Towards Cyber Sensing: Venturing Beyond Traditional Security Events

Artūrs Lavrenovs, Kimmo Heinäaro and Erwin Orye
NATO CCDCOE, Tallinn, Estonia
Arturs.Lavrenovs@ccdcoe.org
Kimmo.Heinaaro@ccdcoe.org
Erwin.Orye@ccdcoe.org

DOI: 10.34190/EWS.20.062

Abstract: Host and network-based events are the backbones of any modern IT monitoring and detection system. The number of lower priority security events is significant and might contain weak indicators of cyber attacks; by combining host and network events with sensor data that are not part of conventional IT security, we are able to elevate otherwise missed events to discover hidden cyber attacks. The sensor data is fed into a situational awareness system which augments traditional alerts. This technique is primarily applicable for critical infrastructure, military, government and large organisations where the adversary is sophisticated enough to bypass existing detection methods. We discuss operational and strategic implications by using this type of sensor. We have implemented these principles in two scenarios tested in cyber exercises. In the first proof of concept we focused on sensor fusion by integrating existing non-IT sensor systems with IT security and correlated the collected data. This enabled the Blue Team to detect well-hidden Red Team attacks against a simulated power grid and counteract them. In the second, we explored a large variety of sensors monitoring individual personnel and their operating environment. Sensors used in this research are categorised into biological, environmental and EM spectrum. Biological sensor data includes heart rate, stress level and brain wave monitoring. Environmental sensors monitor the RF spectrum, CO_2 level, VOC level, temperature, humidity, infrared, ultraviolet, visible light, noise level, proximity and vibration.

Keywords: host-based events, network-based events, IDS, security events, sensors

1. Introduction

Critical infrastructure, government and military networks are under ever-increasing threat of cyber attack. The defence solutions market is experiencing growth and vendors are constantly developing new and more advanced solutions that use cutting-edge approaches like Artificial Intelligence (AI). But most of these solutions rely on traditional sources of data - host and network-based events. As adversaries in this scenario are up to state-level actors, they have the resources to investigate, adapt to and overcome new advanced defences. This calls for a widening of our view and exploring additional sources of data.

It does not necessarily equate to acquiring another new and expensive solution containing both hardware and software components, but rather evaluating what sensor data is already available. Sensors can be viewed from system-centric and human-centric perspectives. The former primarily focuses on the states of the systems, inferring human properties indirectly when possible, while the latter addresses biological data that can be measured directly via dedicated sensors. We refer to the combining of traditional and other sensor data for the purpose of detecting cyber attacks as cyber sensing.

In Section 2 we review existing research addressing human behaviour analysis and training models. In Section 3 we explore data sources acquirable from different sensors. In Section 4 we describe a proof of concept correlating traditional security events with building automation and gathering data from environmental sensors. Section 5 discusses operational and strategic implications of cyber sensing, and Section 6 presents our conclusions.

2. Related work

There are two related streams of work in this field. The first is from a human-centric perspective and the second concerns how computer models can be used to predict human behaviour.

Sensor fusion mimics the human brain combining multiple senses. Lin et al. (2004) propose a neural network architecture to integrate data from several physiological and behavioural sensors to improve reliability and resistance to impersonation (multimodal verification system).

Augmenting cyber-physical systems with sensors monitoring humans results in cyber-physical-human systems (CPHS) (Sowe et al., 2016). CPHS is not investigating defensive or offensive applications researched by us but the safety of humans (Gelenbe et al., 2012) and security of the sensors and produced data.

2.1 Human behaviour analysis

The analysis comes in three stages: human behaviour measurement, modelling and analysis (Carus et al., 2014). In straightforward research, a relatively small number of people are observed for a long time using many sensors. Typical of this research is people living in smart homes, where we build up personal data that can be classified under 'normal behaviour' and from that point, we find a threshold where deviating behaviour can be identified using supervised or unsupervised machine learning. Examples of this are:

- A multi-view setup of cameras used to recognise people's behaviour based on human action recognition with multi-view scenarios and real-time execution (Chaaraoui et al., 2014).
- Analysing relationships between behavioural patterns and energy consumption by correlating home-based activities with electricity use (Chen, Cook and Crandall, 2013).
- Analysis of the behaviour of elderly people in their homes by detecting patterns in sensor data in an unsupervised manner (Rieping, Englebienne and Kröse, 2014).

