Lights on Power Plant Control Networks

Stefan Mehner, Franka Schuster, and Oliver Hohlfeld

Brandenburg University of Technology Cottbus – Senftenberg
{stefan.mehner,franka.schuster,oliver.hohlfeld}@b-tu.de

Abstract. Industrial Control Systems (ICS) are critical systems to our society. Yet they are less studied given their closed nature and often the unavailability of data. While few studies focus on wide-area SCADA systems, e.g., power or gas distribution networks, mission critical networks that control power generation are not yet studied. To address this gap, we perform the first measurement study of Distributed Control System (DCS) by analyzing traces from all network levels from several operational power plants. We show that DCS networks feature a rather rich application mix compared to wide-area SCADA networks and that applications and sites can be fingerprinted with statistical means. While traces from operational power plants are hard to obtain, we analyze to which extent easier to access training facilities can be used as vantage points. Our study aims to shed light on traffic properties of critical industries that were not yet analyzed given the lack of data.

1 Introduction

Insights gained from two decades of Internet measurement research enabled the evolution and optimization of Internet technology—including Internet performance and security. This research has provided many fundamental results on network operation or network traffic that form the basis for network planning or optimization (e.g., the finding of the self-similar character of Internet traffic [9] or the evolving application mix [11]). These efforts focus on the Internet as the largest and general-purpose communication network. In contrast, important application-specific networks that increasingly rely on Internet technologies received less attention from a measurements' perspective yet, e.g. power plants.

The hidden networks. Critical industrial systems such as power plants rely on industrial control systems (ICS) for their operation. These systems are based on proprietary protocols and typically closed. Especially the access to critical infrastructures in the energy and water sector is highly restricted, which limits the potential for conducting measurement studies. As a result, little is publicly known about networks, which limits research potential for enhancements, e.g., to improve their security. While these networks experience an attack vector of increasing relevance [8], the little public knowledge about their properties and functions limits the design of mitigation strategies. By studying traffic-level properties of operational power plants, we make step towards closing this gap.

Related work. Prior work on measurement of *real* infrastructure traffic focused on SCADA networks. While SCADA networks control the interactions

of dispersed assets to enable power and water distribution, Distributed Control Systems (DCS) are dedicated to the control of the local core processes in power plants and water treatment sites [15]. The most recent SCADA work shows that IEC 60870-5-104 is used as the only protocol in the studied infrastructure [10]. In contrast, our work will show that DCS systems feature a much richer application mix. Other studies investigated in water treatment and distribution facilities, a gas utility as well as an electricity and gas utility. They show that SCADA traffic largely differs from Internet traffic given the absence of human users (and thus diurnal patterns) and self-similarity [3]. In later works, the same authors prove that traffic in SCADA networks is periodic and provides a stable connection matrix [2,5]. Since these prior measurements focussed on SCADA infrastructures, we will complete the picture (for the energy sector) by focussing on the DCS part of the energy supply by investigating four operational real power plants.

Our contribution. In this first of its kind study, we shed light on traffic properties of critical yet unstudied type of infrastructure networks: internal (i.e., not intentionally Internet facing) control networks of three power plants and one power plant simulation facility. Our study is enabled by the rare opportunity to capture traffic traces during maintenance slots at three operational power plant sites. Our main objective is to provide a first empirical perspective on these otherwise hidden networks since traffic properties pave the way for controlled simulation and evaluation studies. Our contributions are as follows.

- We show that DCS networks in power plants feature a rich protocol mix that differs by automation layer. This is in contrast to typical SCADA networks that are often realized using a single protocol only. That is, while a recent survey [6] found in every ICS testbed one dominant protocol, we show that in power plant DCS features 'zoo' of protocols on different network levels.
- We show that the proprietary and publicly undocumented ICS protocols used can be identified by applying statistical clustering approaches. These clusterings even work in the absence of payload by analyzing inter-arrival times and header information.
- We finally applied our methods to a dataset from a plant simulator that serves as training facility. Given that measurements in operational power plants can only be conducted in rare maintenance windows, we study if easier to access simulation facilities are an option for measurements.

2 ICS 101

ICS aim at controlling and supervising machines or processes (e.g., coal firing).

