Science Hackathons for Cyberphysical System Security Research

Putting CPS testbed platforms to good use

Simon N. Foley
IMT Atlantique, Lab-STICC,
Rennes, France

Fabien Autrel
IMT Atlantique, Lab-STICC,
Rennes, France

Edwin Bourget
IMT Atlantique, Lab-STICC,
Rennes, France

Thomas Clédel
IMT Atlantique, Lab-STICC,
Rennes, France

Stephane Grunenwald
IMT Atlantique, Lab-STICC,
Rennes, France

Jose Rubio Hernan
Télécom SudParis,
Evry, France

Alexandre Kabil
IMT Atlantique, Lab-STICC,
Brest, France

Raphaël Larsen
IMT Atlantique, LaTIM,
Brest, France

Vivien M. Rooney
IMT Atlantique, Lab-STICC,
Rennes, France

Kristen Vanhulst
Télécom ParisTech,
Paris, France

ABSTRACT

A challenge is to develop cyber-physical system scenarios that reflect the diversity and complexity of real-life cyber-physical systems in the research questions that they address. Time-bounded collaborative events, such as hackathons, jams and sprints, are increasingly used as a means of bringing groups of individuals together, in order to explore challenges and develop solutions. This paper describes our experiences, using a *science hackathon* to bring individual researchers together, in order to develop a common use-case implemented on a shared CPS testbed platform that embodies the diversity in their own security research questions. A qualitative study of the event was conducted, in order to evaluate the success of the process, with a view to improving future similar events.

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems**; • **Security and privacy** → *Security in hardware*; • **Software and its engineering** → *Software development process management*;

ACM Reference Format:
Simon N. Foley, Fabien Autrel, Edwin Bourget, Thomas Clédel, Stephane Grunenwald, Jose Rubio Hernan, Alexandre Kabil, Raphaël Larsen, Vivien M. Rooney, and Kristen Vanhulst. 2018. Science Hackathons for Cyberphysical System Security Research: Putting CPS testbed platforms to good use. In *Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC '18), October 19, 2018, Toronto, ON, Canada.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3264888.3264897

"The hardest single part of building a software system is deciding precisely what to build"
[Frederick P. Brooks Jr., "The Mythical Man-Month", 1975]

CPS-SPC '18, October 19, 2018, Toronto, ON, Canada

ACM ISBN 978-1-4503-5992-4/18/10…$15.00
https://doi.org/10.1145/3264888.3264897

1 INTRODUCTION

Securing Cyber-physical systems (CPS) is non-trivial, requiring techniques and understanding that span a variety of computational and physical components. It is often this diversity that is exploited by attackers: a siloing of expertise on the part of researchers and/or developers can mean that vulnerabilities are overlooked, introduced or interoperate in ways that are not anticipated [14]. CPS testbeds provide platforms that can help understand and investigate research on security techniques, and testbeds that reflect the diversity of CPS components have an important role to play [1, 8]. However, construction and maintenance of 'real-life' testbeds requires expertise [8, 12] which, along with potentially high capital and recurrent costs, may deter their use in research.

Sharing research equipment among research groups is one course of action. This is more than just sharing infrastructure, it is also about pooling expertise and collaboration. The challenge is *how* to enable this research using shared CPS platforms.

In this paper we consider how time-bounded collaborative events, centred around shared CPS testbed platforms, can be used to enable security research. These events [4], such as hackathons, jams and sprints, are increasingly used as a means to bring groups of individuals together, to explore challenges and develop solutions. The primary contribution of this paper is a methodology for a *science hackathon* used to develop a common use-case using shared CPS testbed equipment. In addition to providing a 'real-life' CPS infrastructure, the use-case supports diversity in the research questions it addresses. This is achieved by drawing together security research from a number of individual projects, including diagnostics, resilience, visualization and anomaly detection.

