Security Lancaster | Lancaster University

Security and Resilience of Cyber-Physical Infrastructures

Proceedings of the First International Workshop held on 06 April 2016 in conjunction with the International Symposium on Engineering Secure Software and Systems, London, UK

Awais Rashid, Wouter Joosen, Simon Foley (editors)

Lancaster University Technical Report No: SCC-2016-01

Security Lancaster | Lancaster University

Preface

The world is experiencing a massive growth in connected cyber-physical infrastructures – ranging from IoT-based smart environments to critical infrastructures such as power grids, energy, water and manufacturing systems. The number of connected devices is expected to grow to tens of billions by the year 2020. Very large cyber-physical infrastructures are envisioned which will integrate multiple applications run by a variety of stakeholders within a shared fabric. Examples include future industrial environments, infrastructure monitoring technologies and intelligent transportation systems. In such contexts, thousands of nodes will be deployed and used by a large number of stakeholders to provide a multitude of services. Such shared fabrics will remain in operation for a long time (potentially decades) and the physical composition, the services provided and the stakeholders involved will change with time. This scale of future cyber-physical infrastructures and their dynamic nature in terms of stakeholders, services and physical properties over long time periods poses unique security and resilience challenges.

This one day workshop brought together a community of researchers interested in in areas such as embedded systems, industrial control and SCADA systems, critical infrastructure and IoT settings. The workshop was a "working" meeting and not a mini-conference. The presentations were selected on the basis of short 4 page position papers describing work in progress or testbeds currently in use across the research community. This technical report brings together the accepted papers for the workshop. A working paper, based on discussions at the workshop, is currently in preparation and will be published subsequently.

Details of the workshop programme are available on the workshop web site:

https://sites.google.com/site/serecin2016/

The editors thank the authors and attendees of the workshop for exciting and insightful discussions on a topic critical to security and resilience of current and future societies.

Awais Rashid, *Security Lancaster, Lancaster University, UK*
Wouter Joosen, *iMinds-DistriNet, KU Leuven, Belgium*
Simon Foley, *Department of Computer Science, University College Cork, Ireland*

Contents

Composite vulnerabilities in Cyber Physical Systems *

Pierre Ciholas
Security Lancaster
Infolab21, Lancaster University
Lancaster, LA1 4WA, United Kingdom
p.ciholas@lancaster.ac.uk

Jose M. Such
Security Lancaster
Infolab21, Lancaster University
Lancaster, LA1 4WA, United Kingdom
j.such@lancaster.ac.uk

ABSTRACT

Many recent incidents and attacks on cyber physical systems across various industries demonstrate that the risks emerge from combination of vulnerabilities in different elements as opposed to IT networks and systems where a single vulnerability in an element can compromise the entire system. In the past decade, efforts have been increasingly invested in CPS security research worldwide, sharpening or adapting guidelines, recommendations, standards and tools. However, a gap remains in understanding the vulnerability level as well as the capacities of recovery in case of composite vulnerabilities in such systems. This paper presents a brief literature review followed by a road map for composite vulnerability research in CPS.

CCS Concepts

•**Computer systems organization** → **Embedded systems;** *Redundancy;* Robotics; •**Networks** → Network reliability;

Keywords

Cyber-Physical Systems, CPS, Industrial Control Systems, ICS, SCADA, vulnerability, composite

1. INTRODUCTION

A Cyber Physical System (CPS) is a network of interacting and collaborating computational elements controlling physical entities [19], including sensors, actuators, control processing units, and communication devices. For example, Industrial Control Systems (ICS) are a type of CPS that are widely used to automate and control in domains like electricity, water, oil and gas, transportation, telecommunications, banking, emergency services, etc. [21] [22]. The widespread growth of wireless embedded sensors and actuators is even creating several new CPS for other domains, such as medical devices, automotive, and smart cities [2]; and increasing the role that the information infrastructure plays in existing CPS, such as in the power grid leading to the next generation smart grid [6] [8].

CPS security is paramount because these systems are usually critical infrastructures in which a failure or malfunction has the potential to cause debilitating impact to the society [21]. However, it is only during the last decade that CPS security has received attention. This is partly because the use of the Internet and commodity hardware and software has exposed industry to security vulnerabilities and threats from which it was historically isolated [17]. For instance, the number of recorded vulnerabilities has seen a sharp rise over the last decade. Another reason is the raise in awareness of the potential impact of attacks on critical infrastructures, which has captured the public imagination [17], including the discovery of Stuxnet in 2010 [18], and Flame in 2012 [20].

Despite the increasing amount of research devoted to CPS security and the rise of awareness on its importance, existing research and tools focus mainly on assessing single vulnerabilities or already ongoing attacks leaving a gap on the assessment of composite vulnerabilities within CPS. This is of crucial importance, as CPS are composed of many different components at many different layers possibly cross-cutting administrative domains [4] [13], and CPS complexity is only expected to increase with the advent of the smart infrastructures (smart grid, smart cities, etc.). Therefore, single vulnerabilities of particular components could be combined to give rise to more severe composite vulnerabilities that can be exploited to hack into and take control of this type of systems. For instance, the famous Stuxnet exploited a combination of vulnerabilities in PCs, ICS interfaces, controllers and policies to successfully take control of the target infrastructure [18].

2. STATE OF THE ART

During the last decade, numerous different organisations started studying the security of CPS and ICS.

The European agency ENISA (European Network and Information Security Agency) publishes studies, articles and guidelines to the attention of ICS administrators [7] [8]. Considering the content of their documents, it is obvious that an increasing effort is invested in order to take on the new security challenges that CPS face, without hesitating to question the available techniques and standards. How-

*(Produces the permission block, and copyright information). For use with SIG-ALTERNATE.CLS. Supported by ACM.

WOODSTOCK '97 El Paso, Texas USA

DOI: 10.475/123_4

ever, we could not find in any of their publication special methodology or procedure to understand the effects of different vulnerabilities aggregated in one system.

Another agency that published substantial guidelines and standards related to CPS is the United States NIST (National Institute of Standards and Technology). We note that in most of the NIST documents related to CPS security published in the beginning of the century and before, the guidelines and recommendation were written following the best practices for an information technology network and not an operational security network. This difference is significant and has been clearly formulated in one of their publication that we consider as a milestone in the field of industrial control systems security, the "Guide to industrial control systems (ICS) security" published in 2011 [22].

Several other national agencies also publish interesting documentation such as the CPNI (Centre for Protection of National Infrastructure) [9] or the SANS (System Administration Networking Security) [14]. However, similarly to the early publications of the NIST, the recommendation and guidelines provided in their publications are not designed specifically for CPS. All these agencies and publications have contributed to sharpen and adapt the guidelines, recommendations, standards and tools available.

Despite the increasing amount of research devoted to CPS security and the rise of awareness on its importance, existing research and tools focus mainly on assessing single vulnerabilities or attacks only when they are happening leaving a gap on the assessment of composite vulnerabilities within CPS. Most projects related to Intrusion Detection Systems (IDS) within CPS only focus on detecting already ongoing attacks due to existing vulnerabilities and do not try to detect vulnerabilities before they are already being exploited [1] [10] [3] [5] [11]. Detecting vulnerabilities before they are exploited is of crucial importance, as CPS are composed of many different components at many different layers possibly cross-cutting administrative domains [4] [13], and CPS complexity is only expected to increase with the advent of the smart infrastructures (smart grid, smart cities, smart houses, etc.).

Current methodologies and assessment tools tend to consider each element and layer separately, assuming a perfect separation amongst the different elements composing a CPS. However, the recent CPS focusing on ICS attacks showed clear weaknesses in such approach. The detailed documentation and case studies reveal that numerous targeted CPS were successfully compromised by the attackers using cross-layer and inter-components vulnerabilities combined to accomplish the final objective of the attack. Therefore, we understand that single vulnerabilities of particular components are being combined to give rise to more severe composite vulnerabilities that are or can be exploited to hack into and take control of this type of systems. For instance, the famous Stuxnet exploited a combination of vulnerabilities in PCs, ICS interfaces, controllers and policies to successfully take control of the target infrastructure [18]. Studies outline that conventional verification and validation do not provide satisfying results when applied to CPS.

Consequently, studying and understanding the potential impacts of combined vulnerabilities, developing system-wide assessment tools that consider more than just single vulnerabilities of particular components [15], as well as integrating those tools and evidences with the wider risk management process in Cyber Physical Systems is essential. These gaps are considered to be one of the main problems preventing us from building robust, reliable and resilient mission-critical Cyber Physical Systems [24].

One of the key challenges associated with understanding the complex vulnerability landscape within CPS is related to the diversity of contexts in which they are deployed and used.

3. ROADMAP

We present a road map for future research on composite vulnerabilities within CPS. In a nutshell, the research should start by gathering the current projects related to CPS security, followed by a detailed analysis of the most recent and well documented CPS incidents to further the understanding of how the single vulnerabilities have interacted together to allow the incident to happen. By analysing various CPS incidents involving numerous single vulnerabilities, recognisable patterns and interdependencies types between single vulnerabilities might arise, allowing to formulate a first set of identifiable composite vulnerabilities along with their specifications, severity, impact and possible defense strategies. The research should then analyse the current methodologies available to elicit the single vulnerabilities in a cost-effective way when applied to a CPS. Further research should be conducted on the current methodologies to improve the detection and provide additional information on single vulnerabilities that might be of interest to further the understanding of composite vulnerabilities. This initial research could constitute a solid foundation for the development of frameworks to assess and aggregate vulnerabilities in a CPS. In particular, such framework might produce a set of composite vulnerabilities based on single vulnerabilities of individual components, which could then be used as the input to risk assessment and risk management processes for operational CPS and as the input to improve security at the design and pre-deployment phase of CPS.

3.1 Understanding composite risks

Understanding the impact and potential consequences of vulnerabilities combined, as well as designing a mitigation management solution, particularly in very special systems such as in CPS, requires a solid understanding of risk literature and vulnerability types. Therefore, research efforts should be dedicated to the gathering then studying of the academic and industry specific literature and current projects, in order to have a complete overview of the state of the art. A detailed analysis of the recent well-known and documented attacks or incidents in CPS, such as Stuxnet [18] or Flame [20] might reveal identifiable vulnerabilities interactions or dependency patterns. If such recognisable patterns exist, further study can be done to highlight their specifications, their severity, their impact and, if possible, develop a methodology to reveal them before an incident happens. Understanding the link between literature and real-world application is instrumental in the evaluation of risks involving multiple vulnerabilities. The results of this first phase of research will define the new risks arising from vulnerability combination within the context of CPS. This will form a solid foundation for the development of a framework but also for future research related to CPS.

3.2 Eliciting single vulnerabilities

Composite vulnerabilities are composed of single vulnerabilities. Therefore, it is paramount to be able to elicit single vulnerabilities in a cost-effective way. The different assurance techniques used to detect the vulnerabilities and measure the security level must be evaluated and adapted to the particular context of CPS [23]. Recent studies have reviewed assurance techniques on ICS, which are one type of CPS, providing a better understanding of their efficiency at the different stages of the system development as well as their usage through an ISO 27000 certification process [16]. Further research needs to be conducted to improve the detection of single vulnerabilities through the different assurance techniques when used in the specific context of CPS. Also, studies need to be conducted to see how assurance techniques can be extended to provide extra information about single vulnerabilities that can be important to understand how they may be composed or aggregated together.

3.3 Frameworks for composite vulnerabilities

With a clear understanding of the interconnections and dependencies that single vulnerabilities can have provided by the first phase of research described in section 3.1 and cost-efficient methods to elicit single vulnerabilities as described in section 3.2, it may be possible to start the development of a framework to composite vulnerabilities. The research should then aim to develop methodologies to detect and chain single vulnerabilities within a system, allowing to recognise the different patterns of single vulnerabilities interactions or dependencies they can have to reveal the composite vulnerabilities. These methodologies should be general enough to be applicable to existing or projected CPS with tools allowing to evaluate their severity, impact and possible counter-measures. The efforts in this process should be focusing on having a formal representation of the framework and a clear methodology to go from single vulnerabilities as input and produce a set of composite vulnerabilities, including and accounting for the severity of such composite vulnerabilities in terms of their potential impact on assets. The ultimate aim should be to explore to what extent composite vulnerability detection and analysis could be automated.