SCADA vs. DCS. There is not one type of ICS, but many [7]. Most prominently, both SCADA and DCS systems enable the supervision and control of industrial processes. Studied types of SCADA networks are wide-area networks that are scattered over hundreds of kilometers, e.g., bulk power grids for power distribution [10] or water treatment and distribution facilities [2,5]. In contrast to the widely studied SCADA network, DCS offer integrated solutions provided by a single vendor and are thus often designed for local use—like the power plant

networks studied in this paper. Consequently, the used protocols and traffic patterns in both network types differ and, as a consequence, it is not possible to infer the characteristics of one network type from the characteristics of the other. For example, IEC 60870-5-104 (one of the most widely used SCADA protocols, e.g., studied in [10]) is a telecontrol protocol that does not play any role in DCS. Also, SCADA networks are dominated by a single or few protocols, while—as we will show—the integrated nature of DCS systems yields a much richer protocol mix. The unstudied nature of DCS traffic thus motivates our work.

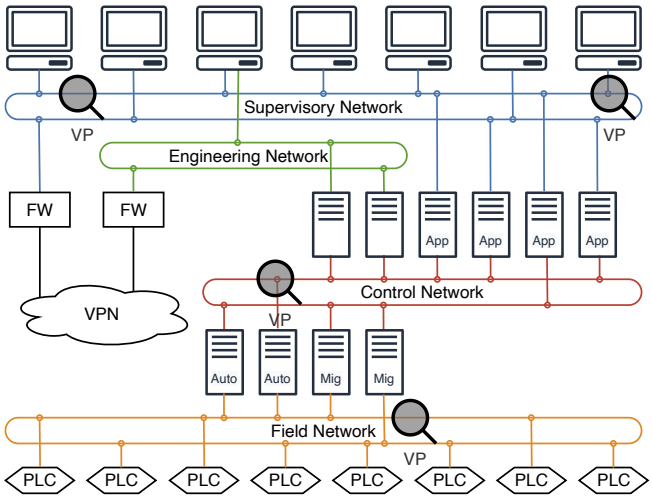

Fig. 1: Network architecture levels of plant DCS with vantage point locations.

ICS Architecture. While ICS differ in their physical extension and associated network characteristics, they are quite similar in terms of network organization. Typically they follow a layered, hierarchical design. The lowest network part, the *Field Network*, consists of the physical level, including sensors and actuators measuring and adjusting basic physical parameters (e. g., temperature, pressure, speeds, feeds). They are connected to programmable logic controllers (PLC). Each PLC controls actuators and/or monitors sensors and hence realizes a low-level control according to an implemented application-specific control loop. In the *Control Network*, inputs from various PLC are collected and evaluated by subsystem-specific servers in order to aggregate all activities corresponding to a subprocess. For power plants, these sub-processes include, for example, temperature monitoring of the boiler. There are also servers that prepare the process data for visualization on the Human Machine Interfaces (HMI) and others that connect legacy systems to the control system. Monitoring and manual adjustment of the physical process is realized in the *Supervisory Network*. Here the operator interacts with the subsystems via HMI that incorporate the data from the Control Network.

Power Plants. ICS in plants use DCS architectures, often realized as large proprietary networks provided by a single vendor. As such, the concrete network

design including its configuration (e. g., addressing schemes) is vendor dependent and differs. We observe three typical server types in plants, which are also depicted in Figure 1. Automation servers (*Auto*) provide so-called automation objects that enable clients to automate well-defined procedures by directly accessing reusable functionality made available by the server. Migration servers (*Mig*) aggregate data from Field-Network communication to representations appropriate for the application servers (*App*) connected to the Control Network. These application servers in turn realize further data preparation for graphical applications, such as the control center's screens, attached to the Supervisory Network. Another characteristic is that these networks are rarely upgraded. In the case of power plants, operators only plan for one network upgrade during the entire lifetime of a power plant, i. e., once installed, the networks run for decades without major changes. This is in stark contrast to classical Internet-networks that are upgraded much more frequently.

3 Power Plant Datasets

Our study is based on packet traces captured at three operational power plants (see Table 1). The traces contain traffic from all physical subprocesses including coal firing, fluid flow, and turbine operation. As our datasets account for two of the leading vendors of control systems that are of widespread use world-wide, they enable us to provide a representative picture of traffic characteristics.

	vendor	level	duration	# packets	# devices
Plant 1	A	supervisory	3.3 h	38 M	39
		control	18.6 h	96 M	44
		field	1.4 h	6 M	52
Plant 2	B	control	54.9 h	17 M	89
Plant 3	B	control	2.7 h	61 M	65

Table 1: Dataset overview

Plant 1. The first dataset was captured in the main process control network of a unit of capacity class of 800/900 MW as part of a two-unit coal-fired power plant. We were allowed to monitor vantage points at all three network levels shown in Figure 1. In this network, process control technology of type A (we omit vendor names due to a non-disclosure agreement) from one of the few major vendors in this field is installed, which is of widespread use in power plants worldwide.