The paper is organized as follows. Section 2 discusses some related work on time-bounded collaborative events and CPS testbed platforms. Section 3 proposes the use of transverse use-cases on shared CPS testbed platforms as a means of supporting diverse research questions. Section 4 proposes the use of a science hackathon as a means to develop transverse use-cases and Section 5 describes a CPS testbed platform that was developed following this process. A qualitative study on the experience of this activity, described in Section 6, has been used to evaluate the process. Section 7 provides discussion and conclusion.

2 RELATED WORK

There is a good deal of published literature on time-bounded collaborative events, ranging from academic studies of public hackathons [16] to accounts by practitioners of their own experiences [6, 9]. An ongoing theme is the tension between the pressures to produce artefacts versus the idealised notion of a hackathon as a free-wheeling melting pot of creativity, ideas, and skills. In an ethnographic study of how participants organise themselves and their development practices, Richterich [16] found that the pressures to produce artifacts can skew the activity away from personal learning and technical depth. On the other hand, Frey et. al. [6] consider how organisational structures can slow innovation and describe their practical experiences of using hackathons as a means to facilitate innovation within the organisation.

The time-bounded event described in this paper is closest, in sentiment, to the science hackathon [9] which exposes researchers to new research challenges and helps them to develop their own research ideas within a broader research landscape. Science hackathons are time-bounded collaborative events where a group of researchers come together to identify research challenges and 'hack' new lines of research. A case in point is the science jam at the ACM Conference on Human Factors in Computing Systems, with the objective to enable small groups of individuals to research, and develop a research poster within a two day timeframe.

Contemporary hackathons that focus on CPS and IoT development are becoming common, although we are not aware of their use in a scientific context for CPS research. Wagner [18] discusses how the Scrum agile methodology might be adapted in the context of a design-sprint for CPS systems, although it has not been tried in practice. Taha et. al. [17] describe an agile development of an open-source CPS testbed that is intended for research and education, and supports simulation and verification of continuous and discrete CPS models. Our transverse use-cases go beyond simulation, with an emphasis on working with real-life infrastructure and platforms, concurring with Green et. al. [8] who highlight that in real-life, CPS infrastructures, such as Industrial Control Systems, comprise a wide range of different equipment and that research testbeds should reflect this diversity. This emphasis of dealing with real-life infrastructure and equipment for training and research is also seen in the development of Capture the Flag style gamification of an ICS testbed platform [1, 7]. The activity described in this paper is not a conventional competitive hackathon; as a combination of jam and hackathon, it emphasises hacking CPS platforms as a means to integrate, demonstrate and explore lines of research.

3 TRANSVERSE USE CASES

The Cyber CNI Chair is an Institut Mines Télécom (IMT) industry chair on the cyber-security of Critical National Infrastructure with an emphasis on security of cyber-physical systems. A collaboration between three geographically dispersed IMT schools and eight industry partners, *targeted* industry use-cases form an important part of the research activity of each doctoral/post-doctoral researcher. With over fifty academic and industry researchers and investigators working together on ten separate, although related, targeted projects, there is a risk that these research efforts become siloed, if not by project, then by school. Furthermore, Intellectual Property (IP) constraints surrounding an industry target use-case may make it difficult to share the results within the chair or to the broader scientific community. Foreground research results may become intertwined with commercially sensitive target use-case technology/background IP that an industry partner wishes to protect.

Building on the research activities across separate projects, *transverse use cases* are proposed as a means of addressing these risks. These use-cases are independent scenarios that are developed around a common cyber-physical system platform, and are intended to build synergy and enable unencumbered demonstration and sharing of research across the chair, as well as to the wider community. The ambition is to build innovative prototypes that span the research projects, while being fail-fast, so as to avoid the sunk-cost-fallacy, as needs be. We consider transverse use-cases in the spirit of Jackson [5, 10]: how do we construct an appealing and thought-provoking cyber physical system scenario in order to explain and further investigate our research?

4 SCIENCE HACKATHONS FOR SECURITY

A science hackathon provides a means to develop transverse use-cases, where researchers come together to explain ongoing research, work on requirements, learn and teach each other about relevant technology and develop rapid prototypes. In the following we describe the process that was developed. Our approach focussed on collegiality and exploration, as opposed to the development-centric and competitive focus of contemporary hackathons.