Another popular research topic is the abundance of data created by mobile phones. In modern cell phones, there are often accelerometers, gyroscopes, magnetometer, GPS, barometers, proximity sensors and an ambient light sensor. An example of an academic analysis of human behaviour using location knowledge obtained from data sent by heterogeneous moving objects can be found in Yus et al. (2014).

Closing in on the effect of human behaviour on cybersecurity, most of the related work focusses on the protection of one's own IT-systems. An example is studying the human factors aspect of cyber defence. Gutzwiller et al. (2015) focus on education, training and situational awareness to enable people to pick up cyber attacks that have been missed by automated tools.

Modelling human behaviour to anticipate insider attacks via system dynamics (Greitzer and Hohimer, 2016) is one of the few papers that looks at cyber effects from the perspective of an insider threat, trying to predict the probability by using a diverse set of data sources from the cyber domain and inferred psychological and motivational factors of a malicious insider activity.

2.2 Training models

Computer-aided learning methods depend on the kind of data that is available. If data for both inputs and outputs is available, supervised learning is often a good approach. If only input data is available, unsupervised learning or artificial intelligence is usually used. In cases where no direct access to modify the output is possible, but some measurement of the quality of an output that follows an input is, then reinforcement learning can be used. In the realm of human behaviour, it is very difficult to predict the outcome with a known set of inputs. Therefore, most research currently focusses on unsupervised learning.

Bass (1999) described how detection systems will fuse data from heterogeneous distributed network sensors to create cyberspace situational awareness. He says 'situational awareness technology is alarmingly primitive relative to protecting our critical electronic infrastructures'.

An analysis of unsupervised activity discovery and primitive sequence information in this field can be found in Seiter et al. (2014). Feature representation of sensor data that does not rely on prior expert knowledge of human activity and using unlabelled data for activity recognisers is examined in Bhattacharya et al. (2014). A behavioural activity recognition model for mini-wearable devices is used to recognise activities of daily living in Hu et al. (2014), human activity recognition from video images in Jalal et al. (2017) and training a neural network to predict whether video frames and audio are temporally aligned to predict sound source localisation and audio-visual action recognition in Owens and Efros (2018).

3. Exploring data sources

This paper is about exploring new ways to detect cyber incidents. A wide variety of sensors is used, without any assumptions about whether they can be related to the subject or not. Data collected from the sensors is correlatable with knowledge and the timeline of Red Team activities in a cyber exercise. It might be possible to detect cyber incidents from improbable sources. Further studies are needed to check which sensors could be used for early warning of an attack and which are merely an indication of an incident already discovered.

Data sources used can be divided into system- and human-centric. System-centric data sources are based on the integration of existing security and automation systems and are further studied in Section 4. Human-centric data sources are based on individual sensors that are described here. These can be divided into three classes: environmental, electromagnetic spectrum and biometric.

3.1 Environmental sensors

The environment of the defending Blue Teams could be monitored with a wide variety of different sensors. The positioning of environmental sensors is critical for accurate measurements. In this study, no assumptions of values have been made and sensors were situated at random locations near the Blue Teams. The following environmental data sources were monitored: vibrations (acoustic noise level and acceleration); atmosphere (air quality, temperature, humidity and atmospheric pressure); and computer use (current consumption).

Acoustic noise level can be measured with a microphone, digitized with an AD converter and averaged over time. Acceleration can be measured using integrated circuits containing miniaturised mechanical structures (microelectromechanical systems, MEMS). An accelerometer can detect movement, vibration and shock. (Scanaill et al., 2013)

Air quality was monitored by measuring CO_2 level, volatile organic compound (VOC) level and alcohol level. Gas content in the air can be measured with miniature semiconductor sensors that contain a heated thick-film metal oxide layer on a base that contains electrodes for measuring the resistance of the circuit. Composition and structure of the metal oxide layer are adapted to the target gas (Scanaill et al., 2013). CO_2 is produced by human respiration and the amount increases with physical activity (Fucic et al., 2012). VOCs are carbon-based chemicals that are considered as indoor air pollutants. They are typically emitted from building materials and furniture (Scanaill et al., 2013). Alcohol level is normally measured as breath-alcohol content from a single person. Temperature can be monitored by measuring voltage over a semiconductor p-n junction under constant current (Scanaill et al., 2013). Relative humidity can be determined by measuring the capacitance of material exposed to air (Sazonov, 2014). Atmospheric pressure can be measured with a piezoresistive MEMS device, where pressure bends material causing changes in resistance (Scanaill et al., 2013).

