2021 YEARBOOK

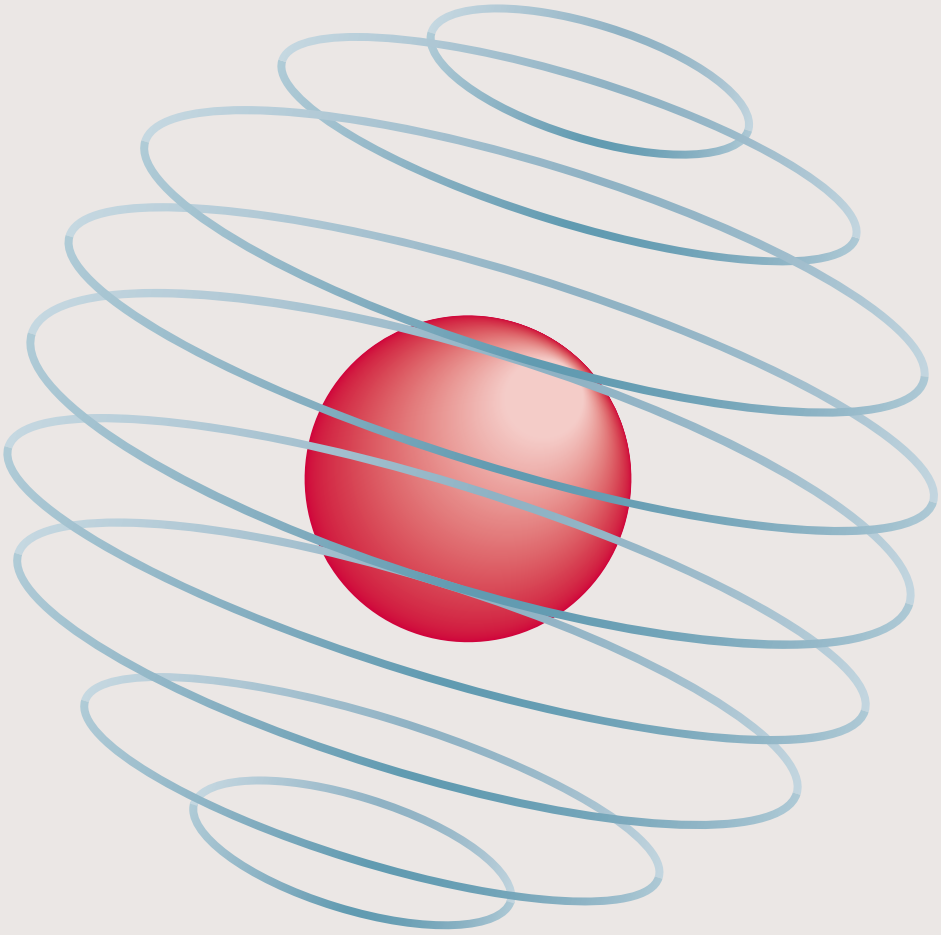

UNIVERSITY OF OXFORD

CENTRE *for* DOCTORAL TRAINING *in* CYBER SECURITY

UKRI Engineering and Physical Sciences Research Council

Contents

Director's Welcome

When we prepared the 2020 CDT Yearbook, I don't think any of us imagined that we would spend the whole of the academic year without once assembling the members of the CDT together in one place. But we have adapted to a new –hopefully temporary – normal way of working, and most of the regular work of the CDT has carried on.

Courses have been `attended' (both live and self-paced on line); seminars have been given, conferences have been attended, theses have been written, submitted, and examined; and even some social events have taken place. We moved our annual showcase event online, and it seemed to be very well-received. Our students have been distributed around the world, though many have continued to work in Oxford. Pandemic restrictions have slowed progress for some, and quite a few have made use of extensions of time (and funding!) made possible by EPSRC and University schemes.

Thanks to vaccinations – not least those made possible by colleagues elsewhere in this University – we can be cautiously optimistic about getting together soon face-to-face. We shall be a smaller group by then, because many CDT students have moved on to life after the doctorate. We kicked off a new alumni network during the year – for our CDT students, and for any other Oxford graduates working in the Cyber Security field, and we look forward to deepening professional relationships as those who have studied here move into positions of leadership across the sector.

This summer marks another turning point as David Hobbs leaves us to spend some time in the USA. He has been pivotal in the practical organisation of the CDT since it started, and has been the Centre Administrator for the last couple of years. It's been a role of many parts, from organising courses to financial management, with a large component of student welfare and support thrown in, as well as yearbook editorship! We will all miss him hugely, but wish him and his family well as they embark on a new adventure.

We also said farewell to Katherine Fletcher, earlier in the year. She has facilitated much of the networking across the University around Cyber Security, and has been supporting the CDT for some time in external liaison, too. She hasn't moved so far, and is now applying her talents to managing research teams in Medical Sciences.

You will see from the pages that follow another snapshot of the huge range of achievements from the CDT students. Theirs is the impact that will make a difference in a world ever-more reliant on good cyber security in all its dimensions.

Andrew Martin
Professor of Systems Security
Director, CDT in Cyber Security

UNIVERSITY OF OXFORD

This is to certify that

Mary Katherine Bigham

Wolfson College

having satisfied certain conditions

The Security of Human Computer Interaction by Speech

and having satisfied all the conditions prescribed by the Statutes of the University on 11 November 2020, was on 6 May 2021 admitted to the Degree of

DOCTOR OF PHILOSOPHY

University Offices,
Oxford
6 May 2021

Submitted Theses

RANJBAR BALISANE

Supervisors: Andrew Martin, Department of Computer Science

Engineering Secure, Usable, and Privacy Preserving Identity Management System using Trusted Computing

Researching systematic approach to enhancing authentication privacy and security using trusted computing.

While a lot of research is focused on enhancing a particular method of authentication such as enhancing hashing algorithms, or ways which fingerprint templates are stored, or finding new and more secure ways to authenticate a user. This research focuses on the underlying architecture and instead of asking the user to comply with complex policies and change with technology, it changes the underlying architecture using the advances made in technology to provider better privacy protection, security and more usable security.

Bio

Ranjbar has a First Class B.Sc. (Hons) in Ethical Hacking & Network Security and M.Sc. with Distinction in Forensic Computing. Prior to joining the CDT in Cyber Security at Oxford, he worked as an eCampus project manager for Soran University, Kurdistan, initiating the first large scale eCampus in Iraq, working in collaboration with LS Cables (LG). He has experience in vulnerability discovery, penetration testing, programming, networking, and protocol design.

Publications

R A. Balisane and A. Martin (2016). Trusted Execution Environment-Based Authentication Gauge (TEEBAG). In Proceeding of the New Security Paradigms Workshop (NSPW), Colorado, USA, 2016. ACM.

Atamli-Reineh, R. Borgaonkar, R. A. Balisane, G. Petracca, and A. Martin (2016). Analysis of Trusted Execution Environment usage in Samsung KNOX. In Proceedings of the Workshop on System Software for Trusted Execution (SysTex), Trento, Italy, 2016.

Mini-Project; OpenSky – Towards Secure Next Generation Air Traffic Communication Protocols

Mini-Project: Trusted Computing versus Identity

JACQUELINE EGGENSCHWILER

Supervisor: Rebecca Williams, Faculty of Law

Non-State Actors and Norms of Responsible Behaviour in Cyberspace

Computer systems and networks have become key determinants for the proper functioning of global markets, political institutions, and societies at large. Given their extensive reach into almost all areas of human activity, their safekeeping has become of strategic importance for a diverse range of actors. The proliferation of offensive cyber operations, such as WannaCry or Petya/NotPetya, has spurred calls for normative measures of restraint, and behaviour-guiding rules of the road.

Despite surging numbers of academic publications pertaining to cybersecurity generally, and norm-making processes specifically, the contributions of non-state actors to global cybersecurity governance efforts have remained under-theorised.

With a view to offering correctives, this thesis examines the roles assumed by non-state actors in global cybersecurity norm formation processes. Specifically, it analyses how, in which capacities, and how effectively non-state protagonists engage in norm cultivation endeavours by surveying nine exploratory case studies, grouped into three stakeholder clusters, i.e. (a) civil society and academia, (b) corporate actors, and (c) expert communities.

Triangulating different qualitative means and methods of data collection and analysis, this thesis suggests that non-state actors have come to exert discernible politico- legal influence over discussions about norms of responsible behaviour in cyberspace. Advancing empirically more informed and varied conceptualisations of the parts played by non-state actors in cybersecurity norm creation projects, this dissertation suggests that their roles can be systematised along the following profiles: (a) knowledge brokers, (b) awareness raisers, (c) norm leaders and cooperation incubators, (d) diplomatic change agents, (e) discussion feeders and gap fillers, (f) implementation assistants and capacity builders, and (g) custom shapers.
The case studies reveal noteworthy variations in how non-state entities seek to shape actor behaviour and realise regulatory effects. The results of this inquiry go to show that non-state actors have to be taken seriously as key contributors to global cybersecurity steering eorts, and that their actions and authority have come to extend beyond advocacy or lobbying.

Bio

Jacqueline Eggenschwiler is a doctoral researcher at the University of Oxford's Centre for Doctoral Training in Cyber Security. Her research looks at the contributions of non-state actors to global cybersecurity norm formation processes and corresponding governance implications. Jacqueline holds degrees in International Affairs and Governance, International Management, and Human Rights from the University of St. Gallen and the London School of Economics and Political Science.

Publications

Jacqueline Eggenschwiler, Ioannis Agrafiotis, and Jason RC Nurse, 'Insider Threat Response and Recovery Strategies in Financial Services Firms' (2016) 2016(11) Computer Fraud & Security 12 (https://perma.cc/L29A-H56C).

Jacqueline Eggenschwiler, 'Accountability Challenges Confronting Cyberspace Governance' (2017) 6(3) Internet Policy Review 1 (https://perma.cc/3KT6- SHGP).

Jacqueline Eggenschwiler, 'A Typology of Cybersecurity Governance Models' (2018) 13(2) St Antony's International Review 64 (https://perma.cc/4AK4-JGUU).

Jacqueline Eggenschwiler and Jantje Silomon, 'Challenges and Opportunities in Cyber Weapon Norm Construction' (2018) 2018(12) Computer Fraud & Security 11 (https://perma.cc/CK9P-CK47).

Jacqueline Eggenschwiler, 'An Incident-Based Conceptualization of Cybersecurity Governance' in Ryan Ellis and Vivek K Mohan (eds), Rewired: Cybersecurity

Governance (John Wiley & Sons 2019).

Myriam Dunn Cavelty and Jacqueline Eggenschwiler, Behavioral Norms in Cyberspace: Can Corporations Make the Digital Sphere Secure? (2019) (https://perma.cc/4GQE-EWVR) accessed 10 December 2019.

Jacqueline Eggenschwiler, International Cybersecurity Norm Development: The Roles of States Post-2017 (techspace rep, April, EU Cyber Direct 2019) (https://perma.cc/7PR4-C72T).

Jacqueline Eggenschwiler, 'Expert Commissions and Norms of Responsible Behaviour in Cyberspace: A Review of the Activities of the GCSC' [2020] Digital Policy, Regulation and Governance (https://perma.cc/4QKH-F2WG).

Jacqueline Eggenschwiler and Joanna Kulesza, 'Non-State Actors as Shapers of Customary Standards of Responsible Behaviour in Cyberspace' in Dennis Broeders and Bibi van den Berg (eds), Governing Cyberspace: Behavior, Power, and Diplomacy (Rowman & Littlefield 2020).

Other:

2017 EuroSSIG Fellow (https://eurossig.eu/eurossig/)

NextGen@ICANN 58, Copenhagen

JOHN GALEA

Supervisor: Daniel Kroening,
Department of Computer Science

Optimizing a primitive based approach to generic taint analysis

Dynamic taint analysis is a fundamental technique in software security that tracks the flow of interesting or suspicious data during the execution of a target application. However, it is also known to suffer from excessively high runtime overhead which limits its overall practicality. The overhead is especially problematic when the taint analysis performed is generic. While this type of analysis is more powerful than standard taint status tracking, due to its support of user-defined custom taint policies, its versatility naturally implies that it cannot be optimized for a single task. Along these lines, this dissertation addresses the challenging problem of enhancing the performance of generic taint analysis.

Our research focuses on taint engines implemented using dynamic binary instrumentation. Essentially, the performance bottleneck of these systems stems from complex taint propagation code that is dynamically instrumented, at the level of machine instructions, into the target application. In order to facilitate the implementation of generic taint analysis, instrumentation is typically achieved by inserting transparent calls, which conveniently invoke propagation routines without disrupting the execution behaviour of the analysed application. However, these calls perform expensive context-switching, and therefore degrade overall runtime performance significantly.

We hypothesize that the overhead of generic taint analysis can be reduced in comparison to the slowdowns incurred by the analysis implemented with the use of transparent calls. This hypothesis is tested by investigating three optimizations. Firstly, call-free generic taint propagation is evaluated as a means to mitigate expensive context switching. Secondly, the idea of vectorized generic taint analysis is explored to determine the effectiveness of using SIMD features to propagate multiple taint labels simultaneously. Thirdly, dynamic fast path generation is proposed so that the execution of ineffective taint propagation can be adaptively elided. Our optimizations are presented with respect to a generic taint analysis driven by user-defined primitives, which act as building blocks for custom taint propagation. Crucially, this primitive based approach enables the user to focus on optimizing core pieces of functionality regarding the desired policy, without the need to delve into the complex internals of the system.

The optimizations are integrated into a generic taint tracker called the Taint Rabbit, and their effectiveness is evaluated on various real-world software and CPU-bound benchmarks. With average speed-ups of around 42% on SPEC CPU benchmarks, and as high as 99% on data compression tools, our results demonstrate the advancement of optimized and generic taint analysis.

Bio

John Galea completed his B.Sc. undergraduate degree in Computer Science and Artificial Intelligence at the University of Malta. He continued his studies and was awarded a Master's degree in Computer Science in 2015.

John concluded reading for a DPhil in Cyber Security at the Department of Computer Science, University of Oxford. His interests revolve around binary analysis, vulnerability discovery and dynamic binary instrumentation (DBI). Some of his research builds upon the DBI engine DynamoRIO, an open-source project which he actively contributes to and helps maintain.

Publications:

Mini-Project: The Verser protocol: Verifying Services of IoT Devices Based on Their Capabilities

Mini-Project: ROPEX: Towards the Circumvention of Data Execution Prevention via Automatic Exploit Generation

JOHN GALLACHER

Supervisor: Joss Wright, Oxford Internet Institute

Online intergroup conflict: How the dynamics of online communication drive extremism and violence between groups

This thesis investigates the mechanisms through which social media may foster and exacerbate intergroup conflict both online and offline, and explores different intergroup and intragroup dynamics which may drive this trend.

This work firstly investigates the prevalence and impact of online communication between opposing groups on offline intergroup relations, demonstrating how intergroup communication does frequently occur online, but can be confrontational, hostile, and fails to improve intergroup relations, instead predicting future offline violence.

Following this, this work then develops a novel approach to detect online hate speech by leveraging datasets from multiple social media platforms in order to train natural language machine learning classifiers. The work then uses these classifiers to explore the influence of extreme outgroup denigration in ingroup discussions on fringe social media platforms popular with the far-right. It demonstrates how exposure to online hate plays a role in shaping individual radicalisation trajectories, increasing the likelihood of offline hate crimes, and driving mutual radicalisation with opposing groups.

Finally, the work explores the effect of hostile manipulation from state actors on the dynamics of online conversations, presenting novel techniques to reveal how manipulation of these discussions increases the polarisation of online conversations from genuine users.

Together these findings shed light on the mechanisms by which the Internet promotes intergroup conflict, extremism, and violence, providing new insight into initial steps that could be used to counter these effects.

Link to thesis repository - https://ora.ox.ac.uk/objects/uuid:9997b0cf-89a9-4ffb-8320-b98cc36845a0

Bio

John was DPhil student within the University of Oxford's Cyber Security Centre for Doctoral Training, graduating in 2021.

Outputs:

Gallacher, J.D., (2021) Leveraging cross-platform data to improve automated hate speech detection. arXiv. https://arxiv.org/abs/2102.04895

Gallacher, J.D., Bright, J,. (2021) Hate Contagion: Measuring the spread and trajectory of hate on social media. PsyArXiv https://psyarxiv.com/b9qhd

Gallacher, J.D., Heerdink, M, (2021) Mutual radicalisation of opposing extremist groups via the Internet. PsyArXiv https://psyarxiv.com/dtfc5

Gallacher, J.D., Heerdink, M. & Hewstone, M., (2020) Online engagement between opposing extremist political groups predicts physical violence of offline encounters Social Media + Society

https://journals.sagepub.com/doi/full/10.1177/2056305120984445

Gallacher, J.D. Radicalisation of the Online Far-Right: Hate Speech Propagation from Mainstream to Fringe Web Platforms. (2020) The European Consortium for Political Research - Advancing Political Science Conference, Innsbruck, Austria

Gallacher, J.D., & Heerdink, M., (2019) Measuring the effect of hostile information operations: a case study of Russian Internet Research Agency interference in online conversations Defence Strategic Communications, 6, 155-198

Gallacher, J.D., (2019) Automated Detection of Terrorist and Extremist Content. Extreme Digital Speech: Contexts, Responses, and Solutions. VoxPol Network of Excellence for Research in Violent Online Political Extremism

Gallacher, J.D., & Fredheim, R., (2018) Division aboard, cohesion at home: How the Russian troll factory works to divide societies overseas whilst spreading pro-regime messages to domestic audiences. Responding to Cognitive Security Challenges, Chapter 5, NATO Strategic Communications Centre of Excellence

Fredheim, R., & Gallacher, J.D., (2018) Robotrolling 3/2018. NATO Strategic Communications Centre of Excellence

Gallacher, J.D., Barash, V., Howard, P.N., & Kelly, J,. (2017) Junk News on Military Affairs and National Security: Social Media Disinformation Campaigns Against US Military Personnel and Veterans. Data Memo 2017.9. Oxford, UK: Project on Computational Propaganda

MANUEL HEPFER

Supervisors: Thomas Powell and Thomas Lawrence, Saïd Business School

Organizational resilience: the case of cybersecurity

Manuel is interested in the resilience of organizations facing adversity. Despite decades of research, key elements of organizational resilience remain poorly understood. The purpose of Manuel's research is to advance the study of organizational resilience.

Manuel works on three research projects.The first research project reviews and integrates literature on organizational resilience published in leading business and management journals. The second research project is an empirical study that explores the cognitive foundations of organizational resilience in the context of cybersecurity, comparing how

three global organizations have responded differently to the same cyberattack. The third research project enhances our practical understanding of organizational resilience by showing how executives can improve their organizational resilience to cyberattacks and capture strategic opportunities before and after a cyberattack occurs.

Bio

Manuel Hepfer is a doctoral student at Oxford University, studying in the Said Business School under the supervision of Professors Thomas C. Powell and Thomas B. Lawrence. His research focuses on the resilience of organizations in the empirical context of cybersecurity.

Manuel's academic background combines the areas of computer science and business administration. After graduating from Reutlingen University (Germany) top of his class with a bachelor's degree in Business Informatics in 2015, he pursued his education at the London School of Economics, where he graduated in 2016 with Distinction in Management, Information Systems, and Digital Innovation. After gaining experience during practical placements in management consultancies (PricewaterhouseCoopers, Audi Consulting) and large corporations (Porsche Financial Services, Robert Bosch GmbH), Manuel started his Doctoral degree at Oxford University.

Publications:

Hepfer, M., Powell, T.C., "How to Make Cybersecurity a Strategic Asset." MIT Sloan management review (2020).

MONICA KAMINSKA

Supervisor: Lucas Kello,
Department of Politics and International Relations

To Retaliate or Not: A Matter of Cyber Risk Perception

This thesis investigates variation between states' responses to hostile cyber operations, focusing on the cases of the United States, Israel, and Estonia. The thesis develops a conceptual framework based on the risk society theory in Sociology, which explains how cultural beliefs and historical experience producing risk aversion or risk acceptance interact with the material environment to produce particular international outcomes in state responses. Specifically, the thesis argues that the reason the United States failed to punish cyber attacks meaningfully is rooted in the dominant risk paradigm, fuelled by a cultural and historical aversion to surprise and uncertainty, which guided the US approach to cyber security competitions. In the case of Estonia, the same concerns about risk impeded a response, although these did not afflict Estonia itself, but NATO. While NATO was reluctant to issue a meaningful response due to the perceived risks of the operational environment, Estonia also chose to not pressure NATO for a response due to a culturally and historically engrained fear of NATO abandonment in a future crisis with Russia. Meanwhile, the reason that Israel exhibited a more robust pattern of punishing response is that policymakers were not constrained by the paradigmatic desire to avoid and minimise risk in international security competitions. Instead, their response calculus was shaped by a historical and cultural focus on a single primary adversary, rather than the operational environment, and a far more risk-acceptant strategic culture.

Bio

Monica is a final year doctoral student, hosted by the Department of Politics and International Relations and supervised by Prof. Lucas Kello. Alongside finishing her DPhil at Oxford, Monica is a postdoctoral researcher at The Hague Program for Cyber Norms at Leiden University – Institute of Security and Global Affairs. She is also a trustee of the European Cyber Conflict Research Initiative (ECCRI) and has previously held research positions at the Centre for Technology and Global Affairs and the Computational Propaganda Project, both at the University of Oxford. Monica's research examines international cyber conflict, particularly states' responses to hostile cyber operations. Monica holds an MPhil in Geographical Research from Cambridge University and a BSc in International Relations from the London School of Economics.

Outputs:

Monica Kaminska, 'Restraint under Conditions of Uncertainty: Why the United States Tolerates Cyberattacks', Journal of Cybersecurity 7, no. 1 (9 March 2021), https://doi.org/10.1093/cybsec/tyab008.

Book chapters:

Monica Kaminska, Fabio Cristiano, and Dennis Broeders, 'Limiting Viral Spread: Automated Cyber Operations and the Principles of Distinction And Discrimination in the Grey Zone', in 2021 13th International Conference on Cyber Conflict: Going Viral, ed. Tatiana Jan□árková et al. (Tallinn, Estonia: NATO CCDCOE Publications, 2021).

Monica Kaminska and Jantje Silomon, 'Tackling the Cyber Skills Gap: A Survey of UK Initiatives', in Cyber Security Education: Principles and Policies, ed. Greg Austin (London: Routledge, 2020).

MARTIN KRÄMER

Supervisor: Ivan Flechais,
Department of Computer Science

Empowering Privacy in the Connected Home – Communal Use of Smart Technologies

The latest wave of internet-connected smart home technologies promises convenience and control over a diverse network of different systems, such as appliances, utilities, and entertainment devices.

Striking the balance between convenience and control proves to be a minefield for product designers. Firstly,

the data needs of these technologies amplify concerns over improper data collection and processing practices, highlighting a power imbalance between users and manufacturers. Secondly, convenience and control favour specific practices of use that manifest in related power differentials among household members. Additionally, devices are sometimes utilised for coercive control or domestic abuse.

These are issues of information and interpersonal privacy that surface in the home. However, due to the rapid evolution of technology, the nature of these issues remains under-explored. To fill in this research gap, we ask: `\rqmain' Because privacy is a concept that invites many different definitions and interpretations, the thesis adopts an exploratory and inductive approach. It approaches the overarching research question in four steps: (1) 34 semi-structured interviews inquiring people about their internet-connected and smart device usage practices; (2) a six-month ethnographic study of six households' experiences with smart home devices; (3) a conceptual framework to position emerging findings for research and design; and (4) four case studies that demonstrate the applicability of this framework to privacy in smart homes. Inductive thematic analysis of interview data provided insights into the ways in which technology use in the home was communal. Building on these insights and their relationship to privacy, we chose a grounded theory approach to analyse and present our ethnographic data. Sensitising concepts from ethnomethodology provided focus and perspectives on the establishment of communal use. Key findings include (1) fluid divisions of labour (planned and unplanned) that contributed to the construction of roles with respect to devices; (2) the ways in which household members' interactions contribute to a sense of normalcy (e.g. appropriate use) and to the management of relationships inside and outside the home; (3) that household members sometimes articulated this normalcy in rules to highlight expectations of use and everyday considerations of privacy. We used conceptual framework analysis to link these insights with salient concepts from existing literature on privacy for smart technologies. The framework offered an additional agentic perspective and sensitising concepts to inform innovation in research and design. The case studies drew on the framework to discuss strengths and weaknesses of research contributions and policy initiatives.

The insights gained from the study offer implications for data protection regulations along with academic debates on interpersonal power imbalances in the home.

Bio

Martin is a final year doctoral student hosted at the Department of Computer Science, University of Oxford and supervised by Prof. Ivan Flechais and Dr. Helena Webb. Martin has spent the past 10 years researching and working at the intersection of people and technology. Not only does he care deeply about the ways in which technology can improve lives, but is also critically aware of the societal impacts this can bring about.

Martin's doctoral research concerns privacy in human-centred-computing. His thesis explores, unpacks, and designs for the social organisation of smart device use in the home to empower privacy in familial (communal) settings. To this end, he has worked with members of the public in ethnographic, interview, and co-design studies. His main contribution is a framework that serves to empower people through product innovation by sensitising designers and researchers to the ethical challenges of communal use and privacy in the home.

Publications:

Chalhoub, George, Martin J. Kraemer, Norbert Nthala, and Ivan Flechais. ""It did not give me an option to decline": A Longitudinal Analysis of the User Experience of Security and Privacy in Smart Home Products." In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1-16. 2021.

Martin J. Kraemer. "Beyond the Individual: Exploring Data Protection by Design in Connected Communal Spaces." In 2020 {USENIX} Conference on Privacy Engineering Practice and Respect ({PEPR} 20). 2020.

Seymour, William, Martin J. Kraemer, Reuben Binns, and Max Van Kleek. "Informing the design of privacy-empowering tools for the connected home." In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp. 1-14. 2020.

Martin J. Kraemer, Ivan Flechais, and Helena Webb. "Exploring Communal Technology Use in the Home." In Proceedings of the Halfway to the Future Symposium 2019, pp. 1-8. 2019.

Martin J. Kraemer,, and Ivan Flechais. "Researching privacy in smart homes: A roadmap of future directions and research methods." In Living in the Internet of Things: Cybersecurity of the IoT-2018, pp. 1-10. IET, 2018.

DENNIS MALLIOURIS

Supervisor: Andrew Simpson,
Department of Computer Science

Finance & Cyber Security: Uncovering Underlying and Consequential Costs of Security Breaches and Investments

Severe security breaches can be costly for publicly listed firms. In order to deploy scarce resources in the most efficient way, information security decision makers in such firms need to conduct holistic economic cost-benefit analyses that take into account all costs associated with security breaches and strategic investments in countermeasures. In addition to obvious direct costs following security breaches and investments, there are also multiple indirect underlying and consequential costs introduced or elevated by such events. In this dissertation, four novel cost elements affecting firms' financing costs that are prone to be overlooked in basic cost-benefit analyses are discussed. Over the course of four empirical studies, the effect of security investments on firm value (via abnormal returns) as well as the effect

of security breaches on cost of equity (via systematic risk), cost of debt (via credit default swap spreads), and information asymmetries (via insider trading-induced abnormal returns) are analysed. These studies reveal four cost elements that can all be characterised as indirect underlying/consequential financial costs, which ought to be considered in holistic economic information security cost analyses.

Analysing a US-centric sample of 202 severe security breaches between 2005 and 2019 revealed that severe security breaches are associated with significantly positive increases of insider trading abnormal returns (hence information asymmetries) and systematic risk (hence cost of equity). However, the examination of CDS spread changes following the same severe breach announcements indicated that CDS spreads (hence default spreads and by implication cost of debt) do not increase significantly in the immediate days surrounding the event, suggesting that costs of debt are subject to different dynamics.

Furthermore, assessing an international sample of 221 security investments, the dissertation reveals positive and negative effects on firms' market value following upon small-scale and large-scale security certificate investments, respectively. These market reactions constitute secondary benefits and underlying/consequential costs that also ought to be taken into account for the purpose of expected net benefit of information security (ENBIS) analyses. Collectively, these insights suggest highly relevant information security-induced implications for cost of capital that need to be recognised in ENBIS calculations.

Given the absence of a formal definition of the cost type underlying/consequential costs and a framework or model to facilitate assessments of such costs, the dissertation also provides a novel nuanced definition of indirect costs of information security and the Iceberg Model of Information Security Costs which will help guide further academic research and practitioners' investment endeavours.

Bio

Dennis started the Oxford CDT programme in 2017. He has an MRes in Management (Distinction) for which he studied at University College London and London Business School. He was awarded the SoM Full Scholarship for the duration of his studies. Priorly, he graduated first in his class with an MSc in Management & Finance from UCL. Dennis also obtained a BA in Management (first-class), has a certificate in financial valuation from Oxford's Saïd Business School (OCVP), and is a qualified management accountant (IHK). Dennis worked on multidisciplinary in-house consultancy projects at Siemens in Germany and the UK, at a technology-driven hedge fund, and in financial research & valuation. Dennis grew up multilingually and speaks English, German, Greek, and French. He represented Oxford University in regional and national competitions, and New College on university-level, in multiple sports. He gratefully acknowledges CDT/EPSRC funding, New College's 1379 Society Old Members Scholarship, New College's Sporting and Cultural Award, and a Willis Towers Watson grant. His research at the CDT explored financial, strategic, and organisational implications of cyber security for publicly listed firms. His research projects analysed underlying and consequential costs of security breaches and information security investments, focusing on firms' financing costs. Dennis submitted his dissertation in June 2021 and will be working in Finance.

Publications:

D.D. Malliouris & A.C. Simpson (2019). The stock market impact of information security investments: the case of security standards. In: The 2019 Workshop on the Economics of Information Security (WEIS 2019).

D.D. Malliouris & A.C. Simpson (2020). Underlying and consequential costs of cyber security breaches: changes in systematic risk. In: The 2020 Workshop on the Economics of Information Security (WEIS 2020).

D.D. Malliouris, A.T. Vermorken & M.A.M. Vermorken (2020). Aggregate insider trading and future market returns in the United States, Europe, and Asia. International Journal of Finance & Economics (in press).

JAMES PAVUR

Supervisor: Ivan Martinovic,
Department of Computer Science

On Space Cyber-Security

This project focuses on cyber-security concerns for satellite systems. On going experimental research includes the investigation of privacy and security properties for modern satellite broadband connections over the Digital Video Broadcasting for Satellite (DVB-S) protocol, and the integrity and authenticity of space situational awareness data for flight control and orbit determination. The project also considers the strategic and political effects of Cyber-ASAT (Anti-Satellite Weapon) capabilities. Longer term, the thesis will focus on providing core security principles and best practices to enable the secure operation of critical space missions. These principles will be derived from experimental research and strategic analysis.

Bio

James hails from Atlanta, Georgia (USA) and holds a BSFS in Science, Technology, and International Affairs from Georgetown University in Washington, DC. He is at Oxford on a Rhodes Scholarship (Georgia and Wolfson, 17). His thesis revolves around the security and privacy aspects of satellites and space-based systems with his most recent research focusing on satellite telecommunications and broadband services.

He has dabbled in cybersecurity through a variety of professional experiences – including functioning as the principle cyber decision maker at a 500 employee non-profit (Students of Georgetown Incorporated). His internship experiences include working as a Reverse Engineer for Embedded Systems at Booz Allen Hamilton, auditing building control and SCADA systems as a contractor for the US General Services Administration, and investigating computer crimes with the US Postal Service's Office of the Inspector General. He has also contributed to telecommunications and

privacy policy research through Georgetown's Software and Security Engineering Research Center.

His language of choice is python, although (with generous use of Google) he is also proficient in C/C++, JavaScript, C#, PHP, and Visual Basic.He enjoys hackathons and CTF competitions, collecting (and sometimes consuming) tea, flying kites, and pretending he knows how to play squash.

Publications:

James Pavur, Daniel Moser, Vincent Lenders, and Ivan Martinovic. 2019. Secrets in the Sky: On Privacy and Infrastructure Security in DVB-S Satellite Broadband. In WiSec '19: Conference on Security and Privacy in Wireless and Mobile Networks, May 15– 17, 2019, Miami, FL, USA. ACM, New York, NY, USA. https://doi.org/10.1145/3317549.3323418

James Pavur and Ivan Martinovic. 2019. The Cyber-ASAT: On the Impact of Cyber Weapons in Outer Space. In 2019 11th International Conference on Cyber Conflict - Silent Battle, May 28-31, 2019, Tallinn, Estonia. NATO CCD COE Publications, Tallinn, Estonia.

James Pavur. 2018. Cyber Security and AI (Seminar). Rhodes Artificial Intelligence Laboratory - Speaker Series. November 07, 2018, Oxford, United Kingdom.

MICHAL PISKOZUB

Supervisor: Ivan Martinovic, Department of Computer Science

Network traffic behaviour profiling

Nowadays, computer networks have become incredibly complex due to the evolution of online services and the rapid growth of the number of smart devices such as smartphones, tablets and laptops. Most of users' information, even the most sensitive ones, are transmitted over the Internet. Unfortunately, due to this phenomenon we also see an increasing interest of malware developers who are able to find and exploit novel vulnerabilities in network devices to carry out their malicious intents. To tackle these threats, network analysts should be aided with advanced techniques to identify malicious traffic in order to guarantee the security of networks.

In this thesis, we aim to reduce the asymmetric advantage of attackers by examining malware detection and classification using flow-level network traffic. Our methods explore the ability to extract network behaviours generated by malware. We further evaluate the challenge of working with limited amount of data offered by flows to detect and classify network traffic of malware. Malicious flows are intertwined with benign ones originating from a production network to simulate the real-world settings. We gather one of the largest network flow datasets of malware in order to evaluate our proposals and show that we can detect unseen malware variants.

Moreover, we explore the behaviour profiling of network hosts in order to identify them on large networks. We extract unique behaviours and show that we can work only with the amount of information exchanged by hosts in order to successfully extract their unique behaviours and hence distinguish them from others. We show that while such an approach could be used for maintenance of networks, it may also be employed as an attack against network-based moving target defence (NMTD) systems, which is followed by countermeasures and guidelines to avoid such scenarios.