The science hackathon runs across three separate events, as depicted in Figure 1. A 2-hour brainstorming session develops a strawman proposal based on a potential target CPS platform. A one-day requirements jam considers this proposal, and develops a shared understanding of platform requirements, which are, in turn, developed and implemented at a 2-day prototyping hackathon.

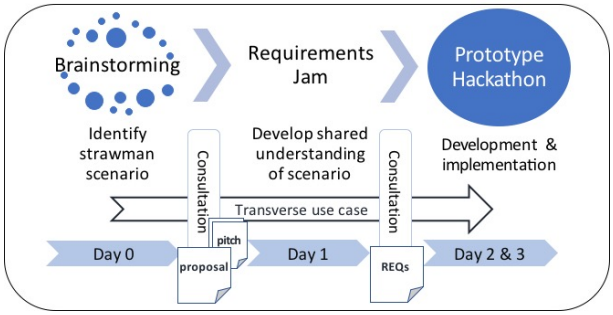

Figure 1: The three stages of the science hackathon

Brainstorming. Following an introduction to the goals of the hackathon, participants spend two hours brainstorming on the available CPS platforms and consider how a platform might form the basis of common scenarios interpreted from their individual research projects. The brainstorming session concludes with a short presentation and discussion of a strawman proposal. After the brainstorming session, a short (2 page) proposal document is prepared collaboratively online and circulated for consultation and feedback.

Requirements Jam. The objective of the one-day requirements jam is to develop shared understanding of the technical requirements for the platform to be developed and how research questions from individual research projects can be cast in this platform. One week prior to the jam, each participant reviews the strawman proposal and develops a (5 minute) presentation slide with their *initial* answers to three questions:

Threat scenario. *What is the threat (related to your research project) that you are focusing on, how do you plan on supporting it in the CPS testbed, and is your approach innovative?*

Technology challenges. *What technical development will be needed on the CPS testbed to implement the scenario? Are there potential obstacles? Can your scenario be implemented with minimal (re-)coding of the target (preferable)?*

Research challenge. *How does this relate to the research questions on your own project? Will the platform development leverage/enable your research work (preferable)?*

Each participant uses this to pitch how their work can contribute to the overall scenario. Pitches are debated, scenarios are synthesized, revised and/or eliminated, and presented and discussed with partners and investigators at the end of the requirements jam.

Following the requirements jam, a requirements document detailing the CPS threat scenario is prepared and circulated for feedback. This is a synthesis of the individual threats identified by participants and cast as a unified scenario for the testbed. Additionally, the document identifies required CPS technologies and development, and the research (project) challenges it is envisaged that the platform will support. The document prioritizes tasks to be carried out, whereby tasks with lower priority can be dropped during prototyping if necessary due to time constraints or feasibility concerns

Prototyping hackathon. The requirements document provides a tentative roadmap for the two day prototyping hackathon. Impromptu groups formed around the requirements, solving problems as they arose. Studies have shown that anticipation of having to prepare and give presentations can interfere with the technical focus of participants in a hackathon [16]. Therefore, since the primary objective was to produce technical artifacts for the transverse use-case, a programme of presentations was avoided and the two-days concluded with a short and informal debriefing session. Documentation and presentations of the results were left for after the hackathon.

5 CASE STUDY

A science hackathon was conducted during the first half of 2018, following the three stages described in the previous section. The brainstorming session (February) identified a Fischertechnik-based CPS testbed as the basis for the strawman scenario. Over the course of the requirements jam (May) and the prototype hackathon (June), the testbed was adapted to be consistent with project research questions. Seven doctoral researchers from seven different projects participated, along with one research engineer who provided domain expertise on the testbed. One Principal Investigator facilitated the process and an Applied Psychologist conducted a study of the experiences of the participants. This section gives an overview of the transverse use-case that was developed.

