



**QUEEN'S
UNIVERSITY
BELFAST**

Enhancing Security and Privacy of Next-Generation Edge Computing Technologies

Hagan, M., Siddiqui, F., & Sezer, S. (2020). Enhancing Security and Privacy of Next-Generation Edge Computing Technologies. In *IEEE Conference on Privacy, Security and Trust (PST): Proceedings* (International Conference on Privacy, Security and Trust (PST)). IEEE . <https://doi.org/10.1109/PST47121.2019.8949052>

Published in:

IEEE Conference on Privacy, Security and Trust (PST): Proceedings

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

Copyright 2019 IEEE. This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Enhancing Security and Privacy of Next-Generation Edge Computing Technologies

Matthew Hagan, Fahad Siddiqui, Sakir Sezer

The Centre for Secure Information Systems (CSIT), Queen's University Belfast
Belfast, United Kingdom

m.hagan, f.siddiqui, s.sezer@qub.ac.uk

Abstract—The advent of high performance fog and edge computing and high bandwidth connectivity has brought about changes to Internet-of-Things (IoT) service architectures, allowing for greater quantities of high quality information to be extracted from their environments to be processed. However, recently introduced international regulations, along with heightened awareness among consumers, have strengthened requirements to ensure data security, with significant financial and reputational penalties for organisations who fail to protect customers' data. This paper proposes the leveraging of fog and edge computing to facilitate processing of confidential user data, to reduce the quantity and availability of raw confidential data at various levels of the IoT architecture. This ultimately reduces attack surface area, however it also increases efficiency of the architecture by distributing processing amongst nodes and transmitting only processed data. However, such an approach is vulnerable to device level attacks. To approach this issue, a proposed System Security Manager is used to continuously monitor system resources and ensure confidential data is confined only to parts of the device that require it. In event of an attack, critical data can be isolated and the system informed, to prevent data confidentiality breach.

Index Terms—Edge Computing, Cloud to Edge, Edge Security, IoT, GDPR, Data Protection, Active Security, Embedded System, Cyber Resilience, Security Micro-architecture.

I. INTRODUCTION

The emergence of connected devices and services that take advantage of embedded computing and connectivity, commonly known as the *Internet-of-Things* (IoT), has clear potential benefits for society. Both consumers and businesses can take advantage of these technologies to enhance and optimise a wide range of activities, including automotive and transportation, healthcare, building management and critical infrastructure operations, in a wide range of ways. When used appropriately, these devices and services can make use of provided or inferred user and environmental data to improve aspects including safety, performance, convenience, reliability and cost. Corporations can also make use of the provided data to enhance the customer experience, make better informed decisions and even discover new business models and market opportunities through use of *Artificial Intelligence* (AI), *Machine Learning* (ML) and *Data Analytics*. It is estimated that the IoT will proliferate to a trillion devices by 2035 [1].

However, where this sharing of data within the IoT brings benefits and opportunities, it simultaneously presents a risk to privacy and security [2]. Large-scale integration and deployment of intelligent devices and related services which deal

with confidential data or critical infrastructure environments pose serious design, supply chain, privacy, security and safety challenges that must be addressed [3], [4], [5], [6]. Current IoT service architectures are vulnerable to attacks and operational failings that, if exploited, may lead to significant losses of data. This would likely cause businesses to fall foul to various international data control regulations, such as the European Union *General Data Protection Regulation* (GDPR), Japan *Act on the Protection of Personal Information* (APPI) and the *California Consumer Privacy Act* (CCPA) [7], [8], [9]. Furthermore, malicious tampering of data may further interfere with data-driven decision making processes that use AI, with potentially disastrous results.

Additionally, the proliferation of data generating devices has vastly increased connectivity, cloud storage and processing demands. The combined worldwide total of IoT data creation and transmission is expected to rise from 216ZB (Zettabytes) in 2016 to reach 847ZB by 2021 [10]. The worldwide total power consumption of data centres, the underpinning backbone of cloud computing, was reported at 416 terawatt-hours (TWh) in 2016, and is predicted treble by 2025 [11]. Such consumption increases have been deemed unsustainable, thus driving the need for optimal approaches to data storage and processing, including edge-based computing. While remaining an integral part of IoT service architecture's core infrastructure, optimised use of edge computing can decrease cloud-based workloads which deal with excessive quantities of unnecessary personal data, while offering low latency results that consume less network bandwidth.