Finally, we propose a novel method of storing network flow data in a domain specific binary file format, which is motivated by the lack of sufficient methods to process large-scale network data on the order of billions of flows. The binary format makes the analyses of methods in this thesis possible, especially when working with the University of Oxford dataset, which contains more than 181 billion flows. We show that our binary format improves the state of the art in terms of storage, while offering faster data processing techniques.

Bio

Michal has been interested in Computer Science (and technology related topics) since his childhood friend introduced him to it at the age of 7. He continued this passion by doing BSc in Computer Science at King's College London and MSc in Computer Science at the University of Oxford. In the Cyber Security CDT in Oxford he completed two projects titled: On The Way To Adaptive Honeypots, and Dynamic Re-Planning For Cyber-Physical Situational Awareness.

Publications

Michal Piskozub, Riccardo Spolaor, and Ivan Martinovic. 2019. MalAlert: Detecting Malware in Large-Scale Network Traffic Using Statistical Features. SIGMETRICS Perform. Eval. Rev. 46, 3.

Michal Piskozub, Riccardo Spolaor, Mauro Conti, and Ivan Martinovic. 2019. On the Resilience of Network-based Moving Target Defense Techniques Against Host Profiling Attacks. In Proceedings of the 6th ACM Workshop on Moving Target Defense (MTD '19).

Mini-Project: On the Way to Adaptive Honeypots

Mini-Project: Dynamic Re-Planning For Cyber-Physical Situational Awareness

ARIANNA SCHULER SCOTT

Supervisors: Michael Goldsmith (Department of Computer Science), Helena Webb (Department of Computer Science) and Harriet Teare (Centre for Health and Law in Emerging Technologies)

Dynamic Consent: a Mechanism for Privacy Control

My research emphasises that the burden of communication must lie with those collecting information (data controllers), rather than those who are asked to provide personal information (data subjects). Informed consent is used in research to protect the members of the public who

agree to take part from being harmed, and hold researchers accountable. Such harms were originally physical (e.g., medical procedures), but now data harms (e.g., identity theft and fraud) must be considered. Two issues with informed consent processes in research are that the process overburdens the participant with information, and introduces consent fatigue which leads to shortcuts being taken to shorten the activity (e.g., saying yes to all options without reading what those options are).

Dynamic consent has been proposed as one way to avoid these problems because it focuses on revocation and engagement: building mechanisms for people to take back their consent at any time, and prioritising communication as a way to inform participants throughout the project. I worked with an online research platform that implemented dynamic consent. I carried out a service evaluation to see if participants were informed and engaged, proposed changes to improve communication, and carried out an intervention study to measure the effect of these changes. I found that communicating relevant information in a way that participants understood was more likely to retain participation over time and created a tool-kit to support others looking to make similar changes.

My work was done within a specific use case but has wider implications. Cybersecurity literature often prescribes user education as a way to mitigate organisational risk, but there are few guidelines as to what such an education looks like, or examples of success. I argue that seeking user input and feeding back on what this input has helped to achieve, or engaging with users is underexplored. Engagement is important but it is empirically underdeveloped - there are few intervention studies demonstrating or evaluating the importance of communication in online research platforms. My work demonstrates that engaging data-practices offer a return on the (seemingly) high administrative and organisational investment required to facilitate such communication (e.g., hosting workshops and coordinating feedback). I gathered conflicting accounts of engagement within my collaborators' study: participants reported very little understanding of overall project aims and took part in good faith while the research team reported that the study design had a strong participant-centric focus. I co-designed a solution to this with both groups, consulting participants and implementing changes as part of a software development cycle with the researchers.

Bio

I was hosted within the Cyber Analytics Group under Professor Sadie Creese, supervised by Professor Michael Goldsmith. My work was co-supervised by Dr Helena Webb, whose focus lies with Human Centred Interaction in the Department of Computer Science (now at the University of Nottingham), and Dr Harriet Teare, based in the Centre for Health and Law in Emerging Technologies which is in the Faculty of Law (now Research Leader at RAND Europe). My DPhil work focused on the conversations medical researchers were having with research participants who had signed up to their research platform. Every day, I am frustrated by the number of websites and applications that pepper me with cookie notices, asking me for my consent. I am overloaded with information and entirely ill-equipped to make an informed choice. The conversations that organisations have with users about the information collected about them and how this data is used needs to change. I found that, in research, the point at which people are asked for their consent (the start of a project) is also an excellent time to ask them how they want their information to be used, and to show them what is being done with the information that others have already provided. There must be ways for people providing data to change their mind, even if it is unlikely that they will do so - such options demonstrate true freedom of choice rather than consent as a tick box exercise. This approach is something that I am using in my post-DPhil engagement, and will take into every project I work on in the future.

I have played an active role within the CDT since joining. In my first year I took three groups of students to compete in the Cyber 9/12 cyber policy competition hosted by the Atlantic Council in Geneva (mixing up cohorts, of course!). The following year, I led and ran the UK's inaugural event. I took CDT friends with me to this event and we organised the best cyber crisis simulation I have ever been to. Unfortunately, I failed to protect them from bullying and abuse of power in the (unpaid) workplace. Since then, I have done my best to serve as an example of positivity, empathy and a can-do attitude. I have written, organised and edited the yearbook - a project that has gone from strength to strength, as you see. Academia can be an incredibly difficult place to build community, as my experience has been that it showcases, celebrates and rewards individualism. As a direct result of this, I have led, supported and encouraged my friends and peers within the CDT in getting together, talking and sharing ideas through conferences, workshops and social events. The work that I have been privileged to be able to spend my time studying, I have folded back into CDT modules and public engagement activities. I am incredibly grateful for my time here and the connections I have made. I very much look forward to celebrating the success of everyone I have been able to meet and learn from.

Publications

Schuler Scott, A., Goldsmith, M., Teare, H., Webb, H. & Creese, S., 2019. Why We Trust Dynamic Consent to Deliver on Privacy. In Trust Management XIII: 13th IFIP WG 11.11 International Conference (IFIPTM) Proceedings, vol. 563, p. 28. Springer Nature.

Schuler Scott, A., Goldsmith, M. and Teare, H., 2018. Wider Research Applications of Dynamic Consent. In IFIP International Summer School on Privacy and Identity Management Proceedings, pp. 114-120. Springer, Cham.

Schuler Scott, A., Goldsmith, M., Teare, H., Creese, S. and Kaye, J., 2018. Dynamic Consent in Cybersecurity for Health. In Int'l Conf. Health Informatics and Medical Systems (HIMS'18). CSREA Press.

Talks

(2020) "RE: I've been forced to sign this and I am not happy", addressing the upheaval in European data-handling (GDPR) and its impact on privacy awareness.

(2020) "Protecting ourselves online: Why we shouldn't take cookies from strangers", an interactive public engagement exercise focused on data use and institutional responsibility.

(2020) "Designing responsible, participatory research", an internal talk to CDT students about designing and implementing co-design methods (inclusive and responsive practice).

(2019) Developed, coordinated and ran "Cybersecurity in Context", a hybrid module designed to teach technical skills and explore causes and responses to cybersecurity events. Lectures, discussion, capture the flag (CTF) exercises and a cyber crisis simulation were used to explore computer forensics, binary exploits, web attacks.

Research impact

(2020) Protecting User Data as a Software Developer, Computer Science

Department, UNIQ Summer School. 2 hours, online, 30 sixth-form students.

(2019) Cybersecurity in Context, Centre for Doctoral Training in Cyber Security. 4 days, in-person, 13 PhD students.

(2018) Cyber Crisis Simulator, Centre for Doctoral Training in Cyber Security. 2 days, in-person, 13 first-year PhD students.

WILLIAM SEYMOUR

Supervisor: Max van Kleek, Department of Computer Science

Re-Thinking Smartness: Designing More Ethical Connected Devices for the Home

Modern smart devices are capable of incredible things: making life easier, more enjoyable, and more secure. But this 'smartness' often comes at the cost of devices harvesting data from the home, constraining how we use them, and changing the ways we relate to each other. My thesis explores what it means for a device to be smart, the ethical concerns that smartness causes, and ways that we might rethink smartness to better support people's needs and values.

In order to better understand the problem, my thesis begins by exploring people's perceptions of what smartness is, and how it manifests across a variety of contexts in the home. Doing so identified concomitant ethical concerns, such as privacy and autonomy, and highlighted the similarities and differences in how these operate across devices. I followed up by addressing two of the identified problems in greater detail, starting with a six-week technology probe deployment designed to give people control over their connected devices by visualising and constraining data flows. Users of the probe found their privacy preferences shifted over the course of the study and showed how, when given the right resources, people can learn and come together to solve privacy problems in the home. My thesis then explores social concerns around devices with voice interfaces (e.g. Alexa), with a survey exploring correlations between trust, anthropomorphism, and relationship development with voice assistants. It shows how people develop relationships with social devices in a similar manner to those between people, raising questions about the potential for social interaction modalities to be used to manipulate. Finally, I bring these lines of enquiry together by proposing the concept of respect as a lens for using standards of interpersonal interaction to evaluate interactions with smart devices. Practical and theoretical perspectives on respectful behaviour from a variety of disciplines are used to link the behaviours of smart devices to previous work on moral theory, agency, social hierarchies, and oppression, laying out how future devices might mitigate ethical problems whilst promoting human development and flourishing.

Bio

William's research focuses on ethical human-computer interaction in the smart home. With a background in computer science, he frequently designs and develops prototypes for use in research experiments. His work also includes the use of speculative and fictional design to probe beyond the edges of what is possible with today's technology.

Publications

Informing the Design of Privacy-Empowering Tools for the Connected Home. W. Seymour, M. J. Kraemer, R. Binns, and M. Van Kleek. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems.

Strangers in the Room: Unpacking Perceptions of 'Smartness' and Related Ethical Concerns in the Home. W. Seymour, R. Binns, Petr Slovak, M. Van Kleek, and N. Shadbolt. Proceedings of the 2020 ACM Conference on Designing Interactive Systems.

Does Siri Have a Soul? Exploring Voice Assistants Through Shinto Design Fictions. W. Seymour, and M. Van Kleek. Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (alt.CHI).

Privacy Therapy with Aretha: What if your firewall could talk? W. Seymour. Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (Graduate winner, SIGCHI Student Research Competition)

IoT Refine: Making Smart Home Devices Accountable for Their Data Harvesting Practices. M. Van Kleek, W. Seymour, R. Binns, J. Zhao, D. Karandikar, and N. Shadbolt. 2019. In Proceedings of Living in the IoT 2019.

Aretha: A Respectful Voice Assistant for the Smart Home. W. Seymour, M. Van Kleek, R. Binns, and N. Shadbolt. 2019. In Proceedings of Living in the IoT 2019

VALENTIN WEBER

Supervisors: Lucas Kello, Department of Politics and International Relations and Joss Wright, Oxford Internet Institute

The Diffusion of Cyber Norms: Technospheres, Sovereignty, and Power

Why have the Chinese and Russian internet sovereignty norms diffused successfully to some countries – thereby undermining the previously dominant US norm of internet freedom – and what are the implications of successful international diffusion of information controls for understandings of sovereignty in International Relations? Starting in the 1980s, internet freedom swept through the world and soon became the leading global cyber norm. This rise was propelled by US goals of increasing economic power and spreading democracy. Many countries that imported internet technologies, saw the economic benefits of the technologies, but with the import of the internet they also adopted loose market and content regulations. However, in the 2000s, the eruption of protests in the Middle East, Latin America, Africa, and across Asia, powered through internet-connected devices, increased the regime threat perception of authoritarian- leaning leaders. As a result,

they started importing technology/techniques associated with internet sovereignty from China and Russia, who had, early-on, built up their information control capabilities. Beijing and Moscow, for their part, have actively promoted internet sovereignty diffusion to geographical areas that they conceived as their spheres of influence – the Commonwealth of Independent States for Russia and the Belt and Road Initiative for China.

This thesis makes several main contributions. It builds a framework to distinguish the three cyber norms it studies: internet freedom, and the Chinese and Russian variants of internet sovereignty. This thesis also devises a neoclassical realist framework of International Relations, which focusses on explanatory variables at the unit level (perception underlying export/import) to explain the international expansion of internet freedom/sovereignty. Additionally, it updates the term of sovereignty to the digital age. Namely, with the diffusion of internet freedom/sovereignty norms, technological spheres of influence emerged, in which the US, China, and Russia have weakened countries' Westphalian sovereignty and shaped their domestic and interdependence sovereignty.

Bio

Valentin is a Research Affiliate with the Centre for Technology and Global Affairs, University of Oxford. Previously, he was an Open Technology Fund Senior Fellow in Information Controls at the Berkman Klein Center for Internet & Society, Harvard University. Valentin is interested in how the cyber domain is changing conflicts and state strategies. His current research focuses on the role of information controls in state strategies.

Outputs

"China's Quest for Foreign Technology: Beyond Espionage." Book Review. Journal of Cyber Policy. April 2021.

"How China's Control of Information is a Cyber Weakness." Lawfare. 12 November 2020.

"Der moderne Überwachungsstaat." In WAS (Bd. 113) – Angst. Leykam Verlag. August 2020.

"Making Sense of Technological Spheres of Influence." Strategic Update. LSE IDEAS. April 2020.

"The Sinicization of Russia's Cyber Sovereignty Model." Net Politics - Council on Foreign Relations. 1 April 2020.

"Studying Information Control Diffusion: An Agenda for Further Research." Open Technology Fund. February 2020.

"Understanding the Global Ramifications of China's Information Controls Model." in Artificial Intelligence, China, Russia, and the Global Order: Technological, Political, Global, and Creative Perspectives. Air University Press. October 2019.

"The Worldwide Web of Chinese and Russian Information Controls." Working Paper. Centre for Technology and Global Affairs. September 2019.

"The Worldwide Web of Chinese and Russian Information Controls." Report. Open Technology Fund. September 2019.

"中国和俄罗斯信息 控制的全球网." Report. Open Technology Fund. September 2019.

"Всемирная паутина российского и китайского контроля за информацией." Report. Open Technology Fund. September 2019.

"Future of Global Competition and Conflict Virtual Think Tank Report." NSI. September 2019.

Vasilis Ververis, Marios Isaakidis, Valentin Weber, and Benjamin Fabian. "Shedding Light on Mobile App Store Censorship." In Proceedings of the 27th Conference on User Modelling, Adaptation and Personalization Adjunct. June 2019. Larnaca, Cyprus. ACM, New York, NY.

Valentin Weber and Vasilis Ververis. "Measuring Censorship on Mobile App Stores." OxPol. 3 May 2019.

"Finding a European Response to Huawei's 5G Ambitions." Norwegian Institute of International Affairs. March 2019.

"A Bold Proposal for Fighting Censorship: Increase the Collateral Damage." Net Politics - Council on Foreign Relations. 31 January 2019.

"Understanding the Global Ramifications of China's Information Controls Model." in AI, China, Russia, and the Global Order: Technological, Political, Global, and Creative Perspectives. White Paper for the Joint Chiefs of Staff. January 2019.

"Linking Cyber Strategy with Grand Strategy: The Case of the United States." Journal of Cyber Policy. 2018.

"States and Their Proxies in Cyber Operations." Lawfare. 15 May 2018.

"The Rise of China's Security-Industrial Complex." Net Politics - Council on Foreign Relations. 17 July 2018.

"Why China's Internet Censorship Model Will Prevail Over Russia's." Net Politics - Council on Foreign Relations. 12 December 2017.

TINA WU

Supervisor: Andrew Martin,
Department of Computer Science

Digital Forensic Investigation of IoT Devices: Tools and Methods

Consumer IoT devices are becoming omnipresent in our homes, their continuous interactions create a rich source of evidence for forensic investigators. We have already seen devices admitted as evidence in criminal cases. However, the investigation of these cases is still in its infancy, with methods and tools only compatible with specific manufacturers. In this thesis we aim to improve the IoT forensic process by developing generic approaches, methods, and tools.

As a first step we undertook an online survey of the digital forensic community (n = 70) on their interpretation of various definitions and a roadmap for future research. This resulted in the development of a concise definition of IoT forensics and demonstrated that research should focus on IoT forensic tools and data acquisition. Next, to identify the gaps in IoT forensic tool development we examined academic publications between 2014-2019 for research-based digital forensic tools. This identified the need for an IoT device identification approach and a tool to aid the analysis of IoT network traffic.

To fulfil the requirement of an identification approach, we proposed our machine learning approach that uses packet header and behavioural features combined with feature selection methods to develop an approach that can accurately and efficiently identify the type of IoT devices on the network.

We then focused on developing a generic acquisition method and proposed the use of Bluetooth Low Energy (BLE), a common protocol used by many smart healthcare devices, and integrated in many other consumer IoT-products. We demonstrated the feasibility of this method

by acquiring health / medical related traces from smart healthcare devices.

In our final study, to fulfil the requirement for a IoT network analysis tool capable of automating the network analysis process we presented and evaluated our tool. Additionally, we analysed the network traffic metadata from 32 IoT consumer devices for multiple forensic purposes. Results showed the following: remote access is unavailable on most devices, the Shannon entropy test is a useful pre-test in identifying unencrypted content, many devices used encryption, but smart cameras are prone to sending unencrypted content when detecting motion or updating and finally we found the majority of IoT traffic was sent to the U.S.

Bio

Tina currently works as a Cyber Security Researcher at Ofgem focused on security research in industrial control systems. She completed her MSc in Forensic Computing and Security at the University of Derby. Previously she worked as an Industrial Research Fellow at the University of Plymouth focused on penetration testing of maritime systems. Before her PhD she worked at Airbus Group as a Research Engineer focusing on research in cyber security and forensics in industrial control systems. Her research interests are in Forensics and monitoring of industrial control systems with a focus on live memory forensics, novel attack detection methods, malware analysis, side channel attacks and the Internet of Things (IoT).

Publications

T. Wu and A. Martin, "Bluetooth low energy used for memory acquisition from smart health care devices," in 2018 17th IEEE international conference on trust, security and privacy in computing and communications/ 12th ieee international conference on big data science and engineering (trustcom/bigdatase), 2018, pp. 1256-1261.

T. Wu and J. Nurse, "Exploring the use of PLC debugging tools for digital forensic investigations on SCADA systems," Journal of digital forensics, security and law, vol. 10, pp. 79-96, 2016.

T. Wu, F. Breitinger, and I. Baggili, "IoT ignorance is digital forensics research bliss: a survey to understand iot forensics definitions, challenges and future research directions," in Proceedings of the 14th international conference on availability, reliability and security, 2019, p. 1–15.

T. Wu, F. Breitinger, and S. O'Shaughnessy, "Digital forensic tools: recent advances and enhancing the status quo," Forensic science international: digital investigation, vol. 34, p. 300999, 2020.

T. Wu, F. Breitinger, and S. Niemann, "IoT network traffic analysis: opportunities and challenges for forensic investigators?," DFRWS APAC 2021

ADAM ZIBAK

Supervisor: Andrew Simpson,
Department of Computer Science

A Success Model for Cyber Threat Intelligence Platforms

Cyber security information sharing is increasingly playing an important role in improving the organisational and national overall security posture. Efforts in the public and private sectors to foster information sharing initiatives have intensified in recent years resulting in the current complex constellation of sharing organisations, forums, platforms and tools, including threat intelligence management solutions. However, despite the growing interest, a number of questions pertaining to the nature of cyber security information sharing and its success remain unanswered.

This dissertation uses mixed methods to consider How can cyber security information sharing be made more successful? In answering this question we draw on practitioners' experiences and collect empirical data to ensure that our conclusions can contribute to practice and theory. As with all evaluation efforts, this dissertation relies on the well-established premise that the success of a process cannot be reasonably assessed without a coherent understanding of the goals it is designed to attain.

Our first contribution therefore provides a nuanced conceptualisation of cyber security information sharing. This includes surveying practitioners for their understandings and attitudes towards different aspects of the process and what it is trying to achieve. It also addresses the disparity among them when it comes to distinguishing between the different forms information sharing can take by proposing a high-level classification of these forms and their objectives.

Our second contribution examines the extent to which the benefits and barriers to successful cyber security information sharing mentioned in the literature are reflected in the attitudes of cyber security professionals. A categorisation of the benefits and barriers is introduced followed by a self-administered survey of practitioners. The results show a degree of inconsistency between theory and practice. It also highlights quality issues as a primary concern for practitioners and a hindrance impacting the success of these efforts.

Our third and main contribution is two-fold. First, we investigate the quality problem further in the context of threat intelligence platforms. Through a systematic review of the literature and a modified Delphi study we identify a set of quality dimensions practitioners employ in determining the quality of the platform's content. Second, we draw together the previous findings, as well as the theories and practices of information systems literature to develop and test a holistic success model --- a framework

for understanding and measuring the key success factors and their interrelationships --- for threat intelligence management platforms.

Bio

After obtaining his bachelor's degree in Computer Science, Adam completed the MSc in IT Law and Management from King's College London (Distinction) where he gained a grounding in the areas of Law which are most relevant to Information Technology. Prior to joining the CDT, Adam worked as an open-source intelligence researcher and Arabic linguist for the International Centre for Security Analysis at King's College London.

Adam's primary research interest is cyber threat intelligence, with an emphasis on evaluating the quality of current threat sharing initiatives and platforms. Adam was the President of the Oxford University Strategic Studies Group (OUSSG) for the 2018/19 academic year.

Outputs

Adam Zibak and Andrew Simpson. 2018. Can We Evaluate the Effectiveness of Cyber Security Information Sharing Efforts?. In Proceedings of the 2018 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (CyberSA 18). IEEE.

Adam Zibak and Andrew Simpson. 2019. Towards Better Understanding of Cyber Security Information Sharing. In Proceedings of the 2019 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (CyberSA 19). IEEE.

Adam Zibak and Andrew Simpson. 2019. Cyber Threat Information Sharing: Perceived Benefits and Barriers. In Proceedings of The 14th International Conference on Availability, Reliability and Security (ARES 2019). ACM.

Adam Zibak, Clemens Sauerwein and Andrew Simpson. Towards a Cyber Threat Intelligence Quality Model for Research and Practice. Submitted to Digital Threats: Research and Practice.

Adam Zibak, Clemens Sauerwein and Andrew Simpson. A Success Model for Threat Intelligence Platforms. Submitted to Computers & Security.

M-P1: When You Lose, Do Not Lose the Lesson: A Case Study on The Sony Pictures Entertainment 2014 Data Breach

M-P2: Conceptual Pipelines for the Access-data Centric Approach to Open Source Intelligence (in collaboration with Horus Security Consultancy)

Goodbye Katherine

In May 2021, our Industrial Liaison Officer, Katherine Fletcher moved to a new exciting role with the Nuffield Department for Public Health. Katherine's support for the CDT and wider Cyber Security Oxford Network helped shape it into the diverse group it now represents. Katherine's enthusiasm for seeking out new network members and the boundless energy in supporting our students will be greatly missed. During her time with us, Katherine supported numerous funding bids, Cyber Security Seminars and DPhil student. She will be a great loss to the team, but we are sure she will be fantastic in her new role.

Joint Student conference

The first cross-CDT symposium took place in February, a collaboration between the Cyber Security, Autonomous Intelligent Machines (AIMS) and Heath Data Science CDTs brought together researchers presenting their work to others outside of their traditional disciplines. Topics included self-supervised learning in medical imaging and the control of information at a state level and its impact on cyber policy. A panel of alumni also shared their perspectives on career development and life after graduation. Events such as the symposium help to connect students with different academic disciplines across the university, enriching the CDT network, building sharing skills and knowledge. Further details can be found here. https://padlet.com/davidhobbs1/c4fql6xdl1s4aphi

CDT podcast update

Arianna Schuler Scott (CDT16) and Claudine Tinsman (CDT18) have been working on a public engagement project: "Proving the Negative", a CDT podcast focussing on the research being undertaken within the CDT and the relevance of that work to society, industry and recent world events.

The aim of this series is to showcase CDT research and share with listeners something new about cybersecurity – the target audience is enthusiasts. This means that it has been produced to be accessible to those who are not doing research. This podcast has been designed to serve several functions – for CDT students, it provides an opportunity to practice interview and media engagement skills. Potential employers might use the series to identify emerging experts and fellow academics may find this a positive example of public engagement to follow.

On the show, co-hosts Arianna and Claudine interview several members of the CDT about their research, breaking these ideas down for a general audience. Interviews have been carried out and PTN is currently in post-production, with contributions from Mary Bispham (CDT14), Sean Sirur (CDT17), Fatima Zahrah (CDT18), Anirudh Ekambaranathan (CDT18), Valentin Weber (CDT15), Julia Slupska (CDT18) and Marcel Stolz (CDT16).

The first episodes are due to be released at the end of August,
so feel free to follow @PTNegative1 on Twitter and watch this space!

Cyber-enabled Foreign Election Interference: The Old, The New, and The Ugly

Arthur Laudrain, CDT18

"Nothing will be left unanswered"[1]. It is in these resolute words that French President François Hollande reacted to serious suspicions of foreign meddling in the aftermath of Emmanuel Macron's election in April 2017. A few days earlier, foreign hackers had exfiltrated and spread online internal documents from the Macron campaign. It quickly appeared the attackers had mixed fabricated documents designed to stir outcry with genuine campaign material. This was reminiscent of the less successful fate of the US Democratic Party leadership during the 2016 U.S. presidential campaign, which also experienced a hack-and-leak operation, also known as doxing.

Although these events diverged in their execution and context, they can both be characterised as foreign electoral interference (FEI). There is now little doubt that the Russian state orchestrated the Democratic National Committee (DNC) leaks in the US[2]. Recent indictments unsealed by the US Department of Justice even suggest that the same team of Russian military intelligence operators in the GRU contributed to doxing operations against both campaigns[3].

These foreign interference events were serious hostile acts. The precise impact on the political outcome of these democratic processes remains disputed[4], but is undeniable in its existence. Recent research suggests Hillary Clinton may have lost up to 75 electoral college votes due to Russian interference. Indeed, the stolen emails showed that the party leadership secretly favoured her over Bernie Sanders, incensing his supporters[5]. In France, the Macron leaks didn't have as much effect for both structural and situational reasons[6], yet they undermined trust in the electoral process. FEI is not a new phenomenon and has been a feature of international relations, typically conducted by great powers of the time, from France and the UK in the XIXth century to the US and the USSR during the Cold War[7]. But the implications of these contemporary practices have lead observers asking whether democracies could even uphold their electoral sovereignty in the XXIst century[8]. Such are the stakes of the issue at hand.

Yet, the response of the targeted states has puzzled observers. In the summer of 2016, President Obama threatened President Putin with undefined retaliation if the Russians kept probing electoral systems and party servers[9]. But it is only after Secretary Clinton was defeated that the Administration enacted diplomatic and other sanctions. If -as it appears[10]- the Obama Administration knew about Russian meddling attempts many months before the election was to take place, why didn't it react more forcefully beforehand? Once Donald Trump took office, the U.S. position became erratic, if not contradictory. On the one hand, Trump publicly rebutted the media's allegations that Russia was behind the DNC leaks. On the other hand, the U.S. Cyber Command conducted operations targeting the Russia-based Internet Research Agency on the day

[1]Triggle, "Hollande Vows 'response' to Macron Hack."

[2]Mueller, "Report on the Investigation into Russian Interference in the 2016 Presidential Election."

[3]Department of Justice, "Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace," October 19, 2020.

[4]For further discussion on this point, see Valeriano, Jensen, and Maness, Cyber Strategy.

[5]Levin, "When the Great Power Gets a Vote"; Levin, *Meddling in the Ballot Box*.

[6]Jeangène Vilmer, "The '#Macron Leaks' Operation."

[7]Mohan and Wall, "Foreign Electoral Interference."

[8]Shimer, *Rigged*.

[9]Isikoff and Corn, *Russian Roulette*.

[10]Ibid.

of the Congress mid-term elections, and made it known publicly at a later stage[11].

Across the pond, despite President Hollande's pledge[12], there was no retaliation to the Macron leaks. Emmanuel Macron, in power soon after, adopted what we previously described as a "red phone" policy, consisting of direct and private communication between the French and Russian governments at a high political level[13] Across the Channel, the U.K. government -in the aftermath of Russian meddling in the Brexit campaign- did not "[seek] evidence of successful interference in UK democratic processes"[14].

We approach these FEI attempts as a consequence of the stability-instability paradox[15], incentivising revisionist states to pursue strategic competition against democratic regimes. The paradox creates the conditions for a deterrence failure, inciting hostile powers to encroach on a democracy' sovereignty, below the threshold where nuclear deterrence would be effective[16].

In this short paper, we first review the existing literature on traditional forms of electoral interference before exploring its contemporary developments, as enabled by technology and social media. This will allow us to better conceptualise the phenomenon.

THE OLD: Historical overview, debates, and definitions

Elections are one of the cornerstones of democratic regimes. It is thus not surprising that political science has centred its attention on both the concept and its practices. Major works have tackled how voters inform themselves and decide for whom to vote[17], elections under authoritarian regimes[18], voter representations[19] and coalitions[20]. Interference into elections is probably as old a practice as elections themselves. Intervention can take many forms, from funding political parties to blackmailing politicians or spreading defamatory stories about a candidate. Notorious instances include France's open opposition to George Washington's re-election bid in 1796[21], and Bismarck' secret support for French Republican leader Gambetta in 1877[22]. The Cold War was an ideal scenario for FEI, with the USA and USSR regularly meddling into political processes of their proxies or even of each other. The US sought to undermine political leaders or parties -among others- in Iran, Italy and Guatemala, while the USSR attempted to influence elections mostly in Europe, including West Germany, Finland and France[23].

Theorisation on FEI, old and new, is nascent. Looking at historical instances, Levin finds that foreign powers tend to intervene if 1) the political landscape of the target is polarised, 2) the intervener has internal support by one of the sides (invitation), and 3) the political alternative would be seriously detrimental to the intervening power ("implacable actor")[24]. He also finds that overt interventions are more likely to succeed than covert ones. Although his tentative case-study on the DNC leaks seems to confirm the applicability of his theory to new forms of FEI, his findings are preliminary.

Overall, these historical instances in their great diversity of time, place, and method still have in common one main goal: Undermining the electoral process and legitimacy of the institutions by acting whether in favour or against a specific candidate or party. There are, however, many different definitions of foreign electoral inference and a few notable unsettled debates among scholars. First and foremost, not everyone agrees on the purpose and scope of FEI. Hansen and Lim, for instance, remain obscure on its purpose ("furthering a particular agenda") but only consider voter influence, and would therefore exclude more

[11]Nakashima, "U.S. Cyber Command Operation Disrupted Internet Access of Russian Troll Factory on Day of 2018 Midterms."
[12]Author translation. Triggle, "Hollande Vows 'response' to Macron Hack."
[13]Laudrain, "France's New Offensive Cyber Doctrine."
[14]Intelligence and Security Committee of Parliament, "Russia."
[15]Seabury, *Balance of Power.*
[16]Kissinger, *Nuclear Weapons and Foreign Policy.*
[17]Alvarez, *Information and Elections.*
[18]Gandhi and Lust-Okar, "Elections Under Authoritarianism"; Levitsky and Way, "Elections Without Democracy."
[19]Darcy, *Women, Elections, and Representation.*
[20]Austen-Smith and Banks, "Elections, Coalitions, and Legislative Outcomes."
[21]Mohan and Wall, "Foreign Electoral Interference."
[22]"Bismarck and Gambetta."
[23]Levin, "Partisan Electoral Interventions by the Great Powers."
[24]Levin, *Meddling in the Ballot Box.*

direct actions such as altering ballots[25]. On the other side of the spectrum, some such as Levin, accept both overt and covert actions[26]. The distinction resides in whether the intervening state acknowledges its actions or remains secretive[27]. For instance, under Levin's definition, an official statement by a foreign head of state or representative could qualify as interference. For Shimer however, the most relevant forms of interference are covert actions against votes of succession[28], as seen recently in France and the US. They are, indeed, the most common and regular. We agree that covert forms of interference are the most dangerous and therefore relevant to study: It is harder to fight against something you are not aware of. However, Russian interference in the Brexit referendum has shown that leadership elections are too restricted a scope for the study of contemporary FEI. If anything, "one in a generation" votes such as referendums are more critical by their long-term repercussions and are bound to attract interest from ill-intentioned foreign powers. Last but not least, Bubeck and Marinov insist on the distinction between process and candidate interference, that is whether the interferer seeks to change the electoral rules or the result[29].

We first consider FEI as a subset of influence operations, that is "the deliberate use of information by one party on an adversary population to confuse, mislead and ultimately influence the actions that the targeted population makes."[30] We define FEI as *the intervention of a foreign power into the electoral ecosystem of a country to influence its political environment, up to the election result*. By electoral ecosystem, we mean not only the election process and infrastructure themselves, but also the media environment and public opinion more widely.

THE NEW: A change of tools and tactics enabled by technology

Although the past five years brought FEI back into the spotlight, only a handful of scholars have published substantive work on its contemporary occurrences as of writing these lines. They typically fall into three distinct yet equally valuable categories. First, we find investigatory accounts, typically of the DNC leaks. They provide invaluable empirical data from elite interviews, an intricate knowledge of foreign-policy and national security bureaucracies, and they place these events within the wider historical context of Great Power competition during the Cold War[31]. Unfortunately, little of such task has been conducted on other recent FEI events, and we must rely on short articles and parliamentary reports for

insider information on the French and British governments' response. Second, we find research investigating the spreading mechanisms of disinformation operations[32] and their effects on voters[33]. Finally, we find works of theorisation, sparse but critical in setting-up the sub-field[34]. Hansen and Lim created a Cyber Voter Interference (CVI) model to investigate domestic variables that improve or impair such operations' efficacy[35]. They notably point the lack of deterrence and compellence framework to guide states' response to potential CVI. In the same vein, Wigell theorises disinformation and electoral interference as part of a wider hybrid interference strategy leveraging vulnerabilities inherent to liberal democracies, such as foreign policy restraint, political pluralism, a free press and an open economy[36].

