From that session, we took several example scenarios: advanced industry, environmental control, future farms and transportation, which served as a guideline for the variety the discussion should address. This turned out to be an important point, due to the fact that many process-centric detection solutions cannot cope with all these scenarios, but rather need to be defined for each use case separately.

A recurring theme in these discussions was identified to be the question of separation between safety and security issues. Purely safety issues should be addressed by control engineers, not by the security community. It was proposed that the security community should avoid involving itself too closely with the specifics of process control, as this leads to several problematic issues. Not only does this lead to a strong dependency on the process, meaning mechanisms need to be customized per process or even per individual factory or production work flow, but the security community specifically needs to avoid building a shadow SCADA system. Finally, effective reaction to security incidents seemed a major open question. Thus, three different general issues were identified at the start of the discussion:

1. combining the constrained environment with standard security management and best practices

2. the availability of semantic knowledge of the process
3. how to effectively and proportionally react to (potential) security incidents

5.3.3 Security Management

The effective deployment of security management has been an active topic in CPS for quite some time. Common problems and example scenarios of failure are often related not to specific attacks on the CPS, but rather relate to poor security management; lack of isolation where appropriate, lack of updates, poor security policies and lack of security awareness. Big examples like Stuxnet are a demonstration of the potential of targeted attacks, but the vast majority of real world compromises are due to old viruses that still affect unpatched systems. This issue strongly relates to the age of the deployed systems, backwards compatibility requirements and the fact that updating often violates guarantees provided by the manufacturer of the system. The working group discussed briefly about why general purpose CPUs/OSs are used, rather than proprietary systems, if generic vulnerabilities are really the cause of the vast majority of incidents. However, it quickly became clear that this was necessary mainly due to cost and tool availability, and it was additionally concluded that proprietary systems could not have offered better security anyway.

5.3.4 Semantic Knowledge

The use of process semantics seems to be an attractive way to further improve detection rates. Because this requires sufficient redundancy in measurements, studying the placement of additional sensors also seems in scope. However, a strong criticism against this is that this is not the domain of security. Rather, safety and process engineers should be responsible for sufficient redundancy in messages and for the detection of safety-relevant events, regardless of whether they come from an attacker or random failures. This view was somewhat controversial, but supported by the fact that incrementally adding process semantics to intrusion detection will lead to an increase in system complexity, and will eventually result in the mirroring of the SCADA system, referred to as a shadow SCADA system. Although some discussion pointed out that mirroring the process monitoring is an approach used by safety and for fault tolerance, the step towards Byzantine fault tolerance was considered far to expensive, due to the fact that $3n + 1$ sessions are required to tolerate n Byzantine faults.

5.3.5 Effective Response

In regular IT security, responses to attacks can often include simply disabling or isolating the affected system, such as a compromised web server, from the network. However, for CPS attacks, this approach is not a reasonable solution, since this response can often lead to significant financial damage. In addition, the attacker can likely exploit such intrusion responses by designing attacks specifically to trigger these responses. In such a case, the response may cause more damage than it prevents

5.3.6 Conclusion and Open Challenges

We identified a fundamental question; is detection of exploits on the physical component, as well as effective placement of additional (physical) sensors, an IT security problem? As raised in the discussion, there are many arguments supporting the use of information about the process to improve detection rates. On the other hand, we need to be weary of the shadow SCADA problem – the security system should not replicate SCADA functionality. In the

end, we agreed that significant disagreement exists on whether process-aware detection is a good idea.

Apart from this result, we also concluded that in reality, many of the research issues we are discussing are not (yet) a problem. Rather, the current challenges faced by the industry are mainly regarding more general IT security challenges, many of which can be found by security consultancy and subsequently solved by standard IT security measures. However, effective deployment of these measures and practices is expensive and challenging due to scaling issues and certification. One potential source of improvement is the certification of processes to change running systems, so that guarantees from manufacturers are preserved while allowing to implement standard security practices to keep systems up to date.

5.4 Cyberforensics

Moderation: Heiko Platzlaff (Siemens – München, DE)
Minutes: Stephan Kleber (Universität Ulm, DE)

Proposed discussion topics included:

- Forensic readiness in ICS
- Machine Learning in forensics (ML)
- Forensics of operational data
- Live (monitoring, SIEM) vs. dead (post mortem) forensics where crucial problems of pro-actively making systems forensic ready are non-disturbance of business processes, volatility, influence on operational data
- Differences of forensics between ICS and conventional IT systems
- Safe and secure storage for forensic data is missing
- Is it about discerning host from network?

5.4.1 Abstract

The working group "Cyberforensics" did pursue the question about how forensic analyses in the ICS context can be made more efficient, or even feasible in the first place. One option is to employ methods of machine learning to automate certain steps of the process. This is especially useful at the beginning of an analysis, by filtering out the data to be investigated manually afterwards. Another discussion is the role of SIEM in forensics to correlate incidents, even with data not present in PLCs any more. Based on that is the ranking of the importance of a certain entity for a successful forensic analysis. Finally the central difference in terms of forensics between conventional systems and ICS can be a chance: Context aware analysis makes it feasible to take semantics of operational data into account. Especially PLCs, however, need to be made forensics ready first, for the best of results.