3.4 Meaningful evaluation methodologies

To test the effectiveness of any frameworks developed, they must be implemented and used in CPS case studies. Conducting experiments to test composite vulnerability frameworks in the wild is challenging as it requires strong collaboration with industry and, even then, it may be difficult to do it in a way that does not impact the operational system in which the experiments will be conducted. During the early test phases, the framework could be evaluated on a testbed, such as the ICS testbed at Security Lancaster [12], but to fully test its effectiveness, an experiment on a deployed CPS will be required such as Smart Grid [6]. The framework should be improved at every test phase accordingly to the results of the experiment. The contribution of this framework in a risk assessment or a risk management processes should be evaluated and documented. Ultimately, these evaluation methodologies can be used to compare the different frameworks to define improvements and future work.

4. REFERENCES

[1] E. J. Byres, M. Franz, and D. Miller. The use of attack trees in assessing vulnerabilities in scada systems. In *Proceedings of the international infrastructure survivability workshop*. Citeseer, 2004.

[2] A. Caragliu, C. Del Bo, and P. Nijkamp. Smart cities in europe. *Journal of urban technology*, 18(2):65–82, 2011.

[3] A. Carcano, I. N. Fovino, M. Masera, and A. Trombetta. State-based network intrusion detection systems for scada protocols: a proof of concept. In *Critical information infrastructures security*, pages 138–150. Springer, 2009.

[4] P. Didier, F. Macias, J. Harstad, R. Antholine, S. Johnston, S. Piyevsky, M. Schillace, G. Wilcox, D. Zaniewski, and S. Zuponcic. Converged plantwide ethernet (cpwe) design and implementation guide. *Cisco Systems and Rockwell Automation*, 2011.

[5] N. Erez and A. Wool. Control variable classification, modeling and anomaly detection in modbus/tcp scada systems. *International Journal of Critical Infrastructure Protection*, 10:59–70, 2015.

[6] H. Farhangi. The path of the smart grid. *Power and energy magazine, IEEE*, 8(1):18–28, 2010.

[7] E. . E. U. A. for Network and I. Security. Protecting industrial control systems - recommendations for europe and member states, 2011.

[8] E. . E. U. A. for Network and I. Security. Smart grid threat landscape and good practice guide, 2013.

[9] C. . C. for the Protection of National Infrastructure. Critical security controls guidance, 2014.

[10] I. N. Fovino, A. Carcano, T. De Lacheze Murel, A. Trombetta, and M. Masera. Modbus/dnp3 state-based intrusion detection system. In *Advanced Information Networking and Applications (AINA), 2010 24th IEEE International Conference on*, pages 729–736. IEEE, 2010.

[11] N. Goldenberg and A. Wool. Accurate modeling of modbus/tcp for intrusion detection in scada systems. *International Journal of Critical Infrastructure Protection*, 6(2):63–75, 2013.

[12] B. Green, B. Paske, D. Hutchison, and D. Prince. Design and construction of an industrial control system testbed. In *PG Net-The 15th Annual PostGraduate Symposium on the Convergence of Telecommunications, Networking and Broadcasting*, 2014.

[13] B. Green, D. Prince, U. Roedig, J. Busby, and D. Hutchison. Socio-technical security analysis of industrial control systems (ics). In *Proceedings of the 2nd International Symposium on ICS & SCADA Cyber Security Research 2014*, pages 10–14. BCS, 2014.

[14] S. Institute. Critical security controls, 2014.

[15] W. Knowles, D. Prince, D. Hutchison, J. F. P. Disso, and K. Jones. A survey of cyber security management in industrial control systems. *International Journal of Critical Infrastructure Protection*, 9:52–80, 2015.

[16] W. Knowles, J. M. Such, A. Gouglidis, G. Misra, and A. Rashid. Assurance techniques for industrial control systems (ics). In *Proceedings of the First ACM Workshop on Cyber-Physical Systems-Security and/or PrivaCy*, pages 101–112. ACM, 2015.

[17] M. Krotofil, A. Cardenas, J. Larsen, and D. Gollmann. Vulnerabilities of cyber-physical systems to stale dataâĂŤdetermining the optimal time to launch attacks. *International Journal of Critical Infrastructure Protection*, 7(4):213–232, 2014.

[18] R. Langner. Stuxnet: Dissecting a cyberwarfare weapon. *Security & Privacy, IEEE*, 9(3):49–51, 2011.

[19] E. A. Lee. Cyber physical systems: Design challenges. In *Object Oriented Real-Time Distributed Computing (ISORC), 2008 11th IEEE International Symposium on*, pages 363–369. IEEE, 2008.

[20] B. Miller and D. Rowe. A survey scada of and critical infrastructure incidents. In *Proceedings of the 1st Annual conference on Research in information technology*, pages 51–56. ACM, 2012.

[21] R. J. Robles, M.-k. Choi, E.-s. Cho, S.-s. Kim, G.-c. Park, and J. Lee. Common threats and vulnerabilities of critical infrastructures. *International journal of control and automation*, 1(1):17–22, 2008.

[22] K. Stouffer, J. Falco, and K. Scarfone. Guide to industrial control systems (ics) security. *NIST special publication*, 800(82):16–16, 2011.

[23] J. M. Such, A. Gouglidis, W. Knowles, G. Misra, and A. Rashid. The economics of assurance activities. 2015.

[24] X. Zheng, C. Julien, M. Kim, and S. Khurshid. On the state of the art in verification and validation in cyber physical systems. *The University of Texas at Austin, The Center for Advanced Research in Software Engineering, Tech. Rep. TR-ARiSE-2014-001*, 2014.

Water Treatment, Distribution, and Electric Power Testbeds for Research in the Design of Secure Interdependent Critical Infrastructures

Nils Ole Tippenhauer
Singapore University of Technology and Design
nils_tippenhauer@sutd.edu.sg

Martín Ochoa
Singapore University of Technology and Design
martin_ochoa@sutd.edu.sg

Aditya P. Mathur
Singapore University of Technology and Design
aditya_mathur@sutd.edu.sg

ABSTRACT

A set of interconnected testbeds have been designed for collaborative research in the design of secure and safe public infrastructures for water treatment, water distribution, and electric power. These testbeds contain distributed industrial control systems to investigate the cascading effects of cyber attacks as well as the effectiveness of attack detection mechanisms. While not all parts of the interconnected system are fully operational yet, the general architecture of the testbeds and ongoing and planned research are discussed.

1. INTRODUCTION

This paper presents a set of interconnected testbeds, and related research, at the Singapore University of Technology and Design (SUTD). The testbeds are used for security research on modern Industrial Control Systems (ICS). The testbeds combine ICS for water treatment, water distribution; electrical power generation, transmission, and demand. The long term objective of the testbeds and the associated research tasks is to transform the process of ICS design by bringing cyber security into the early stage rather than as an add on, i.e., after the system is built. While not all testbeds are yet fully operational, these are designed to reflect realistic ICS as in use today. In addition to the testbeds, work is also underway on complementary tools such as simulation environments [4], attack, and defense modeling tools.

The testbeds and the associate research are motivated by an increase in the attack surface of ICS due primarily to the inclusion of commodity IT infrastructure, such as mainstream operating systems and networking technology. Such attacks could come from inside the system perimeter, as by an employee, or through the network from an outside attacker. In either case, researchers have proposed algorithms for the prevention and detection of attacks. Mechanisms for defending an ICS against attacks have also been proposed and experimented with [1, 2]. However, most of the related work remains theoretical, as access to real-world testbeds to validate results is limited [5, 10].

In this work the general architecture and ongoing and planned research are discussed. In particular, the focus is on how the testbeds cover different settings in terms of distributed vs. centralized control, modern and legacy protocols and networking standards, and physical processes with high or low inertia. Also briefly summarized are experiments performed on an operational testbed, and future plans for general access to the lab for local and remote collaborators.

Organization: The remainder of this work is organized as follows. The overall physical and cyber architecture of the testbeds is presented in Section 2. Preliminary experiments, published elsewhere, are briefly summarized in Section 3. The foreseen research and extended collaboration possibilities are in Section 4. The conclusion and future plans for the use of the testbeds are provided in Section 5.

2. ARCHITECTURE OF TESTBEDS

A brief overview of the testbeds at iTrust[1] follows.

2.1 SWaT- Water Treatment

SWaT [8] is an operational testbed for water treatment producing 5 US gallons/hr of filtered water. In a small footprint of approximately 90 square meters (Figure 1), the testbed represents a small-scale version of a modern water treatment plant found in large cities. The overall testbed design was coordinated with Singapore's Public Utility Board, the nation-wide water utility company, and constructed by a third party vendor. SWaT is used to investigate cyber-attacks and respective systems responses, and to conduct experiments with novel countermeasure designs (e.g., physics-based). SWaT consists of six stages labeled P1 through P6 (Figure 2). Each stage is controlled by its own set of dual PLCs, one serving as a primary and the other as a backup in case of any failure of the primary. Overall, the testbed leverages a distributed control approach in normal operations, where each process stage is individually controlled by the local PLCs. For some process stages, the local control requires state information from other stages. Such information transfer is accomplished by networking the PLCs. Both automated distributed control and manual control are possible via the HMI and SCADA workstation.

Communications: Within each process stage, the main PLC obtains data from local sensors and controls actuators such as pumps and valves (Figure 3). In addition to the actuators, level, flow, and water property sensors across the stages enable the PLCs to monitor the status of the system, and to compute and effect control actions. The local communications between a PLC and its direct sensors and actuators is based on Ethernet-based ring topology using Allan-Bradley's Device Level Ring (DLR) protocol. The ring ensures that loss of a single link can be tolerated without impacting data or control functionality. Across different process stages, PLCs communicate with each other through a conventional Ethernet star

[1] A center for research in cyber security located at SUTD.

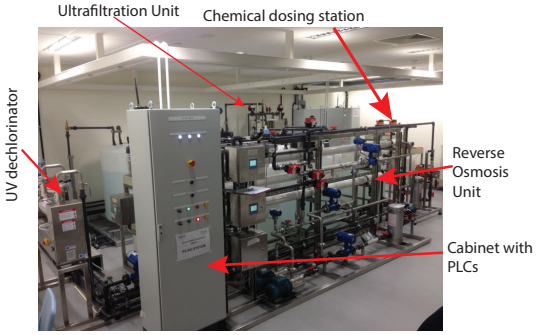

Figure 1: A photo of SWaT. The Reverse Osmosis unit is seen in the front, while the view on Ultrafiltration unit, tanks, and several other components is obstructed.

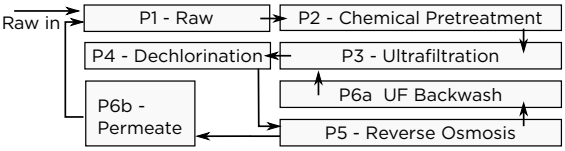

Figure 2: Physical water treatment process in SWaT. Each process stage is locally controlled by a PLC, the overall process forms a continuous loop to recycle water.

topology (the L1 network). It contains a managed Layer 3 switch connecting all 6 process stages, the HMI, SCADA workstation and the historian.

All network communication by PLCs, sensors and actuators in SWaT is using the industrial EtherNet/IP (ENIP) and Common Industrial Protocol (CIP) stack [7]. In ENIP, sensor values or actuator settings are mapped to *tags*. Each tag can be addressed either via a string descriptor defined by the system designer, e.g., MV101 for motorized valve 1 in process 1), or a more direct mapping to bank number and pin number or similar (directly referring to digital/analog pins of a unit's IO panel). Communications among sensors, actuators, and PLCs can be via either wired Ethernet or Wi-Fi links; manual switches allow to change the configuration between the wired and wireless communication.

Stages in SWaT: Stage P1 controls the inflow of water to be treated by opening or closing a valve (not shown) that connects the inlet pipe to the raw water tank. Water from the raw water tank is pumped via a chemical dosing (stage P2, chlorination) station to another UF (Ultra Filtration) Feed water tank in stage P3 where a UF feed pump sends water via UF membrane to RO (Reverse Osmosis) feed water tank in stage P4. Here an RO feed pump sends the water through an ultraviolet de-chlorination unit controlled by a PLC in stage P4. This step is necessary to remove any free chlorine from the water prior to passing it through the RO unit in stage P5. Sodium bisulphate (NaHSO3) can be added in stage P4 to control the ORP (Oxidation Reduction Potential).