Current consumption of mains voltage powered devices (e.g. workstations and servers) can be measured with Hall effect sensors. The magnetic field produced by the current in a wire is measured by measuring voltage over a semiconductor under constant current, and the mains voltage is isolated from the measuring circuit (Scanaill et al., 2013).

3.2 Electromagnetic spectrum sensors

In this study, the aim was to monitor as wide a range of the electromagnetic spectrum as possible. With simple sensors, the following frequency ranges were covered: ultraviolet, visible and infrared light. Radio frequencies were monitored in the range of 0-6 GHz.

Ultraviolet band radiation is produced by the sun or special lamps and can be detected with a photodiode that converts light into electricity with a semiconductor p-n junction. The visible light level can be monitored with a Light Dependent Resistor (LDR), which is a semiconductor that changes conductivity according to the light level (Scanaill et al., 2013). With simple sensors, it is only possible to detect the average level of the band measured. Infrared band radiation is emitted by objects near room temperature. The IR sensor used is a proximity sensor that senses movement in the vicinity of the sensor.

The RF spectrum can be monitored with a Software Defined Radio (SDR). If the frequency range to be monitored is wider than the bandwidth of the SDR, spectrum scanning can be used. Scanning rate (Hz/s) is limited by the hardware and will leave gaps in the time axis for real-time spectrum monitoring. With scanning, it is possible to log all signals detected above a certain threshold amplitude.

3.3 Biometric sensors

Monitoring biometric values of a person in the Blue Team can reveal stress level or changes in behaviour. It can be used to detect cyber incidents as they start to cause problems. In this study, the following biometric sensors were used: heart rate sensor, cortisol level sensor and brainwave sensor.

Heart rate monitoring can be used in sports training or measuring health status and mental stress. Optical wearable sensors (photoplethysmography, PPG) can be made with an LED light source, photodetector and using a signal processing algorithm (Sazonov, 2014). They are available in the form of a wristwatch and are easy to wear in daily activities. Cortisol is a stress hormone, along with adrenaline and noradrenaline. It helps the body to function in stressful situations but is also harmful to health if the stressful situation lasts for a long period (Bonetti, 2014). Brain activity can be monitored with electrodes attached to the scalp (electroencephalography, EEG). The differential signal between electrodes is amplified and filtered to produce waveforms that can also be visualised. Signals can be categorised as gamma (indicating learning), beta (logical thinking), alpha (deep relaxation), theta (sleep) and delta (deep sleep) based on their frequency range (0-100 Hz) (Le et al., 2018).

3.4 Further development

The sensors selected for this study were mostly modern, semiconductor-based, inexpensive and widely available and most were factory calibrated with digital output. Most of the values measured possibly cannot be used to detect cyber incidents, even after further studies. The secondary target of the study was to create a universal sensor platform that can measure and log as wide range of physical values as possible. Measurements that might be useful in detecting cyber incidents, but are not part of this study, include adrenaline level, human body temperature, blood pressure, the consumption rate of beans in the coffee machine and thermal camera imaging of the Blue Team. Collecting and combining data from biometric sensors on all Blue Team members might also reveal new correlations with cyber incidents.

4. Cyber sensing proof of concept

We have implemented cyber sensing of explored data sources into a proof of concept from a system-centric and human-centric perspective. These two perspectives differ significantly therefore we independently developed prototypes for the cyber exercises Crossed Swords and Locked Shields.

4.1 System-centric cyber sensing

System-centric cyber sensing can use pre-existing systems and sensors, thus not necessarily requiring additional hardware components. Every modern building has significant automation controlling the heating, ventilation and air conditioning (HVAC), energy efficiency (e.g. smart lights), safety and security systems. At best, only the building's security system might be integrated with the organisation's cyber defence systems. Usually these sensors can only infer human presence and activity indirectly, commonly by detecting movement, IR radiation or temperature changes.

Critical infrastructure protection has been one of the main components of scenarios in cyber defence exercises for the last few years. Power grids and generation have become classical components of both defence and attack scenarios. And rightfully so, as attacks against power grids have already been seen in the real world and proven highly effective (E-ISAC, 2016).