5.1 A Fischertechnik-based CPS testbed

The testbed used for the hackathon reproduces, at a small scale, how modern industrial control systems are connected to a organization's IT network, as depicted in Figure 2. A virtualized IT network with two VM border routers isolate the network from the Internet and the production network, and two VMs provide a supervision system and an administration system, to control and monitor the PLCs which drive the production line, respectively. A production network connects production line PLCs. Network-enabled Crouzet PLCs support modbus requests from virtualised supervision/administration workstations. An industrial firewall supporting modbus and S7, is configured to enable modbus communications between the PLCs and the supervision and administration workstations. Custom grafcet programs provide automation for the Fischertechnik platform, a machining production system with conveyor belts, DC motors, a pneumatic press, a drill, a mill, a robotic arm which manipulates the processed parts and assorted mechanical and optical sensors. Parts to be processed move from one machining tool to the next, using conveyor belts, pistons and motors. The processed parts are plastic cylinders that are moved by the robotic arm from the end of the line to the beginning so that the process run continuously, as required, without human intervention.

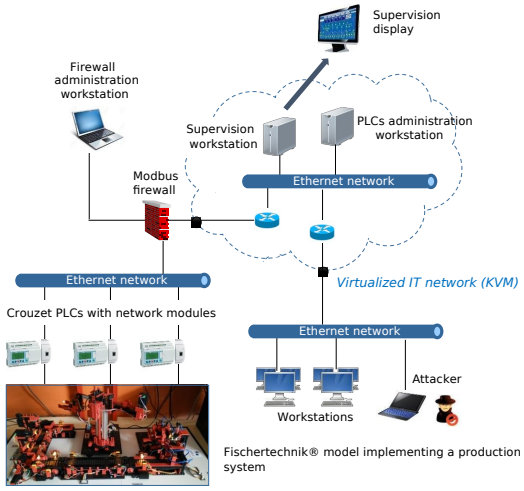

Figure 2: Fischertechnik CPS architecture

5.2 Threat scenario

The requirements jam storyboarded a threat scenario that framed research questions from different projects. It comprised an external attacker exploiting weaknesses in the network perimeter defences and/or an internal attacker with direct access to IT or OT components. The attacker has two objectives: (1) stop a conveyor-belt or release the clamp to halt the production line (easily noticed), or (2) change milling/drilling times in order to reduce finished product quality (more subtle). Both objectives require the attacker to send Modbus packets to at least one Crouzet PLC situated in the OT part of the organization. As only a specialized insider could physically access OT, other attackers need to take control at the OT administration workstation. Two attack vectors were considered: (1) exploit

a vulnerability of the workstation or obtain administrator credentials via a dictionary or password-guessing attack, and (2) disable (physically or remotely) the workstation, forcing administrators to use a less secure secondary rescue workstation.

5.3 Technology challenges

In implementing the threat scenario, testbed changes were necessary in order to enable researchers deploy their tools and investigate their research questions. The need for these changes were identified during the requirements jam, and implemented during the hackathon, along with other unanticipated changes, as they emerged. Further switches were added to the production network, supporting additional hosts, an Intrusion Detection System, network traffic sniffers and hosts to inject network traffic. The PLC grafcets were modified to expose some of their state variables in modbus registers. The Fischertechnik platform was extended to incorporate an interconnection with a second virtualized industrial system, as part of a larger industrial process. In this scenario, the Fischertechnik platform represents a manufacturing plant that builds a replacement component upon its failure in the virtualized industrial system. The state of the virtualized industrial system is modelled in terms of a collection of PLC/Modbus registers (using the EasyModbus library), that the SCADA system monitors and controls the Crouzet PLCs/Fischertechnik platform as appropriate. Virtualization is a good compromise, in term of resources expended during the hackathon, and facilitates making, in a realistic way, more complex scenarios for the research projects that deal with supervision and system modeling.