Although most of the existing literature has focused on conceptualising and identifying electoral interference events, IR scholars are yet to fully embrace the topic. Notably, scholars have given little attention to how we can interpret FEI under the lenses of IR theories. Looking at disinformation through structural and rational realism as well as constructivism, Lanoszka argues that mainstream IR theories are generally ill-equipped to tackle the issue[37]. Further, he argues, disinformation campaigns are unable to shift the balance of power because of major obstacles in the target state, notably scepticism for foreign information, prejudices of the targeted population and state countermeasures. However, he joins Hansen and Lim

[25]Hansen and Lim, "Doxing Democracy."

[26]Levin, *Meddling in the Ballot Box*.

[27]Carson, *Secret Wars*.

[28]Shimer, *Rigged*.

[29]Bubeck and Marinov, *Rules and Allies*.

[30]Lin, "Developing Responses to Cyber-Enabled Information Warfare and Influence Operations."

[31] Isikoff and Corn, *Russian Roulette*; Rid, *Active Measures;* Shimer, *Rigged*.

[32]Ferrara, "Disinformation and Social Bot Operations in the Run Up to the 2017 French Presidential Election"; Jeangène Vilmer, "The '#Macron Leaks' Operation."

[33]Bail et al., "Assessing the Russian Internet Research Agency's Impact on the Political Attitudes and Behaviors of American Twitter Users in Late 2017."

[34]Hansen and Lim, "Doxing Democracy"; Morgan, "Fake News, Disinformation, Manipulation and Online Tactics to Undermine Democracy"; Wigell, "Hybrid Interference as a Wedge Strategy."

[35]Hansen and Lim, "Doxing Democracy."

[36]Wigell, "Hybrid Interference as a Wedge Strategy."

[37]Lanoszka, "Disinformation in International Politics."

when he concedes that social media and computational propaganda could create a "window of opportunity" for hostile powers' interference to have the most domestic effects, e.g. damage a candidate's electoral prospects.

THE UGLY: Democracies' vulnerability

This idea of a window of opportunity is key to conceptualise the difference between historical FEI and its contemporary developments, we argue. As Hansen and Lim have theorised, contemporary FEI have three main components: doxing, disinformation, and trolling[38]. All possess characteristics that make them ideal to seize the windows of opportunity in the target state: direct access to individuals, tailored content, and reactive engagement. Overall, Internet connectivity allows for a more direct, tailored, and reactive outreach to the social layer of the targeted community[39].

These characteristics seem to indicate that, although the primary purpose of these new FEI remains the influence of electoral outcomes, another -maybe more insidious- goal is to sow discord among communities, fracture society using existing divisions on sensitive topics (immigration, race, criminality, inequality) and, eventually, delegitimise the democratic process and result, irrespective of political leanings[40]. This can take place outside of an electoral context (see for instance Russia's operations during the #BLM protests[41]). But political tensions are often higher the closer the electoral stakes.

First, social media platforms allow interferers to push stories directly to people's eyes by bypassing the filter of traditional, editorial-based media[42]. It means people have access to content that is potentially less balanced, less prone to journalistic standards, and with little accountability for false or misleading stories. Second, targeted ad algorithms and online communities of values (around hashtags or groups) allow interferers to craft content tailored to the audience they aim to target[43]. Eventually, this can create a bubble that reinforces existing beliefs and drives communities further apart from each other[44]. Third, the nearly instantaneous nature of social media and the shortening news cycle means that interferers can quickly respond to emerging events and seize that opportunity to exploit heightened tensions or rising political scandals[45]. It is true that doxing operations require time, planning and careful execution. But once the information is acquired, spreading it over open or semi-open platforms is easy to do and happens in a matter of hours.

To better capture and reflect these differences, we introduce the concept of cyber-enabled foreign election interference (CYFI). We define CYFI as *the intervention of a foreign power into the electoral ecosystem of a country to 1) polarise society, 2) undermine trust in democracy, and/or 3) influence election results, by exploiting online mass communications' characteristics of direct access, tailored content and reactive engagement.* The four main tools of CYFI are doxing, disinformation, trolling, and direct ballot manipulation.

CONCLUSION

As a means for states to exploit strategic asymmetry between non-democratic and democratic regimes in the service of strategic competition, FEI deserves more attention in the IR sphere. We cannot ignore the consequences of electoral meddling on the foreign policies of affected states, and thus for international relations.

Scholars of both IR and FEI have already made the parallel between the cyberdomain and tanks during WW1[46], and submarines or aircraft-carriers during WW2[47]: New domains provide new tools and allow for new tactics by changing the tempo of operations. Rather than a radically new concept, contemporary forms of FEI are enabled by the evolution of communication and information techniques. In short, they are cyber-enabled. This brings new challenges: The direct access to a wide range of the population, the ability to craft tailored content and to quickly engage communities as political events unfold make these new forms of influence difficult for states to detect and react to.

[38]As Tenove et al. suggest, we could include direct attacks on electoral infrastructures as a separate category.

[39]The social layer of cyberspace, in addition to the physical and social layers, comprises of both cyber and non-cyber personas, someone's online and offline identities.

[40]Morgan, "Fake News, Disinformation, Manipulation and Online Tactics to Undermine Democracy."

[41]Švedkauskas, Sirikupt, and Salzer, "Analysis | Russia's Disinformation Campaigns Are Targeting African Americans."

[42]Wu, *The Attention Merchants.*

[43]Kaakinen et al., "Shared Identity and Shared Information in Social Media."

[44]Spohr, "Fake News and Ideological Polarization"; Cardenal et al., "Echo-Chambers in Online News Consumption."

[45]Rosenberg and Feldman, *No Time To Think*; Leskovec, Backstrom, and Kleinberg, "Meme-Tracking and the Dynamics of the News Cycle."

[46]Kello, *The Virtual Weapon and International Order*; Levin, *Meddling in the Ballot Box.*

[47]Carson, *Secret Wars.*

Dead Man's Switch: Forensic Analysis Of a Nintendo Switch

Frederick Barr-Smith, CDT18

In direct collaboration with the private security consultancy 'Modux', Frederick co-authored the paper "Dead Man's Switch: Forensic Analysis Of a Nintendo Switch". This paper was part of a research contract with the Defence Science and Technology Laboratory's 'Digital Crime Science' research area and provides tools for processing Nintendo Switches discovered at crime scenes. In particular this research involved the exploitation of the Nintendo Switch and the development of open-source tools for automated reverse engineering and analysis of the data to collect forensic evidence. This paper was published in Forensic Science International and was presented virtually at DFRWS EU 2021 and a guest lecture to present this research at Interpol.

Frederick was also lead researcher on "Survivalism: Systematic Analysis of Windows Malware Living-Off-The-Land" a collaboration with Cisco Talos Intelligence, which has been accepted to IEEE Security and Privacy 2021 (with a notably selective 11.38% Acceptance Rate). This paper's lead researcher was Frederick Barr-Smith and it was produced as part of a collaboration with researchers from Cisco Talos Intelligence.

This paper was also supported by the generous provision of data from Google Cloud Security and VirusTotal. This paper consisted of a large-scale analysis of evasive techniques used by malware, with particular attention paid to state-sponsored malware campaigns. As part of this paper, responsible disclosure was ethically conducted with a wide range of anti-virus firms, to correct susceptibility to Living-Off-The-Land techniques discovered in anti-virus detection algorithms.

Oxford University Competitive Computer Security Society

Sebastian Köhler, CDT18

During the COVID-19 pandemic, working from home (WFH) has become the new normal for millions of employees worldwide to ensure social distancing and reduce the spread of the virus. At the same time, the threat landscape changed and remote working has challenged many system administrators who suddenly had to ensure a secure WFH environment. Due to the relevance of this topic, we focused our last year's events on the security issues and challenges that arose during the pandemic.

To ensure employees can access confidential corporate data while guaranteeing the confidentiality and integrity of the connection, Virtual Private Networks (VPNs) have become more important than ever. This is reflected in the enormous increase in VPN traffic. During the first few weeks of the pandemic, VPN traffic in Europe increased by more than 200% [1]. Unfortunately, the configuration of a VPN is complex and small misconfigurations can put the security of the connection at risk. The two cryptographers Ferguson and Schneier concluded in their cryptographic evaluation of IPsec [3], one of the most widely used VPN implementations, that "[the] main criticism of IPsec is its complexity. IPsec contains too many options and too much flexibility; there are often several ways of doing the same or similar things." Besides IPsec, different VPN implementations, such as OpenVPN and Wireguard, exist. However, due to the widespread deployment of IPsec, one of our sessions focused on the security of IPsec and the most common pitfalls that can occur during configuration. The students had to exploit weaknesses and misconfigurations (Aggressive Mode) in the key exchange (IKEv1) and masquerade a legitimate IPsec client to gain access to a fictitious corporate network via the VPN and capture the secret flag.

Another challenge companies faced during the pandemic was the increased risk of phishing and ransomware. According to a report from KPMG [2], cyber criminals exploited the uncertainty and concerns over the pandemic for their phishing campaigns. In our session about reverse engineering, we demonstrated a detailed and step-by-step static analysis of different malware samples. This allowed the attendees to get to know common tools used for reverse engineering and learn helpful tips and tricks.

Besides the practical sessions and exercises, the society successfully competed in multiple online Capture-the-Flag challenges. Thanks to great teamwork, we were able to complete two events as one of the top 30 teams out of hundreds of participating teams. For a few months, our CTF team Ox002147 was able to defend its place among the top three CTF teams in the UK. While we, unfortunately, lost our place in the top three, we are still ranked among the top 10. With the start of the new academic year, we will continue our practical sessions and prepare to defend our ranking.

As always, we are working tirelessly on new and exciting future events, which will span from secure wireless networks to mobile phone security. We want to take this opportunity to thank the Centre for Doctoral Training in Cyber Security for their continuous support, without which we would not be able to offer such a wide range of events! We cannot wait to see you all again in person! Until then, stay safe and keep hacking!

References

[1] Anja Feldmann et al. "Implications of the COVID-19 Pandemic on the In-ternet Traffic". In: Broadband Coverage in Germany; 15th ITG-Symposium. VDE. 2021, pp. 1–5.

[2] David Ferbrache. The rise of ransomware during COVID-19. https : / / home . kpmg/xx/en/home/insights/2020/05/rise- of-ransomware-during- covid-19.html. 2020.

[3] Niels Ferguson and Bruce Schneier. A cryptographic evaluation of IPsec. 1999.

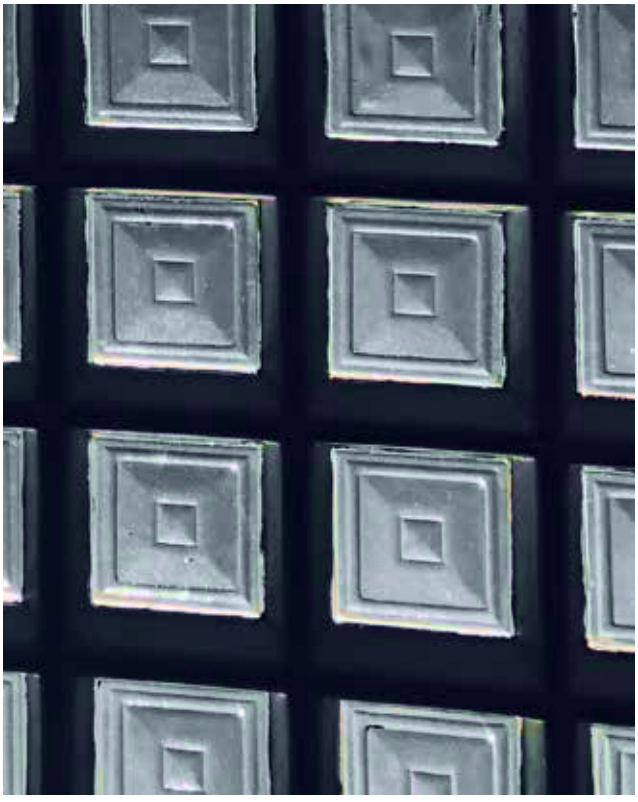

“It did not give me an option to decline”: A Longitudinal Analysis of the User Experience of Security and Privacy in Smart Home Products

George Chalhoub (CDT18), Martin J. Krämer (CDT16) and Ivan Flechais

Despite growing at double-digit rates across the globe, smart home devices still routinely suffer from consumer privacy and security problems. Smart homes today have a dark side, which conjures up images of surveillance cameras and smart speakers that are constantly tracking, listening and watching. Previous work exploring security and privacy in smart homes has mostly consisted of surveys, in-situ design evaluations and interviews.

To address the research gap, we have presented a longitudinal view of smart home security and privacy experiences from the secondary analysis of a six-month ethnographic study of six UK households. We found the experience of managing security and privacy to be inconsistent. We also found that repurposed smart home products introduced negative security and privacy effects.

Methodology

We performed a secondary analysis of a six-month-long ethnographic study of six UK households living in smart homes conducted as part of the “Informing the Future of Data Protection by Design and by Default” project. The study consisted of: (1) planning workshops with participants where they selected smart home products for their home. (2) procuring and providing the chosen smart home products to participants and (3) observing the deployment, installation and use. The data consisted of fieldnotes, photographs, unstructured interviews, and diaries.

Results

Participants expressed concerns over security cameras and smart display cameras, and sometimes taped these up as a result. Similarly, concerns over tracking and microphones were expressed. Privacy concerns were generally managed through consent preferences. However, the experience of consent management was inconsistent. Granting consent was straightforward, while withholding consent caused problems. For example, one household refused to enable ‘Web & Activity’ tracking when setting up their smart speaker. As a result, they weren’t able to play music. Similarly, another household lost the use of their device’s motion detection sensor after using the built-in camera Shutter.Over time, this prompted participants to revisit their preferences and grant consent for services they had rejected.

Furthermore, our results confirm the continued existence

of a well-known problem: password fatigue which was experienced by many participants as they were required to remember an excessive number of passwords as part of their daily routine. We also identified new challenges in access control management and situations where these were not suited to the needs of the households. For instance, while family sharing features were designed into some products, they were not user-friendly and often not used. Instead, participants shared accounts instead of setting up and delegating permissions as they found this much easier.

Implications

Our results suggest that it is the perception of effectiveness of controls that improves the experience of privacy and assurance. Smart home devices must be designed not just to give users controls and choices, but to provide assurance that privacy features are effective.

Our results also show that the life cycle of consent can change over time: users can withhold - grant - revoke - amend consent from ongoing use, new information, or repurposing. We proposed that time and ongoing device use should be considerations for privacy consent management. This would move beyond a model of consent aimed mostly at gaining consent during setup.

Moreover, our results indicate that smart home products were repurposed in five households. Designers should improve their understanding of audiences and contextual uses of smart home products to be able to ground and anticipate how their technologies might be repurposed. This would allow them to accommodate novel applications of smart home technologies and mitigate potential misuse.

Recommendations

We concluded our work with the three recommendations:
First, we recommend that consent management needs to facilitate changes over time, that the experience of withholding consent needs to be improved, and that the experience of making mistakes in consenting to data use should be more forgiving.

We also recommend designers should develop knowledge of the risks and threats of repurposing and improve the transparency of sensitive features (e.g., such as being reminded when sensitive features are enabled).
Finally, we noted that users are more confident in privacy controls that have physical affordances. How to improve the perception of effectiveness of privacy controls that lack physical characteristics is an important research challenge to improving trust.

Acknowledgements

This study was supported by the 2018-2019 Information Commissioner's Office's (ICO) Grants Programme.

Publications arising from work

George Chalhoub, Martin J. Kraemer, Norbert Nthala and Ivan Flechais. "It did not give me an option to decline": A Longitudinal Analysis of the User Experience of Security and Privacy in Smart Home Products. In 2021 CHI Conference on Human Factors in Computing Systems (CHI 2021). ACM. May, 2021.

Re:CONFIGURE – Feminist Action Research in Cybersecurity

Julia Slupska (CDT18)

We present "participatory threat modelling" as a feminist cybersecurity practice which allows technology research to centre traditionally marginalized and excluded experiences. We facilitated a series of community workshops in which we invited participants to define their own cybersecurity threats, implement changes to defend themselves, and reflect on the role cybersecurity plays in their lives. In doing so, we contest both hierarchical approaches to users in cybersecurity—which seek to 'solve' the problems of human behavior—and a tendency in HCI to equate action research with the development of novel technology solutions. Our findings draw highlight barriers to engaging with cybersecurity, the role of personal experiences (for instance of gender, race or sexuality) in shaping this engagement, and the benefits of communal approaches to cybersecurity.

The full Reconfigure report can be found here.:
https://www.oii.ox.ac.uk/reconfigure-report-2020

Attacking cameras using electromagnetic interference

Sebastian Köhler (CDT18)

1. Introduction

Over the last few decades, the underlying architecture of image sensors has experienced a significant shift in technology. Nowadays, two major image sensor architectures exist — Complementary Metal-Oxide Semiconductor (CMOS) and Charge-Coupled Device (CCD) image sensors. Due to the improved semiconductor manufacturing process, the production costs of CMOS image sensors have decreased immensely, while the performance of the sensors increased. As a result, CMOS image sensors have almost entirely replaced CCD image sensors in consumer devices, such as mobile and IoT devices, autonomous vehicles, retail, and surveillance.

However, due to their excellent photometric performance and their capability to capture frames without geometric distortions, CCD image sensors are still used in specific professional and scientific applications [1]. The fields of application range from ground and space astronomy [2], [4] over microscopy [6], industrial automation [13] to military surveillance and defense systems [3], [11]. With the increasing usage of intelligent systems that make safety-critical decisions based on the trusted captured image information, the integrity of the camera inputs has become crucial. Various attacks against camera-based systems compromising the integrity have been demonstrated in academic literature [7], [8], [10], [19]. Since image sensors are optical sensors, the most obvious attack vector is the injection of light. However, injecting light in a controlled way is almost infeasible and only partially possible for CMOS image sensors that implement an electronic rolling shutter mechanism that reads the captured image information row by row, rather than all-at-once (global shutter) [7], [10]. In contrast, CCD image sensors always implement a global shutter inherent to their design. This means fine-grained signal injection attacks using light are not possible. Moreover, a light-based attack requires the line of sight between the adversary and the target camera. Finally, attacks that leverage optical emission tend to be suspicious and easily detected by simple mechanisms. For example, if the frame is suddenly over or under-exposed, an alarm is triggered [14], [15].

In this paper, we overcome these limitations by using intentional electromagnetic interference (EMI). We show that fine-grained perturbations can be injected into CCD image sensors using electromagnetic emanation. While the susceptibility of CCD image sensors against electromagnetic interference has been evaluated in the context of electromagnetic compatibility (EMC) [16], to the best of our knowledge, no research has been conducted from the perspective of an adversary trying to inject fine-grained, controlled perturbations using intentional EMI. Yet, we demonstrate that, due to their architecture, CCD image sensors are vulnerable to signal injection attacks using electromagnetic waves.

2. Threat Model

In this paper, we focus on the evaluation of signal injection attacks against CCD image sensors via intentional electromagnetic interference (EMI).

A. Goal

The overarching goal of the adversary is to spoof the image information captured by a CCD image sensor using intentional electromagnetic interference. More precisely, the attacker aims to manipulate the amount of signal charge captured for specific pixels by inducing a voltage into the image sensor before the signal charge is quantized and amplified by the measurement unit. It is worth mentioning that the adversary is limited in the direction of the manipulation of the signal charge. The electromagnetic waves can only alter the signal charge in the positive direction. This means the adversary can only induce an additional voltage to increase the brightness of pixels but not reduce the captured signal charge to decrease the brightness.

The incentives for an attacker to manipulate the captured frames can vary. For instance, the adversary's intention could be to inject adversarial examples into the frames to cause algorithms used in vision-based intelligent systems to fail. Another goal could be the disruption or distortion of scientific research. For example, the adversary could inject noise into the CCD used for measuring light pollution or into the CCD used in a microscope. Since CCD image sensors are typically used in professional applications, signal injections attacks will always be an issue, independent of the target system.

B. Capabilities

We assume the adversary has access to off-the-shelf equipment, such as software-defined radios, amplifiers, and antennas. Depending on the scenario and the target camera, more sophisticated hardware, for instance, an SDR with higher digital-to-analog converter (DAC) throughput, a more powerful amplifier or a directional high-gain antenna might be necessary. On the other hand, if the attacker only wants to render the image information unusable and does not need to have fine-grained control over the captured frames, any device that can emit electromagnetic waves at the resonant

frequency of the target image sensor is sufficient. We also presume that the attacker has some basic knowledge about digital signal processing to generate the adversarial signal. Moreover, the attack can be executed from outside the line-of-sight and even remotely. Nevertheless, under no circumstances can the attacker access the video output of the target camera. Hence, no synchronization between the attack signal and the readout signal of the image sensor is possible.

3. Signal Injection Attack

Our hypothesis is that, due to their architecture, CCD image sensors are susceptible to intentional electromagnetic interference making them vulnerable to post-transducer signal injection attacks. Normally, a sensor should only react to one specific physical stimulus to which it is intended to respond. In the case of an image sensor, the stimulus is light, or to be more precise, photons. Incident light causes the generation of electronic charges that can be measured and quantized. Yet, the image sensor itself cannot determine whether the signal charge was generated by the photodiode array during the integration period due to the incident light or resulted from electromagnetic interference that coupled onto the circuit. A malicious actor could leverage this fact and emit electromagnetic waves at the resonant frequency of the target CCD image sensor to induce a voltage and subsequently alter the captured image information.

A. Attack Execution

Executing a signal injection attack against a CCD image sensor can be separated into three steps. In this section, we will give a detailed overview of the necessary preliminaries.

1) Signal Generation: In general, any arbitrary data, such as Gaussian white noise, can be modulated onto a carrier wave and induced into the images sensor. However, to demonstrate the possibilities of the attack, we will describe in the following the injection of data in the format of an RGB image. The content of the injected image can be arbitrary. One example would be an image representing an adversarial example. Each pixel of a colored digital image has three color channels — red (R), green (G) and blue (B). The values of the color channels range between 0 and 255, whereas higher values represent higher intensity. In the context of this paper, each pixel of the input image corresponds to one sample S[x,y] of the attack signal, with x and y being the coordinates of the pixel. The amplitude of the attack signal at S[x,y] determines how much additional voltage is induced into a camera pixel. The higher the amplitude, the more additional voltage is induced, resulting in a brighter pixel. The amplitude is represented by the linear luminance Y, which can be calculated using the following equation:

$$Y[x,y] = 0.2126R + 0.7152G + 0.0722B \quad (1)$$

2) Resampling: Ideally, the sample rate of the attack signal should match the readout signal of the image sensor. Usually, the readout signal is determined by the sample rate of the digital-to-analog converter. If the sample rate of the attack signal does not match the readout timings of the image sensor, the injected noise is drifting over consecutive frames and fine-grained control cannot be achieved. To prevent this, the attack signal has to be resampled. If the sample rate of the image sensor cannot be derived from the datasheet, it can be calculated as follows:

$$S = N_{columns} \times N_{rows} \times F; \quad (2)$$

where Ncolumns is the width, Nrows the height and F the frame rate of the image sensor. With increasing frame rate and resolution, the required sample rate of the software-defined radio is increasing too. Depending on the target camera, this might raise the bar for the attacker to inject fine-grained distortions.

3) Transmission: Once the signal has been extracted and resampled, it can be transmitted. Given that the amplitude of the attack signal determines the amount of electronic charge induced into the image sensor, the input image is modulated onto the carrier wave using amplitude modulation. An end-to- end representation of the three attack steps is depicted in Figure 1.

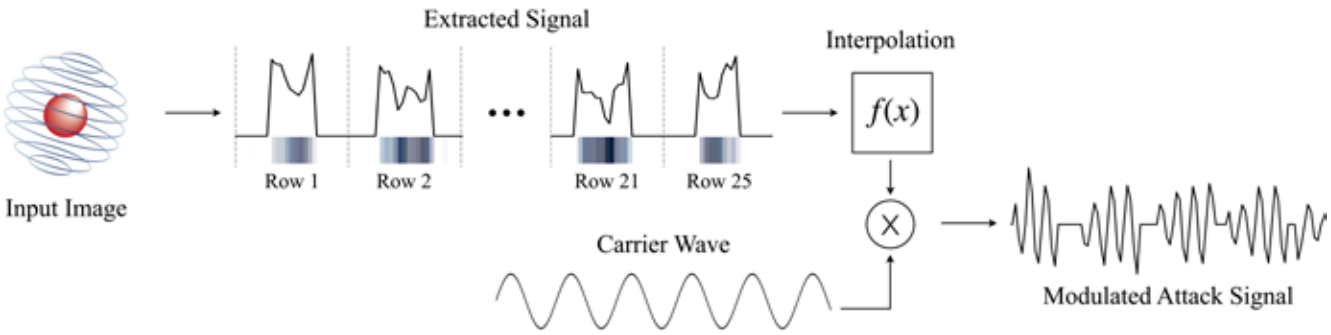

Fig. 1: Overview of the necessary steps to generate a malicious attack signal. First, the signal to be transmitted is extracted from the input image by calculating the luminance Y for each pixel. Second, the extracted signal is interpolated to ensure the different sample rates match up. Finally, the interpolated signal is modulated onto the carrier wave and transmitted via the software-defined radio.

3. Evaluation

We evaluated the susceptibility to intentional electromagnetic interference of two different CCD image sensors. In this section, we describe our method and present the results.

A. Experimental Setup

We validated our hypothesis on two different cameras equipped with CCD image sensors, namely a DFM 25G445- ML and a 420TVL CCTV board camera. The DFM 25G445- ML is a professional GigE color board camera used in a wide variety of applications, for instance, industrial automation, quality assurance, and surveillance [13], and is equipped with a Sony ICX445AQA image sensor [12]. In contrast, the analog CCTV board camera can often be found in older CCTV settings or cheap drones. To ensure that the experiments were not affected by interference from other signal sources, we placed the target camera and the antenna in a closed RF shielded box. At the same time, this prevented uncontrolled radiation of the attack signal, which could otherwise interfere with legitimate communication channels in the tested frequency spectrum or couple to other equipment. Nevertheless, to rule out the possibility that the attack signal is induced directly into the cabling and not into the CCD image sensors themselves, we also performed the attack with the camera switched off. Finally, placing the camera in the EMI shielded box also prevented light from entering the camera and generating a legitimate signal charge, making it easier to detect if the attack was successful or not. Under normal operation, all the captured video frames were almost entirely black. Only some pixels were colored due to the various types of noise, such as readout and dark current shot noise, generated by the camera itself [5], [18]. An overview of the experimental setup and how the different components were interconnected is depicted in Figure 2.

For the generation and transmission of the attack signal, we only used off-the-shelf equipment. More specifically, the signal was generated on a desktop computer running Ubuntu 18.04 and GNURadio 3.8. The PC was connected via Ethernet to an Ettus Research USRP N210, which emitted the signal via a 900 MHz omnidirectional vertical antenna with 3dBi gain. The USRP was equipped with a UBX-40 daughterboard, which provides a maximum output power of 100 mW [9].

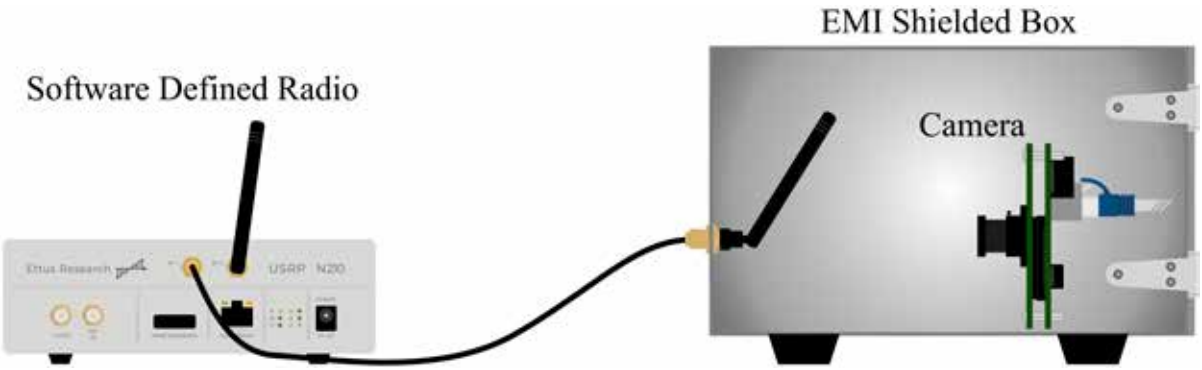

Fig. 2: Experimental setup used for our evaluation. The camera was placed inside an RF shielded box to prevent interference with and from other devices.

B. Carrier Frequency fc

In this section, we describe the empirical approach we used to determine the most effective carrier frequency for the two tested cameras.

1) Signal Generation: We captured video frames while running a frequency sweep with a step size of 1 MHz from 50 to 5000 MHz. We then calculated the Structural Similarity Index Measure (SSIM) [17] between the collected frames. More specifically, for every carrier frequency fc, we collected ten frames, three legitimate frames during normal operation, and seven malicious frames while emitting a sine wave with frequency f = 1 kHz modulated onto the carrier wave. For this experiment, the cameras were placed around 3 cm away from the transmitting antenna and the output power of the software-defined radio was set to the maximum (20.1 dBm).

As previously mentioned, the experiments were conducted inside a shielded and completely dark box. Due to the dark environment, the frames collected during normal operation were almost entirely black. As a result, the SSIM between consecutive legitimate frames was high, meaning they were almost identical. It should be noted that a similarity score of 1.0, i.e., the frames are identical, can never be achieved due to the noise generated by the camera itself. In contrast, for a successful attack, the SSIM between legitimate frames and malicious ones should be as low as possible. The most effective carrier frequency was selected based on the smallest SSIM value. In other words, the frequency that caused the smallest SSIM values induced the most significant perturbations. On the other hand, an ineffective carrier frequency did not induce

any signal charge and led to high SSIM values similar to those measured between legitimate frames.

2) Results: The results of the frequency sweep for both cameras, the DFM 25G445-ML and the analog CCD, are depicted in Figure 3. As the graphs show, the most effective carrier frequency was 190 MHz for the DFM 25G445-ML and 341 MHz for the analog CCD.

Interestingly, while the analog CCD camera was only affected at around 341 MHz, the DFM 25G445-ML was vulnerable at various frequencies. The findings indicate that a malicious signal modulated onto a sinusoidal carrier wave at the appropriate frequency is highly likely to couple successfully onto the CCD image sensor. As a result, an adversary would pick 190 MHz or 341 MHz, depending on the target camera. Due to space constraints and the option to precisely control the camera parameters, such as exposure and gain, which allows us to evaluate the attack under controlled conditions, the rest of the paper will focus on the evaluation and results of the DFM 25G445-ML.

B. Fine-Grained Control

In this section, we show how an adversary can gain fine-grained control over the captured image information. We discuss the prerequisite for such an attack in detail and demonstrate an example attack.

1) Prerequisites: The main prerequisite for such an attack is the right equipment, i.e., a software-defined radio with a high sample rate. Moreover, some details about the target camera are necessary. As described in Section III, to manipulate individual pixels, the sample rate of the attack signal has to match the frequency of the readout signal. Therefore, a rough estimate of the image sensor's sampling rate is required. In case the sample rate is not stated in the datasheet and is publicly available, Equation 2 can be used to approximate it.

2) Method & Result: We replicated the experimental setup as previously described and depicted in Figure 2. The camera was placed 3 cm away from the transmitting antenna, and the image sensor gain was set to 29. We then executed the attack following the steps outlined in Section III. First, we extracted the attack signal from the input image, in this case, the CDT logo. Then we resampled the signal to ensure the sample rate matches the sample rate of the image sensor. Finally, we modulated the signal onto a carrier wave with fc = 190 MHz and transmitted it with the maximum transmission power (20.1 dBm). The resulting frame is depicted in Figure 4.

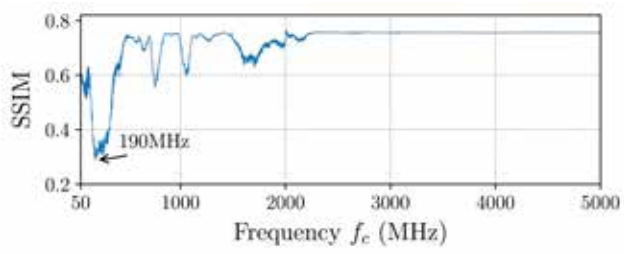

Figure 3 (a) DFM 25G445-ML

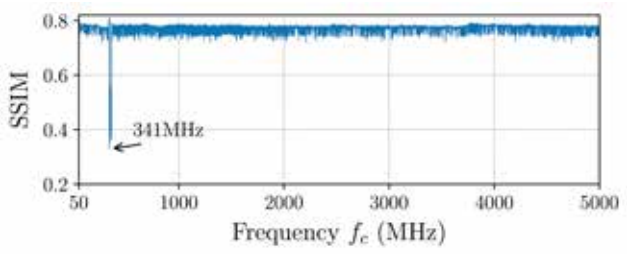

Figure 3 (b) Analog CCTV board camera

Fig. 3: Results of the frequency sweep. The SSIM represents the similarity between the frames captured during normal operation and while a sine wave with f = 1 kHz modulated onto a carrier wave with carrier frequency fc was emitted.

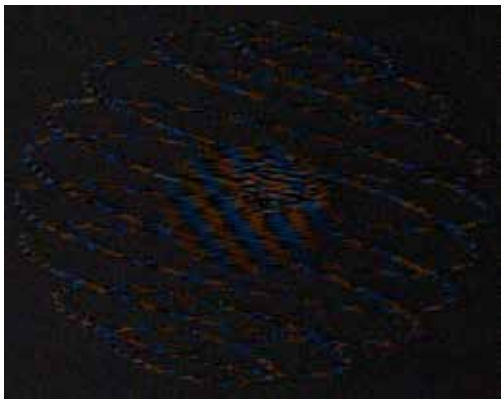

Fig. 4: Example of a signal injection attack, illustrating the fine-grained control an adversary can gain over the captured frame.