For the proof of concept, we used a simple building automation system consisting of access control, smart lights, fire alarm and security alarm. This is representative of what is already commonly available without any additional hardware or software changes in existing building automation systems as long as the status of the sensors can be acquired. Building automation systems typically use measured values for controlling factors such as temperature, and security systems typically use only one-bit information sources (e.g. door opened or closed). In this study, we used the built-in web server of a PLC to expose four one-bit values from the security control of a single room to the IT situational awareness system: security alarm armed, security alarm triggered, safety alarm triggered, lights on. The security systems included CCTV cameras that are accessible from the network, either directly or via a recording device. Figure 1 presents the developed prototype including operator of the engineering workstation and building automation system.

The defended system was a Windows OS computer running SCADA software controlling physical RTU and circuit breakers representing a power grid engineering workstation. In the scenario cyber sensing is defending against engineering workstation compromise that would provide the attacker with initial access which is technique T818 according to the MITRE ATT&CK for ICS Matrix (MITRE Corporation, 2020). The circuit breaker is supplying power to the defender’s cyber command and when opened the power is physically cut bringing the defender’s network down. The workstation is feeding all OS security events into the situational awareness system.

Correlation with the traditional network and host-based events is crucial to the system-centric approach. We integrated these event streams into the exercise situational awareness system Frankenstack (Kont et al., 2017). Low priority events like authentication that otherwise would be lost in the stream of data can now trigger high priority alerts when nobody is actually at the workstation. When these high priority alerts are raised, the adversary has an option to respond by rebooting or terminating potential pivot points that were used to access the engineering workstation, thus countering the Red Team.

The four input bits from the building automation allowed us to define five rules relevant to the scenario and the real world. The monitored trigger is a standard Windows logon security event, which is a low priority event thus not causing any alert. The five implemented scenarios when the alert is raised when the logon occurs:

- The security alarm is triggered: attacker inside the armed location is attempting to access the system.
- The security alarm is armed: attacker has either bypassed physical security system or is accessing the workstation remotely.
- The state of the building automation not updated recently: potential compromise or failure of sensor integration.
- Safety alarm triggered evacuation: attacker might be exploiting chaos caused by the evacuation.
- Smart lights are off: attacker might be accessing the workstation remotely.

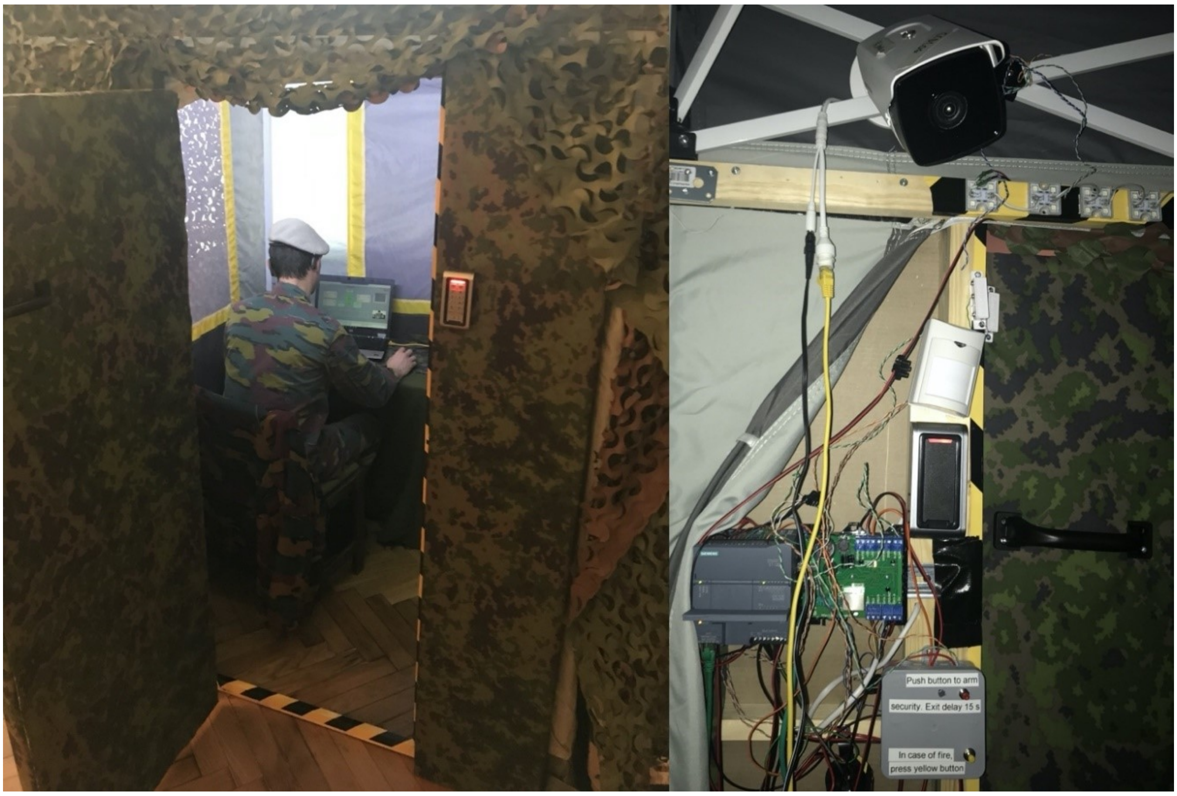
Push button to arm
security. Exit delay 15 s
In case of fire, press yellow button