This paper will detail existing issues surrounding *Cloud to Edge* security, at an infrastructural and architectural point of view. Our proposed System Security Manager approach will then be introduced, that aims to provide system-level segregation of data processing elements within the device, to confine processing and storage of confidential data to secure parts of the device. The issues facing IoT architectures can be summarised as follows:

- Increasing requirements for real-time processing of data.
- Vital decisions taken using AI/ML approaches that require high quality data.
- Consumer awareness surrounding privacy.
- Bandwidth and data consumption costs for cloud & consumers.
- Desire to maintain security of raw data.

- Increase in processing capability of embedded devices.
- Security issues at all architectural levels - desire to keep information footprint as low as possible.

The main contributions of this paper are summarised as follows:

- A proposed shift in IoT services from cloud-centric architectures towards distributed, edge-based processing of data, to facilitate security needs and reduce the quantity of cloud-based transmission and processing of confidential data. Such a shift would improve consumer confidence in next-generation IoT services by reducing the attack surface area from which confidential data may be compromised.
- Defining of active micro-architectural characteristics for securing next-generation edge computing technologies. These characteristics will facilitate runtime monitoring of system's critical resources to detect malicious or anomalous behaviour and initiate active mitigations to ensure safety and security of the edge device and any confidential data it possesses.
- The proposed architectural characteristics provide strong security foundations to cloud-based computing paradigms, ensuring confidentiality and integrity of data produced by the edge device.

II. BACKGROUND

Computing technologies have witnessed significant shifts between centralised and decentralised control, from mainframes to PCs and local networks, to important the more recent centralisation by moving control, data and intelligence of computing systems to the cloud [12]. Due to offering increased flexibility, scalability, reliability, redundancy and computing power, as well as offering greater control to the service provider, the cloud has been predominantly used for significant computing tasks and centralised processing of data, rather than the edge device, which typically are lower powered, and offer less control than a centralised platform.

However, cloud-centric architectures face a number of challenges, particularly as performance, power usage, security and privacy have become increasingly important considerations. Reliance upon third-party vendors for providing key infrastructure components is a significant issue, with a number of high profile attacks having demonstrated major weakness with regard to security and privacy. *Spectre* and *Meltdown*, for example, are two recent highly publicised common processor vulnerabilities that, if exploited, can allow the revealing of memory contents to an attacker [13], [14]. Likewise, use of open source software components allows adversaries direct access to internal code. A vulnerability located within open source software may be of particular value as it may allow exploitation on multiple kinds of systems. The OpenSSL 'Heartbleed' exploit (CVE-2014-0160) [15] and Linux Kernel Copy-On-Write, known as 'Dirty COW', (CVE-2016-5195) [16] are well known and heavily exploited examples. Separately, attacks against communication links and data in transit have been demonstrated, which may cause delay in communication

as well as compromise of privacy, or denial-of-service entirely [17], [18], [19]. Cloud services are further vulnerable to standard operational errors and social engineering attacks that may expose large quantities of data [20]. Some high profile cloud breaches include exposure of unsecured Elasticsearch databases, for example the exposure of 24 million mortgage and credit reports [21].

Figure 1 details a typical *Cloud to Edge* service infrastructure, consisting of the physical cloud infrastructure, such as a data centre and virtualised services, the network infrastructure used for communications, before reaching the local edge devices and their connected sensors, actuators and processors. Alongside are likely security and performance considerations.