This setup of the testbed was used to support the threat scenario described in section 5.2. The fischertechnik process is triggered on failure of a component in the virtualized ICS. This could represent normal wear and tear, or an accident. However, it could also be a result of an attacker interfering with the production processes, forcing early failure. Such complex situations with interleaved safety and security call for more advanced diagnosis methods and are opportunities to illustrate the models developed in the research challenges the platform was designed for. A secondary workstation providing backup administration, configured with an old version of windows with known vulnerabilities, along with weak credentials, provides the (second) attack vector.

5.4 Research challenges supported by testbed

A primary objective of the science hackathon was to develop a CPS platform in which research questions from diverse projects could be demonstrated and explored. By providing distributions for the probability of component attacks and failures, the model proposed by Bourget et. al. [2] can be used to provide real-time diagnosis of security and safety in the testbed. Testbed snapshots of the state of the PLCs, along with alert events from the IDS enable diagnosis of the probability of attack/failure as the threat scenario (Section 5.2) evolves. This can be used to identify the origin of an incident that generates several alerts, compute probabilities of occurrence of future events or evaluate the likelihood of any given event. Such information can be used to decide which is the appropriate response to the incident. These PLC state snapshots and IDS events are also used in another project investigating anomaly detection.

During the implementation, a counter-intuitive observation of the risk model was recognized. The chosen approach was split into two phases. The first phase consists of recognizing the network as well as the Modbus registers. The second phase consists of finding the effective attacks and executing them. Intuitively, the first phase seems to be easier than the second one. In contrast, the recognition phase was much more time consuming than the other phase (according to our implementation). This kind of information is useful especially in the context of risk analysis. Indeed, it allows defining a likelihood of success for each attack (which corresponds to the risk model) in order to carry out the evaluation of the overall risk. In fact, this result highlights the importance of the sensitivity analysis on risk assessment. Moreover, such analysis should take into consideration assumptions on the risk model.

The testbed is a complex arrangement of hardware and software. In its design, various mechanisms help enhance its resilience to attack and failure. For example, network virtualization helps provide defense in depth, while sensor redundancy helps provide fault-tolerance. We use [3] to model deployments of these mechanisms in order to determine, and compare, their collective effectiveness at providing resilience in the testbed. For example, the rescue station is configured with weak credentials that allow an attacker to make a dictionary attack. If these credentials are replaced by stronger ones, how is the overall resilience affected?

Immersive 3D visualization techniques can provide insights into the security and safety of the system that go beyond the more conventional mechanisms and techniques [11]. We are using a 3D virtual environment to simulate testbed configuration, behavior and security events. This kind of simulation can help us to better understand the compatibility of our different models (diagnosis, source and path information, resilience) and to visualize the threat scenarios in a more graphical and interactive way.

6 EXPERIENCE OF THE HACKATHON

"If someone told us, do that, it wouldn't be a hackathon." [Interview extract]

The evaluation aimed to understand team member experience of the science hackathon, with a view to applying insights gained to improve the process in future events. It was decided to conduct semi-structured interviews with individual team members during the second day of the prototype hackathon while the team were still working together. An interview schedule was developed in conjunction with the organiser of the event. Thematic Analysis [15] was applied to the audio recordings. Individual recordings were summarised as transcribed text, emerging themes were identified, and structured the analysis. Transcribed material was anonymised and deidentified, and short verbatim extracts used to illustrate the analysis [13].

An ethical self evaluation was completed. The process of consent was initiated with participants via email, providing team members with information on the proposed evaluation. This was followed by a second email, with an information sheet on method and consent form, for consideration. The proposed data collection was explained verbally to participants on Day 1 of Stage 3; as was informed consent, such as anonymity and that there was no onus to participate. This was an opportunity for questions from the team; questions

and comments were also invited prior to, and following, individual interviews. A consent form was signed prior to each interview by both researcher and participant, and each retains a copy. The forms retained by the researcher are held securely. The interview data comprises a total of 141 minutes, with interview duration ranging from 15 to 30 minutes, with an average of 20 minutes. Themes that emerged from the analysis are discussed below.

