RESEARCH INSTITUTE FOR
SECURE HARDWARE &
EMBEDDED SYSTEMS

RISE ANNUAL REPORT

2019 – 2021

Funded by

in association with
National Cyber
Security Centre

EPSRC
Pioneering research
and skills

CONTENTS

FOREWORD

The last year has been unprecedented in many ways, but it has also demonstrated the continuing threat from hardware attacks and given us reasons to believe that we're making a difference. We have seen new speculative execution attacks, new row hammer variants, and every major CPU vendor has seen a hardware vulnerability of some form.

There is sometimes a tendency to feel like you can trust an FPGA where you can't trust a CPU. However, the FPGA-related vulnerability Starbleed highlighted the complex interplay between hardware, firmware and cryptography.

But even in the year that was 2020, we saw positives for hardware security. Many of the major ISAs introduced new features to tackle long-standing issues, showing that industry is starting to take hardware security seriously.

We've also seen continued progress for the Morello platform, and impressive results from Darpa's FETT bug bounty. It feels like new security paradigms are getting closer and that we're finally starting to fix the legacy. Add to that the continued investment in open-source hardware projects and tools, and you've got a thriving open community of security engineers.

Through initiatives such as RISE, Digital Security by Design, and other Innovate UK programmes the UK Government has demonstrated its commitment to hardware security as a field that is both a concern, and a great opportunity for investment and improvement.

In March 2020, the NCSC published our research priorities for the first time in the form of our Research Problem Book. This set of seven outcome-focused themes set out some of our more significant areas of interest. Through publishing this, we're aiming to inform those looking to conduct cyber security research of where our priorities lie. I'm heartened that all the projects within RISE map to one or more of these themes, and that there are opportunities for hardware security researchers to contribute to nearly all of them.

The field doesn't exist in a vacuum though. As well as RISE, NCSC and EPSRC also run three other Research Institutes: RITICS, VeTTS and RISCS. Last year's EPSRC call for multidisciplinary research across the research institutes shows how important it is to consider the wider environment alongside the hardware. The link between secure hardware and trustworthy software to run on it is clear. However, people use the systems that embody that hardware and software, and we must ensure the overall system is secure and the security burden on the users is minimised. Hardware security can help with that.

Everything I've talked about here reinforces the fact that hardware security and subjects close to it are critical in our long-term goal of making the UK the safest place to live and work online. There are problems to solve, projects to get involved in, and other groups to work with. As the nation and the world continue to reopen there will hopefully be more opportunities for international collaboration as well.

It's an exciting time for hardware security. I can't wait to see where you can take it next.

Dr Ian Levy, Technical Director
National Cyber Security Centre

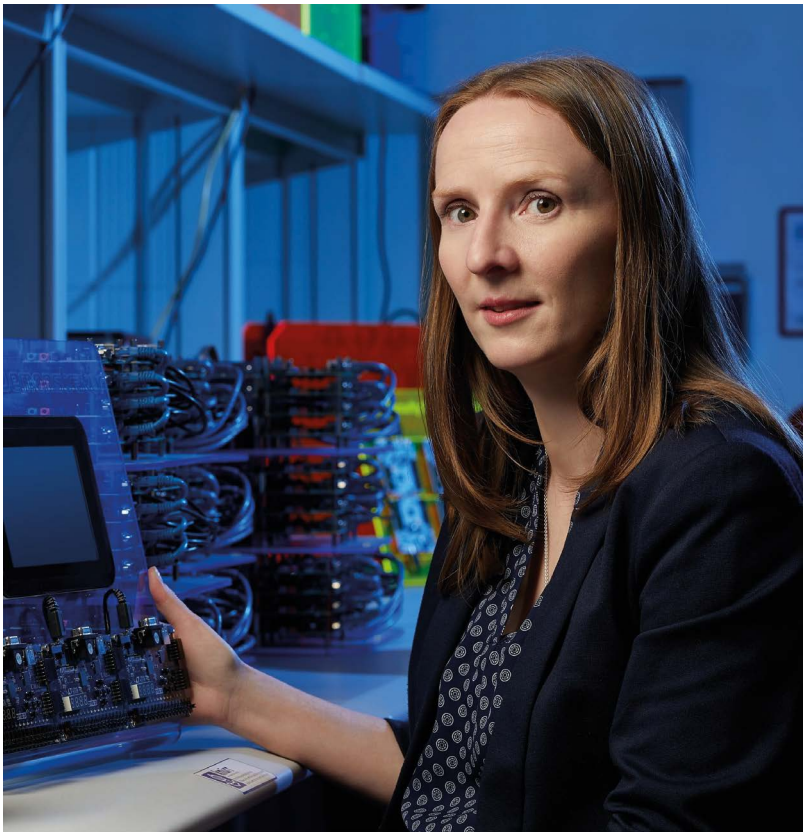

DIRECTOR'S MESSAGE

It has been two years since our last RISE report was published in 2019. Over this time we have made excellent progress across our funded research projects, we kicked-off an international collaboration between the core RISE partners and NTU and NUS in Singapore, and launched a UK competition targeting final year UG/MSc students, sponsored by ARM, to help stimulate the next generation of UK hardware security experts. However, given a key focus of RISE is to grow the UK hardware and embedded systems community by bringing academia and industry together through networking events, the pandemic has made this particularly challenging to deliver effectively.

This report summarises RISE's events and activities since 2019, including our annual conference, and the achievements of the eight RISE-funded research projects. Significant research outputs to date include:

- Plundervolt – an attack developed as part of the University of Birmingham funded project which exploited vulnerabilities with Intel's Software Guard Extensions, leading to errors that could leak secret information such as encryption keys.
- Thunderclap – research by the University of Cambridge team that identified vulnerabilities with USB and Thunderbolt interface standards and which provided security recommendations for hardening systems that were incorporated into the USB 4 standard.
- The Apple Pay vulnerability discovered by the University of Surrey's RISE project which showed that Apple Pay in Express Transit mode if used with a Visa card could be abused to make an Apple Pay payment to any shop terminal, of any value, without the need for user authentication

In 2020 we kicked-off a collaborative project, Secure IoT Processor Platform with Remote Attestation (SIPP), which was funded under the EPSRC International Centre-to-centre call. The SIPP project brings together the core RISE partners, namely Queen's University Belfast and the Universities of Cambridge, Bristol and Birmingham, with leading academics in the field of hardware security and security architecture design from the National University of Singapore and Nanyang Technological University, Singapore, to develop a novel secure IoT processor platform with remote attestation implemented on a RISC-V architecture.

This year we also published a call for Proof of Concept projects, seeking to support the pre-commercialisation of leading-edge technologies arising from RISE-funded projects. The funding is expected to be used to develop an idea through to a stage where a route to commercialisation is clear, either as a spin out, or via licensing or open-sourcing. Two projects were successful, FPGADefender4Clouds – an FPGA Virus Scanner for FPGA Cloud Environments, from Professor Dirk Koc at the University of Manchester and GUPT: A Hardware-Assisted Secure and Private Data Analytics Service, from Professor Markulf Kohlweiss and Dr Michia Honda at the University of Edinburgh.

The semiconductor supply chain has suffered severe shortages over the past two years due to material shortages, the Covid-19 pandemic, natural disasters and other major disruptions. This has led to major supply chain issues in a range of sectors, and in particular in the automotive industry. A small number of foundries in Korea, Japan, Taiwan and China currently dominate the global semiconductor fabrication industry and these nations have plans for further significant investment in this sector to retain their dominance. In addition to this, with the recent high profile attacks against critical national infrastructure, it is not surprising that both the sovereignty and cyber security of the semiconductor supply chain have become significant concerns for many countries.

In August, US President Biden's supply chain review recommended strengthening the US semiconductor manufacturing ecosystem and significant investment in securing the supply chains for critical industries and ensuring the safety and cyber security of products produced within the US. The EC announced the development of a European Chips Act in September with the aim of addressing semiconductor sovereignty concerns in Europe.

It is imperative that the UK also considers the sovereignty and cyber security of its semiconductor supply chain amidst the current geo-political landscape and it is great to see a review of the sector currently underway.

Through the ISCF Digital Security by Design (DSbD) programme, the UK has been investing in projects to help UK digital infrastructure become more secure. The programme has funded ARM to develop a technology platform prototype, the Morello board, which includes a new CPU architecture design that will make processors more resilient to future attacks. DSbD is also funding a range of industry and academic projects that can leverage the Morello board's capabilities.

RISE will continue to play its part in conducting research that addresses security throughout a device's lifecycle, from the initial design and manufacture through to its operational environment. We will also continue to grow the UK's skills in hardware and embedded systems security through our spring school activities and UG/Masters student competitions.

Professor Máire O'Neill, RISE Director
Queen's University Belfast

THE RESEARCH INSTITUTE IN SECURE HARDWARE AND EMBEDDED SYSTEMS

The £5M Research Institute for Secure Hardware and Embedded Systems (RISE), which is hosted at the Centre for Secure Information Technologies (CSIT), Queen's University Belfast, seeks to identify and address key issues that underpin our understanding of Hardware Security. Funded by the Engineering and Physical Sciences Research Council (EPSRC) and the National Cyber Security Centre (NCSC), RISE is one of four cyber security institutes in the UK and aims to be a global hub for research and innovation in hardware security.

RISE aims to address the following research challenges in Hardware Security:

1. **Understanding the technologies that underpin hardware security, the vulnerabilities in these technologies and development of countermeasures.**
    - State-of-the-art Hardware Security primitives: True Random Number Generators (TRNGs), Physical Unclonable Functions (PUFs).
    - Novel Hardware analysis toolsets and techniques.
    - Attack-resilient Hardware platforms, Hardware IP building blocks.

2. **Maintain confidence in security throughout the development process and product lifecycle.**
    - Confidence in Developing Secure Hardware devices.
    - Supply Chain Confidence.
    - Modelling of Hardware Security.

3. **Hardware security use cases and consideration of value propositions.**
    - Novel Authentication, e.g. alternatives to passwords.
    - Secure document viewers.
    - Securing BYOD – attestation, roots of trust.

4. **Development and pull through.**
    - Ease of Development and ease of leveraging best security options.
    - Understanding Barriers to Adoption.
    - Education of Potential User/Developer base.

RISE PROJECTS

The research challenges of RISE are being delivered through a series of projects funded by EPSRC and NCSC. Four projects were funded during the original proposal phase, commencing Nov 2017, and are led by the forming RISE research partners from Queen's University Belfast, the University of Cambridge, the University of Bristol and the University of Birmingham.

SCARV: A Side-Channel Hardened RISC-V Platform.
University of Bristol, Dr Daniel Page.

IOSEC: Protection and Memory Safety for Input/Output Security.
University of Cambridge, Dr Robert Watson, Prof Simon Moore, Dr Athanasios Markettos.

User-Controlled Hardware Security Anchors: Evaluation and Designs.
University of Birmingham, Prof Mark Ryan, Dr Flavio Garcia and Dr David Oswald.

Deep Security: Investigating the Application of Deep Learning in SCA and Hardware Trojan Detection, with the ultimate goal of utilising deep learning. Queen's University Belfast, Prof Máire O'Neill.

A subsequent tranche of 4 projects was initiated in Nov 2018, delivered by the University of Cambridge, the University of Edinburgh, the University of Surrey, and the University of Manchester.

SafeBet: Memory capabilities to enable safe, aggressive speculation in processors. University of Cambridge, Prof Simon Moore.

GUPT: A Hardware-Assisted Secure and Private Data Analytics Service.
University of Edinburgh, Dr Pramod Bhatotia and Dr Markulf Kohlweiss.

TimeTrust: Robust Timing via Hardware Roots of Trust and Non-standard Hardware, with Application to EMV Contactless Payments.
University of Surrey, Dr Ioana Boreanu, Dr Tom Chothia, Prof Liqun Chen.

rFAS: Reconfigurable FPGA Accelerator Sandboxing.
University of Manchester, Dr Dirk Koch.

THE RISE INSTITUTE MODEL

Fulfilling the aims of a global centre for research and innovation in hardware security requires not only world-class research, but also close engagement with leading UK-based industry partners and stakeholders. This additional focus facilitates the accelerated translation of research output into new products, services and business opportunities for the wider benefit of the UK economy.