4. Conclusion

The evaluation in this paper suggests that CCD image sensors are, due to their architecture, susceptible to intentional electromagnetic interference. We demonstrated a signal injection attack against two different CCD image sensors. The presented attack allows an adversary to manipulate the captured frames down to the granularity of single pixels. Although CCD image sensors are not as widespread nowadays and will probably be completely replaced by CMOS image sensors in the future, they are still widely used in professional applications. We conclude that signal injection attacks are a real threat to applications relying on the input from cameras equipped with CCD image sensors.

REFERENCES

[1] Daniel Durini. High performance silicon imaging: Fundamentals and applications of CMOS and CCD sensors. 2019.

[2] Dan M Duriscoe, Christian B Luginbuhl, and Chadwick A Moore. Measuring night-sky brightness with a wide-field ccd camera. Publications of the Astronomical Society of the Pacific, 119(852):192, 2007.

[3] Eric Hagt and Matthew Durnin. China's antiship ballistic missile: Developments and missing links. Naval War College Review, 62(4):87–115, 2009.

[4] Steve B Howell. Handbook of CCD astronomy, volume 5. Cambridge University Press, 2006.

[5] Kenji Irie, Alan E Mckinnon, Keith Unsworth, and Ian M Woodhead. A model for measurement of noise in ccd digital-video cameras. Measurement Science and Technology, 19(4):045207, 2008.

[6] W Gray Jay Jerome. Practical guide to choosing a microscope camera. Microscopy Today, 25(5):24–29, 2017.

[7] Sebastian Köhler, Giulio Lovisotto, Simon Birnbach, Richard Baker, and Ivan Martinovic. They see me rollin': Inherent vulnerability of the rolling shutter in cmos image sensors. 2021.

[8] Jonathan Petit, Bas Stottelaar, Michael Feiri, and Frank Kargl. Remote attacks on automated vehicles sensors: Experiments on camera and lidar. Black Hat Europe, 11:2015, 2015.

[9] Ettus Research. Ubx 40 usrp daughterboard, 2021. https://www.ettus.com/all-products/ubx40/.

[10] Athena Sayles, Ashish Hooda, Mohit Gupta, Rahul Chatterjee, and Earlence Fernandes. Invisible perturbations: Physical adversarial examples exploiting the rolling shutter effect. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 14666–14675, 2021.

[11] Jovan Skuljan. Quadcam–a quadruple polarimetric camera for space situational awareness. In Proc. 18th AMOS Conf, pages 275–285, 2017.

[12] SONY. ICX445AQA Datasheet. https://www.argocorp.com/cam/ImagingSource/common/PDF/sensor/icx445aqa_1.2.en_US.pdf.

[13] The Imaging Source. DFM 25G445-ML. https://www.theimagingsource.com/products/board-cameras/gigecolor/dfm25g445ml/.

[14] Synology. Synology surveillance station, 2020. https://synology.com/en-us/surveillance.

[15] Bosch Security Systems. Fw 6.30 tamper detection, 2016.https://resources-boschsecurity-cdn.azureedge.net/public/documents/TN_VCA_tamper_detect_WhitePaper_enUS_22996235531.pdf.

[16] Robert Wacholc. Investigation into Noise and Stability Effects on CCD and Readout Electronics with Reference to the PLATO Mission. PhD thesis, UCL (University College London), 2019.

[17] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing, 13(4):600–612, 2004.

[18] Ralf Widenhorn, Morley M Blouke, Alexander Weber, Armin Rest, and Erik Bodegom. Temperature dependence of dark current in a ccd. In Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications III, volume 4669, pages 193–201. International Society for Optics and Photonics, 2002.

[19] Chen Yan, Wenyuan Xu, and Jianhao Liu. Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle. DEFCON 2016

Departmental Teaching award for Advanced Security Course

Marcel Stolz (CDT2016)

Marcel was recognised for his outstanding contribution to the Advanced Security MSc course. The award was given based on student feedback and scoring. Marcel reflects on the award as "I like making classes interactive, which I believe was valued by the students. I split them into small groups for some of the problems, where they could discuss their answers and come up with what they considered the best answer as a group and then present it in class. I connected the practical assignments, taught by others for the same course, with the theoretical questions I taught, which I believe was a source of positive feedback.

In addition to this award, Marcel has also recently published several papers and presented his work on platform neutrality; analysis of defence vs intelligence agency interests in states highlighting how Swiss legislation has a (counter-intuitive) de-escalative impact on the cybersecurity dilemma; and the findings from a collaboration with the German think tank SNV, on the German government's cybersecurity ecosystem.

Satellite Hacking: Researching Cyber Space

James Pavur (CDT17)

The number of satellites in orbit is expected to increase by an order of magnitude over the next decade. From weather and geolocation to communications and research, these distant information systems provide critical services that impact billions of lives. Here, at the beginning of a new era in space technology, it is more important than ever to ensure that these platforms are secure.

At Oxford's Systems Security Lab, led by Professor Ivan Martinovic, we are working to study the unique cybersecurity threats and requirements relevant to space technologies. The intention is to identify security gaps that have evolved in modern space missions, determine the underlying causes of these shortcomings, and invent solutions that satellite operators can incorporate to better secure their missions.

Studying SATCOMs

One particularly exciting topic has been our research into the security of modern satellite broadband communications. Satellite-based internet services are a key growth area in the space industry, with companies like SpaceX, Inmarsat, Amazon, and OneWeb betting on satellite constellations as the best way to bring the next billion internet users online. Even today, satellite internet services support millions of customers and businesses. Understanding the security properties and requirements of status quo services can help guide our efforts to design and defend the next generation of satellite broadband.

We began with a series of passive surveys, listening to the radio emissions of 18 satellites in geostationary orbit (GEO). These GEO satellites are located about 30,000 km above the equator. The specific platforms involved in these studies serve customers on five continents, with a combined footprint area exceeding 80 million square kilometers.

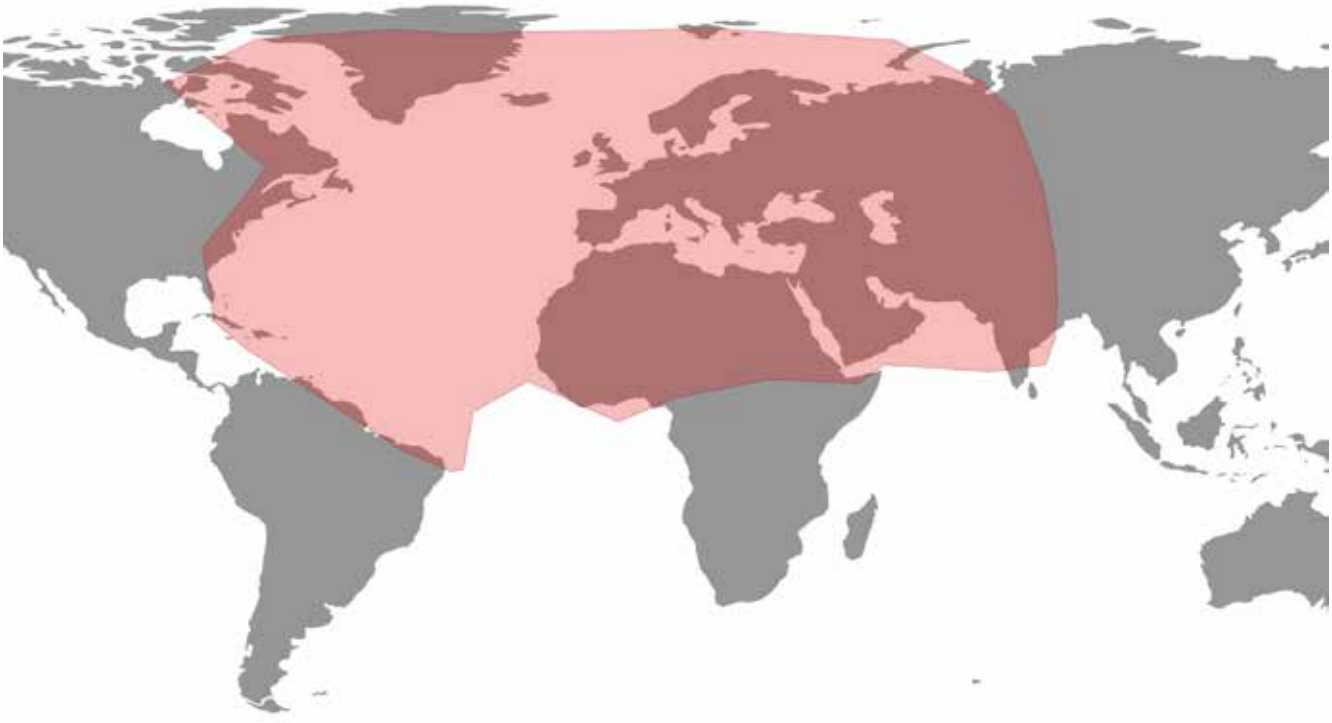

I - The polygon on this map roughly correlates to the combined coverage of the signals received in our study.

Signals Intelligence on a Hobbyist Budget

In exploring these signals, we found that a cyber-attacker could reliably eavesdrop on broadband traffic from dozens of different providers. To make matters worse, they could do so using about £250 worth of widely available home television equipment.

As many satellite internet service providers were not employing over-the-air encryption, this meant attackers could directly observe the internet traffic of satellite broadband customers. Additionally, due to the nature of satellite communications, this attack was virtually untraceable and could be executed over distances of thousands of kilometers.

Looking closer at the contents of these signals confirmed the severity of these findings. We encountered a wide range of data which was inadequately protected. This included consumer traffic, such as SMS text messages from passengers using in-flight Wi-Fi services over the Atlantic. It also included data from governments and some of the world's largest businesses, such as navigational charts destined for cargo vessels in the Mediterranean or login credentials for wind turbines in continental Europe.

When we encounter issues like this in our research, we follow a standard practice known as "responsible disclosure" prior to publication. In our case, this involved reaching out directly to both satellite internet service providers and larger industrial customers to inform them of our findings and make them aware of previously overlooked risks impacting their businesses.

II Partial reconstruction of a JPEG image sent by a maritime engineer encountered during our traffic analysis.

How Does this Happen?

During our responsible disclosure conversations, we learned that many in the industry were notionally aware of the risk of unencrypted wireless communications but had decided to accept it. In part, this was because they assumed equipment to execute these attacks was far more expensive than we found in our own research. However, there were also substantial performance costs to standard encryption approaches – such as the use of end-to-end virtual private networks (VPNs). After some investigation of related work, we learned that the fundamental issues with VPN performance from GEO related to high latency and the TCP communications protocol commonly used for web traffic. In particular, internet service providers had special middlebox tools called performance enhancing proxies (PEPs) which were designed to optimize and speed up a customer's TCP connections.

These optimization tools required the internet service provider to have full visibility into the traffic of their customers so they could determine which packets to optimize. Customers who decided to use a VPN would end up blocking this visibility and would find their connections slow to a crawl.

Building an Open and Actionable Solution

Rather than attempt to convince internet service providers to update their systems to support encryption – a liability which they seemed reluctant to adopt – we worked to invent an approach which would allow customers to encrypt their traffic independently according to their specific needs. Critically, the system had to be comparably performant to unencrypted traffic sent via a traditional PEP.

The ultimate result of this effort was the creation of a hybrid VPN-PEP called QPEP which combines the performance properties of satellite PEPs with a VPN-like encrypted tunneling mechanism. QPEP leverages a modern encrypted transport protocol, known as QUIC, which ensures reliability and reasonable bandwidth exploitation – even in high latency satellite environments.

In testbed simulations, we found that QPEP not only outperforms traditional VPNs, but its design achieves faster page load times than even unencrypted PEPs. Across the Alexa Top 20 list of popular websites, QPEP roughly halves page load times compared to an unencrypted PEP and loads pages over 70% faster than VPN-encrypted connections in the same network. Today, QPEP is freely available as an open-source tool which anyone can download and modify. One advantage to doing this sort of research at Oxford is that we can share information about our solutions freely, without pressure to commercialize or maintain proprietary secrets. This means other researchers can verify and improve on our ideas, paving the way to more secure satellite broadband for everyone. In an industry where encryption solutions are typically closed source and unverifiable "black boxes," we hope that our open approach can rejuvenate innovation around a vital security topic.

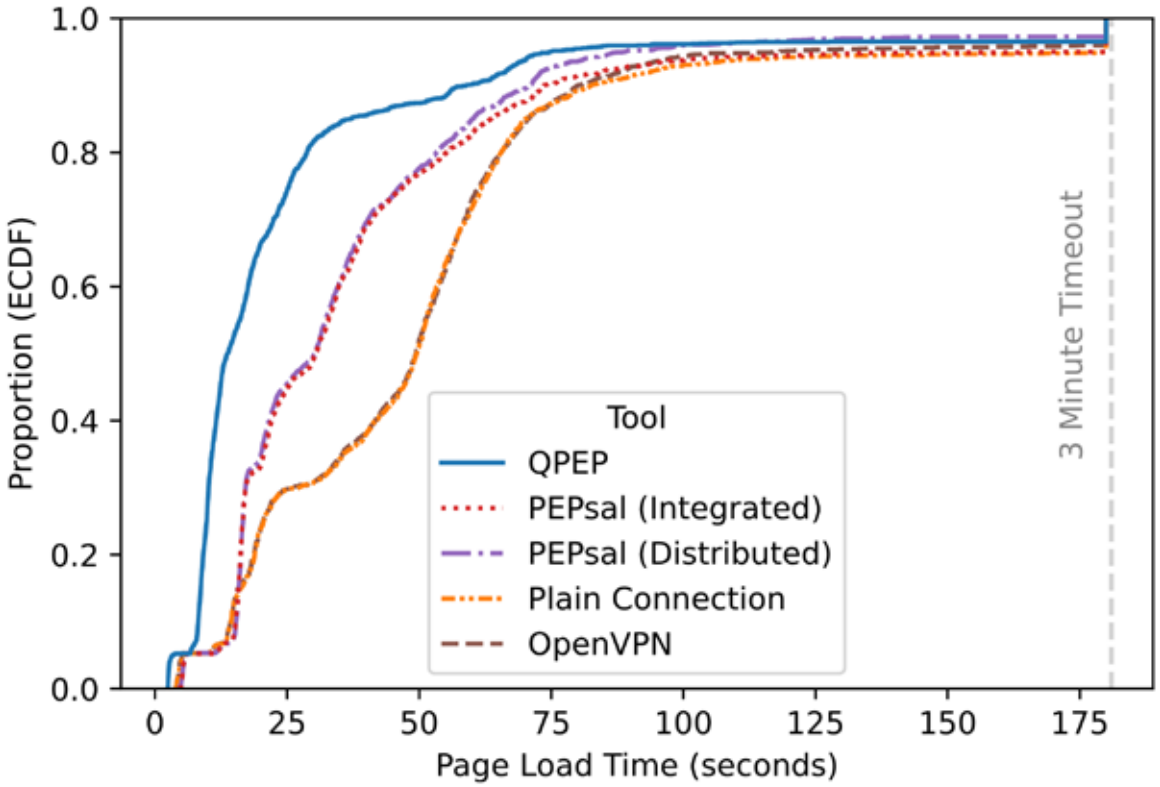

III Distribution of Page Load Times Across Alexa Top 20. PEPsal is an unencrypted PEP and OpenVPN is a typical VPN product.

Going Forward

Our own work on QPEP is continuing as we move towards testing the tool in real-world satellite networks and optimizing its design for use in large networks. We are also considering related topics in securing other satellite communications applications, such as inter-satellite links in Low Earth Orbit (LEO) constellations.

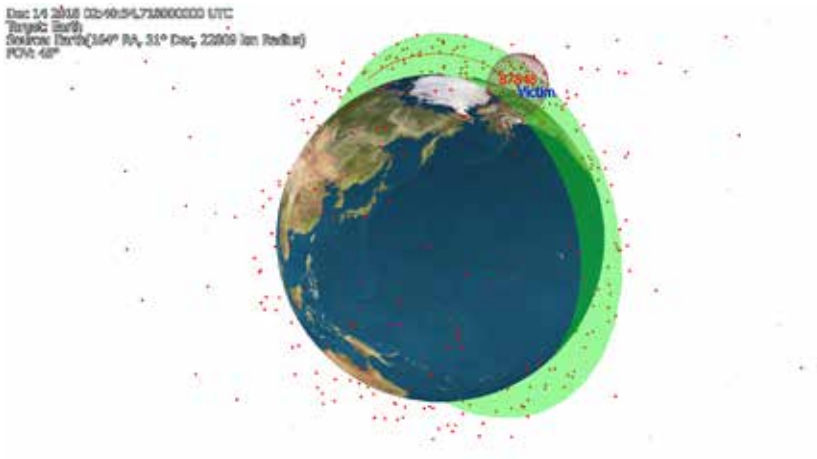

IV Screenshot from our Simulations of Projection Tampering Attacks on SSA Data

Beyond satellite broadband, there are many other security topics of relevance to the space community. For example, we are conducting research into the security properties of space situational awareness (SSA) data which is used to help satellites avoid on-orbit collisions with each other and with pieces of space debris. We've also exploring a variety of other topics, ranging from small satellite platform security to the interaction between cybersecurity and range safety for rocket launches.

Ultimately, space is an exhilarating new frontier for cybersecurity research. Our own exploration of this domain is still in its early days, but we are already finding unusual challenges and opportunities for impactful research. Academic security research on space systems will be critical to supporting the public and private sectors as we leap into the next era of human spaceflight.

Building a smart city 4.0 ecosystem platform: a case study of Jakarta's Super App

Yudhistira Nugraha (CDT Alumni)

In this ongoing pandemic, a smart city has become a key driving factor for accelerating the digital transformation to cope with the Covid-19 crisis. Some emerging technologies are fast becoming a key instrument in a smart city, an increasingly important area in well-applied city services. This article introduces the concept of a smart city 4.0 ecosystem platform that provides new insights on how to translate a vision into reality using Jakarta's Super-App called JAKI (jaki.jakarta.go.id) as a use case. The smart city 4.0 ecosystem platform offers a significant opportunity to advance the understanding of building a smart city with technologies, innovations and collaborations.

What is a Smart City 4.0?

What makes a city smart? There is no clear consensus if one looks for a uniform answer. From a service provider's perspective, a smart city is an innovation related to emerging technologies, such as smart parking, smart home, and smart lighting. Government officials would say that a smart city relates to implementing digital government such as applying for permits or citizen identity cards and making public service less bureaucratic, more effective, efficient and transparent. The city residents would be more pragmatic, expecting a smart city to provide seamless transportation and better living conditions with more job opportunities.

These views do not compete; they are actually aligned because the concept of a smart city should serve added values and provide opportunities to address the city's problems and fulfil the citizen's needs. At its core, the smart city is about the ecosystem and people, people connecting to the people (digital citizens), people connecting to the business (digital economy), and people connecting to the government (digital government). It aims to improve citizens' quality of life, foster economic growth, and promote environmental sustainability by using technologies, innovations, and collaborations.

Our intended outcomes are twofold and interlinked: innovative city and happiness. We believe that the adoption of technology and innovation should lead to the paramount goal; the happiness of our citizens. With various emerging technologies and innovations, such as big data, artificial intelligence, Internet of things, cloud services, and blockchain, our work has considerably upgraded public services such as education, healthcare, transportation, public safety, etc. However, we want to go further. Our smart city concept needs to include fulfilling basic needs such as food, water, and electricity. Only by doing so, we can make sure of public happiness.

The concept of smart city 4.0 aims to improve active participation and engagement of city co-creators, such as city stakeholders and residents, for building a better Jakarta. As a concept, it advocates for stronger collaboration (co-creation) and cooperation between citizens and communities in city development. The city government is

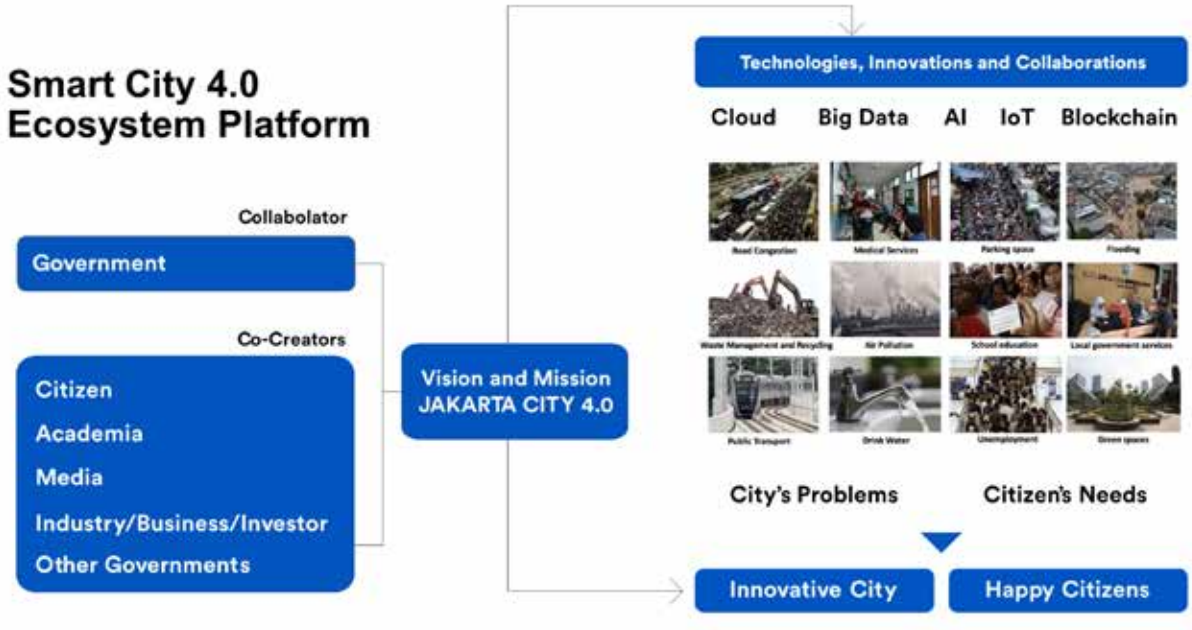

Figure 1: Smart City 4.0 Ecosystem Platform

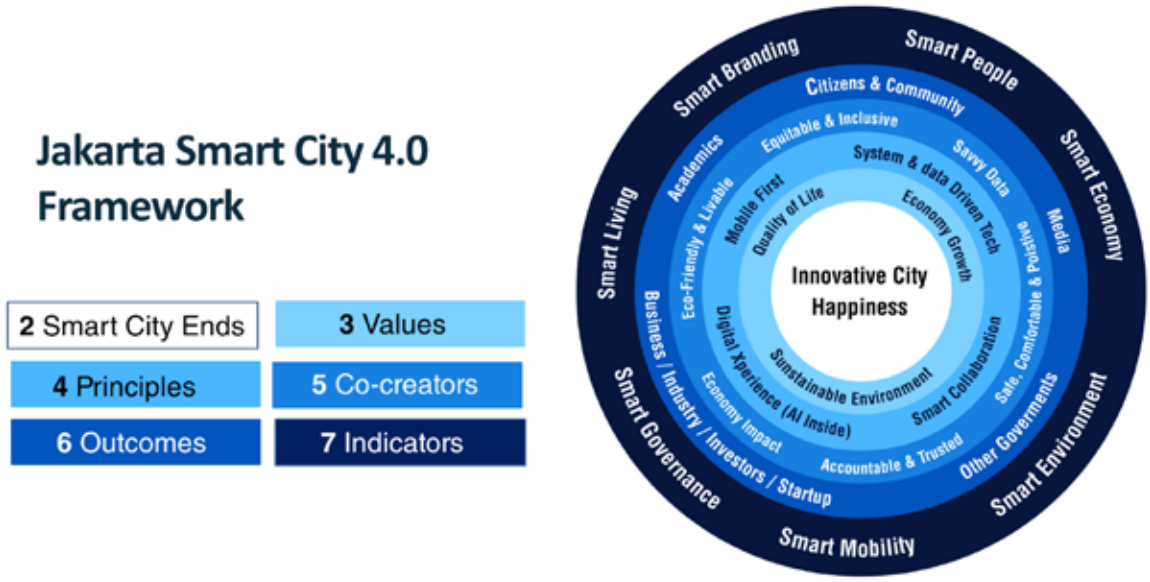

Figure 2: Jakarta Smart City 4.0 Framework

no longer seen as the sole caretaker that has to supply all answers about the needs of the citizens. To do this, the city government has changed gears, turning Jakarta into a more collaborative ecosystem for the citizens, academics, communities, social organisations, start-ups, media, businesses, industries, and other governments, as seen in Figure 1.

We foster this dynamic ecosystem by building Jakarta's Super-App called JAKI. This Super-App has become our government platform to sustain continuous innovation and collaboration to make our version of the smart city a reality.

What We Do in Jakarta Smart City

By carrying out the concept of the Smart City 4.0 Framework as a form of translation and implementation of the governor's vision of "Innovative City, Happy Citizens", Our work has been guided by certain principles to reach our intended outcomes, namely Mobile First, System-and-Data Driven, Digital Experience, and Smart Collaboration, as seen in Figure 2.

Mobile First: We're Doing It for Our citizens

Mobile devices have changed the way people consume information and meet their daily needs. The APJII 2018 survey shows that 93.9% of Indonesians are connected to the Internet daily via smartphones. With JAKI, we want to bring better services to our citizens in the palm of their hands with convenience, effectiveness and efficiency.
As the Jakarta Super-App platform, JAKI has integrated more than 15 features and 28 applications developed by agencies and the public. It provides a single sign-in, which allows a citizen to access all digital services, make payments, and update personal information.

During the pandemic, JAKI serves as a platform for Jakarta's Covid-19 Response via www.corona.jakarta.go.id. It is mobile-friendly, the content is available in two languages (Indonesian and English), and it is also accessible through the website. In almost sixteen months since its initial launch on March 6 2020, the total access to the website has reached more than 42 million hits (data per July 18 2021).

System and Data-Driven: Data-Driven Innovations to Support Smart City

Jakarta has better responses to the COVID-19 pandemic due to smart infrastructure, data analytics, and predictive modelling. Building government as a platform transforms government into a public API where city co-creators can build, connect and interact. JAKI continuously improves the Omni-channel experience of its users to get better and more personalised services. JAKI aims to deliver personalised services and integrate all digital services from birth to death, not only provided by the government but also by our city co-creators. Personalised public services can be developed and adapted to individual circumstances using a Digital ID for Jakarta residents. Hence, various services and data for Jakarta residents can be integrated and customised.

The Vaccination Registration– one of JAKI's most popular features– is the prime example of data utilisation which helps citizens independently register and book their schedule for the Covid-19 vaccination. As of July 18 2021, more than 18 million people have accessed the platform, and about 220,116 people have been registered and vaccinated.
JAKI also allows citizens to report public problems they find across Jakarta, including violations of the Large-Scale Social Restriction. With a picture taken with a smartphone, citizens can submit their complaints through JAKI and will be automatically registered in the Cepat Respon Masyarakat (CRM) Platform. The CRM platform provides transparency by allowing citizens to keep track of their complaints online. As of July 14 2021, there have been 14.163 reports by 7.911 citizens related to Large-Scale Social Restriction. About 13.523 complaints (95,5%) have been solved.

Digital Experience: Organizational Digital Transformation

During the new normal, Jakarta Smart City collaborates with academia and startups to develop a system that will enhance the safety of the citizens using design thinking, system thinking and computational thinking. Understanding citizen needs, solving the problems, seizing the opportunities, and putting the values on livelihood are our digital transformation in building JAKI as a government platform.

We're building a new feature that will promote Digital Health Identity and be integrated into our Super-App JAKI. For example, we provide features in JAKI to support online activities during the new normal such as work from home and learning from home. These features include JakLapor (online complaint system), JakPangan (online food information), JakCLM (online Rapid Test) and our latest innovation, the Vaccination Registration feature to accelerate the vaccination program in Jakarta. JAKI encourages increased digital experience, creating features that work effectively during the pandemic:

JakLapor and JakRespons allow citizens to submit their complaints in minutes. Citizens can also monitor their reports in real-time without having to inquire to the government office.

With machine learning, JakCLM allows citizens to check their risk of getting COVID-19 independently.

Jejak, a digital guestbook, provides QR Code-based supervision of buildings and public premises' capacities in Jakarta. With this feature, the government monitors health protocols and seeks to prevent new clusters of COVID-19.

The Vaccination Registration feature aims to open broader access to COVID-19 vaccines for all citizens. We're also collaborating with the Jakarta Department of Health, The National Armed Forces and private sectors to establish Vaccination Centers around the city.

Smart Collaboration: Collaborating with Co-creators for More Solutions

The pandemic has physically kept us apart from families, friends and colleagues. However, it has also strengthened our bonds. Collaboration has become indispensable to weather this crisis together. At Jakarta Smart City, we are bolstering collaborative projects with government officials, startups, social organisations, and, above all, Jakarta citizens through different collaborative schemes such as the playground, co-develop, and consumer schemes.

For instance, we've been working with the Department of Health of Jakarta in providing the most updated, accurate and trusted data on COVID-19 cases in the capital city. With the Department of Spatial Planning and Land Management of Jakarta, we are developing an integrated map of COVID-19 spread, a map of the controlled zone in Jakarta and an interactive map to support the Large-Scale Social Collaboration (KSBB) program helping those who are in dire need.

We have also worked with the Indonesian Ministry of Communications and Informatics and the Ministry of Health through PeduliLindungi to support vaccination registration in Jakarta. By using public API, it helps citizens independently register and book their schedule for the Covid-19 vaccination.

Jakarta Smart City also collaborates with Harvard CLM Team and Klakklik.ID in developing Kalkulator Covid-19, a self-assessment app powered by Corona Likelihood Metric (CLM) machine learning. With full support from the Jakarta Department of Health, the CLM can also be used as the first screening test to assess the risk of COVID-19 infection of the users.

With Cartenz Lab, we developed a contact-tracing app to support new public behaviour on health, safety and productivity. We also integrated Jakarta Aman and Sekolah.mu into JAKI to keep working and learning from home. All these efforts show that JAKI isn't merely a government-designed app but also a community-designed service that helps people during the pandemic.

As a collaborative ecosystem, at a national level, JAKI won a gold medal at the 2020 Indonesia Entrepreneur TIK (IdenTIK 2020) for the public sector category. We will be representing Indonesia at ASEAN ICT Awards 2021. At an international level, we are delighted for JAKI to be recognised as the WSIS Prizes 2021 Champion for the first time participating in the e-Government category. We will not stop here, considering security by default/design for future improvement; we will continue to develop this super app as a government platform to be implemented globally.

Finally, the year 2020 has become a momentum for Jakarta Smart City, gaining new status as the Regional Public Service Agency (BLUD). This condition will increase our agility, capacity and professionalism in improving public services and providing solutions to city problems. Jakarta stands proud and is transforming into a more innovative and liveable city. We will continue our hard work and ensure the happiness of the citizens of Jakarta.

About the author

Yudhistira Nugraha is a senior member of the Indonesian Government, currently assigned as Director of Jakarta Smart City – Department of Communications, Informatics, and Statistics, the Provincial Government of Jakarta. He has been a civil servant since 2006 and had been working for more than 13 years at the Ministry of Communications and Informatics. His last assignment at the Ministry was as the Deputy Director for E-Government Services in Economic Affairs.
He received a D.Phil. in Cyber Security from the University of Oxford. He also holds a Data Protection Officer Professional University Certificate (ECPC-B DPO) from the European Centre of Privacy and Cybersecurity (ECPC) of Maastricht University, the Netherlands. He also undertakes research and teaching in the area of cybersecurity, privacy, and smart city at Telkom University.

Joint Cyber Security CDT Summer School

Arianna Schuler Scott (CDT16), Anjuli R. K. Shere (CDT18) and Claudine Tinsman (CDT18)

June 2021 saw the inaugural summer school between the seasoned and newly-funded Cyber Security CDTs across the UK. Oxford students Anjuli, Arianna and Claudine were part of the committee tasked with organising and running the (hugely successful!) event. The virtual event ran from 28th – 30th June and was hosted by the Universities of Bath and Bristol. The organising committee also included students from Royal Holloway University of London (RHUL)and University College London (UCL). The event's theme was 'Resilience – Systems, Societies and Threats', showcasing a very great need for multi-, inter- and cross-disciplinary collaboration in the security space. Day one consisted of skills sessions and a cyber crisis simulation, day two offered students the chance to take part in a "data challenge" (further details below), and day three wrapped up with a Societies panel run by our very own Oxford students, keynotes and prizes.

Day 1: Monday 28th June 2021

This summer school was intended to kick off the mid/post-pandemic mingling: Adam Joinson opened the event reminding us that by 2028 there should be about 270 graduated PhDs and DPhils in Cyber Security. Such a significant investment of time and resources, he emphasised, was to create an enormous cohort of experts. These experts will need to network and learn from each other throughout our careers.

Neil Ashdown (RHUL) followed Joinson's opener with a twist on traditional icebreakers, prompting questions such as "What is your favourite film about cyber security?". He also used this as an opportunity to highlight that he chooses to "live free or die hard" every day. Despite Neil's optimism, Die Hard 4.0 did not win the popular vote.

The WannaCry-style Crisis Simulation was set in a fictional country called Alphaland. It was discovered that, given her overwhelming success as President, Anjuli should be given executive power in all scenarios. In protest of the overcast Monday morning in June, and against all odds, the wargame was a fantastic networking event. Attendees demonstrated skills and personal strengths enough to inspire great confidence in global resilience in the coming (post-graduation) decades.

Jason Healey dialled in from New York, the founding director of the Cyber Statecraft Initiative of the Atlantic Council. Perfectly positioned to discuss the implications of international cyber conflict, he described parallels between the wargame scenario, and the dark history of cyber-attacks. This included the ongoing predilection to pay data-ransoms, against expert advice.

After a brief recess, the 'Systems' session was opened by Bruce Schneier, a divisive figure whose presence was particularly exciting as his work is a staple of our first year material. Schneier discussed the pitfalls of Internet security now being "everything security" due to all devices having computing potential. He drew attention to the issue of the Internet not having been designed with security in mind, recounting one of our favourite cyber security news stories – illegal access to a Las Vegas casino via a smart fish tank!