The primary focus of many next-generation intelligent technologies and applications is not limited to human interaction. *Machine-to-Machine* (M2M) interaction is set to grow substantially [22], ultimately leading to even more useful data being generated at the edge, rather than in the cloud. The growing requirements for M2M interaction is shifting the role of edge devices from data consumers, by human users, towards data producers, enabling a wide range of processing capabilities including signal processing, data acquisition, pattern recognition, real-time data analytics and edge inference [23]. A primary reason for this shift is the advancement of embedded technologies and availability of diverse computing architectures, such as heterogeneous multi-core System-on-Chip (SoC). These architectures provide adaptability, flexibility, high performance compute and connectivity to realise different intelligent applications [24] meeting power footprints and form factor at the edge device. This in contrast to early edge devices that only collect and transmit data from sensors to the cloud for data analysis purposes. However this approach of generating masses of data from the physical world, at the edge, can stress the capabilities of cloud computing, due to processing, storage, network bandwidth and latency limitations, leading to data aggregation problems and greater costs [25].

Edge computing is a decentralised, distributed computing paradigm in which the computation is largely or completely performed on distributed nodes. Its aim is to bring memory and computing power closer to the source of activity, allowing technology to directly interact with the physical world [26]. However, just like mainframes and PCs, cloud computing maintains a significant role within next-generation edge computing in terms of a centralised point of access and wholistic data analytics. The enhancements of edge computing capabilities will provide the foundation to realise a range of intelligent and smart technologies, with decision making processes driven by AI and ML inference. This will provide the backbone for a new generation of M2M communication delivering shorter response time, lower latency and improved service availability, enabling a whole new range of computing capabilities. Bringing computation closer to the edge not only decreases communication constraints, but also enables applications to avoid disruption in the event of intermittent or limited network connectivity [27]. A further advantage of edge computing is that it facilitates improved handling of