The key elements within RISE are the academic researchers, an Industry & Stakeholder Advisory Board (ISAB) and the Institute Management team.

The RISE ISAB is chaired by Charles Brookson OBE, and has been created to allow member companies and stakeholders to engage with the research community and to inform funding calls around their real world challenges.

Other functions include:

- Receiving briefings on significant research outputs.
- Identification of research results, which are particularly appropriate for rapid commercialisation.
- Offer pathways to impact, e.g. licensing, spin-out support
- Highlighting shifts in technology or market demand with significance for RISE.
- Informing future RISE research proposal calls
- Helping to build a hardware security community in the UK

The Institute Management team, incorporating leadership and business development, functions to drive forward the development and promotion of the institute to industry and other stakeholders.

You can find out more about RISE and its activities by visiting **www.ukrise.org**

RISE ECOSYSTEM

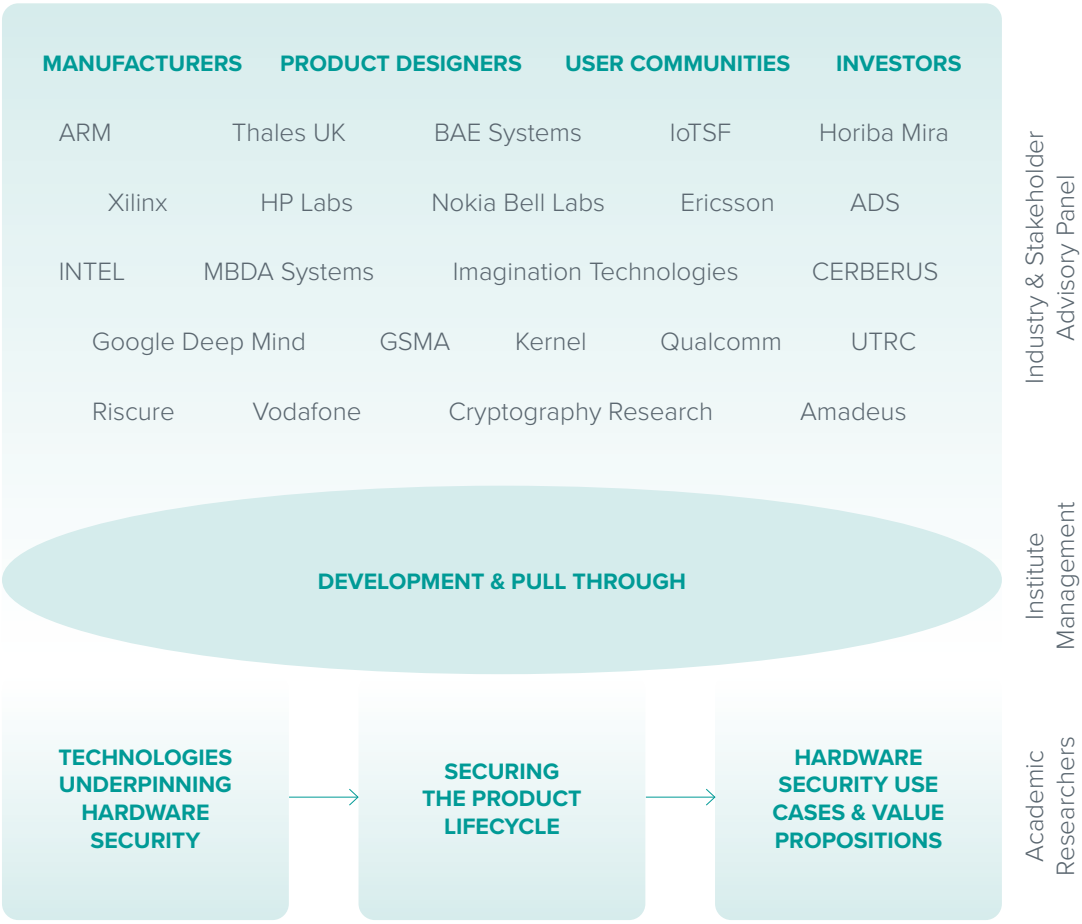

RISE 2019 SPRING SCHOOL

RISE held its 2nd Spring School at the Centre for Secure Information Technologies (CSIT), Queen's University Belfast, on 28 Feb – 1 Mar 2019. Our Spring Schools aim to bring together the hardware security community, from both academia and industry.

Videos of the talks are available at www.ukrise.org/springschool2019

Day One Agenda, Thursday 28th February 2019

09.30 – 10.30	**Registration and Tea/Coffee on Arrival**
10:30 – 10:40	**Opening Remarks**
	Session 1: Importance of Hardware Security
10.40 – 11.20	Martin Dixon, Intel *Opportunities in Hardware Security Research*
11.20 – 12.00	Joe Fitzpatrick, SecuringHardware Millions for defence, not one cent for security
12.00 – 12.40	*NCSC view of Hardware Security Research*
12.40 – 13.50	**Lunch and Networking**
	Session 2: Developing Secure Hardware Devices I
13.50 – 14.30	Ingrid Verbauwhede, KU Leuven *Design methods for hardware roots of trust*
14.30 – 15.10	Samuel Pagliarini, Carnegie Mellon University *Can we build a Trustworthy Billion Transistor Chip?*
15.10 – 15.40	**Break and Networking**
	Session 3: Hardware Security Evaluation
15.40 – 16.20	Emanuel Prouff, ANSSI *Deep Learning for Embedded Security Evaluation*
16.20 – 17.00	Sylvain Guilley, Secure-IC and Télécom-ParisTech *Detection of cache-timing attacks on cryptographic libraries, including post-quantum cryptography*
19.00	**Dinner**

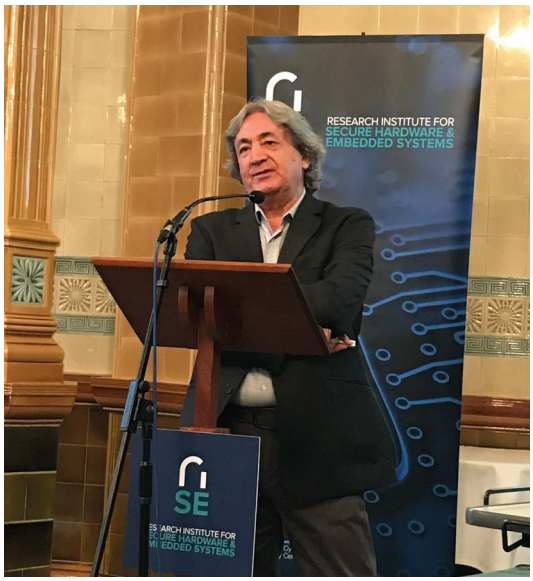

Day Two Agenda, Friday 1st March 2019

08.30 – 10.00	**Registration**
09.00 – 10.00	**Tutorial** Ilhan Gurel, Expert Hardware and Software Security, Ericsson *End to End IoT security*
	Session 4: Development and Pull-Through of HW Security Technologies
10.10 – 10.50	Shahram Mossayebi, Crypto Quantique *Securing connected devices: Story of a deep-tech cybersecurity startup in the UK*
10.50 – 11.30	Jayne Brady, Kernel Capital *Challenges with Commercialisation of Early Stage Deep Tech*
11.30 – 12.00	**Break and Networking**
	Session 5: Advanced Crypto Primitives on Hardware
12.00 – 12.40	Dimitrios Schoinianakis, Nokia Bell Labs *Challenges of Homomorphic Encryption*
12.40 – 13.20	Ayesha Khalid, Queens University Belfast *Physical protection of lattice-Based cryptography – Challenges and solutions*
13.20 – 14.30	**Lunch and Networking**
	Session 6: Developing Secure Hardware Devices II
14.30 – 15.10	David Oswald, University of Birmingham *Trusted Execution in Practice – A Gentle Introduction*
15.10 – 15.50	Simon Moore, University of Cambridge *Thunderclap: Exploring Vulnerabilities in Operating System IOMMU Protection via DMA from Untrustworthy Peripherals*
15.50 – 16.30	Dirk Koch, University of Manchester *FPGA acceleration a boon or bane?*
	Closing remarks

RISE 2019 ANNUAL CONFERENCE

The RISE 2019 annual conference took place at the National Liberal Club in London, on 21st November 2019. The full day plenary program included a keynote from Professor Cetin Koç (University of California, Santa Barbara), lightening talks from early-stage researchers and updates from each of the RISE projects. The conference was well attended, with over 80 participants and concluded with a closed session Industry Stakeholder and Advisory Board (ISAB) meeting. This session proving a useful opportunity for the exchange of views between industry, academic, and government representatives.

Agenda, Thursday 21st November 2019

09.30 – 10.00	**Registration and Tea/Coffee on Arrival**
10.00 – 10.05	**Welcome** Prof. Máire O'Neill, Director, RISE, Queen's University Belfast
10.05 – 10.50	**Keynote** Professor Cetin Koç, University of California Santa Barbara
10.50 – 11.50 RISE Project Updates (15mins per talk to include Q&A)	Simon Moore, Robert Watson, Theo Markettos, University of Cambridge *IOSEC: Protection and Memory Safety for Input/Output Security* Dan Page, University of Bristol *SCARV: A Side-Channel Hardened RISC-V Platform* Mark Ryan, Flavio Garcia, David Oswald, University of Birmingham *User-Controlled Hardware Security Anchors: Evaluation And Designs* Máire O'Neill, Queen's University Belfast *Deep Security: Applying Deep Learning To Hardware Security*
11.50 – 12.15 Lightning Talks – Early Career Researchers and SMEs (4 x 5min talks)	Mahshid Delavar, University of Edinburgh *Quantum Physical Unclonable Functions* Neil Hanley, Queens University Belfast *What Does Security Mean for Approximate Computing* Vashti Galpin, University of Edinburgh *Micro-architecture simulation for verified security and performance* Henry Harrison, CTO at Garrison *Hardsec: using non-Turing-machine logic in FPGAs for security controls*
12.15 – 12.30	John Goodacre, Industrial Strategy Challenge Fund (ISCF) Director *Update on ISCF Challenge on Digital Security by Design*
12.30 – 13.45	**Lunch and Networking**
13.45 – 14.25 RISE Project Updates (10mins per talk)	Dirk Koch, University of Manchester *rFAS - reconfigurable FPGA Accelerator Sandboxing* Simon Moore, University of Cambridge *SafeBet | Memory capabilities to enable safe, aggressive speculation in processors* Pramod Bhatotia, University of Edinburgh *GUPT: A Hardware-Assisted Secure and Private Data Analytics Service* Ioana Boureanu, Tom Chothia, Liqun Chen, University of Surrey *TimeTrust: Robust Timing via Hardware Roots of Trust and Non-standard Hardware – with Application to EMV Contactless Payment*
14.25 – 14.30	**Close**
15.00 – 16.00	**RISE ISAB Meeting**

RISE 2020 ANNUAL CONFERENCE

The RISE 2020 annual conference was held as a virtual event on 30th November 2020. We started in the morning with an invited tutorial session by Ilhan Gurel of Erricson, on the subject of Confidential Computing. The afternoon session included a keynote by Prof. Patrick Schaumont of Worcester Polytechnic Institute on EDA tools for security testing and countermeasure synthesis. Alongside these talks were updates from the RISE projects and updates on entry to academic/industry competition by ISCF on Digital Security by Design (DSbD) Software Ecosystems.