Awais Rashid (University of Bristol) continued the interdisciplinary theme of the conference, talking about the challenges associated with securing large-scale infrastructure and making entire societies cyber-resilient. In particular, he noted that many people's risk thinking patterns are isolated or, at best, sequential, rather than being nuanced understandings of interrelated events (i.e., most people are not security experts). This makes managing complex systems, such as supply chains, difficult. Still, Professor Rashid argued that the difference between solvable challenges and inescapable problems likely depends on whoever is facing them.

Irina Brass (UCL) and Laura G. E. Smith (University of Bath) discussed policy and the role of public influence in establishing regulations for emerging technologies. Laura pointed out how limited research is, into how different aspects of digital environments can impact risks and their associated harms to users. This, she argues, makes regulation difficult. Irina called for technology-neutral, issue-first policy to ensure a focus on values (like privacy), despite any technical hurdles.

The summer school's first session concluded with two sessions - Martín Barrère (Imperial College), talking about the cyber-physical security of critical infrastructure systems. As another event like Stuxnet is increasingly technically possible, his talk was a fascinating look into the avenues for attack of a point of physical vulnerability. Lizzie Coles-Kemp (RHUL) then led a skills session on the ins and outs of navigating academia, including how to build and maintain meaningful and productive partnerships, particularly with external stakeholders. An optimistic note upon which to end day one of the student conference.

Day 2: Tuesday 29th June 2021

Claudia Peersman (University of Bristol) kickstarted day two, discussing the 'Advanced Modelling of Cyber Criminal Careers' project, intended to shed light on the social and economic development of cyber-offenders' careers.

Brit Davidson (University of Bristol) and Lukasz Piwek (University of Bath) gave talks, lending insight into the types of investigations that can be conducted with huge

datasets from the dark web. The audience was then divided into smaller inter-CDT teams for the day's data challenge.

The challenge was to wrangle a database of posts from darkweb forums, creating a visualisation that would be useful for law enforcement purposes. Anjuli was part of Group 1, who carried out a longitudinal history study assessing the time it took for user to start using overtly aggressive language. They created a fully formatted and referenced IEEE-template report on the study in only three hours, and presented their findings. Happily, they won the challenge!

Day 3: Wednesday 30th June 2021

Day three started with a panel discussing resilience (or a lack thereof) in modern society. Arianna organised and delivered this session, from Marina Jirotka and Carolyn Ten Holter's (University of Oxford) joint keynote on responsible technologies and innovation to Charith Perera's (Cardiff University) work on detecting cyber attacks in smart buildings. These were followed by a panel discussion that kicked off with reflections from Klaus Dodds (RHUL) on the recent Integrated Review of UK foreign and security policy, and Jason Nurse (University of Kent) on the role of cyber insurance.

The final summer school session, 'Threats', opened with Kenneth Geers (Atlantic Council) talking about how closely related malware proliferation and geopolitics are. This relationship, he claimed, is often overlooked, due to local counter-intelligence over-investing in controlling particular spaces. This can mean that they don't "see" local criminals acting to attack international targets. This was followed by Leonie Tanczer, Genevieve Liveley and Jenny Radcliffe (Human Factor Security) discussing the importance of acknowledging cognitive biases. Their conversation around the difficulties in negotiating human landscapes when developing cyber security systems and processes could have run for hours but sadly the day had to end! David Balson (Ripjar) delivered an exciting talk on how the line between traditional nation state cyber behaviour and criminal activity is blurring, and Ciaran Martin (former head of the National Cyber Security Centre) delivered the closing keynote.

To keep up with Anjuli's work going forward, follow her on Twitter (@AnjuliRKShere) and check out her recent interview on the Tin Foil Hats Club podcast.

For more information on what Claudine is up to, follow her on Twitter (@ClaudineTinsman) and check out the panel she recently moderated "The dark web: Violence and radicalisation online" for CogX 2021.

Arianna recently submitted her thesis on user engagement as a way to increase participation in online medical research platforms. She is a Senior Analyst in Information Security, currently focusing on mission critical communication and smart factories.

Interesting and informative interactions: why we need to engage with consent.

Arianna Schuler Scott (CDT16)

I am a data protection expert, specialising in communication and consent practices. I am interested in consent as it was designed to protect people but is now often seen as an obstacle (e.g., online data requests, or 'cookies') by both users and designers. Users are asked to make risk calculations and take on liability when sharing data. These requests rarely offer information that is of interest, transfers knowledge or provides opportunity for active participation, so users tend to disengage from this decision-making process. In other words – they don't put too much thought into it. In security, we often blame users for making poor decisions, describing them as a key vulnerability. Choice architects and decision-makers need to engineer environments and tools that inform and support users to engage with these decisions. Users cannot be the weakest link if the "chain" is made of rope.

During my DPhil I collaborated with a study of rare genetic diseases, evaluating and improving engagement within a digital platform asking participants to fill out questionnaires twice a year. Academic analysis, alongside research priorities co-designed with participants and researchers resulted in an intervention making information about the study relevant to those taking part, and easier to understand. I wanted to explore the potential benefits of an informed consent process providing information over the course of a study, rather than simply at the beginning. A key benefit we saw was retained participation over time (questionnaire submissions). This enriched the dataset available to researchers. Engaging with different stakeholders (participants and the research team) throughout design, development and evaluation made this intervention relevant, useful and high-impact as these changes are now part of the live study. What follows is an overview of this work.

Informed consent was established in response to the atrocities of World War II, to protect individuals from unethical medical research practices. As research moves online, fundamental issues with informed consent have become magnified: online consent options tend to overload readers with information (Obar and Oeldorf-Hirsch, 2018), and are designed to coerce the reader into making immediate judgements without understanding the information provided (Millett et al, 2001). As researchers, we design our own informed consent processes, so we are responsible for how different options are presented. We act as "choice architects" (Thaler and Sunstein, 2008). Online informed-consent processes normalise information overload and consent fatigue, which violates the original purpose of informed consent: to protect.

Information overload and consent fatigue are also found in research practice. One solution is to create consent processes that inform people over time rather than all at once, which places a burden of communication exactly where it should be: on researchers rather than their participants. Researchers must architect consent choices in a way that is easy to understand because it is unfair to ask people who are not research experts to understand dense, technical language. Such an approach violates the 'informed' nature of consent.

Someone giving consent makes a decision to trust whoever is asking for that consent, and this decision relies on perceptions of transparency and accountability (Aitken et al, 2016). Even in a sector as highly regulated as research, there are fundamental problems with informed consent practices; participants can rarely recall basic information about a study (Dixon-Woods et al, 2007). What we call "informed consent" is unfit for purpose. Fortunately, researchers can demonstrate transparency and accountability using inclusive and responsive practices. Table 1 provides examples of such practices from my own work:

When sharing information, participants are happy to delegate subject-specific knowledge to domain experts, but prefer to retain overall control of how they take part

Inclusive practice	**Responsive practice**
Online focus groups (4 groups) asking participants for feedback on research direction.	A two-page report thanking participants for their time and describing project progress.
Email requests asking for individual feedback (which would validate research findings).	A 2-minute progress update and graphics distributed via social media.

Table 1: examples of inclusive and responsive practice.

Figure 5: modelling enhanced feedback.

(Nissen et al, 2019). As a cybersecurity specialist, I focus on data protection and privacy. While privacy is perhaps more often discussed within legal spheres, both concepts are cornerstones to informed consent because individuals are often asked to consent (i.e., accept liability) to vague, future data uses that are often the result of their data being shared with other parties.

Researchers may not always know how data will be used in the future, but ethical approval processes demand that this is an early consideration and ask what processes are in place. When data is re-used, research participants want to know (Ludman et al, 2010), and as this is likely to happen after data collection then informed consent processes must account for this. One way to do this is to build data practices that invite participant input (e.g., their experience taking part) and aim to respond (e.g., thanks, and indicate project outputs). Inclusive and responsive practices can be woven throughout research.

My work focused on how Dynamic Consent, which aims to inform participants over time, had been implemented within the project I worked with, the Rare and Undiagnosed Diseases studY (RUDY). Launched in 2014 RUDY focuses on rare conditions, or 'health orphans' (e.g., osteogenesis imperfecta, myeloma and fibrous dysplasia of bone) which often restrict physical and mental capacity, and impact life expectancy (Schieppati et al, 2008). Rare conditions are under-researched, so there is little data to draw from. RUDY is valuable in this regard, and the project is unique because data is entered by participants rather than their clinicians. While RUDY illustrates a shift towards inclusive and responsive practice, in talking to researchers and participants I found that while the research team (n=4) had prioritised 'the participant experience', participants (n=53) could not tell me anything about the study itself. These interviews defined "engagement" as a combination of interest, knowledge-transfer and opportunity to participate.

To find out what participants wanted to know, I ran online focus groups to ask them and discussed these findings with RUDY to see what could be done in the time available. Rather than create a new website, I worked with RUDY researchers on a software development cycle to curate and enhance the information available on the research platform. I ran user tests with participants to assess how engaging these changes were: that they were interesting, transferred knowledge and signposted people towards active participation. Figure 2 shows the website in its current form - the number of participants was highlighted and made more prominent, and a suite of pages was added under "Find out more". Technical terms were made more accessible, e.g., 'different phenotypes' was changed to 'differences between individuals with the same rare disease'. The enhanced feedback intervention provided information about how personal data was used, and consisted of five criteria (see Figure 1)

The enhanced feedback intervention measured the number of questionnaires submitted as a percentage of the number of questionnaires requested (completion rate). General participation increased by 5%, and further, there was a 30% increase in first-time questionnaire completions. In other words, people who signed up to take part in the study were more likely to complete the questionnaires they were asked to do.

In sum, 'engaging' research is interesting, informative and interactive. There should be opportunities for participants to take part in the research process; provide feedback on findings, co-write research papers or consult on an advisory panel. Engaging research is also inclusive and responsive. Inclusive data practices seek input and responsive practices offer feedback. Participatory research relates to seeking input outside of study participation, which improved the validity and relevance of this research. In short, the best work is the result of a two-way conversation. For such conversations to happen, however, researchers have to commit to these principles and develop these practices.

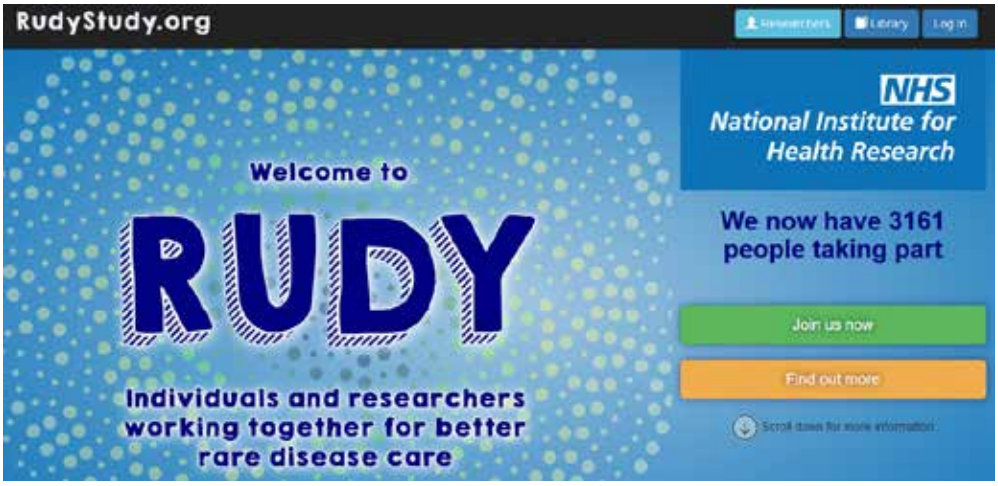

Figure 6: RUDY's project website (available at https://research.ndorms.ox.ac.uk/rudy).

IoT network traffic analysis: opportunities and challenges for forensic investigators?

Tina Wu (CDT14)

As IoT devices become more incorporated into our daily lives, their always on approach makes them an ideal source of evidence. While these devices should use encryption to protect sensitive information, in reality this is not always the case e.g. some expose sensitive data like credentials in cleartext. In this paper, we have conducted an extensive analysis on the communications channels of 32 IoT consumer devices. Our experiments consisted of four main parts; first we carried out a port scan to determine if any ports can be exploited and thus gain remote access. Second, we looked at whether any of the devices used encryption and if not what type of content was exposed. Third, we used the network traffic 'metadata' to identify the destination the data terminated. Finally, we examined the communication between the mobile app and the cloud to see if it can be easily exploited using a proxy server. Our findings show that the majority of devices have remote access unavailable. We found the Shannon entropy test a useful pre-test in identifying unencrypted content. Although many devices encrypted their data, we found several in particular smart cameras would send data in cleartext when they detected motion or during updates. We found the majority of data transverses to the US and stored on Amazon servers with most devices contacting multiple destination. Lastly, we discovered many of the IoT device's mobile apps can be easily exploited using a HTTP Proxy.

1 Introduction

IoT devices and the potential evidence on them become more and more important when working on criminal cases. For instance, a recent criminal investigation involving data from a suspect's Fitbit device proved useful in a murder investigation. Police were able to use the heart rate data that showed a spike at the time of the alleged crime (Buhecker, 2018). On the other hand, a recent survey by Wu et al. (2019) found that over 50%of practitioners did not feel prepared to handle IoT devices and that there is a shortage of tools for IoT forensics. When IoT de-vices are involved in a forensic investigation, it is important to make a decision on how to collect evidence, e.g., memory, in-ternal storage or network layer where this work focuses on the latter approach. Typically, this involves examining the network traffic between the devices and systems it communicates with (e.g., cloud, mobile app) looking for unencrypted information (Servida & Casey, 2019; Kayode & Tosun, 2019). Additionally, if the data is encrypted, the 'metadata' from the network traffic can also be helpful, e.g., the country in which the data resides. Often data is stored in multiple countries creating challenges for investigators when various regulations are involved.

While there has been a significant amount of research on IoT devices from various angles, most existing research does not focus on the forensics implications. Therefore, this article will focus on the following four research questions:

RQ1 Does a device expose ports that allow an investigator to connect / access a device?

RQ2 Do IoT devices utilize encryption when sending/receiving information from the cloud and corresponding App?

RQ3 To which countries do the IoT devices and Apps communicate / establish connections (which is an indicator where data resides)?

RQ4 Do IoT device applications (Apps) utilize encryption when sending / receiving information?

To answer these questions, we examined the network traffic of 32 consumer IoT devices (we bought 17 devices; the remain-ing 15 were part of an existing dataset).

2 Methodology

For our experiment, the following stepwise procedure was used:

IoT device selection: 17 IoT devices were selected which were either wired or had wireless connectivity. In detail, we had: 3 smart hubs, several home automation devices (3 smart plugs, 2 smart bulbs), 6 smart cameras and 3 voice assistants.

Datasets: For our experiments, we used network traffic ob- tained from 2 sources. The first source was data we col- lected from 17 IoT devices connected to our testbed,

the second source was from an existing dataset. Given that IoT devices often come with supporting apps, e.g., for con- figuration, capturing all communication channels required various setups.

Port scanning — After setting up all devices, a port scan was completed to determine which ports are open on the IoT devices. We used Nmap and a quick type of scan for open TCP and UDP ports, using the following command nmap -sS -sU -p 0-65535 [deviceIP] (all ports). This step served as an active approach to analyse the devices. Once the port scan was completed, we tried to connect to open ports using appropriate software, e.g., a browser for port 80 / 443, a SSH-client for port 22, and so on. address, passwords, usernames). As an initial step, the payload was analysed using the Ent2 tool to look for traffic most likely to be unencrypted increasing the chances of finding information. We set the entropy test threshold to 7 (value of the test is between 0-8) to avoid missing unencrypted traffic.

Analysis of metadata. — In addition to the payload, the metadata was utilised where we primarily focused on the location of the connected cloud services. Therefore, a python script was developed to extract the destination IP address, host field and the number of bytes sent to that server from each IoT device, as each IoT device can contact many different servers / destinations. We used the destination IP address to identify the location using GeoIP database and the host address using WHOIS consisted of utilising encryption, HTTP proxy and net- work traffic metadata.

Results: — Finally, we presented the results of the various experiments.

Tool creation: — In parallel to analysing the data, a tool was cre- ated that helped us to more effectively handle the data. The tool is shared and will enable examiners to automatically extract cleartext data and identify where the data termi- nated.

2.1 Port scanning

To answer RQ1, we carried out port scans to identify open ports which may allow to connect to a devices. The result from the scans concludes that the majority of devices used well-known and proprietary ports (port range from 0-49151) we found 8 of the 17 devices used 'upper' TCP/UDP ports which range from ports 49152-65535. 3 devices used TLS/SSL (443), 2 had port 80 open which is typically used to run a HTTP and 1 allowed SSH (22, the Vera Plus hub). The root password to access SSH for this device was written on the hub, so root shell access was easy to gain. The Victure cam exposed a large num-ber of TCP (2) and UDP (19) ports. In contrast to the 2 smart cameras Xiaomi and Yi where all ports were closed. This is beneficial from the security prospective but prevents an inves-tigator gaining remote access to the device and acquiring the filesystem using traditional forensic tools. Often proprietary ports were used to communicate with the app. For instance, the TP-Link devices have port 9999 open in order to control the device using the mobile app.

2.2 Network traffic analysis

We manually analysed all captured network traffic using NetworkMiner[1] and Wireshark. In NetworkMiner we used data to identify the IoT cloud infrastructure. Details about the tool are provided in Section 4.

3 Results

3.1 Utilization of encryption

This section addresses RQ2 wherefore we examined 32 IoT devices and found the majority had secure communication channels especially when the mobile app communicated with the cloud. We examined the unencrypted network traffic for any evidence potentially useful to an investigation. Although the majority of the devices encrypted their content unexpectedly, some devices would send in cleartext video of the detected motion and unique identifiable data during updates. Note, these devices mostly use encrypted channels but performed some actions using unencrypted channels. Many of the devices used secure protocols TLS/SSL. How- ever, we found 9 devices transmitted to the cloud or mobile app with no encryption (HTTP) or partially encrypted (TLS/SSL and HTTP). 7 devices used no encryption between the device-to-cloud and 3 devices used no encryption between the mobile app-to-device.

[1] https://www.netresec.com/?page=NetworkMiner

		Device-to-cloud			Mobile app-to-cloud			Mobile app-to-device		
Category	**Device Model**	**Entropy**	**Cleartext**	**Protocol**	**Entropy**	**Cleartext**	**Protocol**	**Entropy**	**Cleartext**	**Protocol**
SH [1]	Samsung SmartThings hub (v2)(wired)	7.78	✗	TLSv1.2	-	✗	TLSv1.2	-	✗	TLS1.2
	Phillips Hue Bridge(wired) [3,4]	7.72	✓	HTTP/ TLSv1.2	7.99	✗	TLSv1.2	7.81	✗	TLS1.2
	Vera plus hub(wired)[2]	7.87	✗	TLSv1.2	6.24	✗	HTTP/TLS	6.54	✓	HTTP
HA [1]	iBlockcube smart plug	7.74	✗	TLSv1.2	7.22	✗	TLSv1.2	7.45	✗	IPDC
	Amazon smart plug	7.80	✗	TLSv1.2	-	✗	TLSv1.2	7.54	✗	-
	TP-Link plug(HS110)	7.74	✗	TLSv1.2	7.74	✗	TLSv1.2	7.56	✗	TLSv1.2
	TP-Link bulb(LB100)	7.72	✗	TLSv1.2	7.20	✗	TLSv1.2	7.23	✗	TLSv1.2
	LE LampUX	7.87	✗	TLSv1.2	7.28	✗	TLSv1.2	7.91	✗	IPDC
	LiFX lightbulbs†	6.12	✓	TLSv1	-	✗	-	-	✗	-
	iHome†	7.59	✗	TLSv1.2	-	✗	-	-	✗	-
	Nest Protect smoke alarm†	7.67	✗	TLSv1.2	-	✗	-	-	✗	-
VA [1]	Amazon Echo(2nd gen)	7.99	✗	TLSv1.2	-	✗	TLSv1.2	-	✗	-
	Amazon Echo(3rd gen)	7.97	✗	TLSv1.2	-	✗	TLSv1.2	-	✗	-
	Google Home Mini	7.99	✗	TLSv1.3	7.72	✗	TLSv1.2	7.91	✗	TLSv1.2

[1] Smart Hubs (SH), Home Automation (HA), Smart Cameras (SC), Voice Assistants (VA), Smart healthcare and Miscellaneous (M)
[2] On the Vera hub we connected the Aeotec Door/Window sensor Gen5 (ZW120-C), this device uses Z-wave. In order to generate data as none of the other devices were compatible with this hub
[3] On the Phillip Hue Bridge we connected a Phillips lightstrip that uses Zigbee
[†] Devices from the existing dataset

Table 1: The 32 IoT devices used in the experiments, the highlighted entropy values correspond to devices that did not use encryption

To identify unencrypted communication, we started with connections that had an average entropy score of 7 or below these were: Vera plus hub, LIFX bulb 4 , D-Link camera, Samsung camera† , Insteon camera†, Victure cam, Wansview cam, Withings scale† , Withings monitor† and Withings sleep sensor† , the results of the entropy test are shown in Table 1 and 2. While the entropy test correctly detected the devices that used unencrypted traffic, the LiFX lightbulb† used encrypted protocols (TLS/SSL). It was discovered in previous research that the reason for the low entropy of the LiFX lightbulb† was because it was encoded and not in a human-readable format Loi et al. (2017). the cleartext dictionary file to carry out a customised search and conducted a string search (in various encoding) using Wireshark on the network traffic of each device. We looked for device identifiers (e.g. MAC address, serial numbers) and personal information created during setup (e.g. names, email).

3.2 Network traffic metadata

In the following we focus on RQ3 and identify the destinations that the data transverses. Figure 1 shows the flow of traffic to the top 10 countries with the height of the bands corresponding to the number of bytes sent by each IoT device. Overall, the data for 26 of the 32 devices terminated in the US. This is unexpected as the closest Amazon data centre to our testbed (UK) is in Ireland (WikiLeaks, 2018). While trying to establish a reason for our data being sent to the US instead of Ireland, we found that unlike Google and Facebook that provide approximate locations to their data centres, Amazon do not advertise this freely. Furthermore, we found 75% (24) of the IoT devices sent data to multiple destinations, while the remaining devices sent data to a single destination. This is an interesting from an investigative viewpoint, as multiple destinations and jurisdictions, will potentially cause delays in gaining access and securing the data. Also it will also require communication between countries with different legal systems. To gain a better understanding of the type of cloud infrastructure, data was sent to, in this part of the analysis we focused on the most frequently used type and found 21 of the 32 (66%) devices contacted a server belonging to Amazon.

3.3 HTTP Proxy

In this section we used a proxy server to examine the encrypted contents of the network traffic (see RQ4). We found several of the mobile apps allowed a proxy connection while the remaining implemented certificate pinning. We found a case where a device would unexpectedly take snapshots and there was no setting to control this behaviour. On the same mobile app, we found the username and password were sent Base64 encoded. As this mobile app controls several devices, this means we were able to gain access to them all.

2 https://github.com/rsmith-nl/ent (last accessed 2020-03-15).
3 https://dev.maxmind.com/geoip/geoip2/geolite2/ (last accessed 2020-03-15).
4 † Marked IoT devices are from the expanded test scenario.

1 https://www.netresec.com/?page=NetworkMiner
1 Smart Hubs (SH), Home Automation (HA), Cameras (C), Voice Assistants (VA), Smart healthcare and Miscellaneous (M)
2 On the Vera hub we connected the Aeotec Door/Window sensor Gen5 (ZW120-C), this device uses Z-wave. In order to generate data as none of the other devices were compatible with this hub
3 On the Phillip Hue Bridge we connected a Phillips lightstrip that uses Zigbee
† Devices from the existing dataset

		Device-to-cloud			Mobile app-to-cloud			Mobile app-to-device		
Category	**Device Model**	**Entropy**	**Cleartext**	**Protocol**	**Entropy**	**Cleartext**	**Protocol**	**Entropy**	**Cleartext**	**Protocol**
VA [1]	Amazon Echo(2nd gen)	7.99	✗	TLSv1.2	-	✗	TLSv1.2	-	✗	-
	Amazon Echo(3rd gen)	7.97	✗	TLSv1.2	-	✗	TLSv1.2	-	✗	-
	Google Home Mini	7.99	✗	TLSv1.3	7.72	✗	TLSv1.2	7.91	✗	TLSv1.2
SC [1]	TP-Link cam(KC100)	7.99	✗	TLSv1.2	7.97	✗	TLSv1.2	7.81	✗	TLSv1.2
	D-Link cam(DCS-932LB)(wired)	7.82	✗	TLSv1	7.78	✗	TLSv1.2	6.40	✓	HTTP
	Xiaomi cam	7.91	✓	HTTP	7.91	✗	TLSv1.2	-	✗	-
	Yi cam	7.92	✗	TLSv1.2	7.59	✗	TLSv1.2	-	✗	-
	Netatmo Welcome camera†	7.20	✗	TLSv1.2	-	✗	-	-	✗	-
	TP-Link Day Night Cloud camera†	7.41	✗	TLSv1.2	-	✗	-	-	✗	-
	Samsung SmartCam†	6.05	✓	HTTP/ TLSv1.2	-	✗	-	-	✗	-
	Nest Dropcam†	7.78	✗	TLSv1.2	-	✗	-	-	✗	-
	Insteon Camera(wired)†	6.57	✓	HTTP/ TLSv1.2	-	✗	-	-	✗	-
	Victure cam (PC530)	6.06	✗	-	6.45	✗	HTTP	5.74	✗	HTTP
	Wansview cam (Q5)	7.86	✓	TLSv1.2	5.86	✗	HTTP/TLSv	6.68	✗	HTTP
M [1]	Netatmo weather station†	7.40	✗	-	-	✗	-	-	✗	-
	Triby Speaker†	7.59	✗	TLSv1.2	-	✗	-	-	✗	-
	PIX-STAR Photo-frame†	7.48	✗	TLSv1.2	-	✗	-	-	✗	-

[1] Smart Hubs (SH), Home Automation (HA), Smart Cameras (SC), Voice Assistants (VA), Smart healthcare and Miscellaneous (M)
[2] On the Vera hub we connected the Aeotec Door/Window sensor Gen5 (ZW120-C), this device uses Z-wave. In order to generate data as none of the other devices were compatible with this hub
[3] On the Phillip Hue Bridge we connected a Phillips lightstrip that uses Zigbee
† Devices from the existing dataset

Table 2: Summary of findings when looking at unencrypted communication.

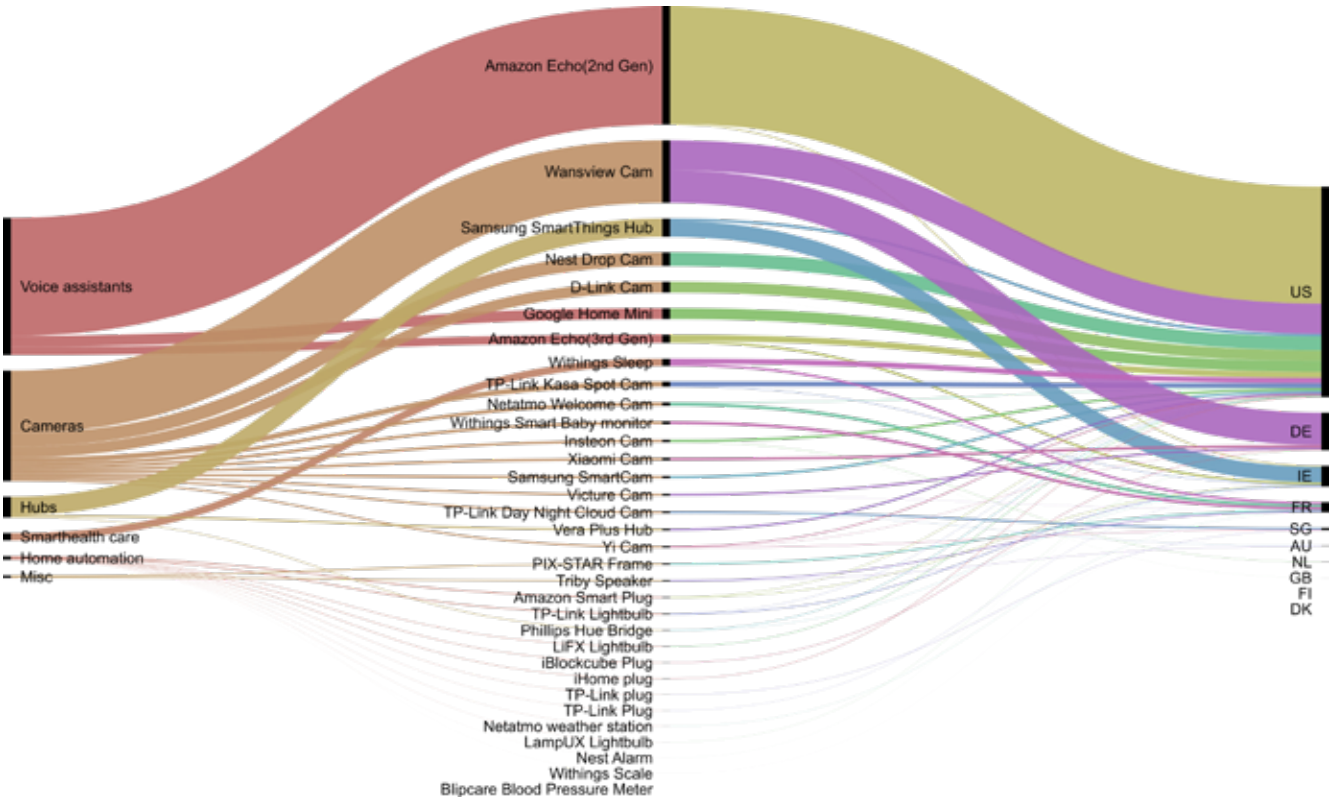

Figure 1: The top 10 data destinations grouped by the number of bytes transmitted by each device to their final destination.

4 Tool creation and Evaluation:

Although our results can be found using separate open source tools, it would require an investigator considerable amount of time to manually extract the data. Consequently, we implemented part of the process in a tool called IoT Network Analyzer, constructed in Python to automate the process. The tool is currently a working prototype and has the following four main features:

Entropy calculation: In order to identify sessions that are in clear text, the Shannon entropy is utilised and calculated over all packets within a session. This calculation is only done on the data portion of the segment; header fields such as ports or IPs of the segment are ignored.

Location of data: The source and destination IP address are extracted to make an assumption of the geolocation. This was accomplished using IP Stack5 which is a "powerful, real-time IP to geolocation API".

5 https://ipstack.com/product

Usage of secure ports: First a list of ports is created with the total number of occurrences for each port. This list is then checked against the pre-defined list of 22 secure ports (??) which can be modified if needed. All this information is then mapped with their corresponding source and destination IP address.

Cleartext extraction: Lastly, our tool tries to extract cleartext. The tool also has the following additional features:

IP addresses and connections overview: A list of the source and destination IP address and their respective number of packets sent.

Packets overview: It analyzes the structure of the different packets and outputs a list of the number of packets including; the total, raw layer, UTF8, byte, ARP, DHCP, DNS, and finally the total amount of packets processed.

Command line interface (CLI): The tool has an interactive CLI based on the cmd26 library. This allows an interactive usage of the tool, data extraction and output capabilities. The tool is available at: https://github.com/Dyvels/ IoTAnalyzer

4.1 Tool assessment

For the evaluation we tested the tool's four main features and the performance of the tool, utilising 5 PCAP files that we knew had cleartext data and another 5 PCAP files with no cleartext data.

Entropy calculation. To evaluate the entropy part, we com- pared our results to a similar tool called Ent and they provided identical results. Our tool has the benefit of automatically calculating the entropy based on the payload whereas Ent requires manual filtering of the IP addresses and extraction of the payloads. This means using our tool an investigator is able to calculate the entropy on the various communication channels at the same time.

Location of data. As we use the external IPStack service for detecting the location, we briefly compared these results with two similar services: IP2Location-Lite7 and Geolite2 (Max- Mind)8 . These have been suggested by Gharaibeh et al. (2017) as they have an accuracy of 80% at country-level. The results are shown in Table 3 and show little difference between the databases, meaning any of these databases would be equally suitable. Note: a method to assess the accuracy of IPStack would be to compare it against a "ground truth" database that has true geographical location for each IP address, however such a database does not exist (Gharaibeh et al., 2017).

To test the accuracy of this feature we compared the results to those of a similar tool T-shark. We anal- ysed the 10 PCAP files using T-shark to filter for the ports and then compared these results to our tool, with both tool showing identical results.

Cleartext extraction. To identify cleartext traffic, we manually examined each PCAP file in NetworkMiner9 followed by a comparison using our tool. NetworkMiner and our

6 https://github.com/python-cmd2/cmd2
7 https://www.ip2location.com/demo
8 https://dev.maxmind.com/geoip/geoip2/geolite2/ Usage of secure ports.

Performance evaluation.

tool iden- tified the same PCAP files containing cleartext information.

The performance test (with respect to processing speed, Central Processing Unit (CPU) and memory usage) was conducted on a system running Windows 10 on a Intel(R) Core(TM) CPU i5-4670K with 32GB RAM. The re- sults from processing the PCAP files of sizes from 500 - 75,000 kilobytes (kbs) and showed that it takes around 1 second to 6 minutes. Table 4 shows that the tool is not resource intensive as it consumes around 60% CPU and has small memory footprint.