Plant 2. This trace was taken at the control level network of one unit of a multi-unit coal-fired power plant of capacity class (in total) of 1000 MW. The control system is from a different vendor also being a dominant supplier in that field. The main activity of the system is realized by virtualized components encapsulating machine-to-machine communication. Since communication within virtual environments does not leave the hosting machines, virtual machine-to-machine

communication is not seen on the wire and a significant proportion of process communication could not be captured by listening at the chosen network switch. That is why the captured process traffic is incomplete. Since virtualization is a new trend in ICS, this trace represents network traffic visible in newer systems.
Plant 3. The third dataset was taken at a third black coal power plant, also at the control level network of the process control system. In contrast, here the control system was shared among both units representing in total 700 MW installed capacity. Here, the applied process control system is of the same type as the one in Plant 2. In contrast, no system components are virtualized and thus more process control traffic is visible.
Some plant network details. We observed that the interaction between the identified network levels is defined by separated networks, each having its own IP subnet. The communication among multiple levels is realized by a certain set of clients and servers that are connected to one or two network levels. In our case, the supervisory network in the studied plants is not intentionally connected to the Internet for remote access. Remote access is a procedure that has also to be manually initiated from inside the power plants and is only performed when necessary. Except for one system that, due to legal regulations, has to regularly report polution measurements to authorities, no data is made available to outside. In all considered plant networks the IP address assignment is static and the only security devices installed are border firewalls to higher-level networks of the owner company or dedicated networks for remote access. The network link speeds are 1 GBit/s. The average packet rates can be derived from Table 1.
Measurement setup. The data was captured during the downtime of the physical generation process during a regular maintenance. In this time, the physical processes were stopped (i. e., no power was produced). However, the operators and the system vendors confirmed that the DCS is running in normal operation, only links to physical actuators are deactivated. For this reason, we assume the application mix studied in this work to be the same as in normal operation. Yet the content of control messages and network load might differ to normal operation (not studied). We captured all traces using `tcpdump` [1] and port mirroring at one or several switches placed in the network part/s stated to the datasets. Vantage points were chosen to capture traffic from servers relevant to plant operation. Not all vantage points shown in Fig. 1 were available at all power plants.
Ethics. The operators granted permission to capture traffic during maintenance intervals in which it was ensured that our setup cannot impair physical processes. All traces contain only machine-to-machine traffic. We do not reveal ICS protocol details or network configurations that could help in attacking these power plants.

4 The Rich Application Mix of Power Plant ICS Networks

A common assumption is that ICS networks are dominated by a single protocol only, e.g., which controls the automation process. Prior work showed that this is indeed the case in a number of industrial settings and identified Modbus TCP / MMS [3,4,5] and IEC 104 [10] as typical protocols used. This originates from

the dedicated design of many SCADA networks and is in contrast to the typical Internet application mix that is dominated by different protocols (see e.g., as observed at an IXP [14]). The existence of only a single control protocol can ease network management and dimensioning. In this section, we show that this common assumption of a single dominant protocol is not the case for all industrial settings. We show that the integrated DCS architecture of the studied power plants yields a more complex traffic composition than commonly believed by actually representing a mix of different application protocols.

4.1 Application Mix of Power Plant 1

Approach. We refer to an "application protocol" as a protocol used to transmit *application* payload—regardless of the underlying transport protocol. We thus omit network control protocols such as ARP and LLC, but consider COTP that also is a MAC-layer protocol. Frequently used dissectors (e.g., by tshark or Zeek) can identify common Internet protocols. ICS networks, however, often use proprietary protocols which are not recognized by such tools. We thus manually inspected all traffic traces and created payload based identifiers for each proprietary protocol. We further identify known Internet protocols (e.g., HTTP/NTP) by the destination port in Zeek logs. All well known protocols are mapped by their name, while the remaining ones are shown by their port number.

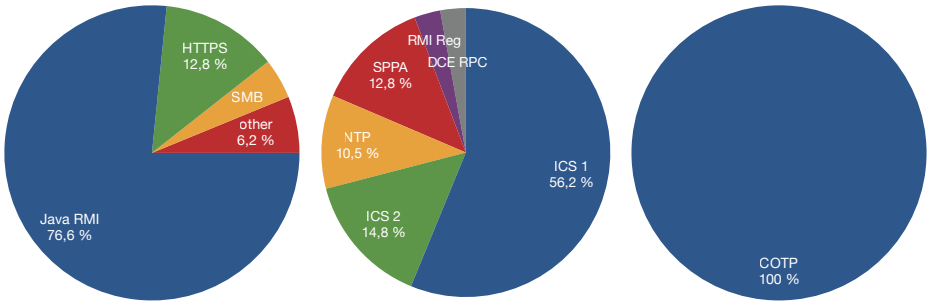

(a) Supervisory Level (b) Control Level (c) Field Level

Fig. 2: Application mix of Plant 1 divided into the hierarchical layers.