Agenda, Monday 30th November 2020

	Online Tutorial
10.00 – 12.00	**Confidential Computing – Ilhan Gurel, Ericsson** Confidential Computing is about preserving code and data confidentiality "in use" by means of HW enforced isolation and memory encryption. This tutorial will cover: • What does Confidential Computing mean? • The need for Confidential Computing • Confidential Computing technologies • Use cases (crypto and key management, attestation, privacy preserving AI/ML, secure bootstrapping, secure containers, secure communication) • Research activities • Relevant open source projects • Hands on demo
	Main Programme
14.00 – 14.05	**Welcome** Prof. Máire O'Neill, Director, RISE, Queen's University Belfast
14.05 – 14.45	**Keynote** Prof. Patrick Schaumont, Worcester Polytechnic Institute, US EDA tools for security testing and countermeasure synthesis
14.45 – 15.25 (5 min talks and Q&A)	**RISE Project Updates** University of Cambridge *IOSEC: Protection and Memory Safety for Input/Output Security* University of Bristol *SCARV: A Side-Channel Hardened RISC-V Platform* University of Birmingham *User-Controlled Hardware Security Anchors: Evaluation and Designs* Queen's University Belfast *Deep Security: Applying Deep Learning to Hardware Security*
15.25 – 15.35	**Refreshment Break**
15.35 – 15.40	Georgios Papadakis, UKRI Innovate UK *ISCF Digital Security by Design (DSbD) Software Ecosystem Competition*
14.25 – 14.30	**RISE Project Updates** University of Manchester *rFAS – reconfigurable FPGA Accelerator Sandboxing* University of Cambridge *SafeBet | Memory capabilities to enable safe, aggressive speculation in processors* University of Edinburgh *GUPT: A Hardware-Assisted Secure and Private Data Analytics Service* University of Surrey *TimeTrust: Robust Timing via Hardware Roots of Trust and Non-standard Hardware – with Application to EMV Contactless Payment*
16.20	**Close**

Videos of the tutorial session, keynote and project updates are available for view on the RISE website at https://www.ukrise.org/rise-conference-videos

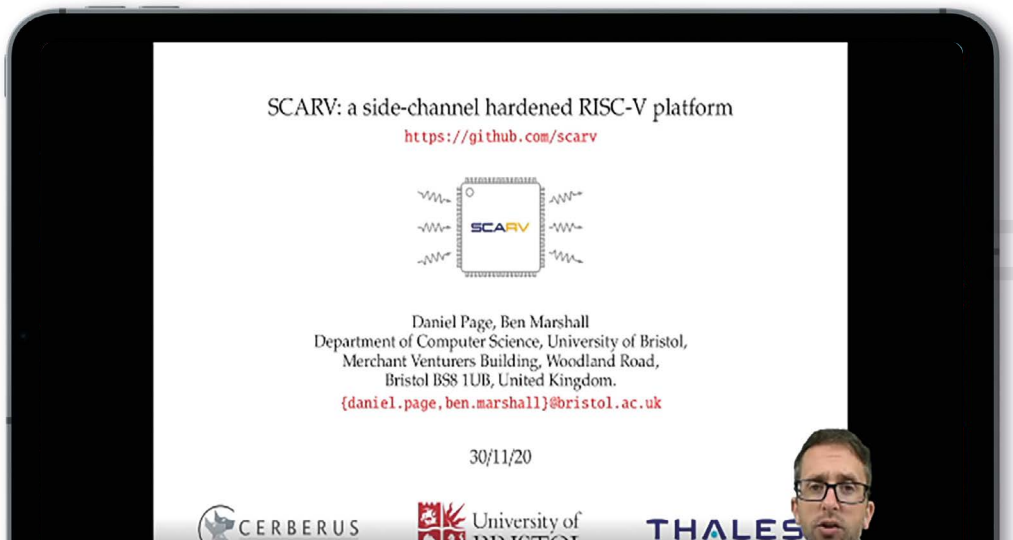

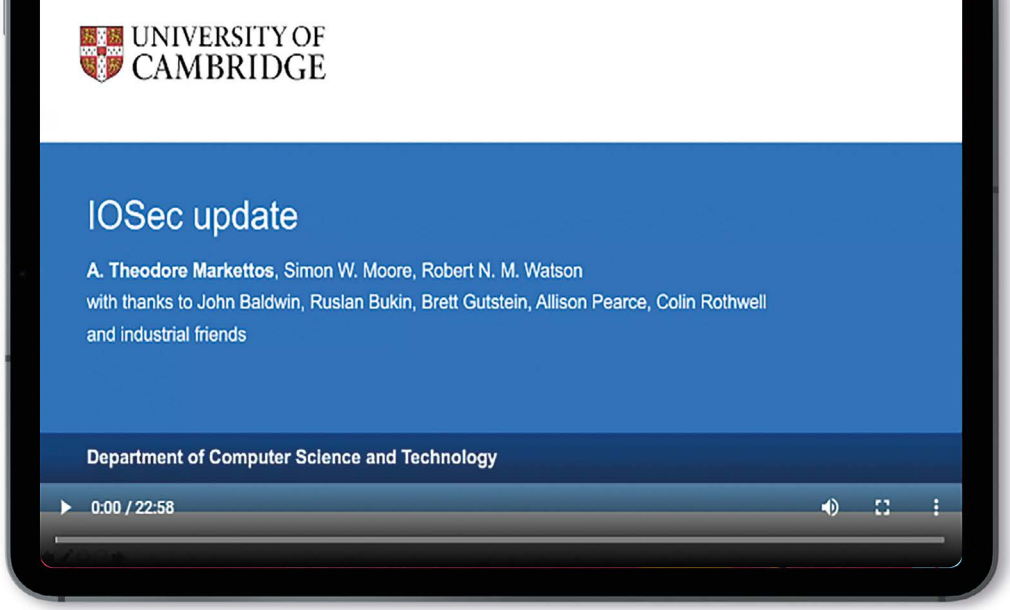
UNIVERSITY OF CAMBRIDGE
IOSec update
A. Theodore Markettos, Simon W. Moore, Robert N. M. Watson
with thanks to John Baldwin, Ruslan Bukin, Brett Gutstein, Allison Pearce, Colin Rothwell
and industrial friends
Department of Computer Science and Technology
0:00 / 22:58


RISE Annual Conference
SE
RESEARCH INSTITUTE FOR SECURE HARDWARE & EMBEDDED SYSTEMS
Tutorial Session: Confidential Computing
Confidential Computing is about preserving code and data confidentiality "in-use" by means of HW enforced isolation and memory encryption. This tutorial will cover:
What does Confidential Computing mean?
The need for Confidential Computing
Confidential Computing technologies
Use cases (crypto and key management, attestation, privacy preserving AI/ML, secure bootstrapping, secure containers, secure communication)
Research activities
Relevant open source projects
Hands on demo
Speaker
Ilhan Gurel
ERICSSON
30th November 2020


Breaking secure enclaves with hardware attacks
Dr.-Ing. David Oswald
User-controlled hardware security anchors: evaluation and designs
University of Birmingham
SE
FutureTPM
INDUSTRIAL STRATEGY
UKRI
UK Research and Innovation
0:00 / 12:05

IMPACT CASE STUDIES

As part of wider RISE dissemination activities, two impact case studies were published.

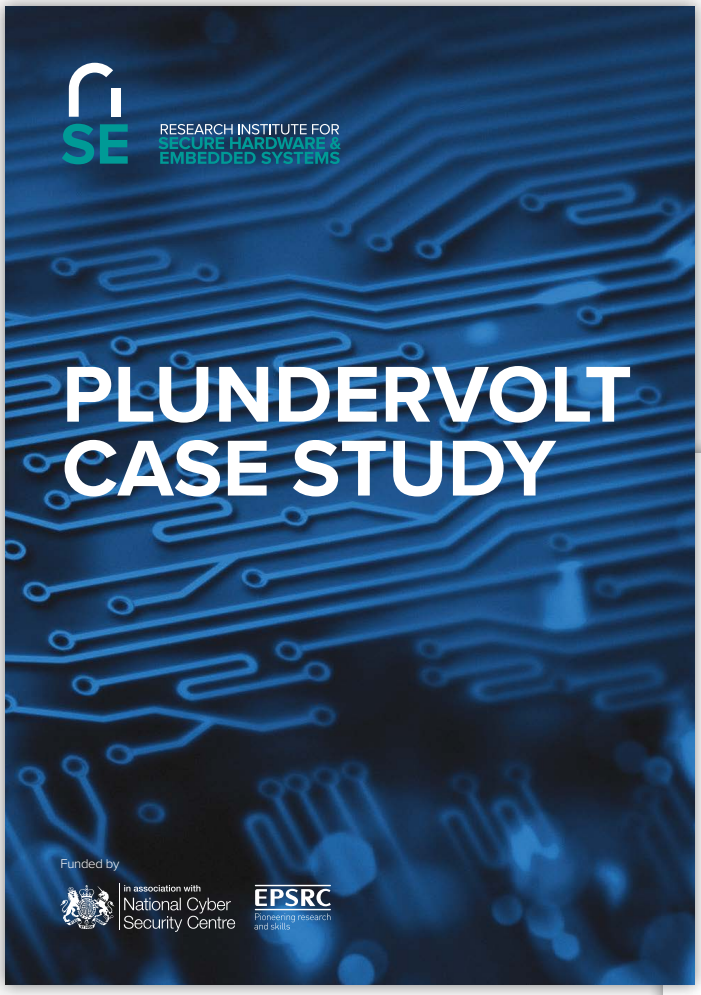

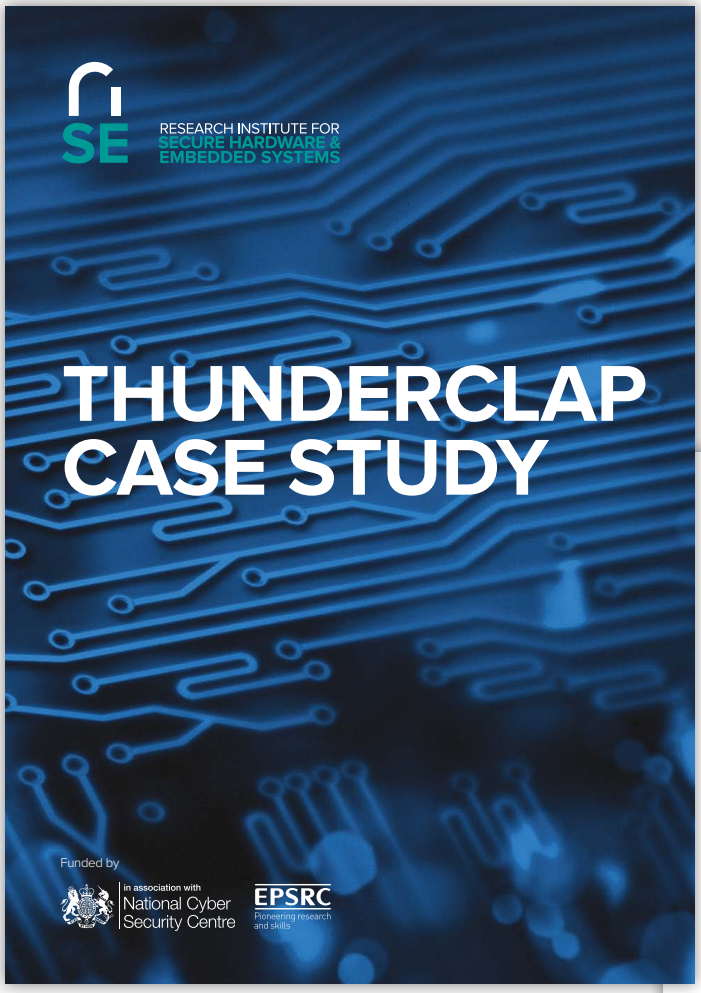

THUNDERCLAP

UK RESEARCHERS IDENTIFY USB SECURITY FLAWS – LEADING TO IMPROVED USB STANDARDS

Security researchers at Cambridge University, members of the UK Research Institute for Secure Hardware & Embedded Systems (RISE), have identified vulnerabilities with USB and Thunderbolt interface standards. The researchers demonstrated that connecting a malicious peripheral device allowed access to secret data and changes to system behaviour.

The vulnerability was found to be applicable to any devices incorporating a Thunderbolt port (Apple laptops and desktops since 2011, some Linux and Windows laptops and desktops since 2016), or devices supporting USB-C, Mini DisplayPort connectors and PCI Express peripherals.

Vulnerability

The attacks exploit Direct Memory Access (DMA) communications, an approach intended to improve the high-speed exchange of data between peripherals and PC memory, without the need for processor involvement.

Although DMA may be secured using an Input-Output Memory Management Unit (IOMMU), it was found that existing IOMMU implementations could be circumvented, or were disabled by default in operating systems to reduce performance overheads.