5.4.2 Machine Learning

There are not many approaches known to us about machine learning (ML) in forensics, and conventional strategies are not applicable. However there are enough raw data points to enable ML, but it is learning with few labels resulting in semi-supervised learning strategies.

Starting points for ML could be to learn signatures for software. This resembles an "intrusion detection system (IDS) post mortem" in a manner of speaking, although this point of view is criticized in the IDS community. On the other hand side effects on log files need to be collected. Finally the algorithm should come up with five to six most important events identifying "very likely intrusions" to be reviewed by the forensic analyst. For an informed decision on the likelihood of a false positive for the current issue, the analyst needs to know the details of the reasoning that lead to the conclusion of the ML algorithm. The approach is per se not false positive free but manual review is necessary nevertheless. This allows for higher false positive rates (10 % and above) than are generally allowed for IDS. The main goal here is to reduce the amount of data to be reviewed by filtering. A good basic property for filtering are timestamps of file changes, i. e., how they correlate to genuine updates.

Under the assumption the majority of systems is clean, this majority can be used as a baseline for heuristically determining that and possibly where something went wrong. The alternative approach to infect and replay a specific instance of an previous attack is not difficult in itself, but laborious for which reason it is only rarely performed.

Filtering using persistence, time and signature filters brings down the number of findings to a small amount of files, i. e., for one incident some tens of files. Then typical modes of operation of malware can be a good starting point for further analysis or filtering. It might get necessary to extend the time frame of such filtering gradually. The algorithm is thought to rate changes by their age and descent deeper into the past until an event corresponding to the trigger of the analysis can be identified. Usually there actually is a time-based trigger restricting the time frame of the search.

Currently ten to twelve filters on different features are enough for a useful reduction of data. The relative importance of features may be determined by ranking from ML. There has been done numerous work on classification of malware that resembles the general ranking problem here.

Encryption of executables in general is a problem for ML, whereas PE headers (Windows executable file headers) cannot be encrypted to remain executable. Even most changes typical for obfuscation and packers are detectable through characteristic changes in the header. In general a mismatch of parameters results. For example a low number of imports for a large executable file size is suspicious. False positives can arise when reverse engineering protection mechanisms are in place, that function similarly to a malware packer.

Filter layers that might be evaluated sequentially for more details potentially are: Network, Files, Host, SIEM.

ML is not a silver bullet, but it is able to lead the way by pointing towards important details. Dynamic analysis in this setting might show active parts of an entity. Forensics may for example identify executables in unusual places in memory. But here the important question to answer is, how to find the boundaries for discerning between "unusual" and "usual" in memory analysis.

5.4.3 The Role of SIEM in Forensics

Live analysis of an ongoing attack is easier and needs less effort than static analysis afterwards. Forensic action is generally regarded as reactive and post mortem, in contrast to IDS. But evasion strategies of attackers on the host might cover actions and prevent forensic analysis. Moreover, resource constrained systems prevent a gapless monitoring during operation. Network traces done by a network data collector like a SIEM could be used to fill those gaps, but the general assumption about SIEM is that it will alert live if something is going on. If this holds, no new information may come from this analysis of SIEM traces. Moreover SIEM

is highly dependent on its configuration.

SIEM may provide an inventory of entities, which may be hosts, network components and the like. Each of those can be assigned a priori with a rating of importance for the operation of the whole system. Based on that information the forensics team may be able to find anomalies more targeted.

In this respect it is questionable whether network and host forensics are two disjunct things. Multiple things may have been happening at different places at the same time.

5.4.4 Correlating Time and Events

As discussed earlier an analyst needs help to be able to make informed decisions about possible findings based on the rating of the filtering and ML algorithm. SIEM already is a system based on expert knowledge and rules derived of that. This information about the set of rules that matched for an incident can be given to the analyst.

Based on time-event and time-time correlation, indirect correlations can be revealed. More can be added by considering other features, like user-IDs involved in a file change event. Going beyond that, causal relations of events can be useful to be inferred automatically. But it is unclear whether this even is possible in full extent. Maybe the typical human approach might help to look at: A human would identify and extract one feature at a time to go after in the system, pointing to important places to do actual forensics at. This might then even lead to the analysis and correlation of an event in SIEM or historical network data to backtrack related events.