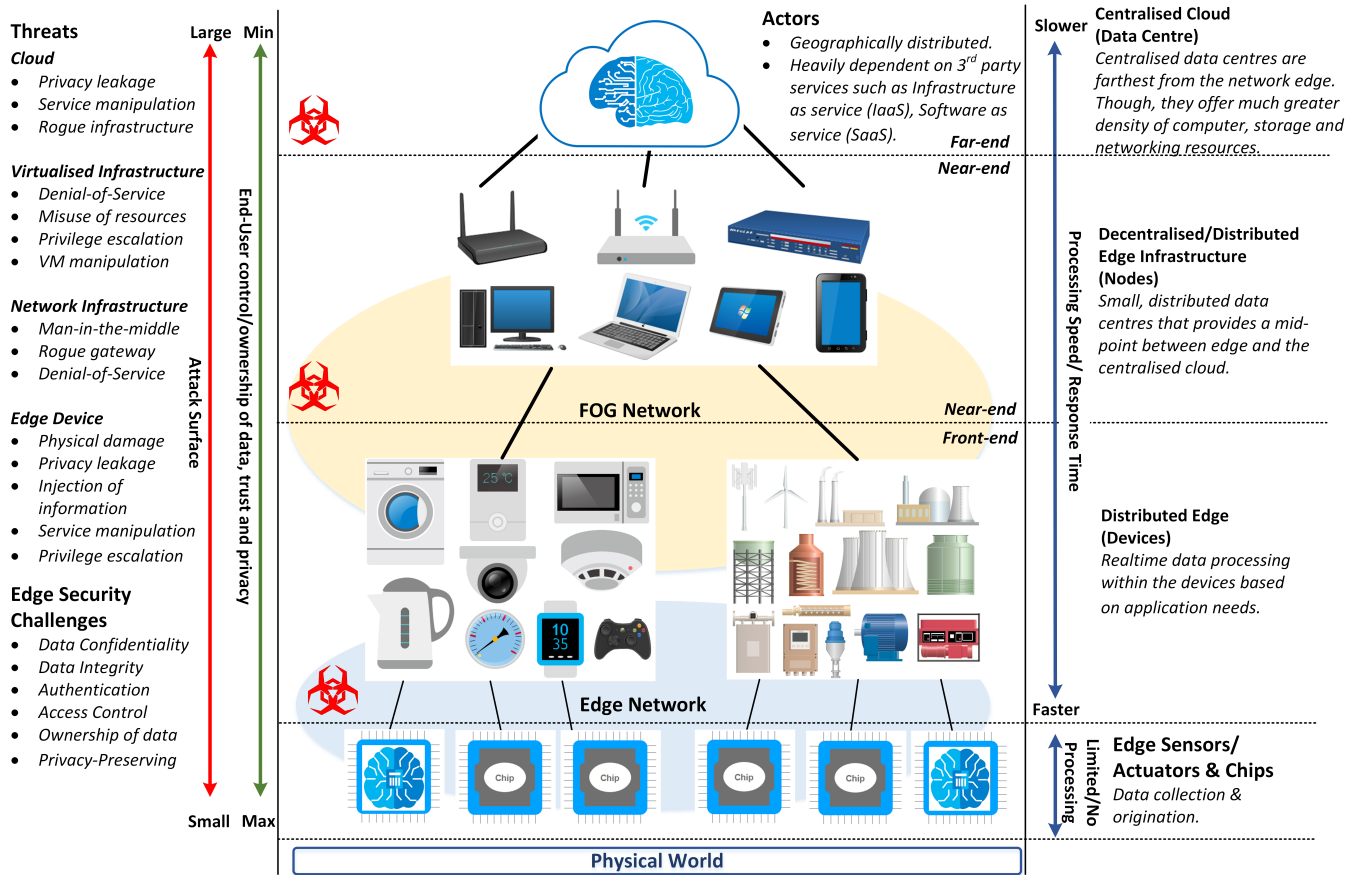


Fig. 1. An illustration of Cloud to Edge infrastructure capabilities, security threats and involved edge security challenges at each layer.

confidential and sensitive data, by processing it at the physical point of generation, at the edge. Subsequently, only processed and anonymised data need be sent to the cloud.

III. THE NEED FOR EMBEDDED RESILIENCE IN NEXT-GENERATION EDGE TECHNOLOGIES

While the localising of data processing greatly reduces the attack surface area from which sensitive data can be attacked, it should be noted these edge devices are prone to many of the same vulnerabilities affecting other aspects of the cloud service architecture, including those mentioned previously. However, edge devices are prone to additional vulnerabilities, due in part to adversaries having full physical access to the device. Side channel analysis, hidden debug ports and boot modes are aspects that may be used to gain additional access to the device [28]. The following are some key issues surrounding embedded edge device security mechanisms:

- A lack of an independent, active run-time security mechanism that can detect threats, malicious activities and protect critical data if existing security mechanisms are compromised to reduce the risk of information exposure or insertion of false data.
- Micro-architectural defence mechanisms offered by security architectures are ad-hoc and passive. They have

been designed to counter specific vulnerabilities or attack only. Open literature has reported examples of where these defence mechanisms have been found vulnerable, attacked and compromised. For example:

- Memory protection* extensions to protect against memory overflow.
- Pointer authentication* to ensure pointer integrity.
- Logical isolation/virtualisation* of resources to avoid side-channel information leakage.
- Chain-of-Trust* security mechanisms to ensure integrity of applications.
- Security architectures rely on building and maintaining a strong *chain-of-trust*. This comprises a series of nested assumptions and as vulnerable as its weakest link. If broken, the security of the complete system is compromised.
- A lack of security standards and protocols which vendors can use to evaluate security of the developed hardware and software components before and after integrating them into the system to ensure secure product development life-cycle.
- A lack of security-aware design and development practices caused by re-using third party hardware and software components, leading to the development of inconsistent and vulnerable solutions.

- Complex hardware-software co-design, security modelling and integration practices, giving rise to vulnerabilities in hardware and software, allowing an adversary to launch attacks.

To approach these mentioned security issues, a new approach is required to ensure security of the underlying data that is handled and processed by the edge device, as well as the service within which it is operating. In order to protect IoT architectures utilising edge processing to handle sensitive data, there is an essential need for an additional layer of defence where critical data is handled, prior to processing. This proposed layer shall complement the existing micro-architectural security mechanisms and provide malicious activity detection and prevention before they become can cause harm or serious damage. This layer can play an integral role in future edge devices handling sensitive data and utilising AI to ensure their trustworthiness within complex M2M environments. This proposed layer will provide an independent, active run-time security mechanism that enable platform-level visibility of the underlying edge device, essential to detect threats and protect the M2M ecosystem.