It is possible to imagine a wide-scale attack making use of compromised devices in the supply chain, such as HDMI/Ethernet dongles, power adapters, or USB-C storage devices. Such attacks would enable remote access of documents, encryption keys, passwords, injection of malicious code etc.

Industry Engagement and Response

Prior to publishing results, researchers contacted affected industry participants to notify them of the vulnerabilities. This led to a series of patches and updates to mitigate against the threats. The hardware platform used to develop the attacks was also shared with industry and is now used in-house for further research and vulnerability analysis.

Since the proposed USB 4 standard incorporates substantial parts of the Thunderbolt specification, including PCI Express DMA transfers, engagement took place with USB 4 standards committee. This led to the incorporation of security recommendations into the upcoming USB 4 standard.

Impact

- Research identified serious weaknesses in existing computer peripheral interfaces.
- Collaboration with industry to disseminate research and understand the threats.
- Industry response with patches and mitigations against the threats.
  - MacOS 10.12.4 – Mitigated specific exploit
  - Windows 10 1803 – Kernel DMA Protections for Thunderbolt 3.
  - Linux 5.0 onwards – Thunderbolt devices assumed less trustworthy by operating system
  - Intel – Published guidance on 'Thunderclap' DMA threats and mitigations.
- Research disseminated via academic papers, website www.thunderclap.io and numerous industry publications.
- Thunderclap FPGA research platform in use by industry for hardening against these vulnerabilities
- Security recommendations for hardening systems incorporated in USB 4 standard.

About RISE

The UK Research Institute in Secure Hardware and Embedded Systems (RISE), funded by the Engineering and Physical Science Research Council (EPSRC) and the National Cyber Security Centre (NSCS) seeks to identify and address key issues that underpin hardware security.

A key focus is to accelerate the industrial uptake of the Institute's research output and its translation into new products, services and business opportunities for the wider benefit of the UK economy.

ukrise.org

RISE PUBLICATIONS

A publications section was added to the RISE website. At the time or writing, the RISE group has published 80+ peer reviewed conference papers, journals, and technical papers.

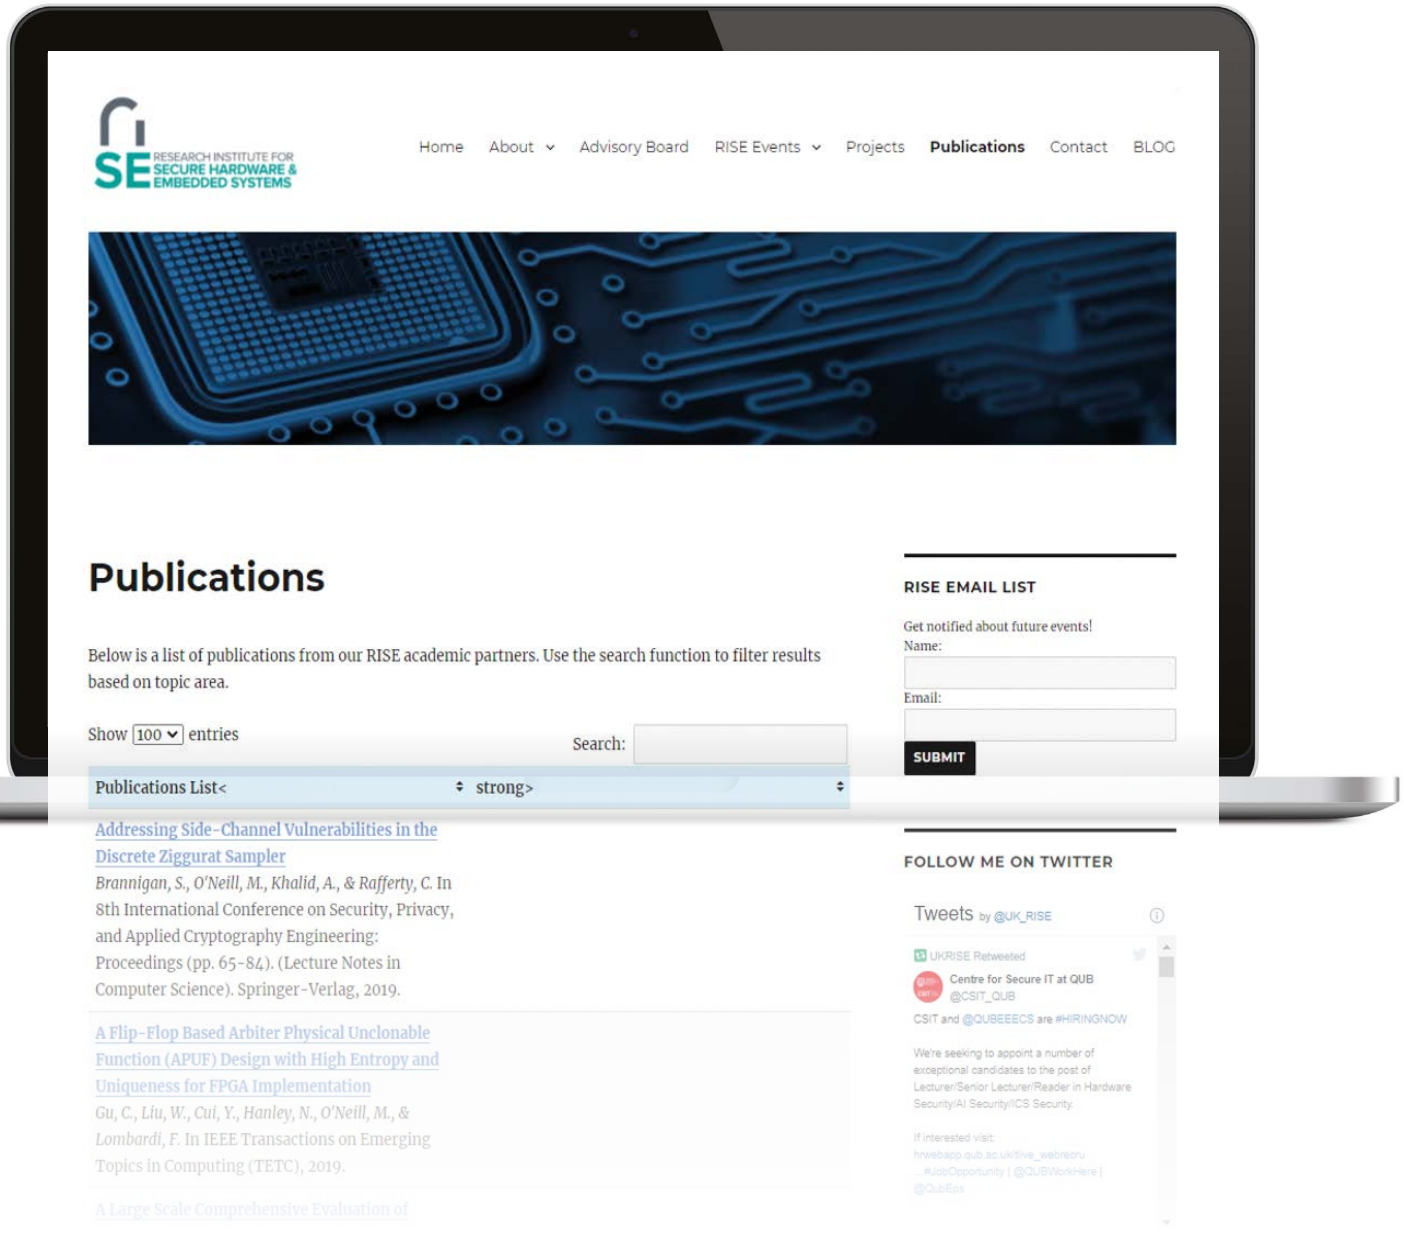

SIPP RESEARCH COLLABORATION

RISE members from Queens University Belfast, the University of Cambridge, the University of Birmingham, and the University of Bristol have been working on the EPSRC funded centre-to centre research programme between RISE and National University of Singapore (NUS) and Nanyang Technical University (NTU). The £1.3 million 3-year programme aims to rethink how security is built into IoT processor platforms and will leverage the strengths of each partner institute to address security across the hardware stack and through remote attestation.

RISE 2020 HARDWARE & EMBEDDED SECURITY COMPETITION

Since its inception, RISE has been working to nurture the Hardware & Embedded Security community. Bringing together both academia and industry through our events, speaking at conferences and disseminating information, RISE has been acting as a nucleating point. With a desire to encourage the next generation of hardware and embedded security professionals, we launched a new competition for final year undergraduate and master's student projects in 2020. The inaugural competition has been sponsored by ARM and we are able to offer some great prizes to promote interest and engage student participation.

RISE PROJECT UPDATES

IOSEC — PROTECTION AND MEMORY SAFETY FOR I/O SECURITY

The IOSEC project was established after the observation that a particularly under-scrutinised part of computer systems is security of the Input/Output (I/O) system. While much work has been done on security mitigations regarding software running on general purpose 'application' CPU cores, the security story when considering peripheral devices had seen much less focus. Peripheral devices today are no longer fixed-function hardware but run substantial vendor firmware, which is at risk of being compromised and triggering malicious behaviour.

We focused particularly on Direct Memory Access (DMA), which is a means by which peripheral devices can access system memory using their own mechanisms, rather than having to go via software running on a CPU. Since the majority of runtime system state is held in main memory, unrestricted DMA allows a peripheral device to read out everything running on the machine, including passwords, encryption keys, secret files, network traffic, etc., and worse tamper with it, injecting malicious code, false data, and so on.

IOMMU Security and Thunderclap

Such attacks were not new, but it was widely believed that a protection mechanism, the Input/Output Memory Management Unit (IOMMU) fixed the problem by blocking malicious DMA from peripheral devices. In particular, MacOS' enablement of the IOMMU in 2012 blocked a slew of prior attacks. We set out to verify if this was true.

One observation we made was that existing attacks had a very simple attacker model; the attacker simply turned up and tried to read or write memory, and this was being blocked by the IOMMU. However, in reality, devices are more complicated and must be given access to some memory in order to do their work. To examine this further, we built a software model of a PCIe network card, to run on a processor on an FPGA. The software model would be passed PCIe messages and respond as the real network card would respond. However, being in software not hardware meant it was much easier for us to modify to add malicious behaviour to the functional network card model (which we extracted from the QEMU emulator). We could then attach the FPGA to a victim machine via a PCIe slot or a laptop Thunderbolt port to explore what a machine would do in the presence of a real malicious device. We have open-sourced this work as the Thunderclap research platform.

We found defences were very poor. Windows made very little use of the IOMMU and was almost entirely unprotected. FreeBSD did, but it was not enabled by default. On Linux and MacOS we were able to see other network traffic such as VPN plaintext. Linux, when the IOMMU was enabled, allocated data for networking packets from the same pool as general kernel data structures which meant that much private kernel data, was exposed to own malicious device. MacOS was the only at the time of study to enable the IOMMU by default, but we found its use somewhat skin-deep. In particular, MacOS network packets contained kernel function pointers that we could subvert to launch a root shell. Finally, Linux was the only operating system to activate a PCIe feature that enabled us to fully bypass the IOMMU and access all of system memory.

We reported these findings to vendors and a series of patches have been forthcoming to address specific vulnerabilities across different operating systems, including advisories from Intel and Microsoft. More generally the community admitted the problem is difficult to solve in the general case, although Intel and others have been taking further measures to address it. Since publishing our work at NDSS in 2019, our Thunderclap research platform has been used by industry to verify some of their mitigations, due to additional development work we undertook to make it a more friendly hardware/software environment to use.

As a result of this work, mitigations have been made mandatory in the USB 4 and Thunderbolt 4 standards, which are now in shipping products.

DMA and CHERI capabilities

While undertaking the Thunderclap work, we became much more aware of the challenges of security in the I/O stack. Performance is critical, and security protections such as the IOMMU that hurt performance are unlikely to be adopted. Furthermore, the area is highly cost sensitive.

Following on from Thunderclap we have been considering whether CHERI capabilities, a promising new fine-grained memory protection primitive originally developed by the University of Cambridge and SRI International, and currently being prototyped in Arm's Morello silicon, provide a tool offering a more rigorous, and potentially more performant, I/O protection mechanism than the IOMMU.