6.1 A process for hacking scientific research

Expectations. The nature of science hackathon, as being a scientific collaboration, had been clear in the information circulated at the outset, and this was reflected in the expectations of the participants, as expressed by the comment: *"I asked colleagues, and they said a hackathon was when you program all day, but the scientific hackathon was explained by Simon, so I understood that it was not about creating a product, but to work with and talk to other people".*

Requirements Jam. The preparatory work required for the jam was reported as not being specific enough, in terms of substance and purpose. For instance, some thought that the questions for the pitch were too broad, and were unsure if their response was as expected. However, when the pitch slides from individual participants were shown at the second meeting, the usefulness of the requirements jam was clear. The preparatory work was seen as an opportunity to clarify individual participation, one participant comments that *".....we had to prepare this part and to write it down, we didn't use it much here, but writing it down made it clearer to us, in our heads".* This was in terms of illustrating the different research perspectives on the same issues, and the identification of shared goals. In another participant's words, *"we can work together in a holistic way, individuals can do parts, then blend".*

The extended time frame. For the team, being together over a longer time frame than usual facilitated cohesion, both professionally and socially. This is illustrated by the remark: *"Usually we have one day meetings, and you cannot talk with people a lot about the work, and the hackathon gives the chance, there is more time to talk to people about their work, about my work, and also not in a professional way, so that is good".*

6.2 Transverse use-case: avoiding research silos

Linking research strands. Participants talked about the collaborative aspect of the hackathon as being interesting, creative, enjoyable and useful. Working with, and learning from, others who have different perspectives and expertise in Cyber Security was engaging. Furthermore, interaction sparked creativity as illustrated by the following comment: *"It gives me ideas, working with everybody, and we produce stuff that we can use for their work, for my work, so I am very pleased".* A benefit for research in the Chair flowing from the opportunity to link disparate research strands was reported, as illustrated by the following: *"Interesting to find out about the diverse knowledge that each PhD student has, for example, good at hacking, good at graphing* [PLC grafcets]. *We have different knowledge, and put together, can create a different point of view, and explain better what they are doing in the Chair."* As the above illustrates, the process of collaboration and making links was something that the individual researchers appreciated and valued.

Positioning individual research projects. While one goal of the hackathon was to foster collaboration in the Chair, additional practical and conceptual benefits were identified for participants' individual work. For instance, being able to develop a practical scenario for their own research, exemplified by the following comments: *"What I wanted was to get an opportunity for a realistic scenario* [for my research]*, and we did it yesterday, so that worked for me"* and also *"the platform is the best way to show people your work and aims."* The conceptual benefit for individual PhD research is the reciprocal input into each others' research, helping to clarify PhD research, as illustrated by this comment: *"Being in the same room for two days and chatting, helped me to think about what I will put in my PhD."*

The function of the transverse use-case. Participants saw the platform as being a useful resource to demonstrate and explain their own, and others', research, and the possibilities it entails. For some, the potential use of the CPS platform is in communicating the meaning of an attack, and its significance. For instance, being able to demonstrate an attack to non-technical people is useful, as the following illustrates: *"It's very interesting to use* [the CPS platform]*, because we can see the attack, we were able to run some stuff on the platform, we started an attack and the platform started moving the wrong way, interesting stuff, and for people that can't understand the technical stuff, seeing the platform moving the wrong way, it's very understandable, and you can say, that is the attack."*

6.3 CPS platform as a communication resource

Conscious of their role as researchers working with industrial partners in the Chair, the participants envisage the platform as a means of facilitating communication to those industrial partners. For example, one participant talked about how they had not previously been successful making convincing prototypes, while at the hackathon, they succeeded with one simple attack scenario. Also, being able to convey to industrial partners what is possible and feasible from a research point of view can be challenging for participants. An example is communicating both the diversity and the broad scope of Cyber Security and, therefore, how the expertise of an individual is limited within the breadth of the domain.