In stage P5, the de-chlorinated water is passed through a 3-stage RO filtration unit. The permeate and rejects from the RO unit are stored in separate tanks. The backwash pump at stage P6 is used to controls the cleaning of the membranes in the UF unit. A backwash cycle is initiated automatically once every 30 minutes and takes less than a minute to complete. Differential pressure sensors in stage

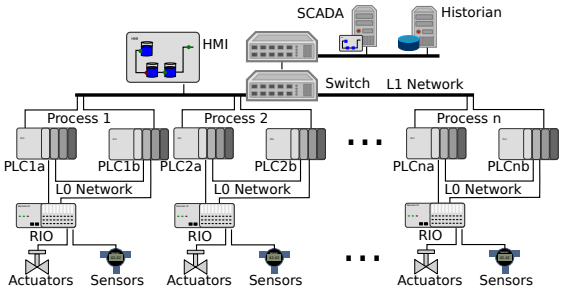

Figure 3: Architecture of the control portion of SWaT. PLC1a, PLC2a,...,PLCna denote PLC controlling their respesctive process stage. Each PLC is complemented by a hot-redundant backup PLC (PLC1b-PLCnb). The PLC communicate with their sensors and actuators through a redundant Ethernet ring (which we refer to as Level 0 network). PLCs communicate among themselves via an Ethernet star topology (which we refer to as Level 1 network).

P3 measure the pressure drop across the UF unit to enable an early backwash cycle if the pressure drop exceeds 0.4 bar.

2.2 WADI: Water Distribution testbed

The Water Distribution and Consumption testbed (WADI) focuses on a replication of a geographically distributed and centrally controlled water distribution network. The physical process (Figure 4) consists of three main stages: supply (P1), primary grid (P2), and return water system (P3). Research at iTrust focuses on P1 and parts of P2, while P3 is only required for more efficient operations (i.e., recycling of "consumed" water). The testbed is set up physically next to SWaT, in a room of approximately 80 m^2 size, and designed for a throughput of 10 US gallons of water per minute.

Physical Process: A major consideration in the design of WADI was the realistic simulation of gravity-based water pressure in the network. To achieve effects similar to the ones experienced in real systems, the water tanks in WADI are mounted at different heights, pipes with varying diameters are used, and booster pumps are available. In addition, a custom leak simulation setup was designed and implemented as part of WADI. The leak simulation allows to divert a defined percentage of water from the main distribution pipe that lowers pressure and volume of available water.

Control System: As WADI is simulates geographically distributed centrally controlled system, it contains a number of remote terminal units and PLCs (NI compactRIO) that aggregate data from local sensors, and transmit that data to a central SCADA system (using the DNP3 protocol). The transmission link for that communication can be switched between (simplistic) Ethernet communication, and 3G-based wireless communication. As such, a range of different attack and defense scenarios can be investigated. In WADI, the PLCs and RTU are connecting to most sensors and actuators directly. Where needed, Modbus/TCP is used for local fieldbus communications.

2.3 EPIC: Power Generation, Transmission, Consumption

The Electric Power and Intelligent Control (EPIC) testbed consists of four process stages: three-phase generation, transmission, micro-grid, and a Smart Home. Each of the stages consists of its own switches, transformers, protection systems and communica-

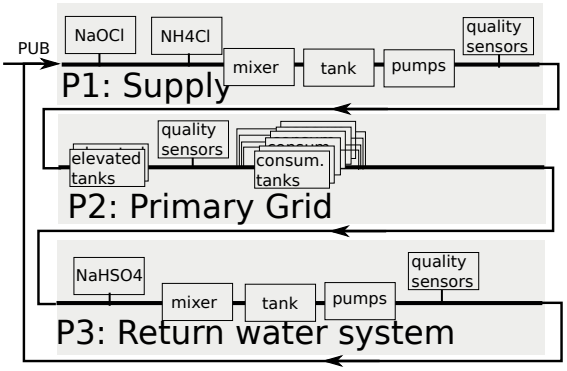

Figure 4: Physical water treatment process in WADI. The supply and and primary grid operations are centrally controlled by a SCADA system. The return water system is not part of the simularion, and was added to allow a continuous loop to recycle water.

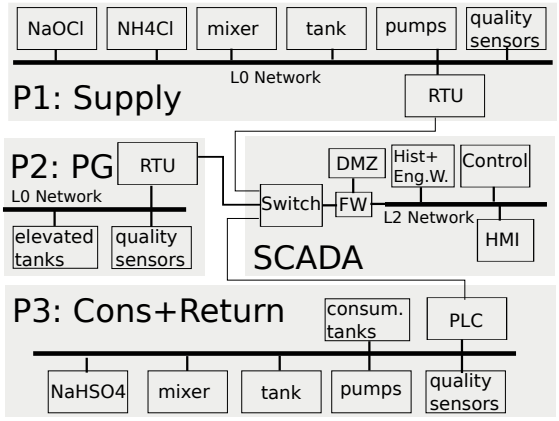

Figure 5: Architecture of the control portion of WADI.

tion systems. The generation, transmission and Smart Home stages represent a typical utility grid network. The additional micro-grid stage consists of photovoltaic (PV) generation, transformers, protection devices, and related battery-based storage.

Physical Process: The generation part consists of three motorized generators powered by the SUTD power grid. In particular, the SUTD grid is used to power M2 and M3, which are mechanically linked to generators G2 and G3 in the micro-grid. Generator G1 is powered when M1 is powered through the Smart Home load bus.

The transmission portion of the design begins from the transformer. Transformers are used to step down the voltage from the distribution levels to transmission levels. While in real systems, electricity is normally transmitted at high voltages in order to reduce losses, EPIC uses lower voltage for safety reasons. For slight changes in load, tap changes in the transformer adjust the voltages.

The Smart Home section consists of various loads of different types, namely resistive, capacitive and inductive. The loads are separated by a bus tie in the Smart Home section. This is to facilitate load shedding as well as assigning emergency loads. A motor, M1 is connected as load to the Smart Home section. This motor can operate the generator, G1 and aid in supplying power to feed the extra loads in the system. The motor can also be used to overload the system and study the characteristics of the system when it overloads.

Control System: In general, the communication infrastructure and control system in EPIC is similar to that in WADI. In particular, it represents a spatially distributed but centrally controlled system. The individual process stages each have one PLC or RTU connected to the central SCADA system. The PLC in the generation process controls the different ways to generate power for the system, and collects data from the sensors and fault detection components. The PLC in the Smart Home section is used to control the variable loads from the SCADA. In addition, the Smart Home has several smart meters connected to the loads, which report the usage to the Advanced Metering Infrastructure (AMI).

3. ONGOING RESEARCH

EPIC and WADI are under construction and scheduled for operation in 2016. Following is a brief recap of some previously published experiments on SWaT, and the current collaboration set-up.

Attack and defense: Since SWaT became operational, a large set of attacks and experiments have been conducted. Initial results have been published on industrial traffic manipulation in the fieldbus [9], the L1 network in SWaT [4], and the impact of attacks on SWaT behavior [3]. As part of those attacks, tools were written to parse and generate the industrial Ethernet/IP protocol traffic. Based on those tools, the attacker is able to manipulate sensor and actuator traffic in real-time. In addition to the network-layer and direct PLC attacks, a series of experiments relating to the physical layer process manipulation and attack detection have been conducted [2].

On-site collaborations: A number of researchers from UT Dallas [9] and MIT [6] have conducted experiments jointly with researchers at iTrust. Security hardware companies such as Checkpoint, ICS^2 and Elbit Systems, as well as others (e.g., Deloitte), have used SWaT extensively for experiments and demonstrations.

4. NEXT STEPS

Envisioned next-steps in terms of research and collaboration possibilities are described next.

4.1 Cascading failures

All three testbeds are connected to the same communication infrastructure, which is configured to allow flexible reconfiguration. The objective of physically interconnecting the testbeds is to explore cascading effects from strategic cyber attacks that target one or more ICS within the connected system. In principle, many of such effects are unpredictable, and will be the subject of research.

Figure 6 depicts the physical interdependence of the testbeds. EPIC will power both SWaT and WADI, so failures in supply will affect water filtering and distribution. Note, however, that this link can potentially be maliciously exploited in an unexpected manner, e.g., when an attacker manages to revert the electricity flow (for instance by injecting accumulated electricity in a battery), there could be consequences to the EPIC testbed. Therefore flows between SWaT and EPIC and between WADI and EPIC are bi-directional. Similarly, SWaT and WADI will be interconnected physically with a water pipe. The normal water flow will be from treatment to distribution, but attackers might revert or prevent that flow.

Cascading effects between SWaT and WADI: The envisioned physical connection between SWaT and WADI will follow the natu-

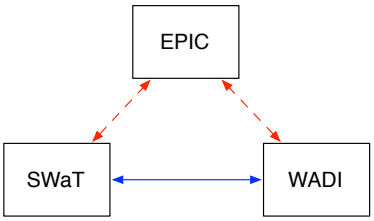

Figure 6: Physical interdependence overview of testbeds. Dashed lines depict an electrical connection, solid lines represent water.

ral water treatment followed by the distribution sequence based on home and industrial demand. Filtered water in SWaT will flow to WADI when demand increases. There are several attack vectors in this scenario. First, attackers might strategically falsify values in sensors in both testbeds causing water to flow from SWaT to WADI with a different volume or pressure than programmed. This could cause harm to both infrastructures in an unexpected manner, since for instance damage can be caused to pumps, pipes and filtering devices if the water pressure is not properly controlled. An intriguing question is to what extent these attacks are bi-directional: must the attacker be able to control both testbeds in order to succeed, or could damages to one plant be caused by controlling the other one?

Cascading effects between EPIC, SWaT and WADI: Failures in power supply to SWaT or WADI caused by an attack, together with strategic falsification of sensors in both ICS, could lead to unforeseen consequences. For instance, if sensors in WADI report a normal water flow coming into the system, but there is actually no water coming in since the SWaT plant has collapsed, the equipment in WADI might suffer damage. On the other hand, if EPIC is maliciously tampered with to cause an over-supply of energy (for instance higher voltage) to SWaT or WADI, while the corresponding sensors are attacked simultaneously, the machinery (i.e., pumps) could be sabotaged. Some research questions are: Can aggressive or subtle alterations in power supply trigger a chain of events that end up in sabotage of either SWaT or WADI assets or normal operation? What are the appropriate attacker models that capture potential (dangerous) attack scenarios? What detection and control mechanisms are required to safely manage system operations during and after an attack, cyber or physical?

4.2 Remote Collaboration

To allow access to all testbeds for external collaborators, a number of virtual machines (VMs) can be reached through a 30Gbit/s lease line and an industrial firewall (Stratix 5900) with VPN appliance (Cisco openconnect). Collaborators with sufficient credentials can connect to the VPN, and then access one of the VMs, and through that directly the testbed networks. It is planned to connect testbeds directly to remote testbeds through suitable network-layer tunnels or similar.

5. CONCLUSION AND FUTURE WORK

A number of ICS security related projects in iTrust are currently using SWaT. Additional projects will soon be launched with WADI and EPIC. Collaborators from organizations within and outside of Singapore have begun using SWaT. To make collaboration easier, it is proposed that remote access to the testbeds be feasible for authorized researchers. Obviously, remote access to physical testbeds with its own challenges such as secure access, visibility into every system component, 24/7 availability, etc.

Currently, the treated water in SWaT is recycled within the treatment process itself. In the near future, the product water of SWaT will also be used as input water for a second testbed, WADI, currently under construction, focusing on water distribution. The interconnection will allow the assessment of impact of attacks, and effect propagation, across multiple testbeds. Similarly, EPIC will allow researchers to test cascading effects in the context of power generation and distribution when linked to SWaT and WADI, which are powered by electricity.

Cascading effects of cyber attacks across multiple ICS is a challenging research problem. The ICS interconnection will also make it feasible to study the impact of multiple simultaneous attacks on two or more ICS. We envision that our testbed will allow to test, fine-tune and validate a variety of defense mechanisms that eventually will produce algorithms and design recommendations for building next generation, resilient ICS.

6. REFERENCES

[1] S. Adepu and A. Mathur. Detecting multi-point attacks in a water treatment system using intermittent control actions. In *Proceedings of the Singapore Cyber-Security Conference (SG-CRC), Singapore*, volume 14, pages 59–74, January 2016.

[2] S. Adepu and A. Mathur. Distributed detection of single-stage multipoint cyber attacks in a water treatment plant. In *Proceedings of the 11th ACM Asia Conference on Computer and Communications Security (in Press)*, May 2016.

[3] S. Adepu and A. Mathur. An investigation into the response of a water treatment system to cyber attacks. In *Proceedings of the 17th IEEE High Assurance Systems Engineering Symposium, Orlando*, January 2016.

[4] D. Antonioli and N. O. Tippenhauer. MiniCPS: A toolkit for security research on CPS networks. In *Proceedings of the First ACM Workshop on Cyber-Physical Systems-Security and/or PrivaCy*, CPS-SPC '15, pages 91–100, October 2015.

[5] A. A. Cárdenas, S. Amin, Z.-S. Lin, Y.-L. Huang, C.-Y. Huang, and S. Sastry. Attacks against process control systems: Risk assessment, detection, and response. In *ACM Symp. Inf. Comput. Commun. Security*, 2011.