To this end we have built another research platform that was inspired by the Thunderclap work. We wanted to model how a peripheral might be structured if it supported capabilities but building a peripheral in hardware only provides a limited data point. Instead, we took an existing dual-core CHERI CPU on FPGA, with one core running FreeBSD. The other core runs a 'virtual device', a software model of an existing peripheral such as an Intel network card, a SATA hard drive or NVMe flash, with a hardware bridge between the two cores to allow FreeBSD device drivers to attach to the virtual device on the second core. Because both models are software, we can reuse the existing CHERI C compiler and toolchain to implement a device that can either use capabilities internally, uses capabilities with its interaction with the host, or both. We are experimenting with different uses of capabilities in the device, device driver and operating system. In this way we explore the security of a range of scenarios more flexibly than if we had started with a hardware implementation.

The initial hardware we implemented this on was the dual-core CHERI MIPS prototype originally begun at Cambridge in 2010. Concurrent with the IOSEC project has been work on a RISC-V incarnation of CHERI, including implementation in a range of 32- and 64-bit processors, and in the LLVM RISC-V backend and FreeBSD/RISC-V. Due to external constraints this also meant almost entirely retooling the FPGA flow. CHERI-RISC-V is now the most supported CHERI FPGA platform, and we are working to port the virtual device model hardware/software across from CHERI-MIPS, to avoid having to maintain the legacy CHERI-MIPS platform. RISC-V is a more solid and forward-looking platform, which enables us to better explore the possibilities across systems-on-chip, including in third-party IP and SoCs.

At the HASP workshop in October 2020, we published a paper describing alternatives for using capabilities to protect DMA. In particular, we considered carefully the use cases and the threat models, which vary depending on what you are trying to protect and what might be compromised. Trade-offs exist based on the extent to which you can modify existing devices, or whether you are only able to add an external wrapper (for example, third party IP cores from a silicon vendor to which you don't have access). We presented a number of structures for interposing on such DMA and are working on implementation in the virtual device model. We aim to show whether such structures can enable a more holistic security viewpoint, as well as whether their overheads can be better than the existing IOMMU-based protection.

SCARV: A SIDE-CHANNEL HARDENED RISC-V PLATFORM

The SCARV (pronounced similar to “scarf”) project sits at the intersection of cryptographic engineering and computer architecture, with broad aims which can be summarised as providing support for efficient, secure execution of cryptography on RISC-V. Existing work can be described in terms of two overarching themes and, up to a point, is supported by one overarching artefact, i.e., an eponymous RISC-V compliant, micro-controller class processor core and associated SoC. As well as integrating and so acting as a testbed for most other work within the project, it has also gained traction outside the project: Kici et al. [8] used it as a test-case for constant-time circuit verification, for example.

Theme 1: Instruction Set Extensions

The concept of an Instruction Set Architecture (ISA) is fundamentally important within the context of computer systems, acting as an interface between hardware, i.e., some compliant micro-architectural implementation, and software executed by it. As a result, the extension of general-purpose ISAs via so-called Instruction Set Extensions (ISEs) to support special-purpose, domain-specific cryptographic workloads, has been and still is an active and important research field. To some extent reflecting by-design support in RISC-V for modularity and extensibility, the first theme has focused on cryptographic ISEs from two related perspectives.

A performance-oriented perspective. Toward the start of the project, we developed the XCrypto ISE for RISC-V; this essentially captured existing work as a structured specification. At the same time, the RISC-V Cryptographic Extensions Task Group (TG) had already started to develop a cryptographic ISE on top of the existing vector extension. To at least some degree, XCrypto helped motivate the TG to establish then drive development of an alternative, scalar extension proposal, e.g., to reflect a need to support cryptography on low(er)-end cores. At the time of writing, the proposal3 is edited by Ben Marshall and due to be ratified in Q4 of 2021; we continue to contribute to it more generally, for example through output related to the design of support for AES [9] (which could be viewed as a RISC-V analogue of technologies such as AES-NI), and development of the first, prototype implementation [10].

A security-oriented perspective. As the interface between hardware and software, an ISA has a potentially important role to play in terms of security; there is increased interest in capturing security-related guarantees in the ISA alongside related to state and computation. Within the context of RISC-V, for example, the proposed Zkt extension captures guarantees around constant execution latency. Under a broadly similar remit, we have produced two main outputs.

First, we developed an ISE [2] which acts to support masked software implementations of cryptography. The central idea is that the ISE captures a suite of “building block” operations used in masking; examples include a Boolean-masked AND (see, e.g., SecAnd, as defined in [1, Algorithm 2]). Use of the ISE yields a result which is both 1) more efficient, because each ISE instruction can replace a significant sequence of ISA instructions, and, crucially, 2) more secure, because common pitfalls with respect to ISA-based implementation of such sequences is avoided.

Second, we developed FENL [4], a specific instance of the general aISA concept introduced by Ge et al. [7]. Their general argument for a “new security-oriented hardware/software contract” imagined the ISA as a less opaque interface by selectively exposing detail in a micro-architectural implementation. FENL acts as a “fence for leakage” by ensuring that an instruction after the fence cannot interact with an instruction before it, and so prevents leakage that may otherwise occur. IIn concrete terms, it allows micro-architectural resources to be selectively used. This, in turn, permits development of software which is resilient against, e.g., power-based side-channel attacks, because any associated countermeasures can be based on more robust assumptions.

Theme 2: Understanding Micro-architectural Leakage

Over the last 20+ years, micro-architectural side-channel attacks [6, 12] have emerged as a significant threat. The majority of such attacks have focused on discrete leakage, e.g., data-dependent variance in execution latency that stems from the behaviour of cache memory. However, analogue micro-architectural leakage can also represent an important consideration: the premise is basically that fine-grained micro-architectural behaviour is observable, for example via power- and EM-based side-channels, and exploitable in associated attacks. Within an increasing body of work which attempts to understand this threat, we have produced two main outputs.

First, we carried out work aiming to better understand the internal design and so leakage characteristics of, e.g., specific ARM Cortex-M0 and Cortex-M3 cores. In [3] we observe that the in-line barrel shifter used to process the "flexible second operand" of ALU instructions will causes interactions between bits; these, in turn, cause leakage which violates assumptions in proofs of security for masking schemes. In [5] we applied side-channel assisted reverse engineering techniques to recover internal details of the Cortex-M3 pipeline. We conclude that, given doing so is possible, it would be useful to simply make them public: this would represent an easier, more reliable route to information that is vital for development of high-assurance software (e.g., that instruments countermeasures against side-channel attack, such as masking).

Second, we developed MIRACLE [11]: this work represents 1) experimental software infrastructure for evaluating micro-architectural leakage, 2) a specific dataset for and exploration of such leakage, spanning 14 different devices, 4 different ISAs, and 4 different vendors, and 3) a web-based interface5 for interactively exploring said dataset. The infrastructure allows characterisation of each device with respect to any leakage effects stemming from sources within the micro-architectural implementation; we use it, for example, to identify and document several novel leakage effects (e.g., due to speculative instruction execution), and scenarios where an assumption about leakage is non-portable between different yet compatible devices.

References

[1] J.-S. Coron, J. Großsch¨adl, M. Tibouchi, and P.K. Vadnala. "Conversion from Arithmetic to Boolean Masking with Logarithmic Complexity". In: Fast Software Encryption (FSE). LNCS 9054. https://doi.org/10.1007/978-3-662-48116-5_7. Springer-Verlag, 2015, pp. 130–149.

[2] S. Gao, J. Großsch¨adl, B. Marshall, D. Page, T.H. Pham, and F. Regazzoni. "An Instruction Set Extension to Support Software-Based Masking". In: IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES) 2021.4 (2021), To appear.

[3] S. Gao, B. Marshall, D. Page, and E. Oswald. "Share slicing: friend or foe?" In: IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES) 2020.1 (2019). https://tches.iacr.org/index.php/TCHES/article/view/8396, pp. 152–174.

[4] S. Gao, B. Marshall, D. Page, and T.H. Pham. "FENL: an ISE to mitigate analogue micro-architectural leakage". In: IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES) 2020.2 (2020). https://tches.iacr.org/index.php/TCHES/article/view/8545, pp. 73–98.

[5] S. Gao, E. Oswald, and D. Page. Reverse Engineering the Micro-Architectural Leakage Features of a Commercial Processor. Cryptology ePrint Archive, Report 2021/794. https://eprint.iacr.org/2021/794.

[6] Q. Ge, Y. Yarom, D. Cock, and G. Heiser. "A Survey of Microarchitectural Timing Attacks and Countermeasures on Contemporary Hardware". In: Journal of Cryptographic Engineering (JCEN) 8 (1 2018). https://doi.org/10.1007/s13389-016-0141-6, pp. 1–27.

[7] Q. Ge, Y. Yarom, and G. Heiser. "No security without time protection: we need a new hardware-software contract". In: Asia-Pacific Workshop on Systems (APSys). 2018.

[8] R.G. Kıcı, K.V. Gleissenthall, D. Stefan, and R. Jhala. Solver-Aided Constant-Time Circuit Verification. https://arxiv.org/abs/2104.00461. 2021. arXiv: 2104.00461 [cs.CR].

[9] B. Marshall, G.R. Newell, D. Page, M.-J.O. Saarinen, and C. Wolf. "The design of scalar AES Instruction Set Extensions for RISC-V". In: IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES) 2021.1 (2021). https://doi.org/10.46586/tches.v2021.i1.109-136, pp. 109–136.

[10] B. Marshall, D. Page, and T.H. Pham. "Implementing the Draft RISC-V Scalar Cryptography Extensions". In: Hardware and Architectural Support for Security and Privacy (HASP). 2020.

[11] B. Marshall, D. Page, and J. Webb. MIRACLE: MIcRo-ArChitectural Leakage Evaluation. Cryptology ePrint Archive, Report 2021/261. https://eprint.iacr.org/2021/261. 2021.

[12] J. Szefer. "Survey of Microarchitectural Side and Covert Channels, Attacks, and Defences". In: Journal of Hardware and Systems Security 3.3 (2019). https://doi.org/10.1007/s41635-018-0046-1, pp. 219–234.

USER-CONTROLLED HARDWARE SECURITY ANCHORS: EVALUATION AND DESIGNS

UNIVERSITY OF BIRMINGHAM

The University of Birmingham is working on different projects linked to RISE. In the following, a summary of these projects is provided.

Remote registration of multiple authenticators

This project is focused on developing protocols that rely on user authenticator tokens, such as the Fast IDentity Online (FIDO) Alliance protocol. These protocols are aimed at easing the burden of remembering authentication credentials for end-users that register into multiple services (called relying parties, RP). However, one limitation of this solution is the lack of robustness: if the user loses their authenticator, they will lose access to the registered services, and fully recovering the access to their accounts might be tedious and cumbersome. In our work, we propose and analyse three different approaches to solve the issue, and discuss their robustness and usability: duplicate authenticators, proxy authenticators, and ring-signature-based authenticators. The goal is that a user can register to a RP with any of their tokens, and this will automatically register the remaining back-up tokens without need to interact with them. We propose and formally analyse protocols for pairing the tokens, registering, and authenticating against RP. We also propose a solution to enable automatic verification for an arbitrary (but finite) number of authentication tokens per user. Currently we have finished the definition and formal analysis part and are preparing the scientific paper for submission.

Using enclaves to improve confidentiality from cloud

An interesting application of a cloud service is Conference Management Systems (CMS). The cloud service provider maintains the infrastructure required to run the service, removing from the user the burden of installing and handling such a system. However, data stored in the cloud can be abused (either intentionally by the cloud service provider or accidentally via a malicious insider or a data breach) to profile participants based on their performance. Sending the data encrypted will drastically limit the functionalities that can be offered by the cloud. On the other hand, using an encryption scheme (such as homomorphic encryption) does not provide the performance nor the flexibility needed for such a cloud service. In this work, we propose a protocol for a privacy-preserving conference management system that leverages an enclave platform at the cloud-side to securely process data, and a shuffling algorithm to hide the links between submissions and reviewers. The system specification and definition of the protocol has been finished, and we are starting the formal analysis part.