Devices	IPStack (IoT Network Analyzer)	GeoIP	IP2location
Withings smart scale	France	France	France
Samsung smartcam	Australia, New Zealand, US	Australia, Cambodia, Germany, Luxembourg, UK, US	Australia, Japan, New Zealand, UK, US
Vera hub	US	US	US
DLink cam	Ireland	Ireland	Ireland
Insteon cam(wired)	Australia, China, Ireland, Japan, Netherlands, UK,US,	Germany, Ireland, Singapore, UK, US	China, France, Germany, Taiwan, UK, US
TPlink cam	Germany, Ireland, Singapore, UK, US	France, Germany, Ireland, Singapore, US	China, France, Germany, UK, US
Amazon Echo	Ireland, UK, US	Ireland, UK, US,	Ireland, UK, US
Google Mini	Mexico, US	Mexico, US	UK, US
Phillips Hue	Germany, Ireland, Singapore, Spain, UK, US	Germany, Ireland, Singapore, Spain, UK, US	Germany, Ireland, UK, US
Xiaomi cam	China, France, Germany, Netherlands, Singapore, Taiwan, UK	China, France, Germany, Ireland, Japan, Netherlands, Singapore, South Korea, Taiwan, UK, US	China, France, Germany, Netherlands, Singapore, Taiwan, UK

Table 3: Countries identified using IoT Network Analyzer and results from two other geolocation databases

Communication Channel	Device Model	Entropy	Protocol	Cleartext data
Device-to-cloud	LiFX lightbulbs	6.12	TLSv1	?
	Samsung SmartCam	6.05	HTTP/TLSv1	MAC address, username, serial number, timestamp and user specified device name
	Insteon Camera(wired)	6.57	HTTP	Port numbers, MAC address, public IP address, unique ID
	Withings smart scale	4.68	HTTP	Weight,height,age, sex etc.
	Triby speakers	7.59	TLSv1.2	Username, serial number, MAC address
	Xiaomi camera	7.91	HTTP	URL and timestamp of captured motion, MAC address
Mobile app-to-device	D-Link camera	6.40	HTTP	JPEG images
	Victure camera	6.45	HTTP	API access key, access token

Table 4: Results from the performance evaluation

5 Conclusion

In this paper we analysed 32 different IoT devices with respect to their network settings to better understand implications for forensic practitioners. The devices that sent data in clear- text were smart cameras and smart healthcare devices. Smart healthcare devices exposed personal data potentially useful for an investigator to identify features of a person of interest. An- other finding was when the Xiaomi camera detected motion, it would send an unencrypted packet containing not only the captured video but also the credentials to an AWS server. One of the forensic challenges discussed in existing work is the storage of IoT data in multiple locations which then leads to different jurisdiction (Yaqoob et al., 2019; Hegarty et al., 2014). In our findings, the majority of data was sent to the US, how- ever, there were also plenty of other countries (see Section 3.2). Note, this was despite our testbed being based in the UK and the 2nd dataset collected in Australia. There would be less le- gal complexity if the data was stored in data centres within the country of origin. From examining the encrypted content, we found an unexpected event where the Kasa app for the TP link devices would take snapshots which could not be controlled. On the same mobile app, we found the username and password used a weak HTTP based authentication that could easily be de- coded. Several other mobile apps accepted proxy connections, however, these were smart plugs and did not have any sensitive data.

9 https://www.netresec.com/?page=NetworkMiner

References

Buhecker, R. (2018). FitBit and Wearables Proving to be a Valuable Tool in Forensics. https://www.secureforensics.com/blog/fitbit-and-wearables-proving-to-be-a-valuable-tool-in-forensics.

Gharaibeh, M., Zhang, H., Shah, A., Ensafi, R., Huffaker, B., & Papadopou- los, C. (2017). A look at router geolocation in public and commercial databases. Proceedings of the ACM SIGCOMM Internet Measurement Con- ference, IMC, Part F131937, 463–469. doi:10.1145/3131365.3131380.

Hegarty, R., Lamb, D. J., & Attwood, A. (2014). Digital evidence challenges in the internet of things. In INC.

Kayode, O., & Tosun, A. S. (2019). Analysis of IoT Traffic using HTTP Proxy. IEEE International Conference on Communications, 2019-May, 1–7. doi:10.1109/ICC.2019.8761601.

Loi, F., Sivanathan, A., Gharakheili, H. H., Radford, A., & Sivaraman, V. (2017). Systematically evaluating security and privacy for consumer IoT devices. IoT S and P 2017 - Proceedings of the 2017 Workshop on Inter- net of Things Security and Privacy, co-located with CCS 2017, (pp. 1–6). doi:10.1145/3139937.3139938.

Servida, F., & Casey, E. (2019). Iot forensic challenges and opportunities for digital traces. Digital Investigation, 28, S22–S29. URL: https://doi.org/10.1016/j.diin.2019.01.012. doi:10.1016/j.diin.2019.01.012.

WikiLeaks (2018). Map of Amazon's Data Centers. https://wikileaks.org/amazon-atlas/map/.

Wu, T., Breitinger, F., & Baggili, I. (2019). Iot ignorance is digital forensics research bliss: A survey to understand iot forensics definitions, challenges and future research directions. In Proceedings of the 14th International Confer- ence on Availability, Reliability and Security (pp. 1–15).

Yaqoob, I., Hashem, I. A. T., Ahmed, A., Kazmi, S. A., & Hong, C. S. (2019). Internet of things forensics: Recent advances, tax- onomy, requirements, and open challenges. Future Generation Com- puter Systems, 92, 265 – 275. URL: http://www.sciencedirect.com/ science/article/pii/S0167739X18315644. doi:https://doi.org/10.1016/j.future.2018.09.058.

Angeliki Aktypi

Supervisor: Kasper Rasmussen, Department of Computer Science

Angeliki holds a Diploma (M.Eng. equivalent) in Electrical and Computer Engineering from the Democritus University of Thrace (Greece, 2014). She received her M.Sc. in Communications & Computer Security jointly delivered by Telecom ParisTech and Eurecom (France, 2014), as a 'VRIka!' Scholar of the French Embassy in Greece. She was enrolled at the Centre for Doctoral Training in Cyber Security, at the University of Oxford in 2016 as a Linacre College student member, funded by the EPSRC. After completing her training year, she joined the Department of Computer Science, under the supervision of Prof. Kasper Rasmussen. During her studies, she has pursued internships at British Telecom (the U.K., 2014), at Thales (France, 2016) and at ICS-FORTH (Greece, 2020), all focused on security research topics.

Her DPhil thesis focuses on the design of secure and decentralised protocols for the communication of entities in the IoT domain and is partially supported by a Russel Group Studentship by British Telecom.

DPhil Thesis: Discovery of and Access to Resources between Entities within IoT Systems

The rapid increase in the use of connected devices (e.g., wearables, smart locks) and cloud applications (e.g., collaboration platforms, big data tools) established trends for a fully programmed and distance manageable lifestyle and shifted the character of the society to be open and network-oriented. The proliferation on connectivity combined with the increase in cyber-crime and the terrorism expansion in cyber space have created an urgent need to rapidly advance our security countermeasures and re-think of traditional approaches. Recognising this need, this DPhil thesis is going to explore and propose the deployment of security infrastructures in connected environments, many times referred to as the Internet of Things. In particular, the objective is to develop secure and resilient to attacks protocols that enable effective discovery and access to resources between entities within IoT systems. Entities include devices (such as sensors, actuators, embedded systems), but also include software-only virtual entities (such as virtual service instances in a cloud or fog computing context).

Publications

SeCaS: Secure Capability Sharing Framework for IoT Devices in a Structured P2P Network, Angeliki Aktypi, Kubra Kalkan and Kasper B. Rasmussen, In 10th ACM Conference on Data and Application Security and Privacy (CODASPY '20). Pages 271–282. ACM. March, 2020.

Unwinding Ariadne's Identity Thread: Privacy Risks with Fitness Trackers and Online Social Networks, Angeliki Aktypi, Jason R.C. Nurse and Michael Goldsmith, In 1st International Workshop on Multimedia Privacy and Security in conjunction with the 24th ACM Conference on Computer and Communication Security (CCS'17). ACM. October, 2017.

Aaron Ceross

Supervisor: Andrew Simpson, Department of Computer Science

Aaron has a background in law and prior to joining the CDT, he worked as a researcher in data protection, privacy, and surveillance at the University of Groningen in the Netherlands as well as the University of Malta. Aaron recently completed the MSc in Computer Science from the University of Bristol in order to expand his technical abilities as well as be able to engage in more multidisciplinary research.

DPhil Thesis: Metrics and models for privacy engineering.

This research proposes to investigate how privacy risks can be more effectively articulated to information system designers. The goal is to be able to establish a link between a systems-level risk analysis and regulatory outcomes. This entails the following questions: (i) what are the challenges faced by information system designers with regards to privacy, as both a broad concept and multilevel values? (ii) what information needs to be made available in order for these challenges to be overcome? (iii) is this information drive measurably effective privacy practices for systems design? This research seeks to match the causes of regulatory liability to specific system design implementation, focusing on being able to accurately link normative regulatory values to a system's operational metrics. This results in measures which may inform organisations of the risks to personal data. One of the means by which empirical data may be gathered for privacy failures within information systems is from events such as data breaches, and other events recorded by authorities. These data may be used as ground truth from which to develop objective, generalisable metrics for privacy risk. This thus provides a more quantitative measures for privacy risk analysis, which hitherto has been largely qualitative, subjective measures. This would therefore better inform the information systems engineering process.

Publications

"Examining data protection enforcement actions through qualitative interviews and data exploration" at the 37th annual conference of the British and Irish Legal, Education and Technology Association (April 2017), and will be published in the International Review of Law, Computers and Technology.

"The use of data protection enforcement actions as a data source for privacy economics" at the 3rd International Workshop on Technical and Legal Aspects of Data Privacy and Security (September 2017), to be published in the SAFECOMP 2017 Workshop Proceedings in Lecture Notes in Computer Science

Mini-Project: Understanding threats to anonymisation: Gap-analysis and research directions

Mini-Project: Exploring Liabilities and Remedies for Data Breach

Munir Geden

Supervisor: Kasper Rasmussen, Department of Computer Science

Munir is a final year DPhil student at the Department of Computer Science, working with Professor Kasper Rasmussen. He has received his BSc in Computer Engineering followed by MSc in Engineering and Technology Management from Bogazici University (Istanbul). Before steering into a research-based career, Munir worked for more than five years as a software engineer in the financial IT industry. He came to the UK to pursue another master's degree in Software Systems Engineering at UCL with a research and security focus which brought him to Oxford to gain a better understanding of cybersecurity by employing an interdisciplinary perspective.

Munir's early security research was on malware analysis. For the identification of malware, he has applied machine learning on both statically and dynamically extracted features from samples. Due to the limitations of those data-driven techniques regarding the detection of novel attacks, his doctoral research has been centred around the attestation of software runtime with the aim of more sound approaches, specifically to address control- and data-oriented attacks in different contexts. His contributions include both designing hardware-based and instrumentation-based detection schemes towards this goal. His recent interest also covers compiler-based security, which redesigns specific compiler optimisations to reduce the program's attack surface deliberately.

Publications

Geden, M. and Happa, J., 2018, October. Classification of malware families based on runtime behaviour. In International Symposium on Cyberspace Safety and Security (pp. 33-48). Springer, Cham.

Geden, M. and Rasmussen, K., 2019, August. Hardware-assisted remote runtime attestation for critical embedded systems. In 2019 17th International Conference on Privacy, Security and Trust (PST). IEEE.

Geden, M. and Rasmussen, K., 2020, December. TRUVIN: Lightweight Detection of Data-Oriented Attacks Through Trusted Value Integrity. In 2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom) (pp. 174-181). IEEE.

Andikan Otung

Supervisor: Andrew Martin, Department of Computer Science

Andikan studied Electronic & Electrical Engineering (MEng) at UCL, where he graduated on the Dean's list. After graduation, Andikan began a career in telecommunications, working in London as a Sales Engineer for Ciena - a multibillion dollar networking-infrastructure vendor. As well as (officially) becoming an inventor, Andikan developed an increasing interest in Security, through his work on the government side of sales operations.

As the more perceptive of readers may have already gathered, Andikan is both technically and commercially minded. His interests reflect this and include: IoT Security, Trusted Computing, DDoS, Cryptography, Multi-factor (including location-based) Authentication as well as commercial strategies enabling the fast creation and adoption of new technologies.

Marcel Stolz

Supervisors: Michael Goldsmith, Department of Computer Science

As a Swiss citizen, I got sensitised to security and defensive matters in my military service. Having become an officer (first lieutenant), I continued my compulsory service in our electronic operations unit. During my Bachelor's and Master's studies in computer science I engaged with topics in LTE (4G) networks and user-friendly design. I gained insights into political work as member and president of the TUX party of the University of Bern, discussing topics such as privacy. Thereafter, I worked with Swisscom, Switzerland's national phone operator, in the Big Data Mobility Insights Squad. My current interest lies in combining my technical knowledge with my interest in politics and experience in military leadership. I explore governance aspects of cybersecurity, both on an enterprise and (inter)national level. I have taken a particular interest in the distribution of responsibilities between technology companies and states, and have developed a theory on the neutrality of global social media platforms.

My main research interests are:

- Platform Governance (eg Social Media Platforms)
- National and Enterprise Cyber Governance
- National and Enterprise Cyber Defence Strategy
- Enterprise Cyber Risk Assessment and Management
- Philosophical Foundations of Democracy in Cyberspace.

DPhil Thesis: Neutrality in Cyberspace

Outputs:

I contributed as an expert through a workshop in 2020 and as a reviewer to the following report by Konrad Adenauer Stiftung (KAS) / Stiftung Neue Verantwortung (SNV) Berlin: Herpig S (2021): Die Beantwortung von staatlich verantworteten Cyberoperationen. Impuls zur deutschen Cybersicherheitspolitik, published by Konrad Adenauer Stiftung and Stiftung Neue Verantwortung Berlin.

Jakobsson AK, Stolz MS (2021): Principled Big Tech: European pursuit of technological autonomy. Published in: Strategic autonomy and the transformation of the EU: New agendas for security, diplomacy, trade and technology. Edited by Niklas Helwig. FIIA Report 67.

Stolz MS (2021): Platform Neutrality–A Solution for the Social Media War?, International Conference on Cyber Warfare and Security. Proceedings of the 16th International Conference on Cyber Warfare and Security ICCWS 2021.

Stolz MS (2021): "Competing Interests of Cyberintelligence and Cyberdefence Activities in Neutral Countries", International Conference on Cyber Warfare and Security. Proceedings of the 16th International Conference on Cyber Warfare and Security ICCWS 2021.

Stolz MS, Creese S, et al. (2020): Cybersecurity Capacity Review Switzerland, Global Cybersecurity Capacity Centre on behalf of the Swiss Federal Department of Foreign Affairs (FDFA), GFCE Code: G0517.

Creese S, Hannigan R, El Kaafarani A, Axon L, Fletcher K, Nagyfejeo E, Schuler Scott A, Stolz M (2020): Foresight review of cyber security for the Industrial IoT, Lloyd's Register Foundation, Report Series: No. 2020.1.

Stolz MS (2019): "On Neutrality and Cyber Defence", Proceedings of the 18th European Conference on Cyber Warfare and Security. Academic Conferences and Publishing International.

Oleh Stupak

Supervisors: Greg Taylor, Oxford Internet Institute and Aleksei Parakhonyak, Department of Economics

Oleh holds MSc in Economics from Paris 1 Pantheon-Sorbonne (France) and BSc, MA degrees in International Economics from Taras Shevchenko National University of Kyiv (Ukraine).

Alongside with academia, he obtained more than six years of experience in private sector. Three of which, Oleh held a CEO position in the self-founded company "thelamp". The company was a projection of his curiosity to the innovative technologies. It provided a full range of IT services and specialised on the tailoring of unique software solution for commercial and governmental purposes.

His MSc thesis in Paris 1 was devoted to the research on DDoS (distributed denial-of-service) attack risk for industries. The model developed during that period is capable of calculating the enterprises' chance of being the subject of DDoS attack considering 25 economic and technical parameters.

Oleh is convinced that the future stands in interdisciplinary approaches and cooperation among schools. His desire for knowledge and world outlook brought him to Oxford's CDT in Cyber Security. Currently, Oleh focuses on the emerging cyber security threats for enterprises. His broad area of interests includes: enterprises behaviour and unfair competition in the information environment, cyber security risks.

DPhil Thesis: The Economics of Industrial Cyber Espionage

Publications

Mini-Project: Unfair Competition in the information environment. Industrial information leakage

Mini-Project: Unfair Competition in the information environment. The DDoS attack

Jack Sturgess

Supervisor: Ivan Martinovic, Department of Computer Science

Jack graduated from the University of Surrey with a BSc (Hons) in Mathematics and an MSc (Dist.) in Information Security. He worked as a software engineer at Accenture for 1 year and at IBM for 5 years; he also co-founded a video games company, while at IBM, that released two titles on PlayStation Network and one on Steam.

An avid traveler, Jack volunteered as a conservationist in Arizona and California after his first degree, camping and working in the middle of nowhere while enjoying fantastic landscapes! His interests also include programming, board-gaming, hiking, and number theory.

Jack's research interests include authentication, biometrics, steganography, and the Internet-of-Things. Being part of the CDT has proven to be an excellent way to study the interweave of these subjects and to understand their implications across other disciplines.

DPhil Thesis: Wearable Authentication Using Inertial Sensors

Publications

Sturgess, J., & Martinovic, I., "VisAuth: Authentication over a Visual Channel using an Embedded Image", Cryptology and Network Security (CANS) 2017

Sturgess, J., Nurse, J. R. C., & Zhao, J., "A Capability-oriented Approach to Assessing Privacy Risk in Smart Home Ecosystems", PETRAS 2018

CDT17 Bios

Thomas Burton

Supervisor: Kasper Rasmussen, Department of Computer Science

Thomas studied an MComp in Computer Science (with Security and Resilience) for four years at Newcastle University. His current areas of interest include secure localisation, misusing localisation schemes for attacks, and mesh networking.

At Newcastle he studied a range of topics from system security and information security and trust to high integrity software development and the challenge of developing highly dependable systems. His third and fourth year dissertation projects covered the security related topics of biometric security and authentication, and cryptographic, specifically zero-knowledge proofs. He has also spent several summers working with a health sciences and bioinformatics group. With them he has worked on a range of projects ranging from website development to algorithm development. During this work, he experienced working with people from a range of academic departments and fields.

DPhil Thesis: Secure Urban Audio Localisation

I am looking at the use of smartphones for urban search and rescue. Specifically I am working on secure protocols for using audio to locate smartphone devices to assist in an urban disaster environment. This type of localisation has a number of potential challenges to overcome as a result of limited low level hardware access, the scale over which the protocols are being used, and how easily an attacker can interfere as the cost of mounting at attack is low. Attackers also have several key advantages, such as, being able to transfer information at the speed of light using normal radio communication which is much faster than the sound used for the localisation.

Selina Yoon Cho

Supervisors: Ivan Flechais, Department of Computer Science and Jonathan Lusthaus, Department of Sociology.

Selina is interested in addressing the socio-technical issues that arise in the intersection of crime and technology. Selina holds an MSc (Distinction) in Information Security and a BA in Economics. Her masters dissertation examined the concept of security through obscurity through its design applications in cryptographic obfuscations and image steganography. Prior to joining the CDT, Selina worked as a security consultant in a plant engineering company. Selina was the President of the Oxford Fintech and Legaltech Society, a platform for fostering technological innovations in finance and law through seminars, hackathons, and research.

DPhil Thesis: Self-governing communities in online game cheating

The business of online game cheating is a multi-million dollar industry, growing evermore robust against the anti-cheat measures put in place by the developers. Despite the prevalence of cheating, there is a lack of understanding in how the cheat resources are managed outside the official gaming scene.

This thesis is an exploratory research into online cheating communities seeking to understand the social and management structures. We design our studies based on the theory of self-governance, and analyse the differing levels of trust and compliance that exist among users. We use a mixed-methods approach combining scraped data from the community websites and interviews with the community members. The findings will be used for forming new theories, and better understanding the nature of the resource management in online fringe communities.

Publications:

S. Y. Cho, J. Happa, S. Creese, "Capturing Tacit Knowledge in SOCs," in IEEE Access 2020.

S. Y. Cho and J. Wright, "Into the Dark: A Case Study of Banned Darknet Drug Forums," in 11th International Conference on Social Informatics, Springer International Publishing, 2019..

Tommaso De Zan

Supervisors: Ewart Keep and Liam Gearon, Department of Education; Andrew Martin, Department of Computer Science.

Tommaso is a DPhil student in Cyber Security at the University of Oxford, where he analyzes policies aimed at reducing the cyber security skills shortage. In particular, he investigates whether cyber security competitions affect students' interest in a cyber security career and how.

In the context of his research, he conducted a six month-traineeship at the European Union Network and Information Security Agency (ENISA) in Athens and continues to collaborate with ENISA on topics related to skills development in the EU. Moreover, he was a visiting student at the Center for International Security at the Hertie School in Berlin.

He is also a Research Associate with the Centre for Technology and Global Affairs (DPIR, Oxford University), a member of the Global Forum on Cyber Expertise, and a steering committee member of the Cyber Youth Initiative (Royal United Services Institute).

Prior to his DPhil, he was an Associate Fellow at the European Union Institute for Security Studies and a Researcher at the International Affairs Institute (IAI). Before joining IAI, he interned at the International Peace Research Institute in Geneva.

He holds a Master's degree in International Relations from the University of Bologna and he was an exchange student at the Josef Korbel School of International Studies and Université catholique de Louvain.

DPhil Thesis: Do competitions affect students' interest in cyber security career? The cyber security skills shortage and public policy interventions

Employers have been lamenting for several years the lack of cyber security professionals in the labour market, the so-called cyber security skills shortage. The shortfall of information security workers means that our data, networks and systems are less secure, which might undermine economic development and national security. Governments have scrambled to redress this trend and have designed and implemented policies to increase the pipeline of cyber security professionals. Among these policies, cyber security competitions have sprouted and received the support of governments and industry alike. Nonetheless, it is generally unknown whether skills shortage policies such as cyber security competitions work and how. This dissertation project investigates the outcomes of these competitions, especially whether they affect students' career interest in cyber security.

Publications:

De Zan T., Giacomello G., Martino L. (forthcoming), "Italy", Routledge Handbook of Global Cybersecurity, edited by Manjikian M. and Romaniuk S. N., Routledge, New York;

De Zan T. (2020), "Future Research on the Cyber Security Skills Shortage", in Cyber-Security Education: Principles and Policies, edited by Austin G., Routledge, New York, https://bit.ly/2UMquJ6;

De Zan T. and Di Franco F. (2020), "EU Cyber Security Skills Development: The Certification of Cyber Security Degrees and ENISA's Cyber Security Higher Education Map," European Union Network and Information Security Agency, https://bit.ly/3dTS7Yc;

De Zan T. (2020), "The shortfall of cyber security competencies in Italy" (Italian), Agenda, Enciclopedia Treccani, Roma, http://bit.ly/2U2jOqb;

De Zan T. (2019), "The Italian Cyber Security Skills Shortage in the International Context", Global Cyber Security Center, Rome, https://bit.ly/2OG9flg;

De Zan T. (2019), "Much Ado About the Cyber Skills Shortage", Net Politics, Digital and Cyberspace Policy Program, Council on Foreign Relations, New York, https://on.cfr.org/2tHwSTx;

De Zan T. (2019), "Mind the Gap: the Cyber Security Skills Shortage and Public Policy Interventions", Global Cyber Security Center, Rome, https://bit.ly/2tpedvs;

Seb Farquhar

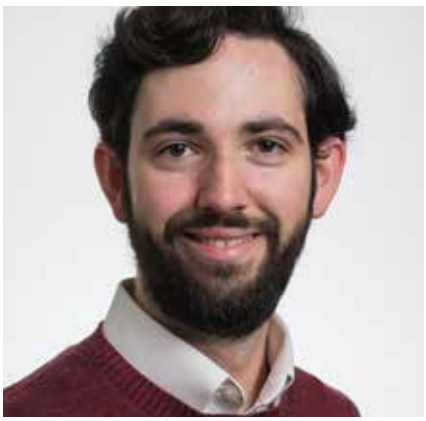

Supervisor: Yarin Gal, Department of Computer Science

Seb is interested in cyber security within machine learning and artificial intelligence. This includes technical work on robust deep learning systems that are able to recognize and safely handle unexpected or adversarial inputs as well as privacy-preserving machine learning and policy questions related to the safe adoption of AI systems. Before beginning this DPhil, Seb worked at the Future of Humanity Institute at the University of Oxford, was Director of the Global Priorities Project, a think tank focusing on global catastrophic risk management, and worked for McKinsey & Co. as a consultant focusing on public sector clients. He has a Master's degree in Physics and Philosophy from the University of Oxford.

DPhil Thesis: Foundations for secure deep learning

In this research project I explore foundations for safe and secure deep learning including:

-Systems capable of detecting anomalies/out-of-distribution behaviour using Bayesian deep learning methods and ensembles.

-Differentially private deep learning systems for learning over extended periods of time or across related simultaneous contexts with managed privacy leakage.

-Deep learning architectures optimized for secure multi-party computation.

Jack Kenny

Supervisors: Catherine Redgwell and Efthymios Papastavridis, Faculty of Law

Jack's background is in public international law. He holds an LL.B degree in Law with European Study and an LL.M degree in International and European Law. Jack is supervised by Professor Catherine Redgwell and Dr. Efthymios Papastavridis at the Faculty of Law. In addition to his DPhil research Jack is a Post-Doctoral Research Fellow at the Hebrew University of Jerusalem on The Prospects for an International Attribution Mechanism for Cyber Operations project. He has several years of research experience including positions at the International Law Programme at Chatham House, the British Institute of International and Comparative Law, and the Amsterdam Centre for International Law. Prior to commencing his doctorate, he held the post of Research Assistant with teaching responsibilities at the Institute for Public International Law, University of Bonn, under Professor Stefan Talmon. Jack has worked on cases involving issues of international and EU law. Current research interests include the relationship between customary law and treaty law, state responsibility, use of force, and the applicability of these topics to new and emerging areas governed by international law such as cyber operations.

DPhil Thesis: Establishing state responsibility for cyber operations

My research considers how existing frameworks of public international law apply to cyber operations. Specifically, the thesis focuses on establishing state responsibility for cyber operations by analysis of applicable primary and secondary rules of international law. The thesis aims to identify and examine primary rules and their application to cyber operations with an analysis of state responsibility and attribution, and discuss possible remedies and forms of restitution.

Klaudia Krawiecka

Supervisor: Ivan Martinovic, Department of Computer Science

Klaudia graduated from the NordSecMob programme in 2017, obtaining a Master's degree in Security and Mobile Computing from two universities: Aalto University and Norwegian University of Science and Technology. Her adventure with computer science began in primary school. She continued to develop her passion during high school and in college. During the second year of her Bachelor's studies, she took an internship in the ICT security office where she was introduced to computer forensics and cyber security fields; in addition, she had a great opportunity to conduct trainings for police officers on NFC technology and the risks arising from its use. She also worked as a Research Assistant at Aalto University in Secure Systems Group. Her research project, which developed into her Master's dissertation, concerned developing SafeKeeper, a system that secures users' passwords on the web. This project received two prestigious awards from the Finnish Information Security Association and the Finnish Computer Science Society.

Furthermore, she actively engage in other activities, including volunteering (e.g., OxfordHack, InspireHer!) and student societies such as Oxford Women in Computer Science Society. Guided by a passion for the spread of education in the area of new technology among children, youth, and adults, and the alignment of this area with opportunities and education for people with disabilities and people socially excluded, she set up a foundation that aims to teach programming and electronics in the form of hands-on workshops.

DPhil Thesis: Authenticating Internet of Things (IoT) devices using out-of-band channels

The amount of Internet of Things (IoT) devices available on the market increases significantly every year. Many such devices are integral parts of smart buildings, which are equipped with modern systems and technologies designed to increase the safety and comfort of their users. Many devices are not equipped with displays or do not allow users to verify their operation. Out-of-band channels such as visual channels (e.g. Augmented Reality) may provide a novel way of authenticating various sensors and give the users an appropriate feedback. The research focuses on designing, implementing, examining and assessing different out-of-band authentication channels and to determine which ones provide stronger security guarantees and fulfil usability, deployability and performance requirements.

Romy Minko

Supervisors: Artur Ekert and Christophe Petit, Maths Institute

Romy's background is primarily in Mathematics, although she also holds a BSc in Chemistry from the University of Melbourne. Romy first became interested in cryptography while at secondary school in Australia and subsequently went on to complete the MSc in the Mathematics of Cryptography and Communications at Royal Holloway, graduating with Distinction. Her research interests lie in post-quantum cryptography and quantum computation; she is currently focussed on supersingular isogeny-based cryptosystems and has previously conducted research in blind quantum computing. Romy is also a 2018 Policy Fellow with the Department of Prime Minister and Cabinet of Australia, where she experienced drafting cybersecurity policy at a government level.

DPhil Thesis: Post-quantum cryptography using multivariate polynomial systems

Multivariate public-key cryptography (MPKC) is one of the four most common branches of post-quantum cryptography, describing cryptosystems based on solving systems of multivariate polynomials. An important step in the cryptanalysis of MPKC system is finding a Gröbner basis for the system. My research focusses on adapting generic Gröbner Basis algorithms to families of multivariate polynomial systems with specific structures. I am currently looking at multivariate linearised polynomials, which have not been studied in great detail.

Additionally I am exploring applications of the HHL quantum algorithm for solving systems of multivariate polynomials, in particular Boolean systems.

Mark Quinlan

Supervisor: Andrew Simpson, Department of Computer Science

Mark Quinlan has gained a mixture of industry experience and academic knowledge prior to starting a DPhil in Cyber Security, his industry experience including BAE Systems where he worked within the cyber field in both commercial and government projects.

His consultancy business designed and built manufacturing resource planning systems, as well as systems security strategy consultancy for companies ranging from British racing teams to Japanese Tier One suppliers.

Mark lives with his partner, and when not pursuing his academic passions he enjoys restoring classic cars, driving karts and said cars, hiking, non-fiction, and enjoying good food and company. Once upon a time Mark was a Dutchman, but has lived in the UK since 2004.

DPhil Thesis: Cyber Continuum; towards a security engineering framework incorporating legacy systems

Mark is looking into privacy and security of Internet of Things devices, and their wider infrastructural landscape. Current work includes a privacy and security analysis of connected cars through the examination of the data- gathering systems of a production vehicle, to ascertain some of the privacy-related threats to which such systems give rise.

Future work: With the lifecycle of an average car being approximately nine years, a connected car has a longer lifespan than the typical IoT device. In addition, it is significantly more likely to be re-sold over its lifetime. When looking at embedded devices across the IoT spectrum, more and more devices are falling outside traditional consumer devices such as speakers and home security, increasingly being used within city infrastructure, and in private and commercial transportation such as cars, the need for security management over longer lifecycles becomes more apparent.

The high-level research objectives are as follows;

1. What is the current state of the literature on the management of legacy embedded systems, and their associated infrastructure?
2. What is the current state of manufacturers providing security updates to their products?
3. What would a theoretical framework incorporating the long-term management of legacy IoT devices look like?

Lonie Sebagh

Supervisors: Jonathan Lusthaus and Federico Varese, Department of Sociology

Lonie is researching the disruption of online criminal trade by various stakeholders from law enforcement to private industry and trading platforms, combining sociology, criminology, and economics perspectives. Her thesis involves three different research methods: 1) laboratory experiments aiming to measure traders' behaviours and responses to operations aimed at eroding trust on online criminal marketplaces, 2) content analysis of reports, news articles, and Blog posts of organisations in different sectors in order to evaluate the way they communicate about the disruption of online criminal trade, and 3) interviews with experts in different sectors to discuss their role in the disruption landscape. The use of mixed methods will allow for the formulation of theory and recommendations about the disruption of online criminal trade by different stakeholders.

Prior to joining the CDT Lonie was a student in Business and Management at the Universities of St Andrews and Edinburgh. Her interest in Cyber Security stems from her work experience in the IT Security department of a private bank in Geneva in 2014, which inspired her to make the transition from Business to Information Technology through a PG Diploma (Distinction) with an emphasis on Computer Security and Critical software engineering.

In her spare time Lonie can be found in her College where she works as a Junior Dean, providing welfare support to fellow postgraduate students.

DPhil Thesis:The disruption of Dark Net Markets – the impact of slander, Sybil, and platform takedown operations

This research will use social laboratory and online experiments, never before used in the field of cybercrime, to better understand cybercriminal networks' responses to Law Enforcement disruption operations in darknet markets. Operations currently performed in the wild will be replicated in controlled settings in order to test their respective effectiveness in lowering product quality and increasing price, therefore rendering these markets less attractive. Participants' behaviours during trading games when subjected to certain disruption operations will inform policies about which operations to use, when to use them, whether to use them individually or simultaneously, and what 'breaking points' Law Enforcement should act upon in order to disrupt these markets more efficiently. The insights from these experiments will be combined with interviews of Law Enforcement agents and forum data where platform users discuss their experiences in these markets.

Presentations:

L. Sebagh, J. Lusthaus, E. Gallo, F. Varese, 2020. Cooperation and distrust in cybercriminal markets – an experimental study of marketplace disruption. In: 2020 Extra Legal Governance Seminar series (ExLegi), May 8, Oxford, England.

L. Sebagh, J. Lusthaus, E. Gallo, F. Varese, 2019. Responses to Slander and Sybil Disruption Operations in Online Criminal Marketplaces – a Laboratory Experiment. In: 2019 European Economic Science Association (ESA) Meeting, September 5-8, Dijon, France.

Sean Sirur

Supervisors: Kasper Rasmussen, Department of Computer Science, Tim Muller, University of Nottingham

Sean has been involved with security and formal methods for the majority of their academic career. In their Computer Science and Physics BSc from Edinburgh, they focused on theoretical computer science (machine-assisted proof/verification, formal languages types and semantics) and security (cryptography and systems security). In physics, they specialised in quantum and statistical physics, with hopes to extend the former into a strong knowledge of quantum computing and with the latter providing a firm basis for information theory. The physics aspect of the degree also involved designing and implementing computable numerical methods and simulations. Sean is interested in the verification of quantum cryptographic protocols in particular, as it includes aspects of nearly every area of their technical interests.

Sean is intent on drawing inspiration for their research from the cyber-security related issues faced by industry and society as a whole. As such, they are currently working on a short research project based on engaging with and interviewing various industry professionals about data protection compliance, with a focus on GDPR. Understanding the efficacy of law and policy as a mechanism to promote safe data and security practices is Sean's primary non-technical focus at the moment.