6 Open Issues

In the final discussion, the seminar led to a number of important conclusions and open research challenges that participants agreed should provide important directions for the future of ICS security research and practice:

- Can Intrusion Detection Systems actually provide better security by becoming "process-aware", i. e., have detailed information about the process that the ICS controls? While this seems intuitive, others argue that everything done in this direction is simply replicating the control system and provides redundancy but not necessarily better security.
- The interaction between safety and security mechanisms is an important aspect and needs further analysis. While today often treated separately, we think that both areas should work more closely together to work on unified mechanisms.
- ICS, also those beyond Critical Infrastructures, should generally have 'last-line-of-defense' monitoring and safety mechanisms that are not connected and not coupled with the potentially attackable ICS. Those mechanisms should provide a ground truth to operators and prevent the system from entering clearly forbidden states.
- Proper reactions to attacks are often very hard to determine for ICS, as a sudden shutdown or disconnection may not be a viable option. ICS security mechanisms like Intrusion Detection and Prevention Systems should therefore be able to provide a flexible reaction to detected security breaches to allow a form of "graceful degradation". So as in the case of safety mechanisms, ICS should enter more robust and fail-safe states when attacks are detected, perhaps to the detriment of efficiency and output of the controlled process.
- More attention should be paid to user interfaces of security mechanisms to allow operators and security experts appropriate analysis and reaction if attacks cause critical situations.

- Security systems should provide more fine-grained output to allow better forensics and proper reaction to incidents.
- We identified a huge gap between ICS security research in academia and industrial practice. While research targets highly sophisticated attacks and countermeasures, many real-world deployments fail because of lack of even the most simple security best-practices. Closing this gap will require a huge effort that should start with identifying which best-practices have to be applied and which do not fit.
- ICSs also pose big challenges for security management because of the huge scale of some installations, the lack of realistic attacker models that would allow one to find the right level of security, and the economic pressure to build cost-effective security solutions.
- In general, diversity and redundancy are good for ICS security. If a large number of ICSs are from a single vendor and use only one brand of devices, attacks and malware can easily spread and create huge damage. It is therefore not clear yet, whether convergence of ICS to a few vendors and standards (in terms of protocols, operating systems, etc.) will provide more benefits to attackers or to defenders.
- The fact that ICSs are often very long-lived installations and that duration of innovation cycles in ICSs is very different from ICT creates huge problems for maintaining ICS security. Well-defined, certified update processes that are guaranteed for the lifetime of ICSs would significantly support security. However, maintaining own ICS software ecosystems also has economic consequences.
- Separation and isolation (like air gaps, virtualization, sandbox, VPNs) are likely the most effective security mechanisms for ICS. As a corollary, this means that multi-stakeholder ICS like power grids are inherently harder to secure, as they require more interfaces between parties.
- While ICS is a very broad term and encompasses a lot of extremely heterogeneous types of systems, participants were confident that the security challenges to be addressed are often very similar and thus that there can be meaningful progress on ICS security in general without the need to divide the field into further sub-disciplines.

We hope that this list can provide beneficial input to the field and pave the way for future research leading to more secure Industrial Control Systems.

We thank the seminar participants for their active participation and the fruitful discussions. Their contributions formed the basis for the findings presented here.

Participants

- Ali Abbasi
University of Twente, NL
- Magnus Almgren
Chalmers UT – Göteborg, SE
- Nils Aschenbruck
Universität Osnabrück, DE
- Gunnar Björkman
ABB – Mannheim, DE
- Damiano Bolzoni
University of Twente, NL
- Alvaro Cárdenas Mora
University of Texas at Dallas, US
- Marco Caselli
University of Twente, NL
- Jorge R. Cuéllar
Siemens AG – München, DE
- Hervé Debar
Télécom & Management SudParis – Evry, FR
- Sven Dietrich
City University of New York, US
- Ulrich Flegel
Infineon Technologies – München, DE
- Dina Hadziosmanovic
TU Delft, NL
- Frank Kargl
Universität Ulm, DE
- Stefan Katzenbeisser
TU Darmstadt, DE
- Richard A. Kemmerer
University of California – Santa Barbara, US
- Stephan Kleber
Universität Ulm, DE
- Hartmut König
BTU Cottbus, DE
- Marina Krotofil
Hamburg University of Technology, DE
- Pavel Laskov
Universität Tübingen, DE
- Michael Meier
Universität Bonn, DE
- Simin Nadjm-Tehrani
Linköping University, SE
- Heiko Patzlaff
Siemens – München, DE
- Andreas Paul
BTU Cottbus, DE
- Konrad Rieck
Universität Göttingen, DE
- Rene Rietz
BTU Cottbus, DE
- Robin Sommer
ICSI – Berkeley, US
- Radu State
University of Luxembourg, LU
- Jens Tölle
Fraunhofer FKIE – Wachtberg, DE
- Alfonso Valdes
University of Illinois – Urbana Champaign, US
- Rens van der Heijden
Universität Ulm, DE
- Alexander von Gernler
genua Gesellschaft für Netzwerk- und Unix-Administration mbH – Kirchheim bei München, DE
- Stephen Wolthusen
Royal Holloway University of London, GB & Gjovik University College, NO
- Emmanuele Zambon
SecurityMatters B. V. – Enschede, NL