Design of a two-tiered TCB

A TCB (trusted computing base) is the set of software and hardware components of a system which form a trust anchor, and upon which the security of the system relies. In this project, we investigate how to split the TCB into two parts, a minimal trusted computing base (M-TCB), and an extended trusted computing base (E-TCB). The M-TCB provides a limited set of functionalities, but the most hardened services; the core functionality of the M-TCB is to protect the platform from strong adversaries and provide the ultimate trust anchor for the system. The E-TCB provides additional security services which need not be protected to quite the same extent. We specify and develop an M-TCB / E-TCB system architecture and a collection of protocols that provide essential trusted computing services against a variety of adversaries with different capabilities: secure boot, platform attestation, firmware A/B updates with rollback protection, and recovery from memory corruption caused by a strong attacker. We have finished the definition of the methodology, and the design and verification of the protocols, and are currently working on the scientific paper.

Hardware devices to support transparent decryption

Transparent decryption has been suggested as a way of achieving auditing fairness and ensuring transparency of access to sensitive information. The core idea is that the sensitive information is encrypted with a key which is held on a special-purpose hardware device, whose firmware restricts the way in which the key can be used. The key can be used to decrypt the sensitive information, but the device allows this to happen only if, along with the decryption request, it receives a proof that the decryption request is present in an append-only public ledger. Interested parties can monitor the ledger to see what decryption requests have been made. The public ledger is maintained by a party which is not required to be trusted, because its actions can be verified. The ledger is maintained in the form of a Merkle tree, and the ledger maintainer routinely and regularly signs and publishes the root tree hash (RTH). This allows anyone that can access

the ledger to verify that the ledger is maintained append-only. Transparent decryption aims to achieve the following security property: If a decryption has taken place, then an interested user can see evidence of that in the ledger, or the interested party can detect that they cannot access the ledger. In this project we propose and study the protocols that must be implemented in order to interact with such a transparent-decryption security device. The formal analysis of this scenario has presented several nontrivial challenges that need to be addressed (e.g., non-volatile memory of the device, unbounded size of the collection of entries in the ledger, etc.). We use and demonstrate the need for advanced features of the verification tools in order to achieve a fully mechanised proof for the scenario. We are on the latest stages of the verification and will proceed to the preparation of the paper shortly afterwards.

Root-of-Trust abstractions for symbolic analysis
In this work, we present a methodology that enables the use of formal verification tools towards automatically verifying complex protocols using roots of trust (RoT). The focus is on reasoning about the overall application security, provided from the integration of the RoT services, and how these can translate to larger systems when the underlying cryptographic engine is considered perfectly secure. The main objective of the present work is to propose a methodology for proving security in scenarios based on services that make use of RoTs, by idealizing the internal cryptographic functionalities of the security device, except those that provide explicit functionalities for the service being offered. We instantiate our methodology in a TPM-based remote attestation scenario. The paper has been published in the International Workshop on Security and Trust Management (STM), co-located with ESORICS 2021. (In press)

Secure IoT processor platform with remote attestation (SIPP)
The SIPP project aims to rethink how security is built into Internet-of-Things (IoT) processor platforms. Firstly, the architectural fundamentals of a processor design need to be re-engineered to assure the security of individual on-chip components. This has become increasingly evident with the recent Spectre and Meltdown attacks. On the upper layer of systems-on-chip (SoCs), hardware authentication of chip sub-systems and the entire chip is crucial to detect malicious hardware modification. Then, at the systems layer (i.e., multiple chips on a printed circuit board), innovative approaches for remote attestation will be investigated to determine the integrity at the board level. Finally, the security achieved at all hierarchical layers will be assessed by analysing physical-level vulnerabilities to ensure there is no physical leakage of the secrets on which each layer relies.

Our current idea is to use CHERI (Capability Hardware Enhanced RISC Instructions) for the SIPP processor. We are currently finishing an FPGA demo of the CHERI Piccolo processor. The aim of the CHERI protection model is to support fine-grained pointer and memory protection within address spaces, and provide primitives to support both scalable and programmer-friendly compartmentalisation within address spaces. After that, we plan to analyse possible vulnerabilities to attacks on the CHERI platform, such as Row hammer.

CHERI-TrEE: Flexible enclaves on capability machines
The CHERI-TrEE decomposes Enclave Execution Systems (EESs) into a set of more basic, orthogonal features. For this purpose, CHERI-TrEE proposes a new EES design that avoids reinventing existing mechanisms on the CHERI-RISC-V capability machine. The result is an EES with novel characteristics like dynamically growing and shrinking enclaves, nested enclaves, sharing of memory between enclaves etc. In this project, we assisted the KU Leuven group with the FPGA implementation of their CHERI-TrEE processor, and the analysis of the hardware in terms of processing speed, area occupation and power consumption.

VoltPillager: Hardware-based fault injection attacks against Intel SGX enclaves using the SVID voltage scaling interface
In this project, we built VoltPillager, a low-cost tool for injecting messages on the Serial Voltage Identification bus between the CPU and the voltage regulator on the motherboard, allowing us to control the CPU core voltage precisely. VoltPillager allows us to mount fault-injection attacks that breach the confidentiality and integrity of Intel SGX enclaves. We present a proof-of-concept key-recovery attack against cryptographic algorithms running inside SGX. VoltPillager attacks are more powerful than recent software-only undervolting attacks against SGX (CVE-2019-11157),

because they work on fully patched systems with all countermeasures against software undervolting enabled. Additionally, we can fault security-critical operations by delaying memory writes. Mitigating VoltPillager is not straightforward and may require a rethink of the SGX adversarial model where a cloud provider is untrusted and has physical access to the hardware.

This project resulted in one paper published about the VoltPillager tool that was presented at the Usenix Security 2021, presenting the first hardware-based voltage glitching attack against a fully-fledged Intel CPU.

Paper URL: https://www.usenix.org/conference/usenixsecurity21/presentation/chen-zitai

SCAvenger

The SCAvenger project concerns the security of machine learning applications. By exploiting memory access patterns and/or timing side channels in Machine Learning libraries, we hope to leak hidden parameters about a neural network (such as weights and biases) during inference which would aid an attacker in stealing or duplicating a network. By creating a replica of a network using these means (as compared to conventional model stealing attacks which work by brute-forcing), we hope to more quickly and accurately produce copies. Given these networks will be an exact match of the original, it should be easier for attackers to craft adversarial examples (AEs) against the original network. Also, we use Side-Channel Attacks (SCAs) in order to scavenge information about a network. Once these vulnerabilities are demonstrated through proofs of concept, we will propose some countermeasures.

Currently, we are working on scaling up some simple proofs-of-concept to more realistic and complex networks. Eventually, we hope to also be able to target inference on GPU hardware. This project is funded by Intel and HP.

Plundervolt: software-based fault injection attacks against Intel SGX

Modern processors offer the user the opportunity to modify the frequency and voltage through privileged software interfaces. With Plundervolt, we showed that these software interfaces can be exploited to undermine the system's security. Therefore, we were able to corrupt the integrity of Intel SGX on Intel Core processors by controlling the voltage when executing enclave computations. This means that even Intel SGX's memory encryption/authentication technology cannot protect against Plundervolt.

In this project, a paper was published in the 41st IEEE Symposium on Security and Privacy (S&P) 2020, with the collaboration of KU Leuven and Graz University of Technology.

Paper DOI: https://doi.org/10.1109/SP40000.2020.00057.

Fill your boots: enhanced embedded bootloader exploits via fault injection and binary analysis

The bootloader of an embedded microcontroller is responsible for guarding the device's internal (flash) memory, enforcing read/write protection mechanisms. Fault injection techniques such as voltage or clock glitching have been proven successful in bypassing such protection for specific microcontrollers, but this often requires expensive equipment and/or exhaustive search of the fault parameters. Besides, their lack of debugging capabilities makes embedded bootloaders notoriously hard to analyse. Therefore, we propose a grey-box approach that combines binary analysis and advanced software exploitation techniques with voltage glitching to develop a robust attack methodology against embedded bootloaders. We showcase our techniques with three real-world microcontrollers as case studies:

1. We combine static and on-chip dynamic analysis to enable a Return-Oriented Programming exploit on the bootloader of the NXP LPC microcontrollers.

2. We leverage on-chip dynamic analysis on the bootloader of the popular STM8 microcontrollers to constrain the glitch parameter search, achieving the first fully documented multi-glitch attack on a real-world target.

3. We apply symbolic execution to precisely aim voltage glitches at target instructions based on the execution path in the bootloader of the Renesas 78K0 automotive microcontroller.

For each case study, we show that using inexpensive, open-design equipment, we can efficiently breach the security of these microcontrollers and get full control of the protected memory, even when multiple glitches are required. Finally, we identify and elaborate on several vulnerable design patterns that should be avoided when implementing embedded bootloaders.

In this project, one paper was published in the IACR Transactions on Cryptographic Hardware and Embedded Systems (vol. 2021, no. 1) with the collaboration of an independent researcher, Qais Temeiza.

Paper DOI: https://doi.org/10.46586/tches.v2021.i1.56-81.

Cutting through the complexity of reverse engineering embedded devices
Performing security analysis of embedded devices is a challenging task. They present many difficulties not usually found when analysing other kinds of commodity systems: undocumented peripherals, esoteric instruction sets, and limited tool support. Thus, a significant amount of reverse engineering is almost always required to analyse such devices. To address these issues we have developed Incision, an architecture and operating-system agnostic reverse engineering framework. Incision tackles the problem of reducing the upfront effort to analyse complex end-user devices. It combines static and dynamic analyses in a feedback loop, enabling information from each to be used in tandem to improve our overall understanding of the firmware analysed. We use Incision to analyse a variety of devices and firmware. Our evaluation spans firmware based on three RTOSes, an automotive ECU, and a 4G/LTE baseband. We demonstrate that Incision does not introduce significant complexity to the standard reverse engineering process and requires little manual effort to use. Moreover, its analyses produce correct results with high confidence and are robust across different OSes and ISAs.

This paper was published in the IACR Transactions on Cryptographic Hardware and Embedded Systems (vol. 2021, no. 3). Paper DOI: https://doi.org/10.46586/tches.v2021.i3.360-389.

DEEPSECURITY: APPLYING DEEP LEARNING TO HARDWARE SECURITY

QUEEN'S UNIVERSITY BELFAST

With the globalisation of supply chains, the design and manufacture of today's electronic devices are now distributed worldwide, for example, through the use of overseas foundries, third party intellectual property (IP) and third-party test facilities. Many different untrusted entities may be involved in the design and assembly phases and therefore, it is becoming increasingly difficult to ensure the integrity and authenticity of devices. The supply chain is now considered to be susceptible to a range of hardware-based threats, including hardware Trojans, IP piracy, integrated circuit (IC) overproduction or recycling, reverse engineering, IC cloning and side-channel attacks. These attacks are major security threats to military, medical, government, transportation, and other critical and embedded systems applications. The DeepSecurity project focuses on two of these threats and investigates the use of deep-learning (DL) in the context of side-channel attacks and hardware Trojans.

Side-channel analysis (SCA) attacks exploit physical signal leakages, such as power consumption, electromagnetic emanations, or timing characteristics, from cryptographic implementations, and have become a serious security concern with many practical real-world demonstrations, such as secret key recovery from the Mifare DESFire smart card used in public transport ticketing applications and from encrypted bitstreams on Xilinx Virtex-4/5 FPGAs. A hardware Trojan (HT) is a malicious modification of a circuit in order to control, modify, disable, monitor or affect the operation of the circuit.

The proposed project seeks to investigate the application of deep learning in SCA and HT detection, with the ultimate goal of utilising deep learning based verification processes in Electronic Design Automation tools to provide feedback to designers on the security of their designs.