Protocol shares differ substantially by level. Figure 2 shows the protocol mix of Power Plant 1 by hierarchical layer. While only COTP is visible at the lowest layer (field), above layers show more protocol variety, as we discuss next.
Field level. We observe one dominant protocol (COTP as identified by Wiresharks dissector) that is served via Ethernet with 100% traffic share. The reason is that on field level point-to-point communication with sensors and actuators (e.g., using Profinet IO to meet real-time guarantees) is used. The sensor data is then directly encapsulated on top of Ethernet as payload using proprietary protocols. At the control and supervisory level, all communication is IP-based

and thus UDP and TCP are the dominant transport protocols. However, both layers differ substantially in transport protocol share. While the supervisory level shows 99.0% TCP communication, the control level shows 40.9% UDP traffic.

Control level. The application mix at the control level features multiple protocols, see the protocol share by packets in Figure 2b. This is rooted in multiple functions that are performed at the control level. First, sensor data is obtained from field level devices and then aggregated and processed by migration and automation servers (see Figure 1). Aggregated data is then prepared for graphical representation (at the supervisory level) by application servers. The communication between the migration and the application servers is based on two protocols. The first is a proprietary TCP-based protocol that we observe on multiple ports, referred to as ICS 1. With a share of 56,2% of the exchanged packets it is the dominant protocol at the control level. The second, referred to as ICS 2, is a UDP-based protocol used to share information between multiple servers via IP multicast. Automation servers, that are responsible for the time-critical automation process, further communicate with application servers by using a proprietary protocol—we call SPPA—on ports 10002 and 10003. In Section 5, we describe in detail our approach of identifying the protocols on the different port numbers. Beyond these proprietary ICS protocols, we also see RPC communication with the supervisory layer using Java RMI (for graphical representation) and further classical Internet protocols such as NTP.

Supervisory level. The communication at the supervisory level is dominated by Java RMI based RPC communication, see Figure 2a again showing the packet share per protocol. Proprietary or classical ICS protocols known in the SCADA domain are absent. Java RMI based RPC communication is used for graphical representation by browser-based thin clients at the supervisory level, which interact with application servers. As in the control level, there is an RMI registry, here on port 1099, with a share of 5.3%. The app mix is dominated by two RMI RPC-based applications each running on a different port with a share of 44.2% respectively 21.4% of the whole trace. In Figure 2a we aggregated all such traffic to 'Java RMI'. Beyond, we see also known Internet protocols such as HTTPS between two servers and SMB for file servers.

Takeaway. *Unlike findings in related works that suggest the presence of only a single protocol in ICS networks, depending on the level, ICS networks can feature richer application mixes comprised of multiple protocols. In terms of protocol complexity, ICS networks at higher layers are thus comparable to typical LANs rather than single protocol networks.*

Bugs can skew the traffic mix. We further observe an unexpected high amount of 68-75% of UDP-based traffic at all levels in this trace. So what is the reason of this? To answer this question we evaluated the application mix and find SNMP to create the high share of UDP traffic. We have consulted the power plant operator and figured out that this is a bug in the control system that unnecessarily pulls status information from network devices. It does not affect the operation of the plant, but disturbs the monitoring of the network components. While similar skews in traffic are normal in Internet-type applications that

get updated often, infrastructure in power plants is typically static, remains in operation for years, and can only be updated during few maintenance windows. Thus changes in the traffic composition can alter network dimensioning and thus might night to be incorporated when simulating networks for planning.

Encryption is uncommon. We observed very little encrypted traffic. The encrypted traffic on port 443 with a share of 12.8% (shown as HTTPS in the plot) was seen only between two pairs of servers in the supervisory level of Power Plant 1. Hence, we observed that encrypted traffic is uncommon, which is in line with similar observations made in other industrial settings [13,4].

4.2 The Application Mix Differs by Power Plant/Vendor

Is the application mix comparable between the different power plants? That is, given that it is challenging to obtain traffic traces from operational power plants, would it suffice to measure one plant and then generalize? We next answer this question by comparing the application mix for the different power plants.

