INFORMATION SECURITY GROUP

REVIEW 2024 - 2025

Index

Introduction

Lizzie Coles Kemp
Professor and Head of Department, ISG

Welcome to the latest edition of the annual ISG Review. As every year, the newsletter showcases the broad range of work that the ISG does – from complex ethnographic work to homomorphic encryption standardisation and everything in between. This ability to offer both breadth and depth in our information and cyber security research and teaching is second to none and makes the ISG a unique place to work and study.

I am particularly proud of our achievements this year; the UK higher education sector is experiencing tough times, and no UK university is immune to the sector's challenges. The ISG is responding with pragmatism, hard work and innovation and, as these articles show, the ISG is still finding the time to excel in research, teaching and external engagement.

This year in the ISG we have been hard at work developing our information security syllabus of the future. In September 2026 we are launching our re-designed MSc programme. We are changing the title to MSc in Information and Cyber Security to reflect how we prepare students for roles in cyber security but that, through its holistic design and interdisciplinary approach, our programme is more than cyber. In our re-design we are giving more focus to AI in cyber, we are mapping our optional modules to contemporary career pathways and introducing a new module on innovation and practice in information and cyber security. Our Coursera programme continues to develop and was awarded with a provisional NCSC certification. This academic year has also seen the introduction of a new masters programme in Applied Data Science and Cyber Security that we co-teach with our colleagues in Computer Science. This programme combines traditional information security teaching with coding and data science skills and is designed to prepare students for information and cyber security that draw on data science techniques.

As the articles in this review show, we have had many research successes. Each of our three sub disciplinary areas (Cryptography, Systems and People and Society) has produced relevant and impactful research outputs. Despite a challenging research funding landscape, we have also made some successful bids for funding with a Transnational Education Exploratory grant award from the British Council for Fauzia Idrees and Waleed Bin Shahid and a grant from the Salient Hub to research AI generated fakes in the community for Andrew Dwyer, Danilo Francati and myself.

The ISG seems to have been out and about more than ever this year. We took part in CyberUK in Manchester, we have given talks across the UK and overseas, published 3 books, and colleagues have been guest speakers on several podcasts. In November last year, we also hosted our inaugural Fred Piper Information Security Lecture and were delighted to welcome the incoming CEO of NCSC and Royal Holloway alumnus, Richard Horne as our inaugural speaker. The Fred Piper Information Security Lecture will be held this year on the evening of 12th November and our speaker will be Siân John, MBE. royalholloway.ac.uk/about-us/events/ai-in-practice-skills-security-and-shifting-realities/

To register for a free place visit royalholloway.ac.uk/about-us/events/

Just after the last newsletter, we said goodbye to Chez Ciechanowicz who fully retired from the ISG after nearly 29 years of service. Thank you, Chez! Many of you will remember Chez as the programme director for the MSc in Information Security. Chez played a significant role in helping to make the ISG what it is today and as we pilot our way through higher education's current difficulties, I am particularly grateful for the work Chez did to make our teaching programme so resilient. I'd also like to take this opportunity to thank Ahmad Salman, Danilo Francati and Darren Hurley-Smith for their contributions to the ISG. These three colleagues also left us this year, and we wish them the very best in their future endeavours.

I'd like to close by saying thank you to you for your on-going support for the ISG. Whether you are someone who only recently came across us or someone who has been involved with us for much of our 35-year history, we are delighted to have you with us.

NEW THREATS TO ACADEMIC INTEGRITY

Chris Mitchell
Professor, ISG

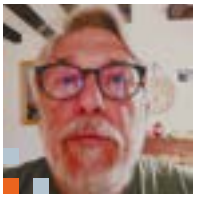

Online academic preprint repositories, such as arXiv and ResearchGate, enable researchers to freely share new research results, helping to speed up and simplify dissemination of new findings and novel insights. However, since the posted preprints are largely unrefereed, readers have to take all claims of new results with a pinch of salt. Of course, this was always the case, even back in the days when printed preprints were circulated by post to friends and colleagues, but the ever-increasing volumes, and the ease of online dissemination, of unchecked material create new threats.