[6] E. Kang, S. Adepu, D. Jackson, and A. P. Mathur. Model-based security analysis of a water treatment system. In *In Proceedings of 2nd International Workshop on Software Engineering for Smart Cyber-Physical Systems (in press; SEsCPS'16)*, May 2016.

[7] ODVA. Ethernet/IP technology overview. https://www.odva.org/Home/ODVATECHNOLOGIES/EtherNetIP.aspx.

[8] SWaT: Secure Water Treatment Testbed, 2015. https://itrust.sutd.edu.sg/wp-content/uploads/sites/3/2015/11/Brief-Introduction-to-SWaT_181115.pdf.

[9] D. Urbina, J. Giraldo, N. Tippenhauer, and A. Cardenas. Attacking fieldbus communications in ICS: applications to the SWaT testbed. In *Proceedings of the Singapore Cyber-Security Conference (SG-CRC), Singapore*, volume 14, pages 75–89, January 2016.

[10] S. Weerakkody, Y. Mo, and B. Sinopoli. Detecting integrity attacks on control systems using robust physical watermarking. In *IEEE 53rd Annual Conference on Decision and Control (CDC)*, pages 3757–3764, Dec 2014.

Testbed Diversity as a Fundamental Principle for Effective ICS Security Research

Benjamin Green, Sylvain Frey, Awais Rashid, David Hutchison
Security Lancaster Research Centre
Lancaster University
United Kingdom
{b.green2, s.frey, a.rashid, d.hutchison}@lancaster.ac.uk

ABSTRACT

The implementation of diversity in testbeds is essential to understanding and improving the security and resilience of Industrial Control Systems (ICS). Employing a wide spectrum of equipment, diverse networks, and business processes, as deployed in real-life infrastructures, is particularly difficult in experimental conditions. However, this level of diversity is key from a security perspective, as attackers can exploit system particularities and process intricacies to their advantage. This paper presents an ICS testbed with specific focus on infrastructure diversity, and end-to-end business process replication. These qualities are illustrated through a case study mapping data flow/processing, user interactions, and two example attack scenarios.

Categories and Subject Descriptors

C.3 [**Special Purpose and Application-Based Systems**]: industrial control systems

General Terms

Security, Human Factors, Experimentation, Design

Keywords

Industrial Control System; ICS; Supervisory Control and Data Acquisition; SCADA; Testbed; Security; Resilience

1. INTRODUCTION

The use of testbeds is essential to understanding and improving the security and resilience of Industrial Control Systems (ICS). The wide spectrum of equipment, diverse networks, and business processes, deployed in real-world infrastructures is particularly difficult to replicate in experimental conditions. ICS broadly spans across three fundamentally different *zones*, with a variety of equipment, skill-sets, and role groups, each summarised here.

- The *manufacturing zone* is where physical process operations take place, prominently built around devices and systems broadly categorised as operational technology (OT), used for monitoring, controlling, and automating process decisions through the implementation of sensors, actuators, and controllers. Observation and manual control of physical processes though human machine interfaces (HMI), engineering workstations, remote terminal units (RTU), data historians, and control servers, is also possible within this zone.

- The *demilitarised zone* forms a boundary between manufacturing zones and enterprise zones, presenting an interface by which data can be captured and stored for further processing. Performing critical functions, devices residing in this zone lean towards conventional information technology (IT), yet have the ability to interact with OT, facilitating remote alarm management, historical data collection, remote desktop access, etc.

- The *enterprise zone* hosts conventional IT devices and systems, further utilising data collected through the demilitarised zone to perform global supervision and long-term strategic planning for the entire infrastructure.

For a more granular view of end-to-end ICS environments, see the Purdue model (figure 1).

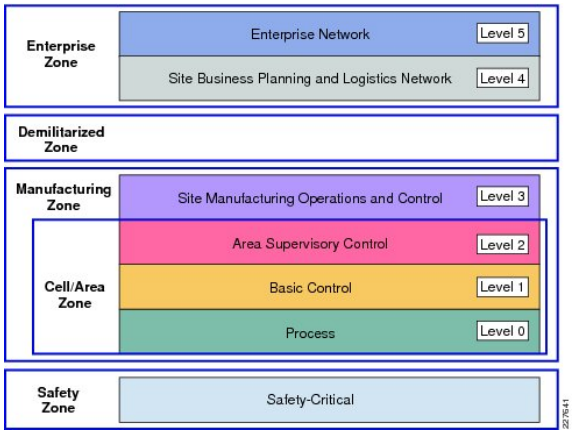

Figure 1: The Purdue Model [2]

The level of diversity in ICS environments is important from a security perspective, as attackers are able to focus their attacks to disrupt operational objects derived from varying ICS zones. First, consider a "Fuzzing" attack exploiting a controller vulnerability; secondly, a controller "Memory Manipulation". The two attacks vary in their outcome. The first could be considered a disruption attack, designed to cause an undesired impact on physical process operations. The second manipulates data used for a variety of objectives, opening several possible end results (reduced efficiency, confusion, disruption, etc.), impacting local and/or remote level (demilitarised zone and enterprise zone) visibility and control of operational processes, and data.

The *modus operandi* of the aforementioned attacks is also significantly different: a "Fuzzing" attack is relatively simple to carry out, and the process (target identification followed by target disruption) can be automatised; "Memory Manipulation" attacks requires an advanced understanding of the target network (data sources and destinations, critical data processing points, redundancy, etc.), and dedicated expertise to intercept and alter a specific data flow.

This paper presents an ICS testbed [5], with a focus on infrastructure diversity, including end-to-end business process replication (section 3). The testbed replicates not only infrastructure found within each zone (field sites, data-centres, and corporate sites), but for each zone proposes alternative forms of equipment, vendors, and protocols. These qualities are illustrated through the introduction of a case study, and two example attack scenarios (section 4). Section 5 investigates potential future work, notably related to the incoming Internet of Things (IoT) evolution.

2. RELATED WORK

Several testbeds have been proposed in the literature for power plants [7], and micro and macro grids [6, 3], based on combinations of real, emulated, and simulated components. Testbeds focusing on water treatment and distribution are less common than power grid oriented testbeds: [1] relies on simulation and virtualisation to reproduce large water infrastructures, favouring the scale and breadth of the testbed over the realism of deploying actual physical components. Some testbeds cover different sectors (power grid, water, gas): [8] is a simulation tool for building various SCADA infrastructures at scale; [9] proposes a combination of physical and simulated components.

The testbeds referenced here are spanning all zones of the Purdue reference model, although they do not always refer to this model explicitly. A variety of attacks has been studied in these environments, including Denial of Service (DoS, either by flooding or specific malicious packets), DNS and routing tables poisoning, traffic sniffing and Man In The Middle (MITM), and malware injections. However, the socio-technical particularities of each zone, and the propagation of attack effects across different environments, represent blind spots in these studies, as they generally focus on specific attacks and/or scenarios.

3. TESTBED ARCHITECTURE

As detailed in [5] the testbed was originally designed with three core factors in mind, flexibility, credibility, and reliability. Since its original conception in 2013, major works have been undertaken to further levels of diversity, and therefore credibility when compared with real-world scenarios. Figure 2 presents the diversity of each zone. Below we describe each zone in relation to the available devices and protocols.

Manufacturing zone: The physical process is built around a set of operational assets (tanks, pipes, pumps, valves, etc.), sensors, and actuators, supporting hard wired electrical signalling, such as 4-20mA, and wireless protocol technologies, such as WirelessHART. The monitoring, control, and automation of physical processes are achieved through a set of sensors, controllers, human machine interfaces (HMI), and network devices. These support the following protocols: S7 over MPI, S7 over Ethernet, DNP3, ModbusTCP, Profibus, Profinet, WirelessHART, OPC, RDP, HTTP, HTTPS, FTP, SFTP, TFTP, SSH, and Telnet.

Demilitarised zone: This zone contains a set of communication devices and servers supporting the handling of all IP based protocols discussed within the manufacturing zone.

Enterprise zone: This zone contains a number of workstations and servers supporting the handling of all IP based protocols discussed within the manufacturing zone.

4. ATTACK SCENARIOS

Figure 3 and table 1 are the output of a case study with a European utility company. Figure 3 provides a greater level of granularity on real-world data flow and processing, and has been replicated within the testbed environment. Table 1 provides a view of some critical role groups, spread across each ICS level [2].

Colour coded to provide basic guidance on the level in which each device resides, figure 3 can be mapped against role groups from table 1. As a risk assessment tool, the creation of data flow/process models, with accompanying role groups, provides a clear end-to-end view of the system.

Figure 3, highlights the complexity of data flow/processing. Operating at an abstracted level, we see the delegation of devices to meet the requirements of specific role groups. However, perhaps of greater interest is the lower-level view presented by the programmable logic controller (PLC). Here we see areas of PLC resources separated and shared based on their functionality. Take "DB2.DBD1" as an example, this is a datablock address, an area of memory allocated for a specific function, in this case storing an input value. It is shared between three system levels (1,2, and 3), used as an input for Historian and RTU data collection.

Where the following sub sections introduce two attack scenarios, developed and applied within the testbed environment, identification of data processing points and user interaction, plays a critical part in the holistic understanding of potential impact.

4.1 Fuzzing

"Fuzzing", is considered to be a blackbox security evaluation technique. Applied to discover software vulnerabilities, Fuzzing randomly mutates well-formed inputs, testing a pro-