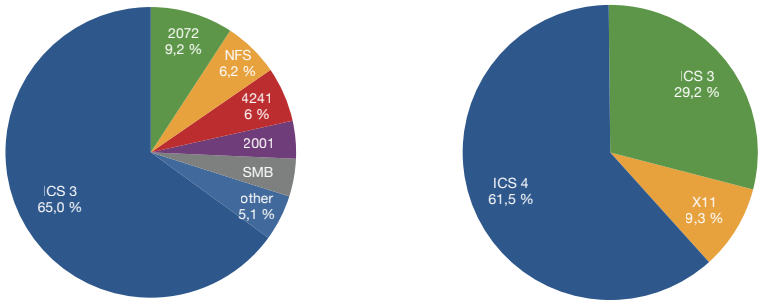

(a) Power Plant 2 (b) Power Plant 3

Fig. 3: Application mix of Power Plant 2 and Power Plant 3

Power plants with different vendor control systems. We begin by comparing the application mix for power plants with different control systems. In our case, the control system of power Plant 1 is by a different vendor than power Plant 2 and 3. In this power plants we have a slightly other network architecture, where supervisory and control level networks are merged. We remark that the operators of that power plants only enabled traffic captures at one vantage point. In direct comparison of the port-based application mix (see Figure 3), we observe different port ranges and number of ports used. At the supervisory level at Plant 1, there is essentially Java RMI traffic running over ports 50000, 50001, etc. At the control level, there are three proprietary ICS protocols that communicate over many ports and represent the largest share of the communication. In contrast, in power plants 2 and 3 we see fewer actively used ports. The (ICS) application protocols also always use the same port number. Further, we see remote procedure call realizations in traffic of both power plant types.

Other application-level characteristics differ. That means, that the port-based application mix differs between power plants that run different control systems. This can be explained by the vendor-driven different design of control systems which usually incorporates at least one vendor-specific protocol whose associated port number(s) individualize the app mix.

Power plants with the same vendor control system. Does running a control system by the same vendor yield comparable application mixes? To answer this question, we compare the application mix of Power Plant 2 and 3—both running a control system from the same vendor. Both plants show a significant share of proprietary TCP-based traffic on port 2010 with 65% respectively 29%. We call this protocol ICS 3 in Figure 3. The largest share at Plant 3, however, is on port 2000 (named ICS 4 in the plot). After consultation with the operator of Plant 2, the ICS 4 protocol is also in use, but we were unable to see it at this VP. In contrast to Plant 2 we also see 9% X11 traffic in the Plant 3 trace. Probably an engineer performed a remote session to servers due to maintenance reasons. Beyond these dominant application ports, other port-level statistics differ. We thus conclude that the two power plants using a control system by the same vendor offer partially comparable application mixes.

5 Towards Understanding the Proprietary ICS Protocols

The ICS protocols used by the studied power plants are proprietary with no documentation provided, not even to the operator. Our study in Section 4 thus relied on a time-consuming manual payload inspection. We therefore now ask if the protocols and their characteristics can be identified purely by statistical means without any a-prior knowledge. We exemplify this study on the control level of Plant 1.

5.1 Clustering Communication by Packet Payload Differences

The payload of a typical ICS protocol mainly consists of few message types and physical (actuator or sensor) values within a well-defined range. Additionally, the nature of automated control loops tendentially lead to recurring physical values transmitted in the payload, generally as well as for each respective protocol. Consequently, it should be possible to identify packets encapsulating the same protocol by comparatively high similarity in the respective payloads.

Approach. We measure differences in packet-level payload by the Levenshtein ratio, which computes the similarity of two binary strings between 0 and 1. To showcase our approach, we extract the payload from the first 10,000 packets for every destination port of interest. Then, using the Levenshtein distance, we compute the similarity for each pair of all packets for every two ports. Finally, we derive the minimum, average, and maximum similarity for every port pair, which is visualized in a heat map as shown in Figure 4.

Payload clustering can trace protocols well. Our analysis reveals that the TCP port ranges 1487 to 3137 and 42239 to 44061 are addressed by communications with highly similar payload (average Levenshtein ratio between 40% and

87%). We call it ICS 1. Concerning the UDP ports 10002 and 10003, which we already know from our prior analysis using the dissector, we identified a average similarity of 74%. We call this SPPA. The port ranges 20202 to 20205 and 20301 to 20304 resulted in an average similarity of 76% and 55%, respectively. On the contrary, e.g., the ports 20202 and 20301 only showed an average similarity of 8%. A look into the trace revealed multicast communication, where both of the ports are either source or destination. Hence, they are considered as related and associated to protocol ICS 2.