Recent joint work with Dr Haitham Al-Sinani, a former ISG PhD student who is now Director of the Department of Cybersecurity and Quality Assurance, Diwan of Royal Court, Muscat, Oman, has revealed how preprint repositories can be used to easily manipulate metrics that are widely used to judge the impact of academic research, and in particular the standing of individual academics. Our interest in this topic arose through our research on applying generative AI in the field of ethical hacking. While monitoring relevant publications via Google Scholar alerts, we encountered a series of papers that appeared to be relevant to our work. However, closer inspection revealed that they were strikingly similar, cited many of the same individuals, and lacked novel content. These observations raised suspicions that the papers were AI-generated and promulgated with the intent of artificially inflating citation metrics.

To understand this, I first need to explain one particularly prominent metric, namely the H-index. The H-index for a researcher is defined as the maximum value h such that the individual has published h papers each of which has been cited at least h times. It attempts to capture both productivity and citation impact in a single number. While the H-index and other metrics, such as the i10-index, are convenient and widely-adopted in academic evaluations, they are also vulnerable to manipulation. This has been known for some time, especially involving the use of highly dubious academic journals which publish large numbers of papers with minimal (and no) peer review. However, using such a route is often non-trivial, as many of the dubious journals make a charge for publishing content, and there is a delay in achieving the desired impact. We appear to have identified a simpler route to inflating an H-index literally overnight. Of course, if an H-index is calculated with care, and dodgy-looking citations are removed, then there is no problem – although even then the efficacy of this metric in evaluating the work of an academic remains highly questionable. However, many people rely on automatic calculations of the H-index, and in particular on the calculation provided by Google Scholar. Google Scholar is a highly valuable and well-trusted online resource for searching the academic literature; as a result I suspect that its H-index calculation is also widely used.

The problem we identified comes down to the fact that preprints posted on online repositories are automatically scanned and used in H-index calculations by Google Scholar. Since some preprint sites, e.g. ResearchGate, appear to perform minimal checking of submitted content, this makes it simple to inflate the citation count for a paper by including it as a reference in an online preprint.

As I already mentioned, we first became aware of this issue through automated Google Scholar alerts, i.e. messages sent by Google Scholar to make researchers aware of new work relating to what Google believes are their interests. We received alerts regarding new papers apparently relating to our work on AI-supported pen-testing. However, on reading these papers we discovered they were clearly at least partly AI-generated and content-free. Moreover, they consistently cited the same set of obscure papers by a small number of authors (one author in particular). After further investigation we found at least 30 of these spurious papers, all uploaded to ResearchGate over a short time period and which had inflated at least one academic's H-index by at least 25. It would appear that Google Scholar indexes all papers which are given a DoI (a widely-recognised unique digital identifier) and ResearchGate submissions are all given such an identifier.

To see how easy it is to achieve H-index inflation, we first created a short 'fake' paper. This paper was created using generative AI tools, and we avoided adding any novel content. We also ensured the paper cited a number of other papers. We avoided citing our own work (to avoid being guilty of inflating our own H-indexes), and largely cited papers that were cited by the spurious papers we had previously found. We also included two clear statements to the effect that the paper had been created using AI and was purely for experimental purposes.

So what happened next? Well, the experiment worked just as expected. Within 24 hours Google Scholar had used the existence of this 'paper' to adjust the citation counts of all the authors we had referenced. This means that anyone wishing to inflate their citation counts (and as a result their H-index) simply needs to have an appropriate number of preprints including the relevant citations uploaded to ResearchGate – or any other preprint site indexed by Google Scholar – and the effect will be almost immediate.

How can this be fixed? Ideally, reputable preprint sites would perform at least a quick check of submitted papers to see if they appear genuine. Moreover, Google Scholar needs to be more selective – perhaps only indexing preprints that have been through some kind of screening process. Arguably, the H-index calculation should only take into account papers that have been through a rigorous review process, which means excluding all preprints from calculations. Until and unless something is done along these lines we should expect to see the Google Scholar citation metrics become increasingly useless.

A more detailed description of our findings is available at arxiv.org/abs/2503.23414

CRYPTOGRAPHY: A VERY SHORT INTRODUCTION 2ND EDITION PUBLISHED

Sean Murphy
(Professor, ISG)
Rachel Player **(shown right)**
(Senior Lecturer, ISG)

The second edition of *Cryptography: A Very Short Introduction* was published in February 2025. It is co-authored by two Royal Holloway colleagues, Prof Sean Murphy and Dr Rachel Player. The book is part of the popular Very Short Introductions series from Oxford University Press, which began in 1995 and now features over 750 volumes.