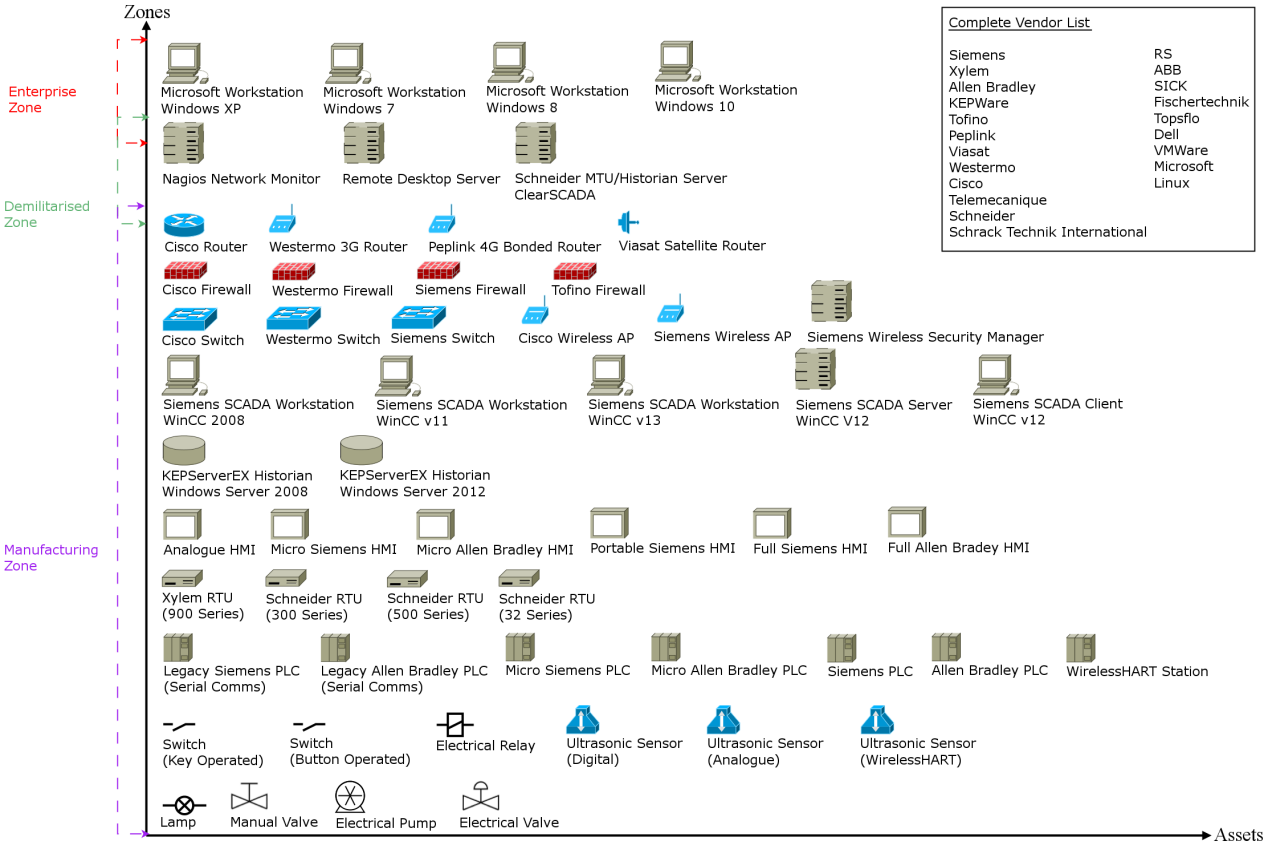

Figure 2: High-level architecture of the testbed

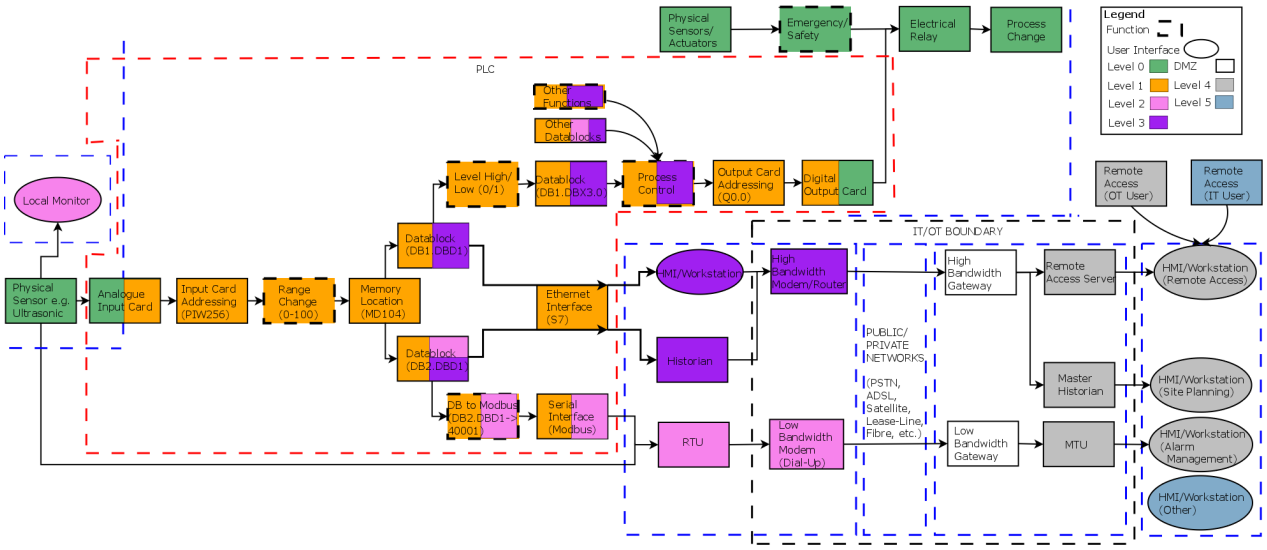

Figure 3: Data flow and processing

grams resilience upon their receipt [4]. Existing works discuss the application of fuzzing to ICS environments [10].

It is possible to conceive a vast array of opportunities where such testing/attacks could be conducted within the testbed

User Roles	ICS Level
Process Control Operators	2,3,4,5
Local Process Managers	2,3,4,5
Regional Process Managers	3,4,5
Regulatory Monitors/Testers	2,3,4,5
Performance Analysts	4,5
3rd Party Contractors	0,1,2,3,DMZ,4,5
Alarm Management Cetre Operator	4,5
Health and Safety Officers	4,5
Home Workers	3,4,5
Support/Maintenance Roles	ICS Level
Electrical Engineers	0,1,2,5
Mechanical Engineers	0,5
Control System Engineers	0,1,2,3,4
Instrumentation Engineers	0,1,2,5
Telemetry Engineers	0,1,2,3,DMZ,4,5
Communications Engineers	3,DMZ,4,5
Information Technology Engineers	DMZ,4,5
3rd Party Contractors	0,1,2,3,DMZ,4,5
Home Workers	3,4,5

Table 1: ICS roles and associated system levels

(section 3). Using figure 3, critical data processing points can be identified as described above. A number of which, where disrupted, could create a cascading effect up the data flow. For example, we applied a standard Fuzzing technique to the PLC, this caused the PLC to fail, requiring a manual reset (power down and up). This is where such attacks can prove devastating to process operations, and remote monitoring. While no level of resilience is configure to provide a replication of the PLC functionality, duplication of sensor data into a local monitor and RTU, directly from the ultrasonic sensor, still provided certain role groups access to critical data, highlighting an undesired change has occurred.

4.2 Memory Modification

"Memory Modification" modifies data stored in memory. Tools such as Snap7 [11] facilitate such modifications, providing an interface with Siemens PLCs over the S7 protocol. Consider the memory location discussed above (DB2.DBD1); once data flows and processing points are understood, selection and modification of memory locations such as this provide an excellent opportunity to cause physical process disruption, and/or inaccurate reporting/alarm management data.

We applied the Snap7 tool to DB2.DBD1, modifying the value beyond its normal operating boundaries. While this had no impact on the physical process, as the data processed through the RTU and historian (monitoring only), it moved up the data flow and into systems residing within the DMZ and Enterprise zones; this created warnings to system users that discrepancies between RTU and PLC derived data points have arisen. However, as the level of complexity found in historian calculations can be high, with reliance on the PLC alone for accurate data, it is possible that if left unchanged for some time, performance based analysis and investment decision could be made based on inaccurate data. This brings us back to the requirement for end-to-end testbed environments, and clear mapping of critical data processing points, with criticality not only identified based on the impact to operational processes, but holistic role group interaction/requirements.

5. FUTURE WORK AND CONCLUSION

In the near future, Internet of Things (IoT) devices are expected to invade a number of industries, including ICS. The extreme dynamism and diversity of the IoT contrast strongly with the slow, monolithic evolution rate of ICS. Our testbed will investigate ICS-IoT interactions through extensions to the existing infrastructure, in particular in terms of wireless technologies and wireless sensors. The diversity showcased in the testbed is also a motivation for automation to replace tedious manual adaptations to all particular devices and environments. Furthermore, formal modelling of system-user interaction and identification of critical data processing points as demonstrated earlier in this paper will be explored as promising and vital parts of our future research.

6. REFERENCES

[1] A. Almalawi, Z. Tari, I. Khalil, and A. Fahad. Scadavt-a framework for scada security testbed based on virtualization technology. In *Local Computer Networks (LCN), IEEE 38th Conference on*, 2013.

[2] P. Didier, F. Macias, J. Harstad, R. Antholine, A. Johnston, S, S. Piyecsky, M. Schillace, G. Wilcox, D. Zaniewski, and S. Zuponcic. Converged Plantwide Ethernet (CPwE) Design and Implementation Guide. *CISCO Systems and Rockwell Automation*, 2011.

[3] G. Dondossola, F. Garrone, J. Szanto, G. Deconinck, T. Loix, and H. Beitollahi. ICT resilience of power control systems: experimental results from the CRUTIAL testbeds. In *Proceedings of the 2009 IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2009*.

[4] P. Godefroid. In *2nd international workshop on Random testing (22nd IEEE/ACM Int. Conf. on Automated Software Engineering (ASE 2007))*.

[5] B. Green, B. Paske, D. Hutchison, and D. Prince. Design and construction of an industrial control system testbed. In *PG Net - 15th Annual PostGraduate Symposium on the Convergence of Telecommunications, Networking and Broadcasting*, 2014.

[6] A. Hahn, A. Ashok, S. Sridhar, and M. Govindarasu. Cyber-physical security testbeds: Architecture, application, and evaluation for smart grid. *Smart Grid, IEEE Transactions on*, 4(2):847–855, June 2013.

[7] J. Hong, S.-S. Wu, A. Stefanov, A. Fshosha, C.-C. Liu, P. Gladyshev, and M. Govindarasu. An intrusion and defense testbed in a cyber-power system environment. In *Power and Energy Society General Meeting, 2011 IEEE*, pages 1–5, July 2011.

[8] C. Queiroz, A. Mahmood, and Z. Tari. Scadasim 2014;a framework for building scada simulations. *Smart Grid, IEEE Transactions on*, 2(4), 2011.

[9] B. Reaves and T. Morris. An open virtual testbed for industrial control system security research. *Int. Jour. of Information Security*, 11(4), 2012.

[10] R. Shapiro, S. Bratus, E. Rogers, and S. Smith. Identifying vulnerabilities in SCADA systems via fuzz-testing. In *Critical Infrastructure Protection V*, pages 57–72. Springer, 2011.

[11] Snap7. Snap7 - Overview. `http://snap7.sourceforge.net/`, 2016.

Experimental Platform for Internet Contingencies

Yannis Soupionis
European Commission, Joint Research Centre (JRC)
Institute for the Protection and Security of the Citizen (IPSC)
Security Technology Assessment Unit
Via E. Fermi, 2749, 21027 Ispra, Italy
yannis.soupionis@jrc.ec.europa.eu

Thierry Benoist
European Commission, Joint Research Centre (JRC)
Institute for the Protection and Security of the Citizen (IPSC)
Security Technology Assessment Unit
Via E. Fermi, 2749, 21027 Ispra, Italy
thierry.benoist@jrc.ec.europa.eu

ABSTRACT

Decentralized Critical infrastructure management systems will play a key role in reducing costs and improving the quality of service of industrial processes, such as electricity production and transportation. The recent malwares (e.g. Stuxnet) revealed several vulnerabilities in today's Distributed Control Systems (DCS), but most importantly they highlighted the lack of an efficient scientific approach to conduct experiments that measure the impact of cyber threats on both the physical and the cyber parts of Networked Critical Infrastructures (NCIs). In this paper we present our novel cyber physical testbed, "Experimental Platform for Internet Contingencies" (EPIC) in support of EU policy making, which can provide accurate assessments of the effects that cyber-attacks may have on the cyber and physical dimensions of NCIs.

Keywords

Cyber-Physical, Cyber security, Networked Critical Infrastructures, Testbed, Emulab, Simulation

1. INTRODUCTION

The Joint Research Centre (JRC)is the European Commission's in-house science service which employs scientists to carry out research in order to provide independent scientific advice and support to EU policy. One such policy is the EU initiative on Critical Information Infrastructure Protection (CIIP) aims to strengthen the security and resilience of vital Information and Communication Technology (ICT) infrastructures. In this context, Joint Research Center (JRC) created *EPIC*, a novel cyber-physical *Experimentation Platform for Internet Contingencies*. To meet the complexity of today's NCIs, EPIC uses an emulation testbed based on Emulab [12, 8] to recreate the cyber elements of a NCI and software simulators for the physical components.

The main testebed requirements, which EPIC fullfils, are:

SERECIN 2016 London, UK

DOI: 10.1145/1235

- **Fidelity**: Experimentation testbeds need to reproduce as accurately as possible the real system under study.
- **Repeatability**: This requirement reflects the need to repeat an experiment and obtain the same or statistically consistent results.
- **Measurement Accuracy**: Experiments should be accurately monitored and measurements should not interfere with the experiment because they might alter the experiment's outcome.
- **Safe execution**: A security experimentation testbed needs to support disruptive experiments with physical processes in a safe manner.

The paper is organized as follows: Section 2 describes the architecture and implementation of EPIC. A set of examples showing the applicability of EPIC to cyber security studies on small-scale and large-scale infrastructures are presented in Section 3. The paper concludes in Section 4.

2. EPIC ARCHITECTURE

The architecture of EPIC suggests the use of an emulation testbed based on the Emulab software [12, 8] in order to recreate the cyber part of NCIs, e.g., servers and corporate network, and the use of software simulation for the physical components, e.g., a chemical process. Fig. 1 provides an overview of EPIC's architecture and experimentation steps, which will be elaborated upon in the remaining of this section.

2.1 Experimentation Software and Hardware

The use of emulation testbeds is a trend that is becoming more popular. One of the most advanced software suites in this direction is Emulab [12]. The name Emulab refers both to a facility at University of Utah and to a software. Nowadays the software is actively supported by multiple universities and there are many private installations throughout the world.