Takeaway. *Our payload based classification scheme could—without any a-priori knowledge—identify all used protocols as compared to our manual payload inspection in the previous section.*

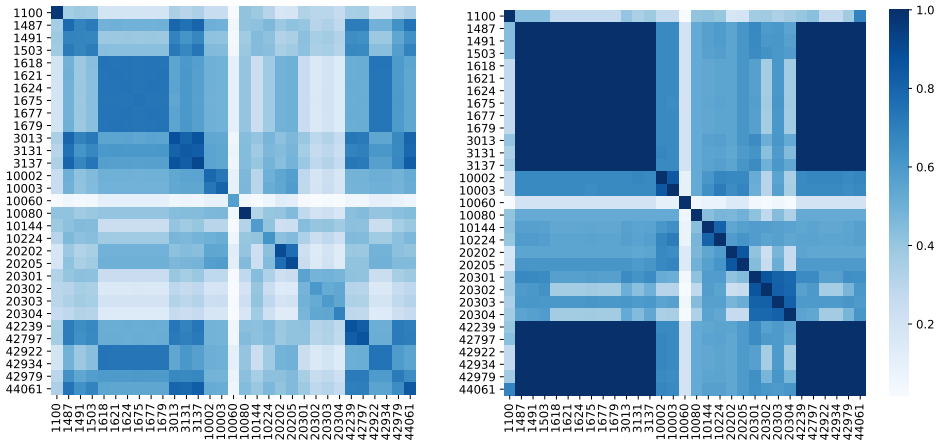

(a) average similarity (b) maximum similarity

Fig. 4: Payload similarity of all ports (TCP and UDP) in the control level trace. The darker the blue color, the higher the similarity of two protocols payload.

5.2 What if We Don't Have Payload?

In public traces usually there is no payload available [12]. Hence, a relevant question is, whether it is possible to achieve similar results by only using available metadata like packet inter-arrival times or packet header information, such as the payload size, the transport protocol used, or unicast/multicast communication? This approach assumes that the same protocols behave similarly with regard to this metadata. If this is the case, no knowledge about the protocol and no time-consuming deep packet inspection would be necessary for a traffic characteristic. As before, we compare the resulting protocol clusters to our previous study.

Approach. We employ the implicit application-level traffic classification approach using the packet size and the packet inter-arrival time distribution by Trivedi et al. [16]. For both metrics, the entire range of values is divided into 50 unequal-sized bins each and normalized. We have adopted this approach and modified it as follows: In contrast to protocols like HTTP, or FTP, industrial

protocols have rather small payloads as often only sensor measurements are transmitted. For this reason, we have chosen other bin sizes as shown in Table 3. The packet inter-arrival times within a flow also differ for industrial protocols, since these have periodic communication with defined cycles. In Table 4 we show the bin sizes we used for our analysis. We normalized the values to the relative share of each bin for both metrics. We consider three cases, payload length and inter-arrival time individually and combined. Furthermore, we added two more metrics: the transport protocol (TCP/UDP) and the communication type (unicast/multicast). To find similar protocol properties, we used two clustering approaches: DBSCAN (with $\epsilon = 0.3$ and $min\_samples = 3$) and k-means (with $k \in \{3..10\}$ cluster).

Metadata clustering does not identify all protocols. Both approaches, the payload clustering and the metadata clustering, look at the same question from two different perspectives. Both clustering approaches find clusters that successfully distinct different protocols, but the results differ slightly. In Table 5 we provide a comparison of the results. Here, the column *Payload similarity* contains the results derived from the payload-based protocol clustering using the Levenshtein ratio, which is used as ground truth, here. The port ranges of ICS 1 protocol were clustered to two respectively three clusters. ICS 2 was correctly clustered by DBSCAN for both metrics. In k-means, ICS 2 was in same cluster as SPPA and DCERPC in most cases. By design DBSCAN was not able to cluster protocols that are using only one or two port numbers, because we set $min\_samples = 3$.

Takeaway. *When payload is unavailable, metadata clustering can identify most protocols. Compared to the payload-based identification, yet at a lower accuracy.*

6 Measuring at a Power Plant Training Facility

Traffic can only be captured in rare maintenance windows of operational power plants. Consequently, we study next to what extent a training facility of a real plant does reflect the traffic observed in the real infrastructure. We consider such a facility as ideal environment for generating representative normal as well as attack traffic. Since this training network is not a productive network, interruptions of the normal network functions do not have any consequences.

Dataset. We capture at a training facility for power plant operators that rebuilds the automation network installed in Power Plant 1. The process control system is of the same vendor as Plant 1, but is an older system version. It only lacks the physical production layer which is completely simulated by software. Further details are shown in Table 2.