IV. CHARACTERISTICS OF ADAPTIVE SYSTEM-ON-CHIP PLATFORM

As mentioned, no active methods exist within embedded micro-architectures to establish or maintain the security of a device once its trust is compromised. This may lead to either the exposure of, or allow modification of confidential data, often without leaving any evidence trail, causing damage to the underlying system and its users.

A. Embedded Security Requirements for Next-Generation Edge Technologies

Considering the vulnerabilities of in-built protections within embedded systems, security functionality should not be limited to *protection* only. The device must *detect* malicious cyber activities and attacks, *respond* by deploying active countermeasures and *recover* the system to maintain critical service operations. The following actions are crucial additional security functionalities required to secure embedded edge micro-architectures:

- **Detection** - The facility to independently monitor valuable system resources and discover activity traits that indicate tampering or compromise.
- **Informing** - This allows for independent informing of decision making elements of the architecture of potentially faulty data or exposure of sensitive data.
- **Mitigation** - This involves the embedded micro-architecture taking evasive actions to avoid negative impacts of compromise. This may include deletion of sensitive data, or disabling of the device.
- **Recovery** - In case of critical operational scenarios, the ability to maintain essential functionalities, such as safety, is vitally important. The ability to disable compromised elements of the device, at physical level, allows secure functionality of the remaining components.

B. Architectural Components to Secure Next-Generation Edge Technologies

Considering the derived security requirements of cyber resilient embedded systems, the following are proposed core micro-architectural components that allow establishing of on-going device activities by continuous monitoring of system resources and activities, keeping track of events to achieve system-level visibility:

- 1) An **Independent Active Runtime System Security Manager**, responsible for *protection*, *detection*, *response* and *recovery* security functions while complimenting existing security mechanisms. It shall continuously monitor system resources, use gathered information to detect benign or malicious system behaviour, respond to detected malicious (system or resource-specific) activities by deploying active countermeasures and recover system back to its healthy state. It is crucial that *system security manager* be physically independent and isolated so its memory resources from the general purpose processor. This physical limiting of attack surface will make the system significantly less susceptible to software vulnerabilities and attacks as was in the case of the TEE, which shares the same physical processor and memory resources with the general purpose processor. Effective realisation of this *system security manager* requires resource-level visibility and monitoring of system's critical components which leads to the second characteristic.
- 2) **Active Runtime Resource Monitors** to observe resource specific behaviours to *detect* malicious activities and report them to the *System Security Manager*. These active runtime monitors are essential as embedded architectures become more complex, with diverse functionalities consolidated into single applications, often involving mixing of sensitive data with non-sensitive data and physical actuation. These active runtime monitors shall generate fine-grained resource specific information which would enable the *system security manager* to articulate, analyse and evaluate system-level behaviours and initiate appropriate *mitigation* and *recovery* strategies. In addition, the gathered information would facilitate continuity of data stream, offering essential information to establish evidence of any anomalous activity.
- 3) An **Active Response Manager** is responsible for implementing *mitigation* and *recovery* requirements of a cyber resilient embedded system that are initiated by the *System Security Manager*. This involves initiating active countermeasures to curtail the detected threat within the system. Moreover, depending on the micro-architecture of the *active runtime resource monitors*, the active response manager can enforce various system-level security strategies, where a compromised resource can be physically isolated from the system. This would allow opportunities to gracefully degrade the system functionality while maintaining essential critical services in next-generation critical infrastructure.

A detailed SoC platform architecture [29], [30] and security modelling approach [31] that realises the proposed characteristics and embedded security requirements has been defined and implemented.