Deep Learning in Side-Channel Analysis

DL has proven to be very effective for image recognition tasks, with a large body of research on various model architectures for object classification. The application of DL to side-channel analysis has shown promising success, with experimentation on open-source variable key datasets showing that secret keys can be revealed with 100s of traces even in the presence of countermeasures. In this project we further improve the application of DL for SCA, by enhancing the power of DL when targeting the secret key of cryptographic algorithms that are protected with SCA countermeasures. We propose a new model, a CNN-based model with a Plaintext feature extension (CNNP) together with multiple convolutional filter kernel sizes and structures with deeper and narrower neural networks. This approach has empirically proven its effectiveness by outperforming reference profiling attack methods such as template attacks (TAs), convolutional neural networks (CNNs) and multilayer perceptron (MLP) models. It was verified using the ASCAD database, a set of databases that provide a benchmarking reference for DL-based SCA research which target a masked AES implementation running on an 8-bit AVR ATMega8151 device. Our model generates state-of-the art results when attacking the ASCAD variable-key database, which has a restricted number of training traces per key, recovering the key within 40 attack traces in comparison with the order of 100s of traces required by previous work using machine learning (ML) approaches. During the profiling stage an attacker needs no additional knowledge of the implementation, such as the masking scheme or random mask values, only the ability to record the power consumption or electromagnetic field traces, plaintext/ciphertext and the key. Additionally, no heuristic pre-processing is required to break the high-order masking countermeasures of the target implementation. This research culminated in a publication in IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES) in August 2020 and a presentation at the CHES2020 conference, held virtually in September 2020.

More recent work has investigated a stacked ensemble model, which trains the output probabilities and Maximum likelihood score of multiple traces and/or sub-models to further improve the performance of CNN-based SCA models and reduce the number of attack traces needed to successfully recover the key to 24.

Deep-Learning based Hardware Trojan Detection

Various functional Trojan detection techniques have been proposed over the past decade. However, approaches based on simulation, side channel analysis (SCA), reverse engineering and logic testing have shortcomings. Both simulation and logic testing have difficulties in generating comprehensive test vectors. SCA approaches usually need a 'golden' circuit and are sensitive to process variation. Moreover, for both the reverse engineering and SCA attacks, the preparation cost of test platforms or the extra overhead of the integration of detection sensors in ICs could make the detection very expensive. Static HT detection techniques, which can check Trojans without the need to run the circuit in design-time, have been proposed to prevent HT insertion before manufacturing and provide timely feedback to the design team. For example, machine learning (ML)-based and neural network (NN)-based HT detection methods have been proposed to detect and prevent HT-insertion at design-time without involving any extra pre-processing or introducing additional overheads. However, in most of these approaches HT related features are directly extracted from circuit designs and they often use engineering intuition to carefully craft the detection model to improve accuracy. Hence, knowledge of the circuits, including the circuit topology, types of components, and types of HTs, is essential for the detection and determines the accuracy of detection results . Furthermore, the crafted detection models often overfit the specific Trojan design used in testing, causing a large performance gap when facing new HT designs.

In this research work, we propose a data-driven HT detection system based on gate-level netlists which requires no prior knowledge of the circuit. The proposed HT detection system provides an extremely simplified detection process without the need for any pre-processing or extra circuit overheads, and it is also effective for various types of circuits. A Natural language processing (NLP) technique is utilized for feature extraction from the circuit netlist for HT detection. To the best of the authors' knowledge, this is the first time NLP has been applied on raw gate-level netlist data for HT detection. Data-driven DL models, namely LSTM and CNN, are utilized for data training based on the extracted features using the NLP algorithm. The results are verified using the Trust-Hub database, an open-source HT benchmarking library. Experimental results show that both the LSTM and CNN DL models achieve good HT detection performance for various Trojan netlists.

This research led to the publication of two journal papers and three conference papers, as follows:

[1] S. Yu, C. Gu, W.Liu, M O'Neill, Deep Learning-based Hardware Trojan Detection with Block-based Netlist Information Extraction, IEEE Transactions on Emerging Topics in Computing, September 2021

[2] M. Xue, C. Gu, W Liu, S. Yu, M. O'Neill, Ten years of Hardware Trojans; A Survey from the Attacker's perspective, IET computer and Digital Techniques, pp 231-246, October 2020

[3] Y. Dou, S. Yu, C. Gu, M. O'Neill, C. Wang, W. Liu, Security Analysis of Hardware Trojans on Approximate Circuits, The 30th edition of the ACM Great Lakes Symposium on VLSI (GLSVLSI), 8-11 September 2020.

[4] S. Yu, C. Gu, W. Liu and M. O'Neill, A Novel Feature Extraction Strategy for Hardware Trojan Detection, IEEE International Symposium on Circuits and Systems (ISCAS), Seville, May 2020.

[5] S. Yu, W. Liu, M. O'Neill, An Improved Automatic Hardware Trojan Generation Platform, IEEE Computer Society Annual Symposium on VLSI, Florida, July 2019.

GUPT: A HARDWARE-ASSISTED SECURE AND PRIVATE DATA ANALYTICS SERVICE

THE UNIVERSITY *of* EDINBURGH

Cyber-physical systems and online services based on “data-driven intelligence” have become an integral part in every aspect of people’s lives nowadays. For that reason, it is imperative that their design complies with strict reliability, security, and privacy requirements and, simultaneously, maintains real-time performance and scalability at a high level. However, the heterogeneity of the infrastructure that they are deployed to (e.g. cloud providers, IoT systems) have proven to be challenging both for developers and users, who are unable to safely assess the benefits and drawbacks of a particular system.

More specifically, when offloading computation to a remote host, we have to trust that they will execute the program they advertise correctly, without tampering with the results, or selling users’ data to third parties. End-to-end encryption addresses some of these challenges, but there are cases where we want the cloud platform to have limited access to the data to perform computation. Using Trusted Execution Environments for these tasks offset some of the trust to the hardware manufacturer rather than the machine operator, by protecting integrity and confidentiality of the programs and data and providing a way for the end users to verify the correctness of the computation, using the mechanism of remote attestation.

As Trusted Execution is already widely used for a variety of industrial applications, it becomes fundamental to accurately model its security guarantees, and how they can be used as a tool to construct more powerful and efficient protocols. Our first approach to bootstrap confidential computation in the cloud resulted in the Steel protocol (PKC 2021), based on work from Fisch et al., 2016. Steel realises the novel cryptographic primitive of “Functional Encryption for Stateful and Randomised functionalities’’ (FESR) and is proven to be secure under the notion of Universal Composability (using the formulation of trusted execution first designed by Pass et al, 2017). Our composable proof allows us to integrate Steel within different protocols that could benefit from the functionality of FESR (such as an upcoming work on Contact Tracing), without requiring a separate proof.

The complexity of interactions between hardware and cryptographic components produces a large attack surface for today’s Trusted Execution Environment implementations; while clearly addressing this kind of bugs at the source is necessary and beneficial, an alternative approach we consider important is to relax our models of Trusted Execution to capture these additional vulnerabilities and build high level schemes to secure the weaker platforms. Two important relaxations are allowing leakage of a trusted enclave’s internal memory and permitting state continuity attacks (rollback and forking). In an upcoming work, we aim to show that, from a weaker enclave formulation that allows both kinds of adversarial behaviours, we can bootstrap the more secure version of Trusted Execution used in the literature.

However, the security guarantees provided by TEEs apply solely to the data that resides inside the volatile memory, protected by hardware. Storage devices in the cloud providers’ facilities, though, are another vulnerable part of the system stack. Therefore, it is quite important to be able to extend the properties that TEEs provide, on the critical, non-volatile application’s data as well.

In recent years, Persistent Memory (PM) tends to acquire its place in the system stack of many cloud vendors. The introduction of PM devices in the storage layer can provide a significant performance boost for the applications, if leveraged to its fullest. Nevertheless, persistent memory, along with its unique characteristics, comes with its own, unique programming model.

Applications should be adapted accordingly, and programmers have to be careful so that they ensure not only the correctness but also the crash-consistency property. This trend urged the need to provide an end-to-end solution to manipulate PM securely. Our proposed system is based on Intel's PMDK, the most widely used framework for PM programming and is built on SCONE, a shielded execution framework which provides high security guarantees to the application, confines its address space inside the secure protected enclave memory and requires no modification on the code. We strive to offer a scalable, fault-tolerant system which ensures strong security guarantees. We also aim to introduce the least possible overheads, so that our system can be used in demanding environments requiring both strong security guarantees and high performance and scalability levels. On top of that, we expose an API, similar to the one that is used by PMDK, so that the applications can be easily adapted and make use of our proposed system without much effort from the side of the programmer. To complement our PM management system, we have also designed a secure network stack based on direct I/O mechanisms (e.g. RDMA) to allow secure communication between the clients and the PM management system.

Lifting the weight of ensuring the reliability, performance (latency/throughput), scalability and security & privacy guarantees off application developers' minds is a task fraught with challenges, to which our research's deliverables offer a significant contribution. In addition, the generality of our approaches and their ability to be integrated seamlessly into existing applications/workloads constitutes solid building blocks for a distributed computing platform suitable for the design, development, and deployment of a wide range of data-driven intelligent applications. Between the theoretical guarantees provided by our modelling and our work strengthening the protections over associated peripherals, we further increase developers' motivation towards adopting TEE-based solutions.

SAFEBET – MEMORY CAPABILITIES TO ENABLE SAFE, AGGRESSIVE SPECULATION IN PROCESSORS

UNIVERSITY OF CAMBRIDGE

During this period, Dr Jonathan Woodruff was employed as the Senior RA working on this project. Jonathan completed the addition of CHERI secure memory protection to the RISCY-OO superscalar processor, which now successfully runs CheriBSD, our CHERI version of the FreeBSD, a full multi-user UNIX-like operating system.

The focus of this research project is mitigation of transient execution attacks on modern superscalar processors. Now that we have our CHERI-RISCY-OO superscalar processor implementation, analysis of known attacks has begun. The table below summarises our initial findings for the Spectre and Meltdown class of speculative execution attacks.

	asm	CHERI asm	C	CHERI C
Spectre PHT – bounds	✓		✓	✓ / ✗
Spectre BTB – oop	✓	✓		
Spectre BTB – ip	✓		✓	✓
Spectre RSB	✓		✓	✓
Spectre STL	✓		✓	✓
Meltdown User-Sup	✗			
Meltdown CHERI	✗			
Meltdown CSR	✗			
Meltdown SpecialCap	✗			

Key:

✓ Produced proof-of-concept exploit

✗ Failed to produce exploit

In collaboration with an Erasmus student – Franz Fuchs – we have reproduced a range of Spectre attacks on the baseline RISCY-OO and our CHERI-RISCY-OO core, which demonstrates that this is an excellent platform to undertake research into speculative execution attacks. For “Spectre PHT – bounds” the initial exploit was mitigated through a combination of our CHERI C compiler replacing pointers with bounds checked capabilities, and the hardware enforcing those bounds. However, given our inside knowledge, we were able to construct an alternative attack that was not inhibited by CHERI bounds checks.

Neither the baseline RISCY-OO nor the CHERI-RISCY-OO cores are vulnerable to Meltdown attacks. “Meltdown User-Sup” is reliant on speculating that TLB checks will pass, but like ARM processors, RISCY-OO ensures that the TLB check passes before issuing memory requests. Similarly, for the “Meltdown CHERI” variant that we added to check that CHERI capability bounds are checked before issuing memory requests. “Meltdown CSR” and “Meltdown SpecialCap” are not current vulnerabilities but could become a problem if we enhance the performance of the processor through further speculative techniques as we explore trade-offs between safety and performance. Results were published CARRV2021: Developing a Test Suite for Transient-Execution Attacks on RISC-V and CHERI-RISC-V

We continue our research into mitigating speculative execution attacks. Our objective is to determine when it is safe or risky to undertake speculative execution, and from this to only undertake safe speculative execution while still achieving excellent performance.

For CHERI-based compartmentalisation and bounds checking, we have formulated the Speculative Capability Constraint (SCC). A CHERI capability system should seek not to issue any data memory request (even speculatively), which is not in the following set of capabilities:

- In the committed state of the register file
- Derivable from the committed state of the register file
- Loadable from memory pointed to the above two groups

We can call this property the Speculative Capability Constraint (SCC). Attacks seek to violate SCC. SCC avoids side-channel mechanism concerns (cache, TLB, BHT ...) by not allowing extra-compartmental values into the microarchitecture.
Several meetings were also held with Microsoft to disseminate the microarchitectural contracts proposed in the SafeBet project, and we hope to be able to collaborate in future SafeBet publications.