The first edition of *Cryptography: A Very Short Introduction* was also produced by the Information Security Group, being co-authored by the ISG's founder, the late Prof Fred Piper, and Prof Sean Murphy. It was published in 2002 and has been a recommended text for students of the MSc Information Security for many years. It is hoped that this new edition will prove similarly invaluable to the next generation of students approaching the topic of cryptography for the first time, as well as appealing to professionals entering the cybersecurity workforce, and to other readers who are curious to learn more about this fascinating topic.

Cryptography has advanced significantly since the publication of the first edition and this new edition brings the reader fully up to date with both important recent academic advances and newer applications like blockchain, secure messaging, and electronic voting. The book is written in non-technical language and does not assume advanced mathematical knowledge, so is designed to be accessible to a broad audience. After introducing symmetric and public-key cryptography, the book provides a detailed discussion of the design of cryptographic algorithms that are secure against quantum computers and the development of cryptographic algorithms with advanced functionalities. The aim is to demystify the art (and science!) of cryptography by tracing its historical use, explaining how it works, and providing examples of its practical use.

The book can be purchased via the Oxford University Press website or from various retailers: global.oup.com/ukhe/product/cryptography-9780192882233

Dr Rachel Player said: "It was a real honour to join as a new co-author for this second edition. The first edition was one of the very first books I read on cryptography, and its friendly style and compact size really helped me to get to grips with the basic concepts without getting overwhelmed. I hope readers find the second edition similarly valuable! I am a huge fan of the Very Short Introductions series, so it is very cool to now have my own book on the shelf alongside other my copies of other books in the series! I was also very pleased to learn recently that there is a plan to translate the book into Japanese!"

Prof Sean Murphy reflected that : "It was a great experience writing the first edition of the book with Prof Fred Piper in 2002. This was about the halfway point between the beginning of open cryptography with Diffie-Hellman and the Data Encryption Standard in the mid-1970s and today, and the landscape of cryptography has greatly changed since then. This second edition aims to update this Very Short Introduction by explaining some of these recent developments in and current uses of cryptography, as well as highlighting some of main issues facing cryptography. I am grateful to have had this opportunity to write this second edition along with my colleague Dr Rachel Player outlining both historical and modern cryptography for a general audience."

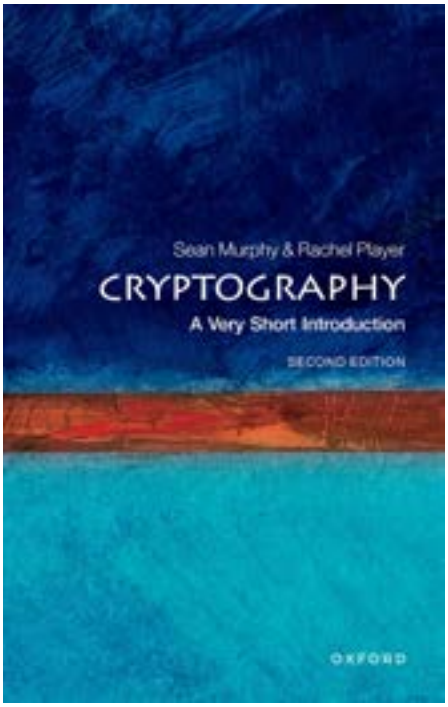

DOING FIELDWORK IN HIGHER-RISK SETTINGS: RESEARCHING THE INFORMATION SECURITY NEEDS OF ACTIVISTS

Mikaela Brough **(picture block left)**
PhD Researcher, CDT in Cyber Security for the Everyday;
Rikke Bjerg Jensen
Professor, ISG **(picture block right)**

In previous editions of this newsletter we have discussed how – and why – ethnography matters to the study of information security in different contexts and among different groups of people. We have illustrated its usefulness in understanding information security practices in otherwise hard-to-access settings; uncovering perspectives, relationships and experiences that often remain hidden or unexplored through other (qualitative) methods. In ethnography, the reliance on first-hand observation for empirical findings and theoretical insights highlights the importance of the researcher's presence within the settings of the group(s) under study, necessitating extended fieldwork. For information security this means long-term explorations of how security is experienced and negotiated between people, how security threats are voiced and felt in different settings, and what expectations are held with respect to security within groups.

Fieldwork in different settings allows us to explore how social structures and relations shape security practices over time, enabling a deeper and inherently grounded analysis of such practices and related security concerns of the group(s) under study.