Sean has worked part-time as a tutor and lab demonstrator for the Computer Security course at UoEdinburgh. Their hobbies include writing music and literature, hiking, dancing and spending time around animals.

DPhil Thesis: The Reputation Lag Attack

Reputation-based trust systems are increasingly common on digital platforms; this includes any model where the trustworthiness of an actor is derived from how other actors judge their interactions with that entity. Common examples include review and ratings systems as often used on electronic marketplaces. In any such system, the propagation of reputation suffers from delays of various kinds e.g. network connectivity, reporting delays and rating-update delays. There is very little existing research on the exploitation of this lag, known as a reputation lag attack. Furthermore, there is no robust theoretical framework for the definition, discussion and evaluation of such attacks. Current research is a first step towards constructing a suitable theoretical framework of reputation lag attacks. Future research includes investigating techniques with which attackers can increase their profit so as to provide an insight into attack patterns. The findings of this future research can then be used to validate and improve upon the framework; construct mitigations against the attack; and to aid investigations into whether the attack is already occurring in the wild and with what frequency.

Publications

Sirur, S., & Muller, T. (2019). The Reputation Lag Attack. In Proceedings of the 13th IFIP WG 11.11 International Conference on Trust Management July 17 - July 19, 2019, Copenhagen, Denmark

Sirur, S., Nurse, J.R. and Webb, H., 2018, October. Are We There Yet?: Understanding the Challenges Faced in Complying with the General Data Protection Regulation (GDPR). In Proceedings of the 2nd International Workshop on Multimedia Privacy and Security (pp. 88-95). ACM

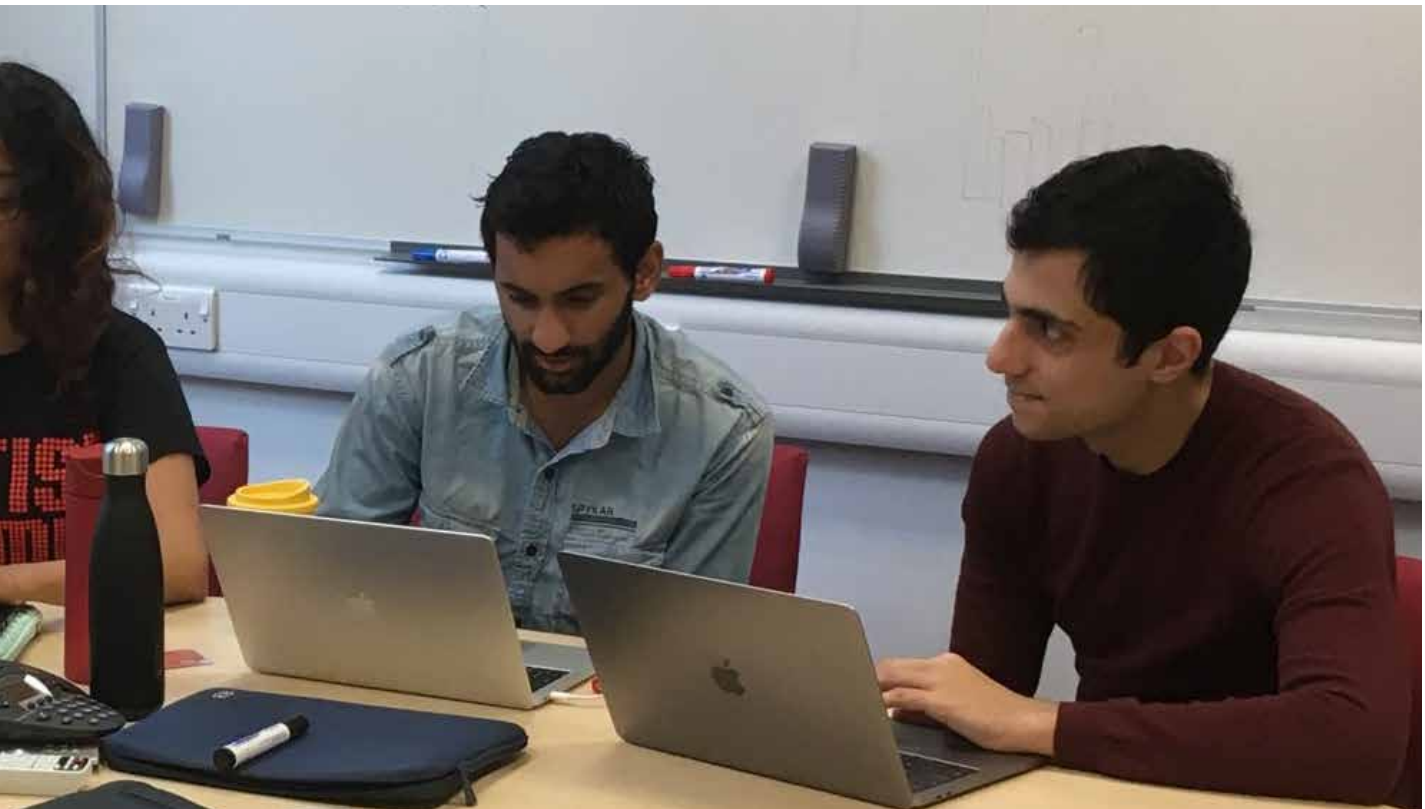

Eva Janečková

Supervisor: Justine Pila, Faculty of Law

Eva is reading for a DPhil in Cyber Security at the University of Oxford within the CDT 2017 cohort. Her DPhil thesis is devoted to the legal implications of computational creativity of artificial intelligence-driven systems deployed in cyber defence. She looks at whether objects generated by those systems can be protected by patents or copyright under currently applicable law. Her thesis will also include policy recommendations focusing on promoting innovation and information sharing in the field of cyber security. In her mini-projects she focused on evolution of data protection regulation in Europe and legal implications of AI-driven creativity in cyber security.

After her graduation from the Law Faculty of Charles University in Prague Eva practiced law in Prague-based law firms as an Associate in corporate, IP/IT, media and telecommunication legal teams. In 2015, her interest in intellectual property law and competition law brought her to the Munich Intellectual Property Law Centre (MIPLC). Eva also dealt with policy coordination in internal market at the Secretariat General of the European Commission in Brussels.

DPhil Thesis: AI-Driven Defence Systems and the Limits of Copyright and Patent Law

This research deals with computational creativity in the context of the current state of the art artificial intelligence algorithms deployed in network defence. It features a legal assessment of whether objects generated by AI-driven systems qualify for copyright and patent protection under the current legal regimes in Europe and in the United Kingdom. In order to examine the eligibility for the legal protection, a human creativity requirement will be analysed, and it will be determined how it is applied to computer-generated objects. This project will conclude with recommendations towards such a legal regime for AI-generated objects which would be ideal for incentivising innovation in the field of cyber security.

Publications:

September 2019 From Creativity Requirements Towards Creativity Tests – a presentation given at the European Policy for Intellectual Property 14th Annual Conference at the ETH, Zurich

Eva Stanková, Human inventorship in European patent law, Cambridge Law Journal 2021, 80(2), 338-365.

Conference talks:

January 2020 AI-generated inventions and creativity tests in patent law - an invited speaker at the European IP Institutes Network Innovation Society (EIPIN-Innovation Society) conference in Maastricht

Henry Turner

Supervisor: Ivan Martinovic, Department of computer Science

Henry comes from Colchester, Essex and holds a MEng from Imperial College London in Computing. His thesis examines security aspects of biometric systems, with a particular focus on voice processing systems and their resilience to realistic attacks, as well as developing ways to better protect users of these systems. During his time at Imperial College London he completed his thesis on security schemes in body sensor networks and facilitating secure communication in embedded medical devices. He has interned as a software engineer at Intel Security (McAfee) working on corporate network monitoring products. Prior to this he also ran a small enterprise publishing iPhone and Android applications during his teenage years, distributing more than 250,000 copies of his applications during the project's lifespan.

DPhil Thesis: Improving the Security of Voice Interface Systems

Voice interfaces have become common on modern devices, and increasingly support complex interactivity, as well as authentication mechanisms to provide personalised (and sensitive) functionality to users. We focus on such voice interfaces, and in particular improving their performance in adversarial situations. We do this through the testing of attacks against such systems and through the development of tools to analyse security properties. In turn these allow us to identify flaws and weaknesses in voice systems.

We then design improved voice interface systems to combat weaknesses and flaws we identify, to improve their overall security properties.

In addition to work on voice interface systems, we intend to try and translate some of our attacks and tooling to other biometric systems, to see if they suffer from similar weaknesses and can be improved in similar ways.

Publications:

Turner, H., Lovisotto, G. and Martinovic, I. Attacking Speaker Recognition Systems with Phoneme Morphing, European Symposium on Research in Computer Security 2019, 23-27 September 2019, Luxembourg

Fredderic Barr-Smith

Supervisor: Ivan Martinovic,
Department of Computer Science

Freddie holds a BSc in Computer Science and Business Management (First) and an MSc in Software and Systems Security (Distinction). He also has held positions in infrastructure and security in multiple sectors.

He also holds several security certifications including Offensive Security Certified Professional (OSCP), CREST Practitioner Security Analyst (CPSA) and CREST Registered Penetration Tester (CRT).

His research mainly is concerned with malware and the techniques malware use to counter and evade antivirus and analysis systems. The objectives of his research being to illustrate and enumerate the various evasion techniques that are used by malware to evade manual and automated analysis. The formation of this research consists of analysis at scale of data and development of proof of concepts of these techniques. This research area also touches upon cybercrime and forensics.

DPhil Thesis: Malware Evasion Techniques

The aims of this research are broadly to illustrate and enumerate the various evasion techniques that are used by malware to evade manual and automated analysis by analysts and automated analysis systems.

These evasion techniques include recognition of malign techniques which have evolved and become exponentially more complex as the arms race of malware development continues.

Analysis of these will include analysis at scale of existing and newly created malware, analysing large datasets from various threat intelligence providers. Additionally there will be creation of proof of concepts that demonstrate innovative techniques within this area.

Furthermore, this research may involve discovery of flaws within antivirus software or other analysis software used to detect and analyse malware. To an extent this flaw discovery can be identified as vulnerability research, with the aim of strengthening the tools used to analyse and protect against malware.

This research will strengthen skills in malware analysis, vulnerability analysis and penetration testing for the researcher, skills necessary to forge a career in cybersecurity and skills of which there is currently a shortage.

This subfield has a constant flow of novel data and has the combined innovation of cybercriminals, a variety of nation state groups and academia contributing to it's rapid development. Therefore analysing this data will contain a large amount of novelty. Additionally contributing new techniques in this vein.

This research will involve active collaboration with a number of private sector companies. Additionally as results of this research may come in the form of detected vulnerabilities within existing antivirus detection and analysis systems, these will help strengthen these systems within both the private and public sector.

This collaboration between academia and private sector is especially important as many of the systems in use by the private sector are only under robust analysis by state actors. Due to this it is important to codify into the academic body of knowledge, the tools and techniques which may only be known or in use by criminal and state actors.

Publications:

Mini-Project: Living Off The Land: Systematic Review of Use of Living-Off-The-Land Technique by Malware

Mini-Project: Onion Optical Illusions: Imitation of Onion Services

George Chalhoub

Supervisor: Ivan Flechais,
Department of Computer Science

George holds a BS in Computer Science from the Department of Computer Science and Mathematics of the Lebanese American University. He obtained his MSc in Computer Science from the school of Electronics and Computer Science at the University of Southampton, in collaboration with Lloyd's Register. He previously worked as a Cyber Security Analyst and is currently an Expert Contributor at Oxford Analytica. His doctoral research is supported by the Information Commissioner's Office and explores the application of user experience (UX) principles in the security and privacy design of smart homes.

DPhil Thesis: The UX of Things: Exploring UX Principles to Inform the Design of Security and Privacy in the Smart Home

Smart homes are under attack. Threats can harm both the security of these homes and the privacy of their inhabitants. As a result, in addition to delivering pleasing and aesthetic experiences, smart devices need to protect households from vulnerabilities and attacks. Further, the need for user-centered security and privacy design is particularly important for such an environment, given that inhabitants are demographically-diverse (e.g., age, gender, educational level) and have different skills and (dis)abilities.

Prior work has explored different usable security and privacy solutions for smart homes; however, the applicability of UX principles to security and privacy design is under-explored. This research project aims to address the on-going challenge of security and privacy in the smart home through the lens of UX design. The objective of this thesis is two-fold. Firstly, to investigate how UX factors and principles affect smart home users and the product design process. Secondly, to inform product design through the development of an empirically-tested data-driven framework for UX design of security and privacy in smart home products.

In the first step, we aim to explore the relationship between UX, security, and privacy in smart homes from user and designer perspectives: through (i) conducting a qualitative interview study with smart home users (n=20) and (ii) analyzing data from a longitudinal study of smart home device adoption and use in households (n=6); and, we plan to explore the role of UX in the design of security and privacy in smart homes through qualitative semi-structured interviews with smart home designers through two rounds of interviews (n=20, n=20).

In the second step, using our exploratory results, we aim to build an empirically-tested data-driven descriptive framework for UX design of security and privacy in the smart home products. To evaluate the applicability of our framework, we are running participatory design workshops with a diverse group of smart home stakeholders. Finally, using our framework, we will extract thematic recommendations supporting security and privacy design practice in smart home products.

By bringing UX design to the smart home security and privacy table, we believe that this project will have a significant impact on academia, industry, and government organizations. Our framework will inform the product design process of security and privacy in this emerging technological area while contributing to scholarly practice.

Publications:

George Chalhoub. The UX of Things: Exploring UX Principles to Inform Security and Privacy Design in the Smart Home. In Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA 2020). ACM. April, 2020.

George Chalhoub, Ivan Flechais, Norbert Nthala, Ruba Abu-Salma and Elie Tom. Factoring User Experience into the Security and Privacy Design of Smart Home Devices: A Case Study. In Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA 2020). ACM. April, 2020.

George Chalhoub and Ivan Flechais. "Alexa, are you spying on me?": Exploring the Effect of User Experience on the Security and Privacy of Smart Speaker Users. In 22nd International Conference on Human-Computer Interaction (HCII 2020). Springer. July, 2020.

George Chalhoub, Ivan Flechais, Norbert Nthala, and Ruba Abu-Salma. Innovation Inaction or In Action? The Role of User Experience in the Security and Privacy Design of Smart Home Cameras. In the 16th Symposium on Usable Privacy and Security (SOUPS 2020). USENIX Association. August, 2020

Anirudh Ekambaranathan

Supervisors: Max Van Kleek and Jun Zhao, Department of Computer Science

Anirudh holds an MSc in Computer Science and Education from Twente University in the Netherlands. His research focuses on applied machine learning in the context of cyber security. Currently he is researching anomaly based intrusion detection systems. His previous research projects focused on Wi-Fi tracking and stylometric linkability in darknet markets. Before joining the CDT he worked as a part-time math teacher in secondary school and had his own cyber security startup.

DPhil Thesis: Understanding Design Features of Family Apps and Design Choices made by Family App Developers

Children have established a significant presence online through mobile devices, leading to an increase in the number of apps designed for children. Mobile apps can provide educational value for children across the globe, giving developers the opportunity to make positive contribution. However, data monetisation remains the main source of income for developers in this space. Targeted ads or game promotions become the norms in the freemium apps, including those used by children. Children not only find them annoying and a waste of time, but also are often nudged to make choices that reduce their personal privacy and leave them more vulnerable to data tracking.

New initiatives are set up to improve children's data protection and online safety. For example, the ICO in the UK is setting up children-specific regulations, and the Federal Trade Commission in the US is calling responses to a review of the Children's Online Privacy Protection Act (also known as COPPA). However, to create proper regulations and incentivise developers to adopt age appropriate design guidelines, we need to understand (1) why responsible design choices are difficult for developers, (2) a better understanding of harms design features in children's apps may cause, and (3) tools for practitioners to effectively assess age appropriateness of children's apps.

Publications:

Mini-Project: Because we cannot control third-party adverts": Understanding Design Choices Made by Family App Developers

Mini-Project: GAN Based Anomaly Detection Using Protocol Conditioning and Feature Corruption

Marine Eviette

Supervisor: Andrew Simpson, Department of Computer Science

Marine is a Radcliffe scholar with interests in Privacy and Cryptography- having undertaken prior research in Post Quantum Cryptography, Cryptanalysis of Historical Ciphers and Online Privacy. Besides this, she previously completed a Masters thesis which sought to help prevent the propagation of Cyber Attacks by utilizing a ZKP verified Distributed Ledger Technology. Her current research interests, however, are strongly rooted in the area of Privacy; where she is choosing to pursue a DPhil project that focuses on exploring how access control models can help to prevent metadata exploitation in the era of big data, in order to aid in protecting user's privacy.

DPhil Thesis: Mitigating the Proliferation and Exploitation of metadata through introduction of Access Control

The implicit nature of consent regarding metadata currently enables companies to extort user's privacy and uncover, otherwise hidden, identity elements, which jeopardises user's sense of control. We find that the scale of metadata has led to online browsing habits shaping an extremely detailed trail of online activities via the passive digital footprints that are unknowingly left behind. Such digital footprints that are being accumulated by means of dataveillance have led to inferences on identity elements that dramatically expose user identity whilst rendering these users powerless in the managing of their personal privacy.

By modelling privacy as an access control problem, we seek to combat inference-driven identity exposure that threatens user privacy. We aim to achieve this through clearly defined policies that target the aggregation of metadata and the consequent inferences. However, at present, there is no one unifying method to privacy and identity management as it relates to this problem. Therefore, whilst it may indeed prove fruitful to reason about access to data objects as a method of protecting individuals' privacy; without further consideration, we are set to be hindered by the unfavourable state of the internet today, as users currently have no control over their own data and consequently cannot control access to it.

The case for ownership becomes increasingly complex when we consider the analysis done on data aggregates; here, companies are analysing user generated data in order to identify patterns and make inferences. These

subsequent processes generate a further set of data, which is often as valuable as the original data, yet this analysis information is automatically owned by the company performing the analysis.

In our work, we hope to reason about data collection and access control from the perspective of users which is decidedly a very challenging task, as users do not typically control access to their collated data. Coincidentally, there have been a number of projects that have begun to investigate methods for users to regain ownership, and consequent control, of their data, through the introduction of decentralised architectures that can be implemented atop of the current internet.

Our research project aims to leverage the solutions to these and similar privacy problems in order to present preliminary ideas of access control modelling as it pertains to metadata and identity management. In line with this, we investigated numerous approaches to similar problems, whereupon we decided to follow a similar route to researchers' design of a framework for secure data collection through the extension of the Category-Based Access Control metamodel.

By utilising Category-Based Access Control we hope to be able to build upon the foundations of access control in order to reason about this complex privacy problem. Our proposed research is of significant value as there is a pressing need for more transparency amongst sharing and management of metadata online.

Publications/ Conferences:

Dec 2020: Towards models for the Mitigation of Metadata Exploitation, 15th IFIP Summer School, on Privacy and Identity Management – IFIPSC2020, Co-author: Andrew Simpson

Jun 2019: Cybersecurity and Intimate Partner Violence, Connected Life Conference 2019, Co-authors: Julia Slupska, Romy Minko, Zhi Tan and Fatima Zahra

Martin Georgiev

Supervisor: Ivan Martinovic, Department of Computer Science

Martin holds a BEng degree from the University of Edinburgh which incorporated an exchange year at University of California, Irvine (UCI). During his stay abroad, he developed an interest in cyber security and published a joint paper with the SPROUT (Security and Privacy Research OUTfit) group at ESORICS. He has various internship experiences ranging from detecting compromised accounts at Facebook to developing innovative systems for the banking industry at Royal Bank of Canada. At Oxford, Martin is interested in authentication using behavioral biometrics. More specifically, he is looking into models for continuous mobile authentication based on touch patterns, their practical applications and privacy concerns stemming from their use.

DPhil Thesis: Behavioural Biometrics for Authentication

The research project aims at investigating the feasibility of using various behavioral biometrics in authentication scenarios. This field of study focuses on establishing uniquely identifying patterns in human activity such as keystroke, mouse and touchscreen dynamics as well as voice, gait and cognitive behaviour. This is typically a non-invasive method for authentication as it does not require users to learn how to operate a particular system or remember unique passcodes and phrases. Furthermore, there are no active steps required for authentication but rather it is a seamless integration with the regular operation of the system. Often authentication systems based on behavioural biometrics can be used as a multifactor safeguard in conjunction with other more traditional cybersecurity measures. Despite being a somewhat well research area with some apparently successful projects currently there are few successful commercial systems employing the technology. The goal of the research is to design, develop and test novel systems for authentication based on behavioural biometrics and close the gap between promising research and practical applications. For instance, developing a continuous authentication model based on phone usage patterns such as touchscreen gestures and gyroscope micro movements in space. Another aspect of the project focuses on identifying problems in past research in the area and some of the reasons it tends to be unsuitable for practical use. For example, the reported sample sizes in some of the studies might not be large enough to accurately represent the population using such systems. Finally, there are unique privacy challenges stemming from the use of highly accurate authentication systems based on behavioural biometrics. One way to maliciously employ this technology is to create unique fingerprints for users which can then be exploited for tracking behaviour and identity throughout multiple non-connected systems. It also might be possible to reveal personal information about users through their behaviour patterns. Gender, age and cultural groups could exert specific traits which might be detectable by the technology described above.

Publications:

Mini-Project: Adversarial noise injection into digital images through electromagnetic interference

Mini-Project: Towards continuous touch-based mobile authentication using neural networks

Hayyu Imanda

Supervisor: Kasper Rasmussen, Department of Computer Science

Inda is a Jardine Scholar at Exeter College. She grew up in Jakarta, Indonesia before moving to the UK to take her BSc in Mathematics from the University of Edinburgh. She then completed her MSc in Mathematics and Foundations of Computer Science at Oxford, where she focused on post-quantum cryptography, with a dissertation on the security of supersingular isogeny key exchange.
After briefly working as a software developer, she went home to Jakarta and took an internship at a consultancy before starting the CDT. In her first year, she co-organised the CDT in Cyber Security conference with Royal Holloway.
She is a 2x full Blue in lawn tennis and the incoming president for the Oxford University Lawn Tennis Club. She feels most at home out in the ocean scuba diving, and deeply cares about wildlife conservation; when possible, she spends time away from Oxford travelling across her diverse home country. You might also find her stranded abroad during a pandemic!

DPhil Thesis: Modelling an Adversary with Privileged Access

Traditional adversary models in cryptography primarily assumes a network adversary with no access to endpoints. We now know that this is far from sufficient. Adversaries with a higher capability, with continuous authorized access to endpoints for example, benefit from the lack of sufficient security design against them.

My research aims to model an adversary with a high level of control over the network, as well as access to privileged information. In particular, I am looking at cases where there exists an asymmetric power dynamic between the adversary and the victim; for example, government whistleblowers, intimate partner violence, and those under targeted surveillance. Due to the power imbalance between the victim and the adversary, previous mitigations to endpoint compromise -- for example, key rotation, wouldn't suffice. New security goals have to be defined, and we design cryptographic mitigations which satisfy those goals.

Publications:

Mini-Project: Mass Surveillance and Isogeny-Based Cryptography: An Introduction

Mini-Project: Location Privacy in Conservation

Jack Jackson

Supervisor: Max Van Kleek, Department of Computer Science

Jack is currently a DPhil Researcher at the University of Oxford, where he resides within the Centre for Doctoral Training in Cyber Security. Jack has previously held roles as a Chief Technology Officer, Principal Technology Consultant, Research Scientist, Cryptographer, and Cognitive Engineer; across a number of Europe's most prolific Startups. For his work, Jack has been honoured with a number of accolades, including a Europe-wide entrepreneurship award. He has also acted as both a keynote and guest speaker at several prolific conferences, including the International Workshop on the Future Perspective of Decentralized Applications, held in conjunction with the 24th International European Conference on Parallel and Distributed Computing - where he also sat as a Program Committee member.
Before joining Oxford, Jack made a name for himself within both the Startup and Blockchain communities across Europe. This stemmed from his researching into privacy-enabling blockchain technologies, where he explored the application of advanced cryptographic mechanisms, such as: homomorphic encryption, differential privacy and secure multiparty computation - to distributed ledger ecosystems. In recognition for his efforts, Jack was invited to act as an Associate Editor at Frontiers Open Access Journal, where he curated his own section on the Fourth Industrial Revolution. In this role, Jack leads a team of seasoned academics, consisting of established professors and postdoctoral researchers.
Whilst at Oxford, Jack has been invited to act as an Expert Consultant on Blockchain Technologies by United Nations (UN) entities. As of 2019, Jack has committed to conducting research into a number of areas, including: cyber insurance, social engineering, and the development and application of deep fake technologies. To these ends, he is utilising his rich and diverse background, which branches across a broad array of areas within: insurance, artificial intelligence, cryptography and distributed ledger technologies.

DPhil Thesis: Deep Phishing

While advances in security and software engineering processes have greatly increased the robustness and resilience of software to cyber attacks, comparable advances in cyber security resilience have not been made

at the human level. This shortcoming can be effectively observed across a spectrum of successful human-centric cyber crime campaigns, such as: social engineering, phishing and psychological operations (PsyOps). Recent developments in the field of artificial intelligence (AI) and increased data collation capabilities facilitated by methods drawn from social engineering, threaten to increase the effectiveness of these attacks. Providing adversaries with the extended capability of scaling their capacity to both act and sound human, underpinned by the information necessary to inform attempted mimicry. Two such advancements include the introduction of deep fake technologies, which have enabled the hijacking of trusted personas at will, by means of impersonation; and the development of open-source intelligence (OSINT) tools, capable of mining publicly accessible information, such as that on social media sites. Understanding the threats these technologies pose in the context of cyber security, especially in the context of enabling targeted social engineering, remains an under-researched area. To these ends, we plan to evaluate existing attack frameworks across each of the aforementioned criminal domains, identifying key aspects which may be automated or enhanced through the assumption of AI. In particular, we aim to explore potential capability extensions within targeted spear phishing campaigns, enabled by the introduction of deep fake technologies into the kill chain; in what we refer to as a Deep Phishing attack. Deep Phishing can best be defined as the AI-facilitated impersonation of an individual, for the purpose of extracting information from a target with which they have sufficient social proximity.

The primary intended outcomes of our research are threefold. First, we aim to extrapolate future attack models given the observable advances in deep fake generation capabilities, and how that might impact more traditional social engineering attack models.

Second, we plan to gain an understanding of the identified emerging threats, through the prototyping of software capable of performing the proposed attacks.

Finally, we intend to explore potential mitigation strategies, including improving target resilience through automated, AI-based red-teaming against individuals.

Beyond this, we wish to analyse how accessible information on an individual (via sources such as social media and data leaks) can be indicative of one's exposure to attack, through data correlation; with the goal of informing future personal information publishing decisions, and next generation user authentication protocols. To these ends, we hope to understand whether fundamental knowledge gaps across users could be addressed using intelligent tutoring (ITS) approaches to personalise and tailor representations with detail appropriate to the user's understanding.

Publications:

Mini-Project: Deep Phishing

Mini-Project: Creating a Pre-Competitive Dataset for Cyber Risk

Talks:

Social Engineering: HSBC Cybersecurity Awareness Week 2019

Deep Fake Technology, CDT Showcase 2019

Deep Phishing: Social Engineering, Redefined as part of Commonwealth Members of Parliament briefing day "Research, Risk and Resilience: Cyber Security and Public Policy", 2020

Sebastian Köhler

Supervisor: Ivan Martinovic, Department of Computer Science

Sebastian is a doctoral researcher in the Centre for Doctoral Training in Cyber Security at the University of Oxford. As part of the Systems Security Lab (SSL), run by Professor Ivan Martinovic, he researches the security of the physical layer of a variety of large and complex systems, such as vision-based intelligent and automotive systems.

He started to specialise on Cyber Security during his undergraduate studies in Computer Science at the University of Applied Sciences Würzburg-Schweinfurt, Germany. Due to his interests in the security of modern cars, he completed his bachelor's degree with a dissertation at the research and development centre of the Dr. Ing. h.c. F. Porsche AG. Before his doctorate, he received a master's degree in Computing & Security and got awarded the prize for the best overall performance on the MSc in Computing & Security for the academic year 2017/18 from King's College London.

In his spare time, Sebastian enjoys improving his knowledge and skills by attending Capture the Flag challenges, hackathons and conferences. As the president of the Competitive Computer Security Society, Sebastian shares his passion for security-related topics by organising and hosting different events, such as workshops, CTFs and talks. Sebastian also likes to contribute to the scientific community by serving as a reviewer. In 2020, he was selected to serve on the Shadow Programm Committee for the prestigious security conference IEEE Security & Privacy, and more recently, he was invited to act as a reviewer for the ACM Journal Transactions on Privacy and Security.

DPhil Thesis: Exploiting Physical Phenomena to Enhance the Security of Cyber-Physical Systems

Large and complex cyber-physical systems, such as autonomous vehicles and industrial control systems, are increasingly relying on the input of a wide variety of sensors. For instance, self-driving cars often use Light Detection and Ranging (LiDAR) in combination with cameras to perceive their environment. In general, a sensor measures a physical quantity such as light, heat and sound. With the ongoing integration of sensors into the decision-making process of cyber-physical systems, the integrity of the measurements is becoming a cornerstone of the correct behavior of the system. However, a sensor cannot validate the authenticity of the measured quantity. An adversary could tamper the measurement by injecting malicious electromagnetic signals into the sensor to spoof a physical stimulus.

In addition, the secure interconnection of those systems to exchange information, such as the sensor measurements, is becoming more and more crucial. A recent example is the communication between an electric vehicle and a charging station. During the charging session, the vehicle measures the State of Charge (SoC) and the battery temperature and reports them to the charging station, which in turn regulates the maximum current. This enables a gentler and faster charging process. This communication has to be secure to ensure that an adversary cannot tamper the communication and spoof the sensor measurements. However, recent research has shown that attacks on the physical layer using electromagnetic waves can bypass security mechanisms on higher layers.

This research project is twofold. On the one hand, we demonstrate signal injection attacks on the physical layer against critical components of cyber-physical systems to impair their correct functioning. For example, we analyse how an adversary can interfere with the charging communication of electric vehicles using electromagnetic waves to interrupt the charging process or even cause irreversible damage. On the other hand, we evaluate defense mechanisms to protect against such attacks. More precisely, we investigate how physical phenomena associated with signal injection attacks in the wireless domain can be facilitated for attack detection and prevention.

Publications:

Mini-Project: Using Structured Demand Manipulation Attacks to Disrupt the Power Grid

Mini-Project: Interrupting the CCS Charging Communication using Radio Frequency Interference

Arthur Laudrain

Supervisor: Lucas Kello, Department of Politics and International Relations

Arthur P.B. Laudrain (@APB_Laudrain) is a Rotary Global Scholar for Peace and a doctoral researcher in cybersecurity at Wolfson College. His interests relate to decision-making, coercion through cyber means and assessing states' military capabilities in cyberspace. He has collaborated with the French Strategic Research Institute (IRSEM, Paris) and is consulting for the International Institute for Strategic Studies (IISS, London). His writing was featured on BBC Science Focus, Lawfare, The Military Balance+ and The Journal of Political Science Education. Arthur previously attended King's College London and Leiden Law School.

DPhil Thesis:Hacks, Leaks and Statecraft: Determinants of Foreign Policy Response to Cyber Coercion

"The nation's actions really matter. A wrong choice could mean irreparable damage. Thus responsible citizens are obliged to fight for what they are convinced is right", said Allison and Zelikow (1999) in their infamous decision-making analysis of the Cuba missile crisis. This bargaining process described well before the advent of the Internet, may explain in part why cyberspace has been mostly a domain of restraint between nations. This is an important question: If we better understand why states exercise restraint in cyberspace, we may also be able to explain cases of escalation. Scholars of cyber international relations have raised numerous hypotheses as to why we witness such restraint. Most rely on rational assumptions with the state as a unitary point of focus. With this research project, I suggest that by investigating key players and small group dynamics during the foreign decision-making process, I may be able to contribute to the explanation of restraint as a foreign policy outcome. Competing objectives and perceptions of individuals as well as group interactions are especially critical in presidential democratic systems such as France and the US. Even more so in matters of national security. Although group decision-making is a valued theory and has been applied on numerous occasions to foreign policy decisions in the US, responses to cyber coercion remain a largely uncharted territory. By investigating the policy and decision making processes in response to instances of cyber coercion in the US (DNC leaks, Sony hack) and France (Macron leaks, TV5Monde hack), I should be able to explain why both states responded the way they did (or did not). Such findings may modestly contribute to the explanation of state restraint that is puzzling the cyber IR community.

Writings and Publications

Trey Herr, Arthur P. B. Laudrain & Max Smeets, "Mapping the Known Unknowns of Cybersecurity Education: A Review of Syllabi on Cyber Conflict and Security", Journal of Political Science Education, 28 Feb 2020.

Arthur P. B. Laudrain, "5G and the Huawei controversy: is it about more than just security?", BBC Science Focus, 21 Mar 2020.

"France's 'strategic autonomy' takes to space", International Institute for Strategic Studies (London), 14 Aug 2019.

"France's New Offensive Cyber Doctrine", Lawfare (Washington DC), 26 Feb 2019.

Conferences

"The State, its Institutions and Processes: Applying Decision-Making Models to French Cyber Security and Defence", Bridging the Gap Workshop, Columbia SIPA (NYC), 11 – 12 Nov 2019.

"The Paris Call for Peace and Security in Cyberspace: A Year Later", Global Governance of AI and Cyber Security Panel, Bonn International Security Forum, 1 – 3 Oct 2019.

"Military Cyber Operations: Comparative Approach of France and the UK", National Research Agency Cyber Studies Seminar, Université de Bordeaux, 6 Jun 2019.

Matthew Rogers

Supervisor: Kasper Rasmussen
Department of Computer Science

Matthew is a 2018 Rhodes Scholar with a degree in software engineering from Auburn University. His experience includes work with Dynetics, an engineering firm, doing malware analysis and reverse engineering APT malware. Additionally he has spent time at the Defense Digital Service, bolstering their cyber capabilities. He has spoken at several conferences on malware analysis and cyber security education. His research focuses on creating cheap intrusion detection systems for serial data bus networks, primarily J1939. From this he hopes to build out mission assurance research for critical transportation and military system.