Figure 1: System-centric cyber sensing proof of concept

Initially, the Red Team has no knowledge of any non-traditional cyber defence systems being used. Through acquiring documentation or triggering the alert and therefore the defender’s response, it becomes clear that the attack has to be specially crafted. Since the defender’s actions are manual, an option is to craft a payload that executes so quickly that it will cut the power off, no matter if it is detected.

A stealthier option is to acquire additional information through compromising the building’s CCTV, allowing the attacker to identify when the operator is in the control room but is not monitoring the displays. In this case, an alert will not be raised and it demonstrates the main weakness of the system-centric approach. It is impossible to get precise information about the operator; only their presence is inferred, but it is not known what they are doing. The operator may be monitoring the displays, taking actions that are directly related to the input of the screens or being distracted and paying attention to their personal phone. The third option is to attack the sensors and the building automation itself to disable the sensor readings, change them to hide the attack or trigger an avalanche of false alerts.

System-centric cyber defence has been demonstrated to be simple, cheap and effective as long as the attackers do not know it is being used. Security rules can be reduced to one or more sensor status bit correlations with one or more host or network-based events using a sliding time window. A set of static rules defining specific defence scenarios can suffice without the need for AI.

4.2 Human-centric cyber sensing

Human-centric cyber sensing focuses on measuring the biological properties of humans directly or through the properties of the environment around them. The volume of generated data is significantly higher than in the system-centric approach as not only might many sensors be necessary, but these sensors will have a wide range of values instead of a single bit and the sampling rate can be high. In this case, creating rules manually becomes challenging and might require AI. If the human-centric approach evolves, it can be expected that there will be a limited number of widespread universal solutions containing sets of different sensors in a single unit. Therefore, we developed one such integrated unit to collect data in various exercises and simulated scenarios.
Environmental values measured change relatively slowly, but biometric sensors and RF spectrum produce a lot of data. Commercial high data rate sensors were deployed to collect data with timestamps (Figure 2).

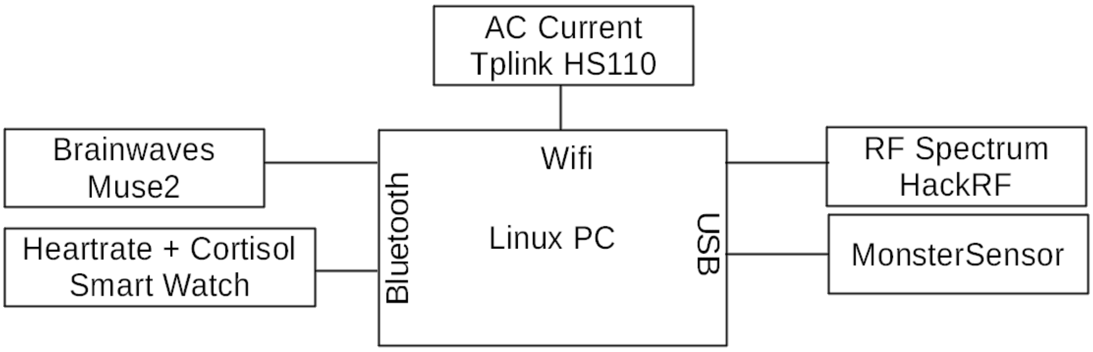

Figure 2: Integration of all human centric sensors

For slow-varying environmental data, we created a multi-sensor logging device called MonsterSensor (Figures 3 and 4). It is based on an Arduino computing unit and low-cost sensor modules which are integrated with an RTC and memory card module for timestamps and logging. The device can either be used autonomously to log values on the memory card or could be connected to a PC for real-time monitoring.