We have developed in our laboratory a testbed using the Emulab architecture and software. By adopting Emulab in EPIC, we can automatically and dynamically map physical components, e.g., servers and switches, to a virtual topology. In other words, the Emulab software configures the physical topology in a way that it emulates the virtual topology as transparently as possible. This way we gain significant

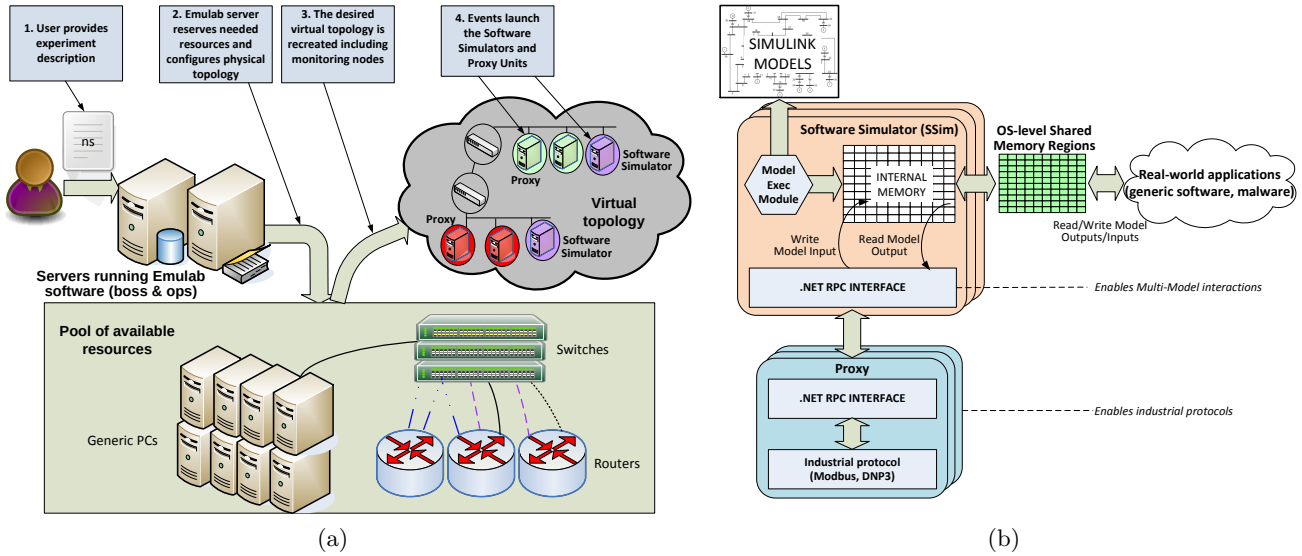

Figure 1: Architecture of the EPIC testbed: (a) Architectural overview and experimentation steps; and (b) AMICI, including software simulators (SSim) and Proxy units.

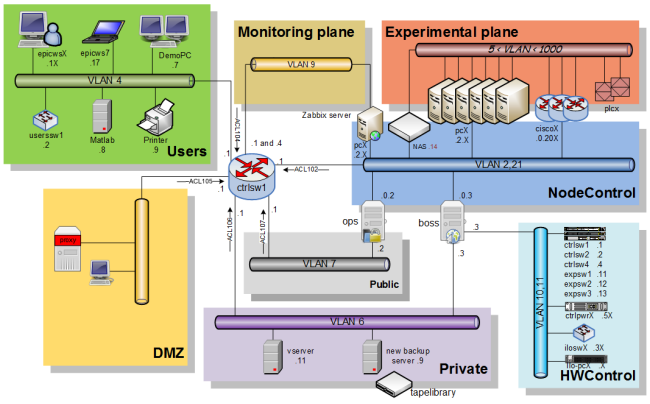

Figure 2: EPIC's logical architecture, which includes experimental and control infrastructure.

advantages in terms of repeatability, scalability and controllability of our experiments.

The basic EPIC architecture consists of two control servers, a pool of physical resources that are used as experimental nodes (generic PCs, routers or other devices) and a set of switches that interconnect the nodes (see Fig. 2) . The Emulab software provides a Web interface (see Fig. 3) to describe the steps that define the experiment life cycle within our testbed (see Fig. 1(a)):

1. First we need to create a detailed description of the virtual network topology, *the experiment script*, using an extension of the Network Simulator (NS) [6] language. The use of a formal language for experiment setup eases the recreation of a similar setting by any other researcher who wants to reproduce our results. In the experiment script we enumerate similar components as different instances of the same component type. This way, predefined *system templates*, e.g., a Linux server disk image, can be easily reused and automatically deployed and configured.

2. Experiments are instantiated by using the Emulab software, which automatically reserves and allocates the physical resources that are needed from the pool of available components. This procedure is called *swap-in*, in contrast to the termination of the experiment, which is called *swap-out*.

3. Furthermore, the software configures network switches in order to recreate the virtual topology by connecting experimental nodes using multiple VLANs. Finally, before the testbed is released for experimentation, the software configures packet capturing of predefined links for monitoring purposes.

4. Experiment-specific software, e.g., simulators, may be launched automatically through events defined in the NS script, or manually, by logging in to each station.

The installation of EPIC at the Joint Research Centre for the emulation, consists today of approximately 200 nodes, which are massively interconnected with two stacks of networks switches and 4 routers. In addition carrier-grade routers and industrial control hardware and software are available as experimental resources(see Fig. 4).

2.2 Physical Systems

For the physical layer EPIC uses simulation, since this provides an efficient, safe and low-cost approach with fast and accurate analysis capabilities. In EPIC we use generic PCs with multitasking OSs to run the real-time software simulation units. Our choice to use Simulink Coder to produce the simulators, although it has major advantages, im-

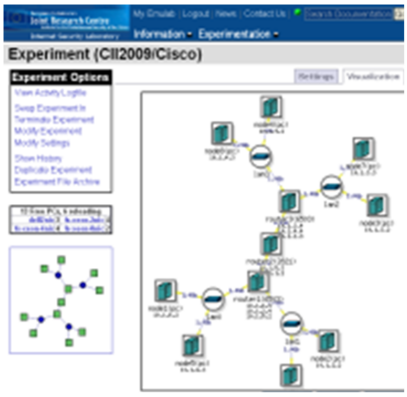

Figure 3: EPIC's web interface.

poses several constraints to the simulated models. An important aspect in this sense is the choice of the model execution rate, also known as the *simulation step*.

The main role of the simulation element (Sim) is to run the physical process model in real-time (SEE fIG. 1(b)). This is done by coupling the model time to the system time in such a way to minimize the difference between the two. Models are constructed in Matlab Simulink from where the corresponding 'C' code is generated using Matlab Real Time Workshop. These are then integrated using an XML configuration file that is flexible enough so that researchers do not need to modify the code. The generated code is then executed in real time and interacts with the real components of our emulation testbed [4].

2.3 Control Logic

The testbed currently provides an array of simulated, emulated and real implementation possibilities for building control strategies at different levels. It recreates control center-specific logic by leveraging real industrial HMI software. The HMI monitors the physical process and forwards operator commands to control devices. Subsequently, a simulation-based approach for testing operational decisions is also available through the *SSim* unit (see Fig. 1(b)).

At the hardware control level the testbed provides real PLCs as well as the possibility to run emulated control code in real-time. The execution engine for emulated control code is embedded into the *Proxy* unit. Control code is provided in the form of an external binary, which is loaded and executed at run-time. The emulated control code can interact with *SSim* units, i.e., with physical processes, through Modbus or RPC calls. Additionally, the testbed supports the execution of control code provided as binary library by the *SSim* unit, in which case the interaction with the physical process model is done through function calls.

3. TYPICAL EXPERIMENTS

The EPIC platform can efficiently:

- recreate realistic network topologies and conditions for example, delay and loss characteristics of wide-area network (WAN) links of the Internet infrastructure.
- study, in a controllable manner, a wide range of disruptions such as:
  - host and link failures,
  - BGP hijacking,
  - distributed denial of service attacks (DDoS), and
  - integrity attacks (e.g. replay).
- implement a wide range of physical process models. Here we mention small-scale processes such as Bell and Åström's oil-fired 160MW electric power plant [1], which is based on the Sydsvenska Kraft AB plant in Malmö, Sweden, and the Tennessee-Eastman chemical process [3], which is also based on a real process, but the authors have introduced slight changes in order to protect the identity of reactants and products. EPIC also enables experimentation with railway systems, based on the train models proposed by Rios and Ramos [7]. These take into account several aspects of real transportation systems such as weight, speed, acceleration, deceleration, and power consumption. A major advantage of EPIC is its ability to provide experimentation capabilities with a wide range of power grid systems. For this purpose we have adopted the IEEE suite of power grid case systems [11], which are based on Matlab open-source libraries, i.e. MatPower [13] and MatDyn [2]. We mention the 9-bus test case, which is the Western System Coordinating Council's (WSCC) 3-machine 9-bus system, and the 30-bus, 39-bus and 118-bus test cases.

Figure 4: EPIC's (a) servers, switches and routers for the network emulation and (b) PLC's for real physical nodes.

Here is a representative experiment [9], which explores the consequences and propagation of disruptions in a scenario involving three critical infrastructures: the ICT infrastructure, the power grid and the railway system. We consider

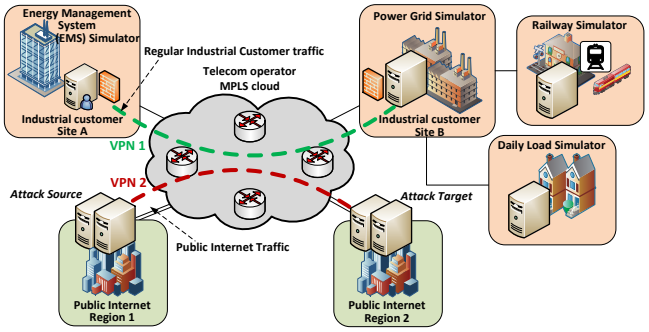

Figure 5: Experimental setting involving ICT infrastructure, electrical grid and railway system.

the hypothetical scenario of a cyber-attack and specifically a Distributed Denial of Service (DDoS) attack, that is causing a severe telecommunication service degradation which propagates across several Critical Infrastructures. We recreated the typical architecture of an installation in which the power grid is controlled remotely (Fig. 5). Here, *Site A* runs a simplified model of an Energy Management System (EMS) [10] to ensure voltage stability. The EMS continuously monitors and adjusts the operational parameters of the power grid model running at *Site B*. In this scenario we adopted the IEEE 39-bus New England system that includes a total of 39 substations together with 10 generators. The daily load imposed to our system derives from real data [5] and the intervention of the EMS is required to keep the grid stable. In order to illustrate the propagation of disturbances to other CIs we used a simple railway system model [7] that includes 10 trains. The train model takes into account several aspects of real transportation systems such as weight, speed, acceleration and deceleration.

During normal operation trains start by increasing their speed up to a maximum value and continue with the same speed until they reach the next station. Once the attack disrupts the normal operation of the grid, sub-nominal voltage levels trigger circuit breakers and disconnect the railway system from the electricity grid. The loss of power supply brings obviously all trains to full stop, which is an issue not only for the passengers, but also for the railway (additional cost). The severe risks that are involved , should be taken into consideration and specific protective measures should be implemented.

4. CONCLUSIONS

Cyber attacks constitute one of the main threats to critical infrastructures. Combining modelling and simulation with experimental activities, the JRC studies the cyber-security threats that cyber-physical systems like smart grids face.

The JRC Experimental Platform for ICT Contingencies (EPIC) is a network test-bed specifically designed to support the execution of repeatable and safe cyber-security experiments in a fully controllable experimentation environment. Moreover it has been extended in order to allow researchers to conduct real-time experiments with simulated physical systems tightly coupled with real cyber systems. It should be stated that currently EPIC is not accessible to other researchers, unless they are cooperating/collaborating and invited at the JRC premises. This access policy is going to change in the future.

Concluding, EPIC has proved its value as a modern scientific instrument, which can provide valuable insights into the disruptive effect of cyber attacks on physical processes.

5. REFERENCES

[1] R. Bell and K. Åström. Dynamic models for boiler–turbine alternator units: data logs and parameter estimation for a 160MW unit. *Lundt Institute of Technology, Report TFRT–3192*, 1987.

[2] S. Cole and R. Belmans. Matdyn, a new matlab-based toolbox for power system dynamic simulation. *Power Systems, IEEE Trans.*, 26(3):1129–1136, 2011.

[3] J. Downs and E. Vogel. A plant-wide industrial process control problem. *Computers & Chemical Engineering*, 17(3):245–255, 1993.

[4] B. Genge, C. Siaterlis, and M. Hohenadel. Amici: An assessment platform for multi-domain security experimentation on critical infrastructures. In *Critical information infrastructures security*, pages 228–239. Springer, 2012.

[5] M. Manera and A. Marzullo. Modelling the load curve of aggregate electricity consumption using principal components. *Environ. Model. Softw.*, 20(11):1389–1400, Nov. 2005.

[6] ns 2. The Network Simulator. http://www.isi.edu/nsnam/ns/. [Online; accessed July 2013].

[7] M. A. Ríos and G. Ramos. Power system modelling for urban massive transportation systems. *Infrastructure Design, Signalling and Security in Railway*, pages 179–202, 2012.

[8] C. Siaterlis, A. Garcia, and B. Genge. On the use of Emulab testbeds for scientifically rigorous experiments. *IEEE Communications Surveys and Tutorials*, 15(2):929–942, 2013.