Yet, ethnography is costly; both in time and, well, money, and sometimes also in terms of mental health, physical safety and so on. Extended fieldwork in distinct and often unfamiliar settings usually means spending months away from friends and family, removed from established support structures and social networks. But doing fieldwork is also enriching, both academically and personally. Here, we share some insights from our own experiences of conducting long-term fieldwork in higher-risk settings with activist groups over the last year. All fieldwork took place outside the UK and, indeed, outside Europe, yet, at this point we refrain from naming the specific settings for our own safety and for the safety of our interlocutors. [i] Instead, we focus on the practicalities of researching the information security needs of activist groups at times of heightened security threats.

Mik's research is a multi-sited ethnographic study on how environmental social movements ensure the security of their information in different, yet politically volatile contexts. Rikke's research explores the information security practices of anti-government protesters and activists in one setting marked by extensive State violence, larger-and smaller-scale protests and frequent confrontations between national security forces and protesters. We have each spent several months in these settings over the last year. Since January 2024, Mik has conducted five months of fieldwork in one national context, which includes a one-month scoping trip and four months of full fieldwork. Rikke's fieldwork with the activist groups under study has, to date, lasted just over six months, including a two-week scoping trip and six months of full fieldwork spread across two trips (four months and two months). These examples also give a sense of how we structure our ethnographic fieldwork: conducting a shorter scoping trip to establish connections, building rapport with initial interlocutors and gatekeepers and developing site-specific security and safety protocols, before carrying out the full fieldwork. While this is an approach that has worked well for us given the settings we work in and the fact that our individual research projects are supported by grants, [ii] we recognise that this might not be the right approach for every ethnographer, and that many will conduct even longer fieldwork.

This is a snippet from Mik's fieldnotes from observing an adversary mapping activist workshop in a rural community (edited for anonymity):

As I move around the room, I notice some sheets filling up quickly while others remain largely blank. Some groups list specific individuals (e.g. [name redacted]) while others opt for broader categories (e.g. "military", "pastors"). Across all maps, the 'friendly but weak' quadrant is the fullest, followed by the 'hostile but strong' quadrant. The maps show a discordant mix of NGOs, POs, local government offices, local politicians, names of local families, union leaders, and fellow activists, with no strikingly obvious pattern in how actors are placed across regions (e.g. some groups consider local government offices adversarial; others consider them allies). At this point, I ask the participant beside me why her group has classed their decided allies as uniformly weak and their adversaries as strong, leading to a detailed story from her town in [Location A] about enforced disappearances related to the protest of a mine.

Of course, conducting ethnographic fieldwork in higher-risk settings as a researcher at a university is not always straightforward. It means following institutional processes and protocols, including often lengthy applications to the Research Ethics Committee (REC) detailing the research design and methods, participant engagements and protection mechanisms as well as complex risk assessments carried out in close consultation with institutional health and safety officers. Yet, while such institutional processes are both necessary and helpful in thinking through potential risks and mitigations relevant to a specific setting, extended fieldwork is often unpredictable and requires continuous re-assessments and often re-adjustments during the course of the research. Therefore, considerations and decisions about day-to-day interactions and activities also remain with the fieldworker. This is certainly our experience. Being in the field, with the people whose practices we study, often bring to light tensions and sometimes conflicts within and between groups, as Mik's fieldnotes here also exemplify. This can make being an observer in such spaces both uncomfortable and challenging: How do you position yourself? How many details do you capture? How are you perceived by different interlocutors? How does your presence shape discussions and practices? How much do you reveal about internal tensions, organising structures, actions and so on, when publishing your findings?

The higher-risk contexts of our fieldwork also raises questions about our own safety: How much do – and can – you share about your experiences, and with whom? Does your presence put interlocutors at risk? How safe are you in this setting? A snippet from Rikke's fieldnotes taken during a large-scale protest in 2024 bring some of these to light (edited for brevity and anonymity):

While many of these protesters might never have protested before, I observe how they have quickly become seasoned protesters with their own approach, establishing check-in protocols and creating their own groups. I'm in the street, surrounded by thousands of people. I see a lorry coming through the crowd, with protesters on the roof and hanging off its sides as they cheer, dance and sing. There is euphoria, happiness and a sense of 'revolutionary relief'. As one protester tells me: "I am so happy. This is not just the poor, but this is everyone coming together. People have had enough." The teargas today is really harsh, piercing, and even though I'm masked up and try to shield my eyes, it makes my eyes water. The atmosphere is defiant, joyful and unified [...] In those moments, the sense that the future would be different, that they had made history, dominated. I'm in a crowd of protesters cheering, news of people having been killed and thoughts about how the security forces will respond beginning to creep in. It is no longer a protest, maybe it never was, but the start of a revolution. At least these are my thoughts in this moment. One of my interlocutors who I have been separated from rings me to ask where I am. We decide to meet at one of our check-in points.