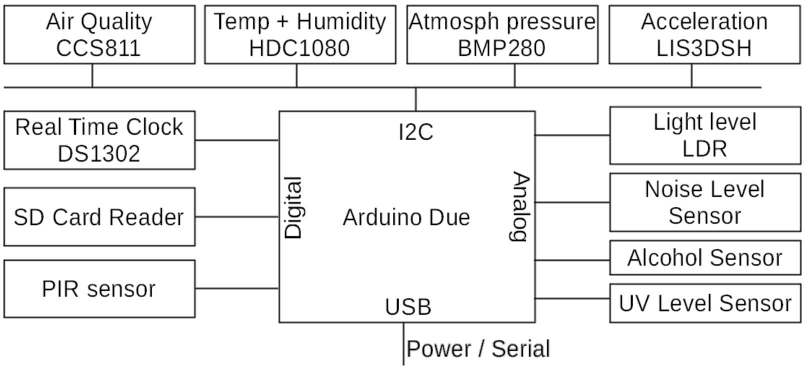

Figure 3: Multi-sensor logger (MonsterSensor) architecture

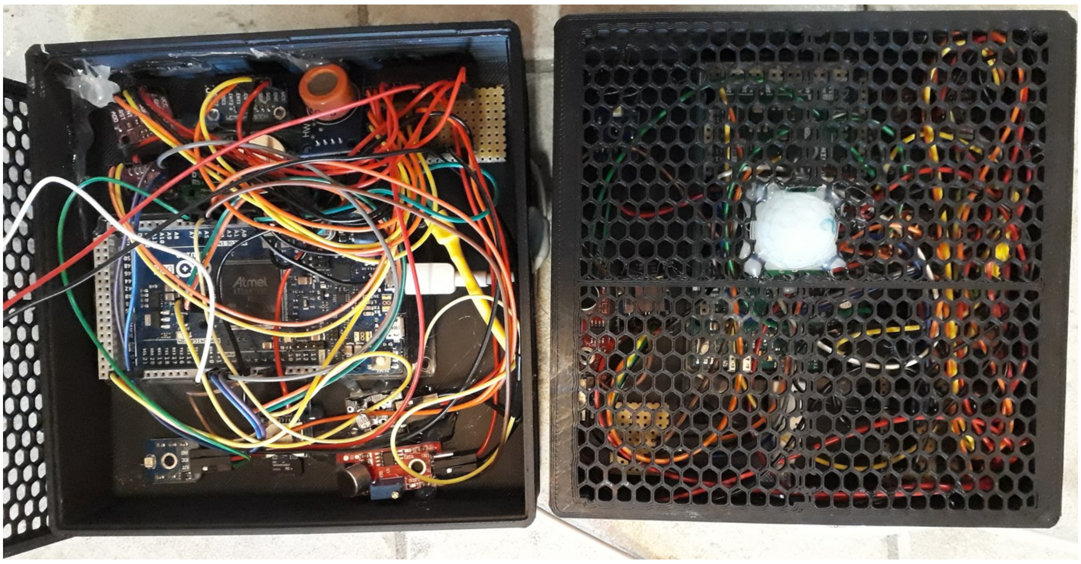

Figure 4: MonsterSensor prototype construction and housing

5. Operational and strategic implications

It is compelling to have out-of-band information, because there is almost always a person involved with cybersecurity, even when a lot of the processes are automated. Having this information adds to operational advantages by identifying and limiting risks or increasing opportunities for offensive cyber actions. The people who will be 'sensed' are mostly employees, but can be end-users, state officials, management or intruders. The monitoring of human behaviour adds to the knowledge we currently have on the effect that individuals have on cybersecurity.

Some infrastructures or smart buildings already have a lot of system-centric sensors installed. Centralising and processing all the data that is already produced by those sensors will provide insights into the correlation between human behaviour and cybersecurity.