[9] Y. Soupionis and T. Benoist. Demo abstract: demonstrating cyber-attacks impact on cyber-physical simulated environment. In *ICCPS'14: ACM/IEEE 5th International Conference on Cyber-Physical Systems*, pages 222–222, 2014.

[10] T. Tuan, J. Fandino, N. Hadjsaid, J. Sabonnadiere, and H. Vu. Emergency load shedding to avoid risks of voltage instability using indicators. *Power Systems, IEEE Transactions on*, 9(1):341–351, feb 1994.

[11] University of Washington - Electrical Engineering. Power Systems Test Case Archive. http://www.ee.washington.edu/research/pstca/, 2012. [Online; accessed July 2013].

[12] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An integrated experimental environment for distributed systems and networks. In *Proc. of the 5th Symposium on Operating Systems Design and Implementation*, pages 255–270, 2002.

[13] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas. Matpower: Steady-state operations, planning, and analysis tools for power systems research and education. *Power Systems, IEEE Transactions on*, 26(1):12–19, 2011.

A Microgrid Testbed for Interdisciplinary Research on Cyber-Secure Industrial Control in Power Systems

David M Laverty, Mats-Robin Jacobsen, Xiaodong Zhao
Queen's University Belfast
School of EEECS, Ashby Building
125 Stranmillis Road, Belfast, UK
+44 (0) 28 9097 4651
david.laverty@qub.ac.uk
mjacobsen02@qub.ac.uk
xzhao06@qub.ac.uk

Kieran McLaughlin, Ivo Friedberg, Rafiullah Khan, Sakir Sezer
Centre for Secure Information Technologies (CSIT)
Queen's University Belfast
+44 (0) 28 9097 1890
kieran.mclaughlin@qub.ac.uk
ifriedberg01@qub.ac.uk
rafiullah.khan@qub.ac.uk
s.sezer@ecit.qub.ac.uk

ABSTRACT
This paper describes a smart grid test bed comprising embedded generation, phasor measurement units (PMUs), and supporting ICT components and infrastructure. The test bed enables the development of a use case focused on a synchronous islanding scenario, where the embedded generation becomes islanded from the mains supply. Due to the provisioned control components, control strategy, and best-practice ICT support infrastructure, the islanded portion of the grid is able to continue to operate in a secure and dependable manner.

CCS Concepts
• Computer systems organization → Embedded and cyber-physical systems → Sensors and actuators • Security and privacy → Systems security • Hardware → Power and energy.

Keywords
Testbed; Smart Grid; secure control; microgrid; synchronous-islanding

1. INTRODUCTION
'Smart Grid' is a wide encompassing term, but applications falling into this category generally refer to the application of modern Information Communication Technology (ICT) to solve challenges and constraints in the operation of traditional electrical energy infrastructure. Whilst there are many decades of experience of the application of ICT at transmission system voltages, where there was a clear business case for dedicated utility own communications, there is now momentum behind the application of such applications at lower distribution level voltages. At these voltage levels, the business case for dedicated communications is not clear, and it is generally necessary to make use of public telecoms infrastructure. While these novel methodologies are needed to allow existing infrastructure to fulfill a changing role, the operation of critical control algorithms on ICT infrastructure, and in particular public ICT infrastructure, introduces the risk of cyber-attacks that can cause physical damage.

The relationship between the traditional safety domain and the cyber security domain is a major challenge for the acceptance of 'smart grid' solutions. It is a challenge that cannot be sufficiently addressed by the necessary application of existing solutions from either domain [13]. In the CAPRICA project the authors investigate how control algorithms can be designed to be resilient in the face of cyber-attacks. Further, the authors develop ICT solutions that make use of the best practice in the traditional ICT domain.

Of particular interest to electrical energy networks at present, notably across Europe and in particular the authors' home network on the island of Ireland, is the issue of the integration of renewable electricity generation. At the time of writing, the all-Ireland power system frequently meets demand with as much as 50% of supply coming from non-synchronous machine generators, predominately made up of wind generation [1]. The issues that arise, including reduction in system inertia and changing fault levels, was not foreseen when the architecture of the electrical infrastructure was designed in the 1950s/60s/70s when centralize bulk generators supplied by fossil fuels was the norm. Hence, there is a requirement for novel methods of power system operation to enable the system to function with this and higher levels of system non-synchronous penetration (SNSP).

The authors' have an interested in a state of power system operation known as 'synchronous islanding' [2]. Conventionally, power system islanding is a mode of operation to be avoided, as is the case in the authors' complimentary research work in islanding protection [3]. However when an island is appropriately controlled, the dangers of uncontrolled islanding are avoided and there is a benefit to the utility in terms of reduction in customer minutes lost, and generation is kept online avoiding the need for it to be disconnected, shut down, and restarted.

In this paper the authors present a physical electrical system testbed, representing a microgrid containing embedded generation, that has been built at Queen's University Belfast (QUB). The testbed was originally constructed to demonstrate the feasibility of synchronous islanding of a single machine, then adapted for multi-machine islanding, and then again to create a testbed on which to investigate the interdisciplinary domain of cyber-secure industrial control systems (see Section 2). The authors will describe the detailed testbed setup in Section 3 together with possible cooperation opportunities. Section 4 will highlight the authors' current research that is performed on the testbed before concluding the paper in Section 5.

2. THE USE CASE
Due to the geographical location of the resource, large scale renewable electricity sources are predominantly located in isolated areas. This presents a challenge to the successful integration of renewable electricity generators, since such remote areas tend to be weakly electrically connected. That is there is no access to the bulk transmission grid, rather power must be transferred at lower voltages across the distribution network. Issues arise including that of voltage rise, power quality and flicker, and constraints and

bottlenecks caused by the capacity of the infrastructure to accommodate reverse power flows. Amongst these challenges is the issue of islanding.

Islanding is the scenario in which an embedded generator (that is, a small generator connected to the distribution network) finds itself operating without connection to bulk utility dispatched generation. This is often caused be a fault upstream of the generator being cleared (disconnected) by the correct operation of protection systems (circuit breakers open). See Figure 1. When a generator operates in an island, there is the possibility for it to supply utility customers outside of the generator owner's premises. This presents a risk to the other customers due to lack of regulation of voltage and frequency, a risk to utility personnel restoring the fault (who would assume the downstream side to be de-energized) and risks the destruction of the embedded generator is reconnection is made out of synchronism with the mains supply.

A generator must be connected to the mains under a process called 'synchronization', during which the voltage of the generator is matched to that of the mains, the frequency is matched to that of the mains, and the phase angle of the voltage is matched to that of the mains. Should any of these properties be incorrect, the generator cannot synchronize due to risks of severe damage to the equipment and personnel safety. For these reasons, islanding is forbidden on most systems.

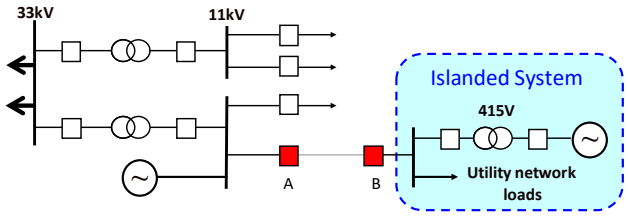

Figure 1: Example Islanding Scenario

The incumbent forms of islanding detection in use today are Rate-of-Change-of-Frequency (ROCOF) and Vector Shift. A thorough technical description of these technologies is outside the scope of this paper, but in summary these technologies are no longer fit to serve on modern grids. Both technologies were designed for use on grids with low penetrations of embedded generation, and known deficiencies that could lead to nuisance tripping (that is, disconnecting a generator when there was no need to) had virtually no impact on the wider system. In the present day, the widespread use of embedded generation means that these technologies must be desensitized otherwise nuisance tripping would be frequent, and the potential for one nuisance trip to cause further 'cascade tripping' and loss of system supply is now of concern. Consequently, novel methods of determining when a generator has become disconnect from the mains (known as 'loss-of-mains' or islanding) is an active research domain [3].

A contrary viewpoint to disconnecting a generator once islanding has been identified is that if one were to control the generator so that it was held artificially in synchronism with the utility grid, then the risks of out-of-synchronism reconnection of the generator are mitigated. Issues of power quality are intrinsically addressed, and those of personnel safety would be resolved through staff training.

Synchronous islanding is made possible by the application of time synchronized phasor measurements, known as synchrophasors. A phasor represents the amplitude, frequency and phase angle of a voltage or current on a power system. Synchrophasors are phasors acquired by a device disciplined to the UTC timebase (typically using a GPS receiver), such that measurements can be taken over a wide geographical area and then usefully compared.

In a synchronous island, the embedded generator operates a control loop to match the properties of a phasor acquired at the generator terminals to a phasor representing a reference location. A suitable reference location would be a highly interconnect substation.

Prior work has considered the physical aspects of the synchronous islanding problem, and identified useful technical design parameters. For example, the system shall fail if the phase angle diverges beyond ±60° [4], or if the telecoms delay becomes longer than 300ms [2].

A current parallel activity is investigating a physical demonstration of multi-generator synchronous islanding, which necessitates a further supervisory control loop with telecommunication between generators. This was previous presented as a simulation exercise [5], where the ICT was considered ideal to allow emphasis on the physical challenge. The physical demonstrator must address the role of ICT.

The quality of the phasor measurements themselves are important which has prompted the authors to develop their own phasor measurement unit. The original version of this work has been open sourced with further details in [7]. A new version, which has been constructed in a modular manner specifically to enable interdisciplinary collaboration, is in the final stages of development. This new version the OpenPMU is split into three modules; data acquisition, signal processing, and telecommunications. This allows subject experts to contribute to their module, with a simple API between each module. This yields a technology platform that is highly suitable to develop further 'smart grid' applications in addition to the PMU function.

In this work, it is important to consider that the present standard governing PMU telecommunications, IEEE Std. C37.118.2 [6] provides no intrinsic security mechanisms, instead passing responsibility for security to the network. The modular design of the PMU used in this work allows state-of-the-art security mechanisms to be designed intrinsically into the measurement technology.

This work investigates the link between the cyber domain and the physical domain. High quality measurement technology allows measurement error to be discounted so that small signal analysis of control loop response can yield information as to the state-of-health of the telecoms network, potentially providing a mechanism to identify unauthorized intrusion or manipulation of the sensor data.

To investigate these challenges, the authors have implemented a testbed that allows for the evaluation of novel control algorithms for various Smart Grid use cases from a control and an ICT perspective. Currently the authors are focusing on the synchronous islanding problem but the capabilities of the testbed are not limited to this use case.

3. TESTBED DESCRIPTION

The laboratory setup consists of one load bank, two generator sets, a transmission line model, two synchrophasor measurement units (which can be standard commercial equipment or they could be OpenPMUs described in [6]) and a tie to the main utility grid. A simplified schematic is shown in Figure 2. One of the synchrophasor measurement units is connected to the islanded area, and the other to the main utility grid providing a reference angle for the controller of the islanded area to steer the islanded area towards. Alternatively one could skip the main utility grid synchrophasor unit and provide a stream of historical synchrophasor

measurements that would act as the reference for the setup. Figure 3 shows an overview of the laboratory setup.

The laboratory test setup contains two generator sets which will be used to mimic the behaviour of hydro-electric generation. These generator sets are built by Scott & CO. A 7.5 bhp 1500 RPM DC machine acts as the turbine providing mechanical torque, coupled to a 3 phase 5 kVA 2 pole pair synchronous machine acting as the generator.

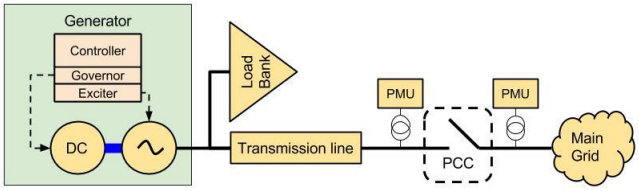

Figure 2: Schematic of laboratory setup of single machine synchronous island

The generator sets have two main ways of providing controllability to the laboratory setup; adjusting the torque provided by the DC machine effectively changing the active power output of the generator and controlling the excitation current in the synchronous machine which will lead a change of reactive power produced or drawn by the generator.

In order to control the generator sets in the laboratory environment a mini-computer, a Raspberry Pi [9], is provisioned to act as a controller for each set (Fig. 4). The Raspberry Pi functions as the turbine governor and the automatic voltage regulator, controlling the torque of the generator set. It controls the generator set in such a way that the response mimics the response of the power plant it represents, in this case a hydro power plant.