V. CONCLUSION

Significant improvements in edge computing performance have brought about possibilities for complex processing to be performed at the source of data collection, instead of the cloud where it is typically performed. Alongside, issues surrounding data protection legislation alongside the real-time performance and resource consumption challenges of cloud computing have further boosted possibilities for distributing processing capabilities to the edge devices. However, such a process would not be without issues, particularly in terms of security of critical data or processes that rely on receiving accurate information. This paper has presented some of the security challenges and requirements, in light of international data protection regulations. Embedded security requirements have been derived from these challenges to improve the resilience of M2M systems. The paper establishes a strong need for embedded cyber resilience, due to an existing lack of active detection, response and recovery security functionalities within existing embedded security systems. This is realised through the proposing of runtime monitoring and system-level visibility of resources activities, with active response functions to enhance, maintain and ensure secure operation of intelligent technologies during the life cycle of the device.

REFERENCES

- [1] P. Spark, "White Paper: The route to a trillion devices: The outlook for IoT investment to 2035," ARM, Tech. Rep., 2017. [Online]. Available: <https://community.arm.com/iot/b/blog/posts/white-paper-the-route-to-a-trillion-devices>
- [2] A. Ukil, J. Sen, and S. Koilakonda, "Embedded security for Internet of Things," in *Proc. IEEE National Conference on Emerging Trends and Applications in Computer Science (NCETACS)*, March 2011, pp. 1–6.
- [3] V. Sharma *et al.*, "Security, Privacy and Trust for Smart Mobile-Internet of Things (M-IoT): A survey," *CoRR*, 2019. [Online]. Available: <http://arxiv.org/abs/1903.05362>
- [4] S. Ravi *et al.*, "Security in Embedded Systems: Design Challenges," *ACM Trans. Embed. Comput. Syst.*, vol. 3, no. 3, pp. 461–491, Aug. 2004.
- [5] N. Aphorpe *et al.*, "Spying on the Smart Home: Privacy Attacks and Defenses on Encrypted IoT Traffic," *CoRR*, vol. abs/1708.05044, 2017. [Online]. Available: <http://arxiv.org/abs/1708.05044>
- [6] D. N. Serpanos and A. G. Voyiatzis, "Security Challenges in Embedded Systems," *ACM Trans. Embed. Comput. Syst.*, vol. 12, no. 1s, pp. 66:1–66:10, Mar. 2013.
- [7] Council of European Union, "Council regulation (EU) no 2016/679," 2016. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [8] Personal Information Protection Commission, Japan, "Amended act on the protection of personal information," 2016. [Online]. Available: https://www.ppc.go.jp/files/pdf/Act_on_the_Protection_of_Personal_Information.pdf
- [9] California Office of Legislative Counsel, "Assembly bill no. 375: 'the california consumer privacy act of 2018'," 2018. [Online]. Available: https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=2017_20180AB375
- [10] Cisco, "Cisco Global Cloud Index: Forecast and Methodology, 2016–2021 White Paper," Tech. Rep., 2016. [Online]. Available: <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html>
- [11] Tom Bawden, "Global warming: Data centres to consume three times as much energy in next decade, experts warn," Tech. Rep., 2016. [Online]. Available: <https://www.independent.co.uk/environment/global-warming-data-centres-to-consume-three-times-as-much-energy-in-next-decade-experts-warn-a6830086.html>
- [12] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, "Internet of Things (IoT): A Vision, Architectural Elements, and Future Directions," *Future Gener. Comput. Syst.*, vol. 29, no. 7, pp. 1645–1660, Sep. 2013.