Human-centric sensors such as heart rate monitors, accelerometers and pedometers are used widely nowadays. Smartwatches are built to collect, store and process data. However, the information that these devices produce is private and not legally accessible for analysis. Installing human-centric sensors or getting access to data of personally owned devices will have a deep impact on privacy and ethics.

The questions that need an answer, and where cyber sensing can be a breakthrough:

- Is there someone in the room?
- Are they authorised to be there?
- Is the person paying attention to their work?
- Is the person taking action if an event occurs? When the event happens, are there correlations to inputs to the system?
- Are there behavioural differences among people that are correlated with cybersecurity?

One should at least be cautious, because this information could also be used to target individuals rather than optimising a process.

5.1 Defensive perspective

The most straightforward use of this information would be from a defensive perspective. It is clear that looking at sensors that are under one's own control is easy to achieve. Achieving correlations between successful cyber attacks and behaviour of people will increase awareness and lead to improved cybersecurity. This could be achieved by better procedures, leadership, technology and knowledge development.

With enough data and good learning algorithms, it is possible to analyse human behaviour and find weak spots in a system or individual. Since in cybersecurity the attacker only has to find one weak spot and the defender has to protect every part of the system, this knowledge will help in establishing priorities to mitigate the threats.

Redeveloping Section 4.1 power grid scenario in a human-centric sensing requires addressing wide range of open questions:

- What would be the sensors and the setup?
- Where would data go? How it would it be transmitted?
- What would be the simple rules and what would be the AI?
- What happens if adversaries target sensors? What if they target data? What if they target AI?
- What if one targets private data from wearables? What if one feeds data into their own AI and find terminally ill employees even before they know it?

5.2 Offensive perspective

Target monitoring, although not the intended use, can be accomplished if patterns of human activity are known when the attacker is conducting cyber attacks. Using this information it makes it possible to attack the target when human defences are at their lowest.

If sensors are implemented from a defensive perspective and generate data when certain sensor output values are reached, it is beneficial to analyse the attack surface generated by those sensors. If an adversary knows what actions there will be if a certain sensor or group of sensors reaches a certain value, these sensors become targets themselves. This means that using the information from those sensors will require additional security measures.

State actors, private companies and criminals may, for different reasons, be interested in this kind of information. States that infiltrate the 'sensor system' below the threshold of detection and use AI on real-time data will learn a great deal and can eventually also cause a lot of damage by using this information.

5.3 Insider threats

An insider threat, also mentioned in *Modelling Human Behaviour to Anticipate Insider Attacks* (Greitzer and Hohimer 2016) is a very difficult threat to mitigate. Analysing human behaviour with the use of sensors will increase the probability of detecting this behaviour, although it will probably be necessary to combine other sources of information.

6. Conclusions

Host- and network-based events drive traditional monitoring and detection systems. The amount of generated log data is tremendous and manually defined rules are reaching the limits of what is possible to define. To address this, researchers and solution vendors are applying AI. In this research, we explore what could be the next step after the currently used defence systems augmented with AI reach their limits.

The next logical step would be to introduce additional data sources that are relevant in the security context. We have presented our exploration of sensors that could be used as these sources. We have defined two distinct approaches; system-centric and human-centric sensing. We covered each of the approaches by developing a proof of concept using it.

System-centric cyber sensing has already been shown to be feasible today without additional hardware, using sensors that have been installed as part of different systems, primarily building automation. The security rules can still be defined manually as many relevant sensor statuses can be reduced to a binary value. In this approach human presence is inferred, meaning activities of humans cannot be measured directly, which is a major drawback. When a sophisticated adversary is aware that such a defence exists and of its relevant rules, they can craft scenarios by-passing the rules. We have implemented both the defence and by-passing scenario for the simulated power grid.

Human-centric sensing has proved to be a much more complex topic. The possibilities for collecting data are endless. We expect that these multi-sensor units, covering a wide range of inputs, will gain traction. We developed this type of universal multi-sensor prototype and are gathering data in cyber exercises as it is clear that this approach requires AI. Sensors generate too much data to be reduced to simple manual rules handling a few binary values.

Cyber sensing is highly promising and could very likely be the next leap in cyber defence advancement. We predict that cyber sensing will follow the same pattern of development as current IT defence systems. Solutions will start from simple system-centric sensing, and as they develop and reach their limits, it will start to transition into human-centric sensing relying on AI.