Application mix differs substantially. Contrary to our expectations the application mix is completely different in both systems. At the supervisory level there is no RPC based communication, but other TCP-based traffic at 10 different ports with a share of about 10% each. Using our payload based classification approach from Section 5.1 we were able to identify most of this ports to the same protocol, which we will call ICS 5 in Figure 5a. The major part of the

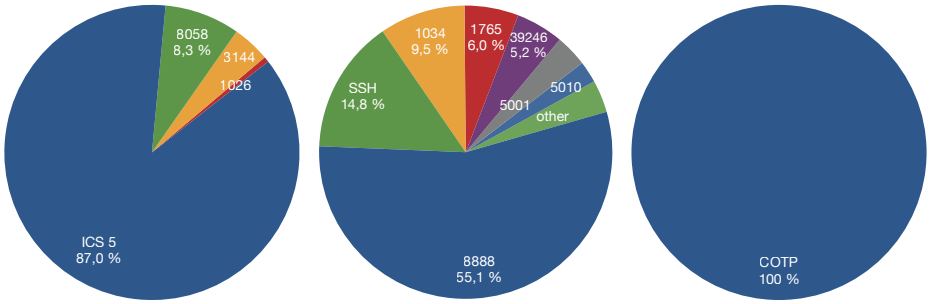

(a) Supervisory Level (b) Control Level (c) Field Level

Fig. 5: Application Mix of power plant training facility

control level traffic consists of one flow on port 8888. Apart from this we see some other proprietary traffic and SSH communication. Obviously the current system version as used in Plant 1 is in most parts a new development. Another reason is that the simulator acts as training facility for Plant 1 and thus tries to completely mimic the control functionality, but it doesn't implement all the power plant functionality. Just as on the field level in Power Plant 1, we also see exclusively COTP in this trace.

Takeaway. *The traning facility shows a completly different protocol mix differs and it might not be sufficient to only capture traffic at simulation facilities to understand real-world ICS networks.*

7 Conclusion

Access to networks that control critical infrastructures, especially DCS, is highly restricted and thus little is publicly know about its properties. We took the rare opportunity to capture traffic traces during maintenance slots at three operational power plant sites and one training facility. Our main objective is to provide a first empirical perspective on these otherwise hidden networks. Unlike prior work that studied SCADA networks, we show that DCS networks in power plants feature a rich protocol mix that differs by automation layer. Any evaluation (e.g., simulation) of plant network traffic must account for these traffic mixes. We further show that the proprietary and publicly undocumented ICS protocols used can be identified by applying statistical clustering approaches—with and without payload. These clusterings even work in the absence of payload by analyzing inter-arrival times and header information. We finally applied our methods to a dataset from training facility that replicates power plant 1. Given that measurements in operational power plants can only be conducted in rare maintenance windows, we study if easier to access simulation facilities are an option for measurements. We observe that the resulting application mixes differ substantially. With this paper we thus aim to shed light on a relevant but not yet studied type of network.

Declarations

Acknowledgement. Franka Schuster acknowledges funding by the German Federal Ministry of Education and Research (BMBF) grant WAIKIKI (funding reference number: 16KIS1198K).
Authors' contributions. This study has been solely conducted by Stefan Mehner (main author) on a previously captured dataset as part of his PhD thesis. The study design was developed by Stefan Mehner and Oliver Hohlfeld. All authors contributed to the discussion and writing of the paper.

References

1. Tcpdump and Libpcap: https://www.tcpdump.org
2. Barbosa, R.R.R., Sadre, R., Pras, A.: A first look into scada network traffic. In: 2012 IEEE Network Operations and Management Symposium. pp. 518–521 (April 2012). https://doi.org/10.1109/NOMS.2012.6211945
3. Barbosa, R.R.R., Sadre, R., Pras, A.: Difficulties in modeling scada traffic: A comparative analysis. In: Passive and Active Measurement (2012)
4. Barbosa, R.: Anomaly detection in SCADA systems: a network based approach. Ph.D. thesis, University of Twente (4 2014). https://doi.org/10.3990/1.9789036536455
5. Barbosa, R.R.R., Sadre, R., Pras, A.: Exploiting traffic periodicity in industrial control networks. International Journal of Critical Infrastructure Protection **13**, 52 – 62 (2016). https://doi.org/https://doi.org/10.1016/j.ijcip.2016.02.004, http://www.sciencedirect.com/science/article/pii/S1874548216300221
6. Conti, M., Donadel, D., Turrin, F.: A survey on industrial control system testbeds and datasets for security research (2021)
7. Galloway, B., Hancke, G.P.: Introduction to industrial control networks. IEEE Communications Surveys Tutorials **15**(2), 860–880 (2013)
8. Hemsley, K.E., Fisher, D.R.E.: History of Industrial Control System Cyber Incidents. Idaho National Laboratory (2018)
9. Leland, W.E., Taqqu, M.S., Willinger, W., Wilson, D.V.: On the self-similar nature of ethernet traffic. SIGCOMM Comput. Commun. Rev. **23**(4), 183–193 (Oct 1993). https://doi.org/10.1145/167954.166255, http://doi.acm.org/10.1145/167954.166255
10. Mai, K., Qin, X., Ortiz, N., Molina, J., Cardenas, A.A.: Uncharted networks: A first measurement study of the bulk power system. In: Proceedings of the ACM Internet Measurement Conference. p. 201–213. IMC '20, Association for Computing Machinery, New York, NY, USA (2020). https://doi.org/10.1145/3419394.3423630, https://doi.org/10.1145/3419394.3423630
11. Maier, G., Feldmann, A., Paxson, V., Allman, M.: On dominant characteristics of residential broadband internet traffic. In: ACM IMC (2009)
12. Mathur, A., Tippenhauer, N.O.: Swat: A water treatment testbed for research and training on ics security. 2016 International Workshop on Cyber-physical Systems for Smart Water Networks (CySWater) pp. 31–36 (2016)
13. Ndonda, G.K., Sadre, R.: A two-level intrusion detection system for industrial control system networks using p4. In: 5th International Symposium for ICS & SCADA Cyber Security Research 2018 5. pp. 31–40 (2018)