DPhil Thesis: Securing J1939 Systems through Minimal System Modifications

Since the early 2000s millions of industrial systems have taken their existing Controller Area Network (CAN) Bus infrastructure and added a software standard, J1939, to simplify communication between the different electronic control units (ECUs) controlling the vehicle. The standard was initially designed for ground vehicles, but is now common place across agriculture and forestry equipment, military vehicles, marine vessels, power generators, and much more. While this is useful for industrial systems, the underlying infrastructure is still CAN, a serial data data bus protocol with no authentication, or effective security mechanisms. For the last decade academia and enthusiasts continually showed hacking an automobile is possible with access to the CAN Bus, even going as far as remotely gaining access. The J1939 standard only simplifies the hacking process by removing the need for reverse engineering the proprietary CAN messages of consumer automobiles. Not only does this simplify hacking single vehicles, it makes attacks agnostic to installed ECUs, enabling non-targeted attacks across fleets of heterogeneous vehicles.

We propose using the J1939 standard for defensive purposes. Instead of relying purely on header and timing data we can analyze the data field, making sense of the 8 bytes of data previously left untouched for practicality's sake. We begin this research with 2 premises: we can only add a single device to the J1939 Bus without modifying any existing ECUs, and we cannot have any false positives. Modifying every installed ECU is expensive, and discourages future firmware upgrades, effectively discouraging security. False positives generally risk alert fatigue, causing true positives to go unnoticed. The safety critical systems typically found running J1939 are too valuable for any level of false positives to be acceptable. To test for false positives we run our IDS against real truck data.

For this research we built a state-based rules framework which compares arbitrary J1939 data fields, adjusted to their real values. In doing so we created over 40,000 rules, 10,000 of which require some level of training to maintain system knowledge across system reboots. With these rules we are able to detect an attacker transmitting fixed-rate messages (e.g., every 100ms) across the bus if a legitimate ECU is already transmitting it. We apply this same timing based security guarantee to non-fixed-rate messages by ensuring the conditions for that message, such as a diagnostic trouble code message, being sent are met. These conditions come from fixed-rate messages, and so provide the same security guarantee. Additionally we ensure the attacker is unable to prevent an existing ECU from speaking short of physically removing it from the CAN Bus, an action that requires far more advanced physical access than traditionally seen in automotive hacking. This work falls within the EPSRC engineering research area, and was done in collaboration with Shift5 Inc. Future work will being the areas of Incident Response using the J1939 standard, using side-channel mechanisms for defensive purposes, and using a hueristics approach on the J1939 data field.

Talks:

Speaker at the Association of Old Crows 55th Symposium panel on "Preparing EMS Superiority" : "Electronic Sheepdogs : Providing the Hacker's Mindset to Everyone"

Publications:

Mini-Project: Incident Response and Prevention Recovery in J1939

Mini-Project: A State-Based Rules Framework for Serial Data Bus Networks

Yashovardhan Sharma

Supervisor: Ivan Martinovic, Department of Computer Science

Yashovardhan has been fiddling with computers ever since he was a toddler. A penchant for testing the limits of what was possible given a computer system made him naturally inclined towards computer security.

He completed his MPhil in Advanced Computer Science from the University of Cambridge and his BTech (Honours) in Computer Science and Engineering from IIIT-Delhi. Having worked on projects ranging from Healthcare to Artificial Intelligence to Human Computer Interaction, his focus recently has been on the area of Privacy and Security, with an emphasis on Cryptography. During his sojourn at Oxford he hopes to design and build trusted systems that leak minimal information and are reasonably resistant to compromise.

The interdisciplinary and collaborative nature of the CDT is a key reason for Yashovardhan to have come to the "other place". Born and raised in India, he is also known for his ability to enjoy non-spicy food.

DPhil Thesis: Analysing the Safety of Collision Avoidance Protocols in Aviation

Collision avoidance protocols for autonomous and semi-autonomous vehicles form the backbone of a safe and reliable global transportation system. In the case of aviation, the Traffic Alert and Collision Avoidance System (TCAS) is responsible for ensuring the safety of aircraft and reducing the chances of mid-air collisions. However the safety and efficacy of TCAS is yet to be analysed from a security perspective- especially given the current and ever-increasing levels of air-traffic with regard to a protocol conceived nearly two decades ago.

This research project aims to model and analyse the constraints specified by TCAS that are required for safe operation. The goal is to determine whether TCAS still meets its operational goals with regard to collision avoidance, and further investigate whether it is vulnerable to malicious attack. The novelty of our research methodology is that we use real-world data (aircraft transponder messages) for our analysis of TCAS's safety, rather than relying on simulations or statistical testing. This allows us to glean accurate and up-to-date information regarding the usage of TCAS around the world. Based on our results, we are then in a unique position to understand the potential risks posed by TCAS, the consequences posed by them, and most importantly - possible methods of remedying them.

Publications:

Mini-Project: Analysing the Safety of Collision Avoidance Protocols in Aviation

Mini-Project: Exploring the Impact of Availability in Secure Enclaves

Anjuli R. K. Shere

Supervisors: Andrew Martin, Department of Computer Science and Jason Nurse, University of Kent

Anjuli (@AnjuliRKShere) is an analyst, writer, and researcher, with experience of journalistic and security-related investigations. While attending 'Particle Summer School' at CERN, she was inspired by the scientific progress created by global collaboration. She has since studied a BA (Hons) in Politics and International Relations at the University of Nottingham, and spent a year gaining a Certificate in Social Sciences and Humanities at Sciences Po, Paris.

Alongside her master's degree in Science and International Security in the Department of War Studies at King's College London, Anjuli began reporting on current affairs for the New Statesman. She specialised in strategic security threats posed by emerging technological concerns, and wrote her dissertation on the extent to which machine learning could protect the NHS from cyber-attacks. Her professional endeavours also include working as the conference and research analyst for the Association for International Broadcasting.

During the first year of her doctorate in Cyber Security at the University of Oxford, Anjuli co-organised and emceed the CDT conference on Cyber Espionage and returned to her work as an intelligence analyst on Channel 4's award-winning fugitive simulation, 'Hunted'. She also conducted cross-disciplinary research projects within the faculties of Law and Computer Science, covering open-source intelligence, data protection legislation, state surveillance and emerging technologies.

Currently, Anjuli's thesis research aims to create a framework for mitigations of Internet of Things threats to the free press, for transcontinental news organisations to integrate into their cyber security strategies. Her case study countries are the UK, USA, Australia and Taiwan.

DPhil Thesis: How can physical, legal and virtual threats from novel Internet of Things devices affect press freedom in the UK, USA, Australia and Taiwan, and how might these threats be mitigated?

The existence and maintenance of a free press can be used as a barometer for the state of a democratic society, as public access to factual information about powerful people and organisations is key to an educated electorate. Therefore, attempts to curtail transparent, accessible and free journalism can be seen as threats to a branch of the critical national infrastructure of a democracy, and thus to the democratic state itself. This has potential implications on an individual human rights level and in terms of international relations and security.

My DPhil research aims to comprehensively investigate and document what journalists and media organisations are doing, procedurally and technologically, to protect themselves against innovative and well-resourced attackers. I would compare this with the recommendations of experts from a variety of backgrounds, including academic, governmental and non-governmental. The objective is determining the extent to which the current protections (and the systems by which these are chosen and updated) are effective against contemporary and anticipated threat models - particularly emerging technological and legal threats – and how these protections might be improved for my chosen case studies and for news organisations in democracies more broadly.

Publications:

Jun 2020: Selected to present "Now You [Don't] See Me: How have the GDPR and a changing public awareness of the UK surveillance state impacted OSINT investigations?" (my Mini-Project 1 findings) at the Surveillance and Society conference in Rotterdam (postponed due to COVID-19)

Jun 2020: "Reading the investigators their rights: A review of literature on the General Data Protection Regulation and open-source intelligence gathering and analysis" (my Mini-Project 1 literature review) is to be published in the SCR peer-reviewed publication affiliated with New College/University of Oxford, The New Collection, http://mcr.new.ox.ac.uk/journal/

May 2020: "Securing a Free Press inside a Networked Panopticon: The case of the Internet of Things" (my Mini-Project 2 report) was published at The 5th European Workshop on Usable Security (EuroUSEC 2020), https://eusec20.cs.uchicago.edu/eusec20-Shere.pdf

Jan 2020: "Growth of privately held data increases risk of espionage" (co-written with Neil Ashdown) for Jane's Intelligence Review, https://www.janes.com/images/assets/638/94638/Growth_of_privately_held_data_increases_risk_of_espionage.pdf

Nov 2019: "AIB November Meeting Report," in AIB Media Freedom Initiative Briefing and Update for Members, https://aib.org.uk/Media-Freedom/AIB-Media-Freedom-Initiative-report-191119.pdf

Oct 2019: "5th Annual Inter-CDT Conference: Cyber Espionage Report," in University of Oxford CDT in Cyber Security Yearbook 2019, https://www.cybersecurity.ox.ac.uk/site-resources/uploads/2020/02/2019-yearbook_web.pdf

May 2019: Centres for Doctoral Training in Cyber Security: University of Oxford & Royal Holloway Cyber Espionage Conference Report: 2nd-3rd May 2019, https://pure.royalholloway.ac.uk/portal/files/33922755/Cyber_Espionage_Conference_Report_.pdf

Dec 2018: "ASAP90 Conference 27th-28th September 2018 Report" for the Association for International Broadcasting, https://aib.org.uk/asap90-conference-27th-28th-september-2018/

Sep 2018: ASAP90 Conference Guidebook & Radio Taiwan International corporate social responsibility pledge, https://asap90.rti.org.tw/wp-content/uploads/2018/09/CONFERNCE-WITH-NOTES-20SEPT2018-0926.pdf

New Statesman (profile: https://www.newstatesman.com/writers/321610)

Julia Slupska

Supervisors: Gina Neff, Mariarosaria Taddeo, Joss Wright, Oxford Internet Institute and Max Van Kleek, Department of Computer Science

Julia Slupska (@jayslups) is a doctoral student at the Centre for Doctoral Training in Cybersecurity and the Oxford Internet Institute. Her research focuses on how cybersecurity concepts and practices can address technologically-mediated abuse. She is also exploring how feminist theories and methodology—such as participatory action research and the ethics of care—can improve cybersecurity. Previously, she completed the MSc in Social Science of the Internet on the role of metaphors in international cybersecurity policy. Before joining the OII, Julia worked on an LSE Law project on comparative regional integration and coordinated a course on Economics in Foreign Policy for the Foreign and Commonwealth Office. She also works as a freelance photographer.

DPhil Thesis: Design Justice in Security Architectures

Feminist theorists have long argued that gendered security problems, such as domestic abuse, are "individualized" and taken out of the public and political domain (Tickner 2004; Walby et al 2014). Unfortunately, the emerging field of cybersecurity risks recreating these dynamics by omitting or dismissing gendered technologically-facilitated abuse (or "tech abuse") such as stalking, surveillance, and image-based abuse (or "revenge porn") from the threat models that shape where researchers investigate challenges to security (Slupska 2019).

The project is based on the following research questions/objectives:
RQ1: How can cybersecurity practices better serve the targets of tech abuse?
RQ2: How can feminist theory and praxis improve cybersecurity research and practice?

On the basis of these two research questions, I plan to develop a feminist approach to cybersecurity which draws on feminist critiques of security studies (Enloe 1989; Cohn 1987; Tickner 2004), feminist technoscience (Wajcman 2007) and the emerging 'design justice' model for technology design (Constanza-Chock 2018). I will start by conducting a set of empirical studies, which will form the basis of a normative political theory for cybersecurity. These empirical studies may include:
- co-designing an "abusability" test for smart devices or image sharing platforms with survivors, tech abuse experts, and conventional cybersecurity experts
- follow-up interviews with co-design workshop participants to explore contrasting understandings of security and strategies for approaching tech abuse
- interviews with product managers exploring how abusability could become incorporated into industry practice
- participatory action research in the form of feminist digital security trainings

This project will use innovative co-design methodologies which have only rarely been applied to cybersecurity. Following feminist approaches to knowledge creation and the emerging 'design justice' model for technology design (Constanza-Chock 2018), people's individual experiences and individual positionality may help to expand how cybersecurity researchers do the work of threat modelling and usable security design. Rather than dictating what threats citizens should be worrying about, this project will develop a model for eliciting and listening to citizens' concerns to expand the scope of threat modelling in cybersecurity. This process will also create pathways for citizens to engage in shaping future research directions for cybersecurity: ones that are grounded in the lived experience of those who are traditionally excluded from discussions of cyber- or information security.

Inspired by Marwick and Boyd's (2018) call for projects that discuss more diverse conceptualizations of "the user" or the subject, I will use collaborative, participatory, and creative practices to address cybersecurity challenges in the UKRI-funded "Reconfigure" citizen science research project. Participatory security design avoids the assumption that security of the individual will follow from technical security and ensures that actors who may ordinarily be marginalized have their perspectives taken into account (Heath et al. 2018). It incorporates 'situated' knowledge and practices (Haraway 1988) so that information security can be studied in a grounded way.

Publications:

J. Slupska and L. Tanczer. (forthcoming) "Intimate Partner Violence (IPV) Threat Modeling: Tech Abuse as Cybersecurity Challenge in the Internet of Things (IoT)." In Technology-Facilitated Violence and Abuse – International Perspectives and Experiences. Emerald Publishing.

J. Slupska. "Safe at Home: Towards a Feminist Critique of Cybersecurity", St. Anthony's

International Review, Summer Issue (2019), no. 15: Whose Security is Cybersecurity? Authority,

Responsibility and Power in Cyberspace. Available at SSRN.

J. Slupska and M. Taddeo. "Generative Metaphors in Cybersecurity Governance". Yearbook of the Digital Ethics Lab Yearbook (2019).

J. Slupska, R. Minko, Z. Tan, F. Zahra and M. Eviette. "Cybersecurity and Intimate Partner Violence"

Map the System Research Competition, published in Yearbook of the Centre for Doctoral

Training in Cybersecurity (2019).

D. Chalmers and J. Slupska. "The Regional Remaking of Trade and Investment Law." European

Journal of International Law, Volume 30, Issue 1, (2019), https://doi.org/10.1093/ejil/chz004.

N. Maroun and J. Slupska. "International LGBT Leaders Take the Stage" Public Diplomacy Magazine (2014).

Public Engagements:

"Safe at Home: Towards a Feminist Critique of Cybersecurity", 2020 International Studies Association, cancelled due to COVID-19 global pandemic

"Reconfigure: Feminist Action Research in Cybersecurity" 2020 Human-Computer Interaction (CHI) Direct Action Workshop, cancelled due to COVID-19 global pandemic

"War, Health, & Ecosystem: Generative Metaphors in International Cybersecurity Policy" (05 November 2019), Hague Conference on Cyber Norms 2019 – Best Paper Award.

"'We're All Happily Married Here!': Intimate Partner Violence as a Cybersecurity Issue" (03 October 2019), Royal Holloway Information Security Group Seminars

"Designing IoT Security for Intimate Threats," (21 May 2019), London IoT Meetup Group

"Towards a Feminist Critique of Smart Home Security Analysis" (9 March 2019), Oxbridge Women in Computer Science Conference

Security Analysis" (9 March 2019), Oxbridge Women in Computer Science Conference

Claudine Tinsman

Supervisors: Max Van Kleek, Department of Computer Science and Rebecca Williams, Faculty of Law

Claudine has a BA in Political Science from UC San Diego. During her time in California, she worked in San Diego city government and at a large immigration law firm. She holds a Master of Law (MLaw) in Legal Issues, Crime and Security of Information Technologies from the University of Lausanne. Her master's thesis examined the potential implications of treating intelligent agents as legally liable actors.

Her DPhil research aims to design effective and customisable subjective harm mitigation implementations that enable users to safeguard and promote their mental wellbeing.
She currently serves as assistant to the papers chairs for the ACM Human Factors in Computing Systems Conference 2021 (CHI2021).

DPhil Thesis: Curating Contact and Conduct in Online Spaces

Many people spend a significant amount of their daily lives communicating with one another online. Children and teenagers are significantly more likely than adults to engage socially online, enabling them to communicate with anyone, anywhere from a very young age. These users may be unaware that they are harming others by their words and actions online, while those exposed to such antisocial behaviour online may feel unable to remove themselves from situations detrimental to their mental wellbeing.

Research in psychology and the social sciences has shown that negative experiences with online contact can have deleterious effects on users' mental health and may in turn encourage certain types of antisocial behaviours.

While all users can and should be concerned about who they engage with online, this project focuses on groups of users who are likely to have specific concerns about exposure to online content and contact, such as parental figures, teenagers, and children. This project has the potential to both help users who exhibit antisocial conduct to engage in more prosocial behaviour, and to prevent others from being subjected to harmful behaviours.

At its core, this research seeks to develop tools that afford users greater control over the contact they experience in online spaces. It combines research from psychology, learning sciences, linguistics, and computer science in order to explore automated solutions that can identify specific threats within the context of online interactions in real time. Specifically, pattern analysis of word use in conversations will be used to assess whether a user's online conversation displays antisocial characteristics. These methods will be combined with machine learning to create tools that can be deployed as conversations unfold.

The ultimate purpose of this research is to provide two complimentary measures to make online communication safer: On the one hand, it will produce tools that provide real-time educational interventions to individuals whose conversations display traits of specific types of conduct, such as cyberbullying. On the other, it will provide users seeking to protect themselves and those under their care (e.g. parents and children) with a notification system when conversations escalate to levels that the users deem undesirable.

Talks:

Presented DPhil research to approximately 40 Commonwealth MPs attending a conference on cyber security at the Oxford Martin School.

"Social Engineering Attacks", Cyber Security Awareness Week, HSBC (20/05/2019).

Appeared on the Big Questions Podcast (Oxford Sparks): Appeared as a guest on the show to explain cyber security concepts to a non-specialist public audience (27/03/2019).

Publications:

Mini-Project: Defining Personal Data under the GDPR: Challenges for Organisational Cyber Threat Intelligence Sharing

Mini-Project: A Universal Security Rating System for Mobile Apps: Preventing Today's Fraud From Becoming Tomorrow's Nightmare

Fatima Zahrah

Supervisor: Michael Goldsmith and Jason Nurse, Department of Computer Science

Fatima received a BSc (Hons) degree in Computer Science from the University of Bradford. Fatima's research interests focus around online hate and investigates how online platforms are strategically used by cyber criminals and hate groups. Her work combines insights drawn from social sciences and uses methods from computer science, including natural language processing, machine learning and social media analysis.

DPhil Thesis: Investigating the Cross-platform Behaviours of Online Hate Groups

Online hate thrives globally through self-organized, scalable clusters that interconnect to form robust networks spread across multiple social media

platforms, countries and languages. Despite efforts from law enforcement agencies and platform developers to remove or limit such content, online hate ideologies and extremist narratives are still being linked to several crimes around the world. The networks formed by hate groups have proven to be remarkably resilient and have increasingly shown to migrate across various platforms and networks, maintaining and oftentimes expanding their connections in the process. Previous research in online hate has generally focussed around one particular platform, even though there is sufficient evidence showing that hate groups often strategize the usage of different online platforms in order to circumvent current monitoring efforts. This research will aim to bridge this gap by investigating how online hate groups make use of multiple platforms to propagate criminal and extremist content. More specifically, it will involve a cross-platform-analysis of the behaviours of such hate groups in order to better understand and detect networks of organised hate. This project will be conducted with particular consideration of the following research questions:

RQ1: How can the current online hate research landscape be advanced by considering through exploring several online platforms?
RQ2: How do hate groups adapt their behaviour on different platforms?
RQ3: How is information transferred and shared by hate groups across platforms?
RQ4:How can we model online hate detection and analysis across various platforms?

Through this, the research aims to determine how multiple online platforms are strategically used by hate organisations, and produce more efficient hate detection and analysis methods. The findings from this will then be used to aid the development of a web interface as a tool for law enforcement agencies to detect, analyse and help remove criminal hate.

Publications: Forthcoming:

Zahrah, F., Nurse, J.R.C. and Goldsmith, M., September 2020. #ISIS vs #ActionCountersTerrorism: A Computational Analysis of Extremist and Counter-extremist Twitter Narratives. To be presented at the 2nd Workshop on Attackers and Cyber-Crime Operations, co-located with IEEE EuroS&P 2020.

Presentations:

Slupska, Julia, Romy Minko, Zhi Tan, Fatima Zahra and Marine Eviette. "Cybersecurity and Intimate Partner Violence" (2019) Map the System Research Competition. Also presented at Connected Life Conference 2019.

Life after the CDT

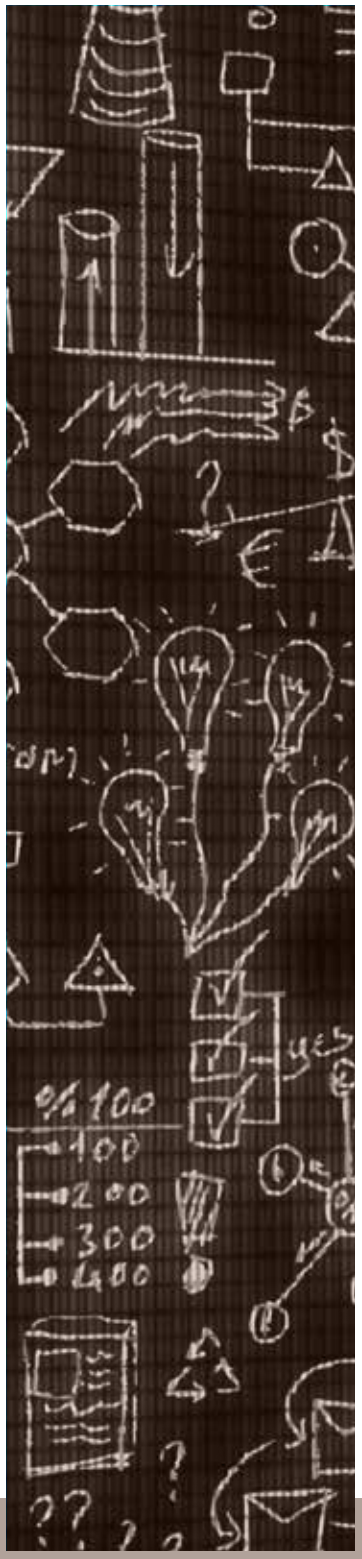

Bushra AlAhmadi (CDT Alumni)

March 6 2020, was my viva date, the day I worked years towards finishing my Dphil and celebrating. I have planned the post viva for years; I planned to travel around the UK, learn new hobbies and take some time off. But, unfortunately, sometimes things don't go as planned. Ten days post viva, a lockdown was announced in the UK due to Covid-19. So, after months of staying home, social distancing and waiting for flights to open, I travelled back to Saudi Arabia in July 2020, where my career post DPhil commenced.

I am now an Assistant Professor at the College of Information and Computer Science at King Saud University. In my first year as a professor, courses in the university were taught remotely. I taught two Bachelor courses, Information Security and Digital forensics courses and supervised a bachelor graduation project. In this project, the students designed and implemented CyberSafe, a mobile application that alerts parents if their child is a victim of cyberbullying by automatically detecting bullying words using Machine Learning. I am also part of and head several committees in the college.

I am also a research visitor at the University of Oxford System Security Lab headed by Professor Ivan Martinovic, where I have continued working on the outcomes of my Dphil into publications. One of these outcomes was the qualitative study on Security Operations Centres, in which many of The CDT industry partners have generously participated. This paper, titled: 99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms, has recently been accepted to the prestigious USENIX Security Conference.

I continued my involvement in community initiatives such as Google Developer Community, where I was chosen to be a Google Developer Expert in Machine Learning. I give talks to the community specifically on Responsible AI and AI+Security. I am also an active member of the Saudi Information Security Association (Hemaya), where I head the association's initiative to empower women in cybersecurity and promote online child safety in Saudi Arabia. For example, we partner with public and private entities to provide training and mentorship for women in cybersecurity to get them into cybersecurity leadership positions. I represented the association in several conferences; the latest was MENA Information Security Conference, where I talked about "Artificial Intelligence for Cyber Security: Are We There Yet?"

I have always been interested in entrepreneurship. Therefore, I am part of CyberME Studio, where we focus on building cybersecurity startups. We are now up to 4 startups that focus on building the next generation of cyber products. I also work as a cybersecurity consultant to the President of the General Court of Audit and other cybersecurity companies in Saudi Arabia. I am honoured that my work with these entities contributed to Saudi Arabia's commitment to cybersecurity which ranks second in the global Cybersecurity Index, issued by the International Telecommunication Union (ITU).

I am very grateful for my time in the university, for my supervisor, Professor Ivan Martinovic, for all faculty, administrators and colleagues who helped prepare me for the career I have today. But, to be honest, the reality of finishing my DPhil hasn't sunk in yet, maybe when I return to Oxford soon and my name is finally called in the graduation ceremony! :)

Security Alumni Network Launch

In April we officially launched the Security Alumni Network, a multidisciplinary networked aimed at connecting Oxford alumni working in the broad security field irrespective of the original academic discipline undertaken. This reflects the aims of the Cyber Security Oxford Network by bringing seemingly disparate fields together to consider Cyber Security issues from novel perspectives.

The inaugural event, chaired by Professor Andrew Martin welcomed cyber security specialists, Professor Ciaran Martin, CB (Blavatnik School of Government) former CEO of the National Cyber Security Centre and Graham Ingram (our university Chief Information Security Officer). The speakers discussed the current Cyber Security landscape and of particular relevance now, protecting Oxford's leading medical research during the current pandemic.

The network has grown to approximately 350 members following this event and a committee formed by a the members is planning a further online event in late summer 2021 and a full in-person in early 2022. If you would like to join this network, please visit https://www.cybersecurity.ox.ac.uk/alumni for further details.

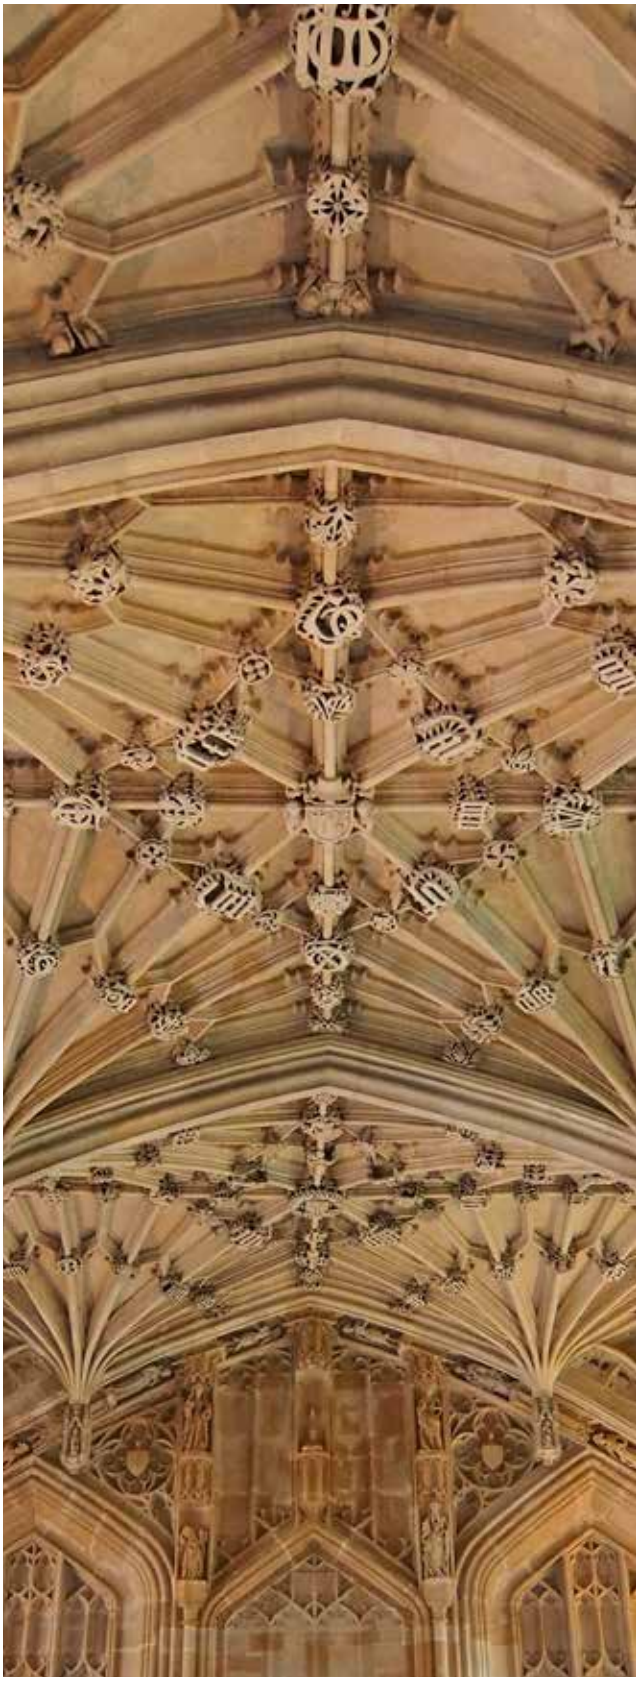

Alumni News

Eman Alashwali (CDT-15) is now an Assistant Professor at King Abdulaziz University (KAU), Saudi Arabia. She recently worked as a lead author of "Saudi Parents' Security and Privacy Concerns about their Children's Smart Device Applications", which is currently under review and will be published soon.

The Alumni network is continuing to expand – please see **www.cybersecurity.ox.ac.uk/alumni** for further details

The CDT Team

Andrew Martin – CDT Director

Professor of Systems Security, Department of Computer Science

An Oxford graduate, Andrew worked as a Software Engineer at Praxis in Bath, where he first encountered some of the challenges of information security and secure systems engineering in the late 1980s. After a DPhil back in Oxford, he escaped to the other side of the world to be a Research Fellow at the Software Verification Research Centre in the University of Queensland. Eventually the excellent weather and relaxed way of life got the better of him, and so he returned to the UK, and entered his current post in 1999.

The core of his research interest here has been in the security in distributed systems. Mostly of late that's been explored through looking at applications of hardware-based security controls – often described as Trusted Computing technologies - particularly as applied to cloud, mobile, and embedded applications (now known as the Internet of Things). His research group has been looking for the architectural elements and design patterns necessary to make trusted clouds and secure IoT a reality. These ideas have the potential to transform how we think about distributed systems and the security of the information they process.

Michael Goldsmith – CDT Co-Director

Senior Research Fellow, Professor, Department of Computer Science

Michael Goldsmith is a Senior Research Fellow at the Department of Computer Science and Worcester College, Oxford. With a background in Formal Methods and Concurrency Theory, Goldsmith was one of the pioneers of automated cryptoprotocol analysis. His research work has investigated a range of Technology Strategy Board and industrial or government-funded projects ranging from highly mathematical semantic models to multidisciplinary research at the social-technical interface. He is an Associate Director of the Cyber Security Centre, Co-Director of the newly launched Centre for Doctoral Training in Cybersecurity and is active in the IAAC Academic Liaison Panel.

Lucas Kello – CDT Co-Director

Associate Professor of International Relations, Director of the Centre for Technology and Global Affairs, Department of Politics and International Relations

Lucas serves as Director of the Centre for Technology and Global Affairs, a major research initiative exploring the impact of modern technology on international relations, government, and society. His recent publications include The Virtual Weapon and International Order (Yale University Press), "The Meaning of the Cyber Revolution: Perils to Theory and Statecraft" in International Security, and "Security" in The Oxford Companion to International Relations (Oxford University Press).

Joss Wright – CDT Co-Director

Senior Research Fellow, Oxford Internet Institute

Joss Wright is Senior Research Fellow at the Oxford Internet Institute, where his research focuses on the analysis of information controls and their global development, and on the design and applications of privacy enhancing technologies.

Joss' work focuses on interdisciplinary approaches to the measurement and analysis of technologies that exert, subvert, or resist control over information. He has a particular interest in bridging the gaps between technically-focused analyses of security and privacy technologies, and their broader social and political implications.

In addition to his work on internet censorship, Joss also co-directs the Oxford Martin School's Programme on the Illegal Wildlife Trade, in which he researches the trade in illegal and unsustainable wildlife products online.

Katherine Fletcher – CDT Industry Liaison Officer

Katherine is the Coordinator of Cyber Security Oxford, the University-wide network of cyber security researchers and practitioners. Her role in the CDT is to help connect students to the wider community of Oxford researchers, and to support matchmaking for research projects with industry or other external partners. Katherine has over 10 years' experience as a Project / Programme manager largely based in Oxford, specialising in large-scale, multidisciplinary research projects spanning academia and industry. Recent experience includes managing research projects in biomedical/computer science (linking pharma industry and academia), open-source software development projects (academic data management) and cybersecurity (multiple business sectors and academia).

Katherine received a BA in International Relations from William Jewell College (Liberty, Missouri, USA; 2001), and an MA in Global Political Economy from the University of Sussex (2004).

David Hobbs – CDT Centre Administrator

After 8 years with the CDT, I sadly leave this amazing community for a new adventure in the USA. Rather than include a bio, I wanted to use this space to thank my colleagues within the CDT, the students and the leadership of our Centre. It has been an absolute pleasure to spend my days surrounded by fiercely intelligent, modest and inspiring people. Our students are among the kindest, most considerate people I have met, all with the very best intentions hoping to change the world for the better. The CDT's Director and Co-Directors demonstrate the very best in leadership, supporting, motivating and offering an opportunity for personal and professional growth. I will look back on my time with great fondness and I can't wait to see what this group of incredible people achieve in the future. Thank you for letting me be a small part of it.

Centre for Doctoral Training in Cyber Security
Dept of Computer Science
Wolfson Building
Parks Road
Oxford
OX1 3QD

cdt@cybersecurity.ox.ac.uk
01865 610644

UNIVERSITY OF OXFORD

CENTRE *for* DOCTORAL TRAINING *in* CYBER SECURITY

UKRI Engineering and Physical Sciences Research Council