References

Bass, T. (1999) Multisensor Data Fusion for Next Generation Distributed Intrusion Detection Systems, Iris National Symposium.

Bhattacharya, S. Nurmi, P. Hammerla, N. and Plötz, T. (2014) Using unlabeled data in a sparse-coding framework for human activity recognition, *Pervasive and Mobile Computing.*

Bonetti, B. (2014) *How to Stress Less: Simple ways to stop worrying and take control of your future*, Capstone.

Candás, J.L. Peláez, V. López, G. Ángel, M. Álvarez, F.E. and Díaz, G. (2014) An automatic data mining method to detect abnormal human behaviour using physical activity measurements*, Pervasive and Mobile Computing.*

Chaaraoui, A.A. Padilla-López, J.R. Ferrández-Pastor F.J. Nieto-Hidalgo M., and Flórez-Revuelta F. (2014) *A vision-based system for intelligent monitoring: Human behaviour analysis and privacy by context*, Sensors, Switzerland.

Chen, C. Cook, D.J. and Crandall, A.S (2013) The user side of sustainability: Modeling behavior and energy usage in the home, *Pervasive and Mobile Computing.*

Fucic, A. Jalali, S. and Pacheco-Torgal F. (2012) *Toxicity of Building Materials*, Woodhead Publishing.

Gelenbe E., Gorbil G. and Wu F. (2012) Emergency Cyber-Physical-Human Systems, *21st International Conference on Computer Communications and Networks*, Munich.

Greitzer, F.L. and Hohimer, R.E. (2011) Modeling Human Behavior to Anticipate Insider Attacks*, Journal of Strategic Security.*

Gutzwiller, R.S Fugate, S. Sawyer, B.D and Hancock P.A. (2015) The human factors of cyber network defense, *Proceedings of the Human Factors and Ergonomics Society.*

Higson, S. (2012) *Biosensors for Medical Applications*, Woodhead Publishing.

Hu, L Chen, Y Wang, S. and Chen, Z. (2014) b- COELM: A fast, lightweight and accurate activity recognition model for mini-wearable devices. *Pervasive and Mobile Computing.*

Kont, M., Pihelgas, M., Maennel, K., Blumbergs, B. and Lepik, T. (2017) Frankenstack: Toward real-time Red Team feedback, IEEE Military Communications Conference (MILCOM), Baltimore, MD, pp. 400–405.

Le, D. N. Van Le, C. Tromp, J. G. and Nguyen, G.N. (2018) *Emerging Technologies for Health and Medicine*, Wiley-Scrivener.

Lee, R.M Assante, M.J. Conway, T. (2016) Analysis of the Cyber Attack on the Ukrainian Power Grid white paper e-isac, https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf, accessed on 16 Jan 2020.

Lin S.H., Mak M.W. and Kung S.Y. (2004) *Biometric Authentication: A Machine Learning Approach*, Prentice Hall.

MITRE Corporation (2020) Engineering Workstation Compromise, https://collaborate.mitre.org/attackics/index.php/Technique/T818, accessed on 16 Jan 2020.

Owens, A. and Efros, A.A. (2018) Audio-visual scene analysis with self-supervised multisensory features. *Lecture Notes in Computer Science* (including subseries *Lecture Notes in Artificial Intelligence* and *Lecture Notes in Bioinformatics*).

Rieping, K. Englebienne, G. and Ben Kröse (2014), Behavior analysis of elderly using topic models, *Pervasive and Mobile Computing.*

Sazonov, E. (2014) *Wearable Sensors*, Academic Press.

Scanaill, C.N. and McGrath, M.J. (2013) *Sensor Technologies: Healthcare, Wellness and Environmental Applications,* Apress.

Seiter, J. Amft, O. Rossi, M. and Tröster, G. (2014) Discovery of activity composites using topic models: An analysis of unsupervised methods, *Pervasive and Mobile Computing.*

Sowe, S. K., Zettsu, K., Simmon, E., de Vaulx, F., and Bojanova, I. (2016). Cyber-Physical Human Systems: Putting People in the Loop. *IT professional*.

Yus, R. Mena, E. Ilarri, S. and Illarramendi A. (2014) Sherlock: Semantic management of location-based services in wireless environments, *Pervasive and Mobile Computing*.