TIMETRUST: ROBUST TIMING VIA HARDWARE ROOTS OF TRUST AND NON-STANDARD HARDWARE

UNIVERSITY OF SURREY

TimeTrust is about using hardware-roots of trust to build (nearly or fully) backwards-compatible counteractions to attacks based on forging proximity and/or forging time checks in cryptographic applications and/or systems. We focus on contactless payments made with bankcards and smart devices (i.e., phones), as our main use-cases. We are centred on a formal treatment, but we also have a significant degree of practical assessments.

TimeTrust's main achievements to date are as follows.

Contribution 1 (Financial Crypto 2019). We showed that if readers are dishonest then Mastercard's relay-resistance protocol cannot work, as the readers can cheat w.r.t. timing measurements and declaring these within the protocol. We proposed a series of new EMV-compatible payment protocols, stopping such timing-dishonest readers. These are based on Mastercard's RRP, and are called PayBCR and PayCCR, depending on whether the bank or the card audits the reader to ensure honest behavior.

See Tom Chothia, Ioana Boureanu, Liqun Chen, "Making Contactless EMV Robust Against Rogue Readers Colluding with Relay Attackers". Financial Cryptography 2019: 222-233

Contribution 2 (AsiaCCS 2020). We create the first cryptographic model for dishonest proximity-checking provers, in a context of authentication protocols and a series of distance-based attacks. This is much wider than payments, in both applications' domain as well as threat model. Yet, we use it to show that our PayBCR and PayCCR (see contribution 1) are cryptographically secure against relays even if readers are dishonest w.r.t. timing.

See Ioana Boureanu, Liqun Chen, Sam Ivey, "Provable-Security Model for Strong Proximity-based Attacks: With Application to Contactless Payment"s. AsiaCCS 2020: 87-100

Contribution 5 (CCS 2020). Together with INRIA and Consult Hyperion, we implement PayBCR in a mix of software and hardware, and test it in practice. In this same work, we are the first to implement Mastercard's RRP. We test both protocols against honest and dishonest readers. We show, for the first time, that Mastercard RRP can be soundly and robustly used in practice, and protects against phone-based relaying (e.g., app-level relaying). Moreover, we show that our enhanced-security PayBCR, protecting against dishonest timing measurements, also can be used soundly in practice, and it is not too computationally costly. Finally, we propose a new Dolev-Yao model for proximity-checking with dishonest readers, which allows for multiple cards and mobility (e.g., cards/phones being moved from one location to another); this model is shown correct at the theoretical level and mechanised in ProVerif to check a few payments protocols, including PayBCR.

See Ioana Boureanu, Tom Chothia, Alexandre Debant, Stéphanie Delaune, "Security Analysis and Implementation of Relay-Resistant Contactless Payments". CCS 2020: 879-898

Contribution 4 (CSF 2021). We produced a refined and precise cryptographic model for proximity checking. We showed that this level of refinement and precision in the attacker model is useful: it leads to exhibiting tens of new attacks in proximity checking. As a first, we mechanised a cryptographic model for proximity-checking in the cryptographic prover called EasyCrypt. As a first, we proved the computation security of Mastercard RRP against relaying, in EasyCrypt.

See Ioana Boureanu, Constantin Catalin Dragan, François Dupressoir, David Gerault, Pascal Lafourcade, "Mechanised Models and Proofs for Distance-Bounding", CSF 2021

Contribution 5 (under submission). We experimentally show a series of new relay-based, man-in-the-middle and fraudulent-payment attacks on mobile contactless-payments' apps, affecting primarily one mobile-app provider (e.g., Samsung, Google, Apple) and card type (e.g., Visa, Mastercard, Amex, etc). This affects both under-the-limit and over-the-limit payments. We show that mobile contactless-payment apps and card types other than the ones in our attack, are largely safe: i.e., our attack does not apply to them. Aside, we also show that Visa's relay-counteraction mechanism does not work, on standard mobile phones, whereas Mastercard's may not be useable (i.e., have an issue linked to usability vs security). We explain why each of these issues is the case and propose patches for each, as well as a general patch. We prove the insecurity of the systems considered. and the security of the patches formally using Tamarin.

RFAS: RECONFIGURABLE FPGA ACCELERATOR SANDBOXING

The University of Manchester

We started our first work with building up an infrastructure consisting of ring oscillators and high-accuracy time-to-digital converters to allow us experimenting at GHz signal rates for DPA and power-hammering attacks. The work on oscillator design continued over the whole reporting period to build faster oscillators and all (possible) oscillator variants, with the goal to get a holistic view on self-oscillating circuits implemented on FPGAs.

In parallel, we worked on more traditional security aspects, including the physical implementation of modules into bounding boxes [3] (Task 3 in the proposal). The proposed module implementation flow ensures that reconfigurable accelerator modules will be physically implemented into strict module bounding boxes, which have isolated fences in between. The module encapsulation of this newly built static system satisfies the requirement verification for the Isolation Design Flow (IDF) of the FPGA vendor Xilinx.

Moreover, we worked on memory protection by incorporated the IOMMU of an ARM SoC to monitor all accesses issued by any hardware accelerator [1]. We also examined details of the memory subsystems in complex ARM SoCs, including scheduling policies on memory busses [2].

We implanted a Bitstream Virus Scanner (Task 4 in the rFAS work programme) [4]. This is a unique protection scheme in which FPGA bitstreams are scanned against hardware vulnerabilities before configuring them to an FPGA. An FPGA bitstream will be parsed (by our tool BitMan) to rebuild a flat netlist. Then a graph search engine will search for several specific virus signatures which could harm the system. Current virus scans include combinatorial feedback paths, short-circuits, wire-tapping, antenna signals and high-fanout nets, which allow us detecting all reported threats targeting FPGAs at the electrical level.

The scanner was extended to characterise the threat-potential from general netlists by quantifying the glitch potential of a given FPGA design (its bitstream) [5]. This considers both the wiring of the netlist and the logic functions (e.g., an XOR produces more glitches than an AND-gate). A broader study of our virus scanner was published at ACM TRETS [8].

We decided to deploy these circuits in our lab and on an Amazon Web Services F1 instance, and we found that we were able to crash an F1 instance by draining more power than what the F1 instance could sustain. We discussed the issue with AWS and the FPGA vendor Xilinx. As an outcome, we designed custom design checker scripts that can mitigate the most severe security issues. We also received USD 60K in cash and boards worth over £ 20K to support our research. Our findings had been published at TCHES [10].

We ran experiments to look at the time it took us to receive the same instance without crashing and with crashing the FPGA instance and we implemented PUFs for the AWS FPGAs, as instances are otherwise fully virtualized and there is no serial number, MAC etc. This work is further interesting as it allows a reverse engineering of the AWS scheduling policies.

We fine-tuned the glitch amplification experiments to maximize both speed of switching and the number of elements (logic primitives and routing resources) that can be used for power hammering.

We found that power hammering reduces throughput on FPGA PCIe links. Experimental results show that by using moderate levels of power hammering (~30w), throughput decreased by about 2GB/s (12%) for large packet sizes. Slowing down PCIe links could be used to cause race conditions or to implement a covert channel between the FPGA and a host CPU.

In collaboration with Nele Mentens (KU Leuven, Belgium) and Ahmad Reza Sadeghi (TU Darmstadt, Germany), we developed a trusted FPGA environment that solves two problems; Firstly it uses our FPGADefender virus scanner to ensure a cloud service provider (CSP) that a user bitstream will not be malicious, and secondly, it ensures the user IP protection by configuring an FPGA only with encrypted configuration bitstreams. This extends mechanisms that we know from the software world (like Intel SGX or ARM TrustZone) to be used for FPGA execution. This work was accepted and presented at FCCM 2021.

To complete our understanding of dynamic power consumption, we characterized the power consumption of routing resources as a function of the signal toggle rate driving these wires. The goal is to find the modes of operation that come with the most severe threat potential.

Related to the previous aspect, we investigated the maximum waste power generation of primitives operating in synchronous mode. This is important as this is currently not flagged as potentially malicious by FPGA design tools. However, first experiments have shown that just rapidly toggling all flip-flops of an FPGA can already crash a chip (a datacentre FPGA provides a few million user flip-flops).

The latest activities in this reporting period included building a virus scanner that could work in an embedded system and research on I/O primitives. We added an extra scan to our virus scanner that detects ring-oscillators that are built through I/O primitives of an FPGA. While a user should normally not have direct access to I/O pins (like in a CPU system, a user should not have direct access to peripherals), AWS is providing this for the I/O pins used to connect to external memory. This is a security flaw and we added this to our virus scanner to protect against the design of ring-oscillators (which are the base to mount information leakage, fault injection or general power-hammering attacks. However, more research is needed to judge if miss-configurations could be a threat (e.g., by setting wrong IO standards).

Outreach Activities

- HiPEAC ACACES summer school in Italy poster presentation discussing the rFAS project ideas ACACES website: http://acaces.hipeac.net/2019
- Keynote at ReCoSoC 2019 in York; presenting rFAS and FPGA hardware security research in Manchester https://www.recosoc.org/pages/program.html
- Discussing FPGA hardware security issues with Ericsson
- Attending the Xilinx Security Working Group (XSWG) event in Munich 2019
- Demo on attacking FPGAs at FPGA 2020 [4]
- Demo at FPL 2020 Demo Night showing attacks and our virus scanner [6]
- Discussions with Amazon AWS and Xilinx on our findings on cloud FPGAs
- Talk “The Future of FPGA-Acceleration in Cloud and Datacenters” at the FCCM 2020 workshop
- Invited talk “Operating Systems for FPGAs - Why and How?” FPL 2020 LEGaTO workshop
- We provided Topic (an embedded design house in the Netherlands) with a test circuit that is creating specific power-hammering patterns on FPGAs. They used that design to test and characterize the power supply circuit of their latest FPGA module (to fulfil a customer request to prove the robustness of their system)

Publications

[1] K. Pham, K. Paraskevas, A. Vaishnav, A. Attwood, M. Vesper and D. Koch, “ZUCL 2.0: Virtualised Memory and Communication for ZYNQ UltraScale+ FPGAs”, FSP workshop, 2019

[2] K. Manev; A. Vaishnav and D. Koch, ”Unexpected Diversity: Quantitative Memory Analysis for Zynq UltraScale+ Systems”, FPT, 2019

[3] K Pham, A Vaishnav, J Powell and D Koch, “A Self-Compilation Flow Demo on FOS–The FPGA Operating System”, Demo at the 30th FPL, 2020.

[4] K. Matas, T. La, N. Grunchevski, K. Pham and D. Koch, ”Invited Tutorial: FPGA Hardware Security for Datacenters and Beyond”, ACM FPGA 2020

[5] K. Matas, T. Minh La, K. Pham and D. Koch, ”Power-hammering through Glitch Amplification – Attacks and Mitigation”, IEEE FPT, 2020

[6] T. La, K. Matas, J. Powell, K. Pham and Dirk Koch, “A Closer Look at Malicious Bitstreams”, Demo night contribution at FPL 2020

[7] T La, K Matas, K Pham and D Koch, ”Securing FPGA Accelerators at the Electrical Level for Multi-tenant Platforms” PhD forum at FPL 2020

[8] T. La, K. Matas, N. Grunchevski, K. Pham and D. Koch, ”FPGADefender Malicious Self-oscillator Scanning for Xilinx UltraScale + FPGAs”, ACM TRETS 13, 3, 2020

[9] S. Zeitouni, J. Vliegen, T. Frassetto, D. Koch, A. Sadeghi and N. Mentens, “Trusted Configuration in Cloud FPGAs”, 29th FCCM, 2021

[10] T. La, K. Pham, J. Powell and D. Koch, “Denial-of-Service on FPGA-based Cloud Infrastructure-Attack and Defense”, TCHES, #3, 2021

RESEARCH INSTITUTE FOR
SECURE HARDWARE &
EMBEDDED SYSTEMS

CONTACT DETAILS
W: www.ukrise.org
E: info@ukrise.org.uk
T: +44 (0) 28 9097 1771
@UK_RISE