Using the platform, and ideas from the hackathon, for communication with industrial partners is envisaged as being a reciprocal process. As such, having feedback is important to the participants, as the following explains: *"Having an industrial partner involved would be good, especially for feedback."* The hackathon is seen as a way of linking research to the real world. The CPS platform, and how it was used in the hackathon, are seen as a means of improving communication between researchers and industrial partners.

6.4 Autonomous collaboration

When asked a general question about what was the best part of the hackathon, participants responses focussed on the enjoyment of working in a team, and being challenged, while not under pressure to produce a integrated end-to-end artifact, as such. As PhD researchers generally working on their own within a supervisory framework, the hackathon provided an opportunity for them to be able to interact with peers, to be curious about others' work, and to learn about it, even if not understanding everything. The following comment exemplifies the sense of pleasure of the team work:

"I enjoyed the challenge of the hackathon, as Simon presented it, having to mix all the work, very interesting for me, the discussions were very interesting, I started the hackathon with great enthusiasm because of the idea of the scientific hackathon, I wanted to do something, talk with others, work in a group, because in a PhD you don't work with colleagues, so it was very interesting to work on the hackathon, new ideas, why not?"

The freedom that the unstructured nature of the prototyping hackathon facilitated was a very positive aspect of the experience. The lack of pressure to produce a specific deliverable meant that the team felt able to try ideas and discard them quickly. The following extract from one interview explains this:

"Thought it might be managed, and it wasn't, which is good, given an area to play, we felt free to do the best we could do, we didn't have a specific aim to reach, so whether it's a success or a small success, we don't feel much pressure about it, we feel free to try to things, if they work, good, if they don't work, try something else, this is good."

7 CONCLUSION

This paper described how a science hackathon, centred around testbed platform development, can be used to investigate security scenarios that reflect the diversity and complexity of real-life cyber-physical systems in the research questions that they address. Unlike a conventional hackathon/Capture the Flag event, the activity was coordinated across time-bounded collaborative events: brainstorming, requirements jam and prototype hackathon. Driven primarily by the PhD students, the autonomous and non-competitive nature of the event was beneficial, an observation consistent with other studies of time-bounded collaborative events [6, 16]. Ameliorating the risk of siloing research was also addressed effectively, as the transverse use case provides a broader context in which to understand and relate individual research challenges.

The testbed platform, and what was developed in the hackathon, was conceptualised as a means of communicating with others about research, such as industrial partners, and those with less technical domain expertise. As such, its role in improving communication is envisaged as being bidirectional, with the suggestion from participants that it could be useful as a way, not alone to explain their research, but also as a way for industrial partners to provide feedback on that research. This finding is a good use of the testbed, one of the main goals of the event.

Remedying the uncertainty expressed by participants around their expectations for the hackathon highlights an interesting area of tension. With a conventional hackathon, expectations are generally well understood: attack or defend. However, expectations are less certain in a science hackathon, as was evident in interviews. While introducing more organization may help provide clarity, it can work against innovation [6]. This motivated the requirements jam as a means to help set expectations, while encouraging autonomous collaboration. While being broad in its questions to the extent that individuals were initially somewhat unsure about how to respond, the jam did, nevertheless, provide clarity to individuals on the potential for them to contribute. As such, this was a valuable learning experience for participants, providing an opportunity to develop a sense of their own research in relation to that of others; how collaboration could work as a process, as well as for future research. Related to this is the uncertainty around the degree of the expected outcome, specifically, whether the conceptual collaborative work was required to culminate in a functioning artifact in order for the hackathon to be deemed a success. While providing additional guidance and direction, other than the level of advice that was given, would have removed some of the uncertainty experienced, it may have done so at the expense of the independent learning gained in this collaborative context. How best to achieve the best balance in this area of tension is an area for future research.

Acknowledgements. This work was supported the Cyber CNI Chair of Institute Mines-Télécom which is held by IMT Atlantique and supported by Airbus Defence and Space, Amossys, BNP Paribas, EDF, Orange, La Poste, Nokia, Société Générale and the Regional Council of Brittany; it has been acknowledged by the French Centre of Excellence in Cybersecurity.