The authors have chosen to add a Pulse Width Modulation (PWM) add-on card to the Raspberry Pi to allow the control of both the 'Eurotherm DC machine controller' and the excitation current for the synchronous machine. The 'Eurotherm DC machine controller' requires a smooth voltage input from 0-10 V in order to function correctly so a simple operation amplifier low pass filer circuit was required to smoothen out the PWM signal generated by the Raspberry Pi add-on card. The operation amplifier also amplifies the voltage by two effectively increasing the output voltage from 0-5 V to 0-10 V.

In order to control the generator shown in Figure 3 a hardware controller was built. The hardware controller utilizes a 'Raspberry Pi 2' as the main computation unit and a micro controller, the Teensy 3.2 [8], as an input/output device that allows the Raspberry Pi to send set-point adjustments to the generator and receive feedback required by the automatic voltage regulator. The Raspberry Pi runs a Python code that controls the generator set such that the generator set mimics the operation of a hydro-electric generator set. This is implemented using time constants to the control input providing the generator set with longer response time.

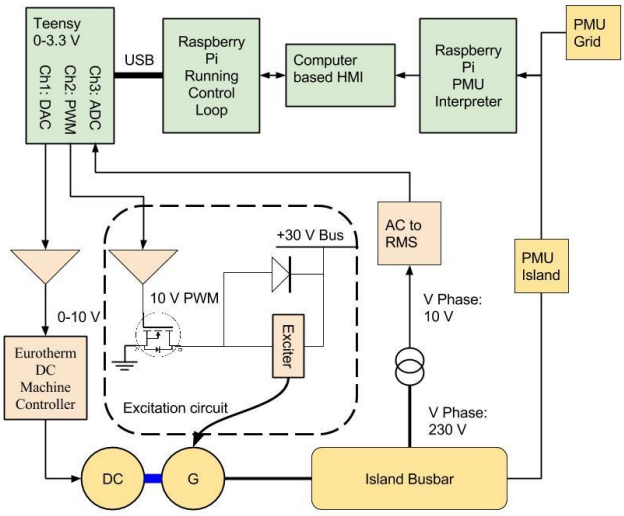

Figure 3: Overview of controller implementation

4. A MODULARIZED PMU

It has been found that present commercially available PMU equipment has unsatisfactory performance. The development of novel real-time control in this work is reliant on the quality of the sensing and feedback apparatus. It was identified early in the work that it would be necessary for the authors to build upon their existing work on the 'OpenPMU' and produce a high quality instrument. The authors' have modularized the functions of the OpenPMU into three major distinct functions. These are data acquisition, signal processing, and data representation (telecommunications), as depicted in Figure 4.

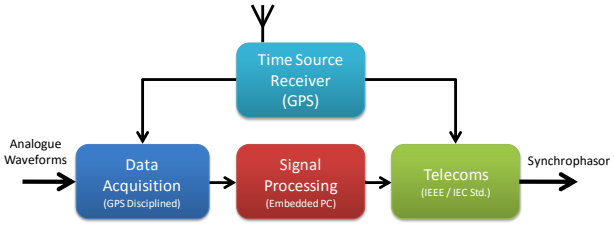

Figure 4: The major components of the OpenPMU system

By modularizing the functions of a PMU, this allows subject experts to focus their efforts on the area of PMU technology to which they can best contribute. The data acquisition stage requires a subject expert in ADC devices, hard-real-time programming environments and time transfer. The signal processing stage is essentially a mathematical problem that can be offloaded to a powerful CPU; subject experts in phase estimation algorithms can apply their expertise here. The telecoms component requires a subject expert in network communications, historians/databases and security.

The interface between each module is achieved using XML datagrams, transmitted by UDP. When two functions appear on the same machine, the operating system's local loopback address is used. All data is presented with human readable tags, and in ASCII format. This is bandwidth inefficient, but allows for extensive interoperability and reduces the learning curve for new contributors. Bandwidth efficient can be improved by adopting JSON in favor of XML. An example of a datagram containing raw waveform sample data is presented in Figure 5. The waveform was

acquired at 12.8 kHz, 16-bits, then formatted into Base64. This section shows 10 ms worth of sample data.

```
<?xml version="1.0" encoding="UTF-8"?>
<OpenPMU>
    <Format>Samples</Format>
    <Date>2015-06-30</Date>
    <Time>10:18:07.460</Time>
    <Frame>23</Frame>
    <Fs>12800</Fs>
    <n>256</n>
    <bits>16</bits>
    <Channels>3</Channels>
    <Channel_0>
        <Name>Belfast_Va</Name>
        <Type>V</Type>
        <Phase>a</Phase>
        <Range>275</Range>
        <Samples>
        VWxJVG5CaVYyeHpXVmhKWjFwWE5XcGlNbEp3WW0xaloyTXlUbT
        lhVnpGc1kzbENNR0ZIUmpCS1IxWjFXVEk1YTFwVFFtbGhWelZv
        WTI1cloxcEhSakJaVTBKcFpWTkNNR050Vm1oa1IyeDFXbmxDY0
        dSRFFuVmtWekZzWTIxc2FsbFhlSE5sVTBKb11tMVJaMlJJU21o
        aWJrNXpXVmhTY0dKdFkyZGhXRkZuWVZjMU1HSjVRbWhKUjBwb1
        l6S1ZaMDVxVVdkamJWWjNZMjFXZWxwWE4=</Samples>
    </Channel_0>
</OpenPMU>
```

Figure 5: Example XML Datagram

This modular platform has potential as a Smart Grid development platform beyond PMU devices, with applications in Smart Metering, Demand Side Management and Power Quality in development.

4.1 Secure Communication Framework

To ensure secure and reliable communication, the developed communication framework is based on IEC 61850-90-5 [10]. The protocol stack for IEC 61850-90-5 standard is depicted in Figure 6 that highlights it evolution from substation automation standard i.e., IEC 61850. The synchrophasor measurements are transmitted using Sampled Values (SV). SV is a stream based messaging protocol designed for high speed sharing of information across the system for time-critical applications.

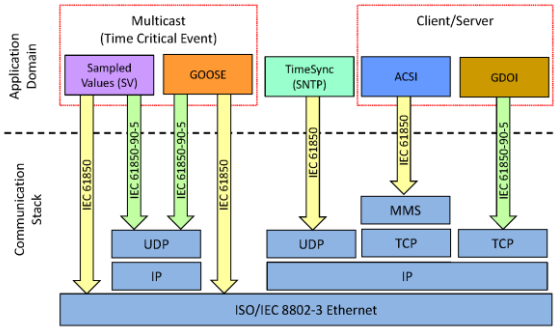

Figure 6: Protocol stack of IEC 61850-90-5

The SV messages are protected by a security mechanism known as Group Domain of Interpretation (GDOI) [11]. GDOI provides enhanced security and protection against man-in-the-middle, connection hijacking, replay, reflection and denial-of-service attacks. It is a group key management protocol that relies on ISAKMP for secure authentication of communicating peers over an insecure wide-area network. The key establishment mechanism of our communication framework is based on Diffie Hellman exchange which is one of the most successful protocol for public key cryptography [12]. The GDOI mechanism assigns Security Association (SA) to communicating peers based on information stream, destination Ethernet or IP address and content/service type. Further, each SA has a certain validity and is refreshed periodically to achieve best possible protection against cryptanalysis. IEC 61850-90-5 along with the distinguishing features of GDOI makes it the well suited communication framework for protecting the critical infrastructure involved in the synchrophasor applications.

4.2 Increased Accuracy of Phasor Measurement Units for Grid Control

As a core part of the phasor measurement units, the data acquisition block in Figure 4 is specially designed to utilize the widely available GPS signal. As illustrated in Figure 7, it consists of four components: a BeagleBone black development board as the main processing unit, a GPS receiver providing the required time information and synchronization, a phase locked loop (PLL) tracked on the UTC time and producing the data sampling clock, and finally an analogue to digital converter digitalizing the input phasors.

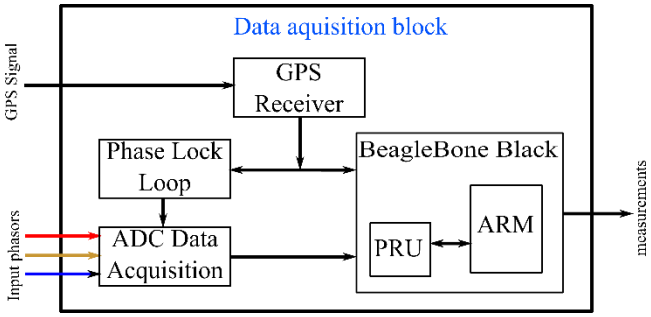

Figure 7: Data acquisition block structure

The GPS receiver provides a pulse signal every second, which can be used to synchronize the start point for ADC sampling regardless of the local time. Then the PLL accepts the 1 pulse-per-second (PPS) signal, and generates a desired sampling clock for ADC. It should be noted that due to the relative low frequency of 1 PPS signal, it has a high requirement on the supporting circuits, which are carefully designed to satisfy the standards.

The unique feature of BeagleBone Black with one ARM processor having two additional Programmable Real-time Units (PRU) inside makes it ideal for task of phasor measurements. It is designed that the main ARM CPU would be gathering sampled data in a higher interval (normally 10 ms) and pack it in one XML frame for the following signal processing block in Figure 4. The PRU acts as a middleware between ARM and ADC, by buffering the digital data from a high sampling rate of 12.8 kHz to a lower rate of 0.1 Hz.

5. Conclusion

This paper has presented a test-bed that enables the prototyping and validation of a several core features necessary for future smart grid systems. A full suite of technologies have been implemented, as necessary for provision of complete end-to-end smart grid real-time control strategies. The presented test-bed incorporates a number of novel features based around the modular design for phasor measurement, including state-of-the-art security mechanisms for secure communications across a wide area network. A key benefit of the presented system is the ability to experiment with the full stack of functionality, spanning from the physical –electrical– layer, through standards compliant communications layers, right up to the application layer, where smart grid control functions can be implemented and tested.

6. ACKNOWLEDGMENTS

This work is funded by the EPSRC CAPRICA project (EP/M002837/1).

7. REFERENCES

[1] Eirgrid, “Wind Generation & System Demand Data”, Eirgrid Plc, Dublin, Ireland, Online: www.eirgrid.com, February 2016

[2] Best, R.J., Morrow, D.J., Laverty, D.M., Crossley, P.A.. Synchrophasor Broadcast over Internet Protocol for Distributed Generator Synchronization. IEEE Trans. Power Delivery. Vol. 25, No. 4, October 2010, pp. 2835-2841.

[3] Laverty, D.M.; Best, R.J.; Morrow, D.J., "Loss-of-mains protection system by application of phasor measurement unit technology with experimentally assessed threshold settings," in *Generation, Transmission & Distribution, IET* , vol.9, no.2, pp.146-153, 1 29 2015

[4] Best, R.J.; Morrow, D.J.; Crossley, P.A., "Effect of loading, voltage difference and phase angle on the synchronisation of a small alternator," Electric Power Applications, IET , vol.3, no.6, pp.531,542, November 2009

[5] Best, R.J.; Morrow, D.John; Laverty, D.M.; Crossley, P.A., "Techniques for Multiple-Set Synchronous Islanding Control," in *Smart Grid, IEEE Transactions on* , vol.2, no.1, pp.60-67, March 2011

[6] IEEE Standard for Synchrophasor Data Transfer for Power Systems," in IEEE Std C37.118.2-2011 (Revision of IEEE Std C37.118-2005) , vol., no., pp.1-53, Dec. 28 2011

[7] D. M. Laverty, R. J. Best, P. Brogan, I. Al Khatib, L. Vanfretti, and D. J. Morrow, “The OpenPMU Platform for Open-Source Phasor Measurements,” *IEEE Trans. Instrum. Meas.*, vol. 62, no. 4, pp. 701–709, Apr. 2013

[8] “Teensy 3.2 Information.” [Online]. Available: https://www.pjrc.com/teensy/. [Accessed: 1-Nov-2015]

[9] “Raspberry Pi Information.” [Online]. Available: http://www.raspberrypi.org/. [Accessed: 17-Apr-2014]

[10] IEC Standard 61850-90-5: Use of IEC 61850 to transmit synchrophasor information according to IEEE C37.118, IEC/TR 61850-90-5, Edition 1.0, May 2012.

[11] B. Weis, S. Rowles and T. Hardjono, "The Group Domain of Interpretation", in IETF Request for Comments: 6407, October 2011.

[12] E. Rescorla, "Diffie-Hellman Key Agreement Method", in IETF Request for Comments: 2631, June 1999

[13] Wei, D., Lu, Y., Jafari, M., Skare, P. M., & Rohde, K. (2011). Protecting Smart Grid Automation Systems Against Cyberattacks. Smart Grid, IEEE Transactions on, 2(4), 782–795.