- [13] P. Kocher *et al.*, "Spectre Attacks: Exploiting Speculative Execution," *CoRR*, vol. abs/1801.01203, 2018. [Online]. Available: <http://arxiv.org/abs/1801.01203>
- [14] M. Lipp *et al.*, "Meltdown: Reading Kernel Memory from User Space," in *27th USENIX Security Symposium, USENIX Security*, Aug. 2018, pp. 973–990. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity18/presentation/lipp>
- [15] I. Ghafoor, I. Jattala, S. Durrani, and C. M. Tahir, "Analysis of OpenSSL Heartbleed vulnerability for embedded systems," in *Proc. IEEE International Multi Topic Conference 2014*, Dec. 2014, pp. 314–319.
- [16] A. P. Saleel, M. Nazeer, and B. D. Beheshti, "Linux kernel os local root exploit," in *2017 IEEE Long Island Systems, Applications and Technology Conference (LISAT)*, May 2017, pp. 1–5.
- [17] J. Qiu, L. Gao, S. Ranjan, and A. Nucci, "Detecting bogus BGP route information: Going beyond prefix hijacking," in *Proc. IEEE International Conference on Security and Privacy in Communications Networks and the Workshops*, Sep. 2007, pp. 381–390.
- [18] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking Bitcoin: Routing Attacks on Cryptocurrencies," in *Proc. IEEE Symposium on Security and Privacy (SP)*, May 2017, pp. 375–392.
- [19] N. J. AlFardan and K. G. Paterson, "Lucky Thirteen: Breaking the TLS and DTLS Record Protocols," in *Proc. IEEE Symposium on Security and Privacy, SP*, May 2013, pp. 526–540.
- [20] L. H. Newman, "Microsoft Email Hack Shows the Lurking Danger of Customer Support," *Wired*, Tech. Rep., 2019. [Online]. Available: <https://www.wired.com/story/microsoft-email-hack-outlook-hotmail-customer-support/>
- [21] D. Olenick, "24 million credit and mortgage records exposed on Elasticsearch database," SC Magazine, Tech. Rep., 2019. [Online]. Available: <https://www.scmagazine.com/home/security-news/data-breach/24-million-credit-and-mortgage-records-exposed-on-elasticsearch-database/>
- [22] GSMA (Organisation), "Cellular m2m forecasts: Unlocking growth," Tech. Rep., 2015. [Online]. Available: <https://www.gsmaintelligence.com/research/?file=9c1e1fdff645386942758185ceed941>
- [23] A. Al-Fuqaha *et al.*, "Internet of things: A survey on enabling technologies, protocols, and applications," *IEEE Communications Surveys Tutorials*, vol. 17, no. 4, pp. 2347–2376, Fourthquarter 2015.
- [24] W. Wolf, A. A. Jerraya, and G. Martin, "Multiprocessor System-on-Chip (MPSoC) Technology," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 27, no. 10, pp. 1701–1713, Oct. 2008.
- [25] S. Sarkar, S. Chatterjee, and S. Misra, "Assessment of the Suitability of Fog Computing in the Context of Internet of Things," *IEEE Transactions on Cloud Computing*, vol. 6, no. 1, pp. 46–59, Jan. 2018.
- [26] W. Shi *et al.*, "Edge Computing: Vision and Challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, Oct. 2016.
- [27] G. Lewis *et al.*, "Tactical Cloudlets: Moving Cloud Computing to the Edge," in *Proc. IEEE Military Communications Conference*, Oct. 2014, pp. 1440–1446.
- [28] A. Kliarsky, "Detecting Attacks Against The Internet of Things," SANS Institute, Tech. Rep., 2019. [Online]. Available: <https://www.scmagazine.com/home/security-news/data-breach/24-million-credit-and-mortgage-records-exposed-on-elasticsearch-database/>
- [29] F. Siddiqui, M. Hagan, and S. Sezer, "Embedded policing and policy enforcement approach for future secure IoT technologies," in *Living in the Internet of Things: Cybersecurity of the IoT - 2018*, Mar. 2018, pp. 1–10.
- [30] F. Siddiqui, M. Hagan, and S. Sezer, "Pro-Active Policing and Policy Enforcement Architecture for Securing MPSoCs," in *2018 31st IEEE International System-on-Chip Conference (SOCC)*, Sep. 2018, pp. 140–145.
- [31] M. Hagan, F. Siddiqui, and S. Sezer, "Policy-Based Security Modelling and Enforcement Approach for Emerging Embedded Architectures," in *31st IEEE International System-on-Chip Conference (SOCC)*, Sep. 2018, pp. 84–89.