REFERENCES

[1] D. Antonioli, H.R. Ghaeini, S. Adepu, M. Ochoa, and N.O. Tippenhauer. 2017. Gamifying ICS Security Training and Research: Design, Implementation, and Results of S3. In *Workshop on Cyber-Physical Systems Security and PrivaCy*. ACM.

[2] Edwin Bourget et al. 2018. Probabilistic Event Graph to Model Safety and Security for Diagnosis Purposes. In *IFIP WG 11.3 Conference on Data and Applications Security and Privacy*. Springer, LNCS 10359.

[3] Thomas Clédel et al. 2018. Towards the evaluation of end-to-end resilience through external consistency. In *10th International Symposium on Cyberspace Safety and Security*. Springer LNCS.

[4] Anna Filippova et al. 2017. Hacking and Making at Time-Bounded Events: Current Trends and Next Steps in Research and Event Design. In *Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 363–370.

[5] S.N. Foley. 2017. Getting Security Objectives Wrong: A Cautionary Tale of an Industrial Control System. In *Security Protocols XXV - 25th International Workshop*. Springer LNCS 10476.

[6] Frank J. Frey and Michael Luks. 2016. The Innovation-driven Hackathon: One Means for Accelerating Innovation. In *Proceedings of European Conference on Pattern Languages of Programs*. ACM.

[7] Francisco Furtada, Lauren Got, Sita Rajagopal, Elaine Cheong, and Ericson Thiang. 2017. *S3-17: SUTD Security Showdown (Event Report)*. Technical Report. Centre for Research in Cyber Security, Singapore University of Technology and Design.

[8] Benjamin Green, Anhtuan Le, Rob Antrobus, Utz Roedig, David Hutchison, and Awais Rashid. 2017. Pains, Gains and PLCs: Ten Lessons from Building an Industrial Control Systems Testbed for Security Research. In *Workshop on Cyber Security Experimentation and Test, (CSET), Vancouver*. USENIX.

[9] D. Groen and B. Calderhead. 2015. Science hackathons for developing interdisciplinary research and collaborations. *eLife* 4 (2015).

[10] M.A. Jackson. 1989. Getting It Wrong: A Cautionary Tale. In *JSP & JSD: The Jackson Approach to Software Development;*, John Cameron (Ed.). IEEE CS Press.

[11] Alexandre Kabil et al. 2018. Why should we use 3D Collaborative Virtual Environments for Cyber Security?. In *IEEE Fourth VR International Workshop on Collaborative Virtual Environments*.

[12] A. P. Mathur and N. O. Tippenhauer. 2016. SWaT: a water treatment testbed for research and training on ICS security. In *2016 International Workshop on Cyber-physical Systems for Smart Water Networks (CySWater)*. 31–36.

[13] D. C. O'Connell and S. Kowal. 1995. Basic Principles of Transcription. In *Rethinking Methods in Psychology. Part II, Discourse as Topic*. Sage Publications.

[14] Olgierd Pieczul, Simon Foley, and Mary Ellen Zurko. 2017. Developer-centered Security and the Symmetry of Ignorance. In *Proceedings of the New Security Paradigms Workshop*. ACM, 46–56.

[15] J. Richie and J. Lewis (Eds.). 2003. *Qualitative Research Practice: A Guide for Social Science Students and Researchers*. Sage Publications, London.

[16] Annika Richterich. 2017. Hacking events: Project development practices and technology use at hackathons. *Convergence: The International Journal of Research into New Media Technologies* (May 2017), 1–27.

[17] Walid Taha et al. 2016. Acumen: An Open-Source Testbed for Cyber-Physical Systems Research. In *Internet of Things. IoT Infrastructures (Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering)*. Springer, 118–130.

[18] Stefan Wagner. 2014. Scrum for Cyber-physical Systems: A Process Proposal. In *Proceedings of the 1st International Workshop on Rapid Continuous Software Engineering (RCoSE 2014)*. ACM, New York, NY, USA.