14. Richter, P., Chatzis, N., Smaragdakis, G., Feldmann, A., Willinger, W.: Distilling the internet's application mix from packet-sampled traffic. In: Mirkovic, J., Liu, Y. (eds.) Passive and Active Measurement. pp. 179–192. Springer International Publishing, Cham (2015)
15. Stouffer, K., Pillitteri, V., Lightman, S., Abrams, M., Hahn, A.: Guide to Industrial Control Systems (ICS) Security. NIST Special Publication 800-82 (2015)
16. Trivedi, C., Trussell, H.J., Nilsson, A.A., Chow, M.Y.: Implicit traffic classification for service differentiation. Tech. rep., North Carolina State University. Center for Advanced Computing and Communication (2002)

A Appendix

A.1 Power Plant Training Facility Dataset

vendor	level	duration	# packets	# devices
A	supervisory	1.1 h	4 M	25
	control	19 h	20 M	28
	field	2.3 h	1.6 M	11

Table 2: Dataset overview of power plant training facility

A.2 Bin Sizes Used for Protocol Clustering

0	5	10	15	20	25	30	35	40
45	50	55	60	65	70	75	80	85
90	95	100	115	130	145	160	175	190
205	220	235	250	300	350	400	450	500
600	700	800	900	1000	1500	>1500		

Table 3: Bin sizes used for the study to divide the TCP/UDP payload

0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	2.0
2.5	3.0	3.5	4.0	4.5	10.0	25.0	50.0	>50	

Table 4: Bin sizes in milliseconds used for the study to divide the packet inter-arrival times within a flow

A.3 Payload Similarity and Clustering Results

port	Payload similarity	DBSCAN p	Kmeans p	DBSCAN i	Kmeans i	DBSCAN pi	Kmeans pi
1100	Java RMI	-	A	-	A	-	A
1487	ICS 1	A	A	A	B	A	B
1491	ICS 1	A	A	A	B	A	C
1503	ICS 1	A	A	A	B	A	B
1618	ICS 1	B	B	A	B	B	B
1621	ICS 1	B	B	A	B	B	B
1624	ICS 1	B	B	A	B	B	B
1675	ICS 1	B	B	A	B	B	B
1677	ICS 1	B	B	A	B	B	B
1679	ICS 1	B	B	A	B	B	B
3013	ICS 1	C	A	B	C	C	C
3131	ICS 1	C	A	B	C	C	C
3137	ICS 1	C	A	B	C	C	C
10002	SPPA	D	C	C	D	D	D
10003	SPPA	-	C	-	D	-	D
10060	DCERPC	-	D	-	D	-	D
10080	Java RMI	-	E	-	E	-	E
20202	ICS 2	E	D	D	D	-	D
20205	ICS 2	E	D	D	D	-	D
20301	ICS 2	E	D	D	D	-	D
20302	ICS 2	E	D	D	D	-	D
20303	ICS 2	E	D	D	D	-	D
20304	ICS 2	E	D	D	D	-	D
42239	ICS 1	B	B	A	B	B	B
42797	ICS 1	B	B	A	B	B	B
42922	ICS 1	B	B	A	B	B	B
42934	ICS 1	B	B	A	B	B	B
42979	ICS 1	B	B	A	B	B	B
44061	ICS 1	B	B	A	B	B	B

Table 5: Comparison of results from payload similarity analysis as well as the clusterings using payload-length (p), inter-arrival times (i) or both metrics (pi) for DBSCAN and Kmeans (n=5 cluster) clustering approaches