Fieldwork involves engaging with others where they are, over longer periods of time, building relationships and making complex interdependent commitments. It is therefore also important to us that our work has the potential to not only advance scientific knowledge but to benefit the people we engage with in ways that they define, following a model of engaged and reciprocal scholarship. [iii] This approach is particularly important when working with groups facing hardship and/or in contexts marked by extractivist dynamics and legacies. Reciprocity involves both conversations about knowledge production and practical exchanges. During our fieldwork, we have both run several in-person workshop sessions and training days with different activist and community groups. We have produced written materials such as a short handbook and guidance documents that cover topics such as secure messaging, file storage, location tracking, passwords, phishing, social media visibility and encryption basics. The content of these materials was shaped by the specific security concerns of the groups under study as gleaned from our observations and conversations during the fieldwork. These forms of knowledge sharing and reciprocity shape our research as much as interviews and participant observations.

[i] In ethnography, 'interlocutor' refers to a fieldwork-based relationship, a person with whom (and about whom) we conduct research. This is different to 'gatekeeper', who is someone that might act as a point of access to specific settings and/or interlocutors.

[ii] Mik's research is supported through the EPSRC Centre for Doctoral Research in Cyber Security for the Everyday (https://www.royalholloway.ac.uk/cdt/) and Rikke's research forms part of the EPSRC-funded project Social Foundations of Cryptography (social-foundations-of-cryptography.gitlab.io/about)

[iii] See the piece 'Then what? Reciprocity in information security research' in this newsletter, written by Jessica McClearn and Rikke Bjerg Jensen.

THE EVOLUTION OF EXCELLENCE: FROM THE SCC (2002-2024) TO THE 'TRUSTWORTHY, EMBEDDED AND AUTONOMOUS SYSTEMS CENTRE' (TASC)

Professor Konstantinos Markantonakis
Professor, ISG and Director of TASC

After more than 20 years of impactful research, commercialisation, and student-led innovation, the Smart Card and IoT Security Centre (SCC) at Royal Holloway is entering a new chapter. To better reflect our evolving research focus and strategic ambitions, we are rebranding as the **Trustworthy, Embedded and Autonomous Systems Centre (TASC)**. This transition builds on the SCC's strong track record in trusted execution environments, mobile device forensics, and cyber-physical system security. Our broadened mission now spans secure autonomous platforms, embedded AI, telecommunications, digital identity and payment systems, and media broadcasting. While our origins lie in smart card technologies, secure elements, and mobile security, these foundational areas now underpin a wider vision for trustworthy computing in real-world systems. TASC will continue to serve as a hub for research excellence, impact, and student engagement within the ISG. We invite collaborators and partners to explore our work at scc.rhul.ac.uk — soon to reflect our new TASC identity.

Building on this transition, in 2025, the Centre continues to play a vital role within the ISG, driving forward innovative research, real-world impact, and meaningful student involvement. Our work remains firmly rooted in addressing pressing cybersecurity challenges through technically ambitious, high-value initiatives. This year's update showcases select advances in our exploration of secure hardware-software interaction and trusted computing environments.

As part of our ongoing research efforts, we recently contributed to two significant publications — one exploring runtime integrity through control-flow attestation, and another introducing a novel framework for secure, distributed machine learning.

Our paper 'Control-flow attestation: *Concepts, solutions, and open challenges*' [1] offers a detailed survey of control-flow attestation (CFA), a method combining control-flow integrity (CFI) and platform attestation to verify runtime execution in computing systems. CFA counters attacks like return-oriented and data-oriented programming by ensuring execution adheres to authorized control-flow paths. Published in *Computers & Security*, a leading international journal in the field of cybersecurity, the paper reviews over 30 CFA schemes from 2016 to 2024, highlighting key features such as event measurement types, verifier-prover models, and trust anchors like TEEs (Trusted Execution Environments) and CPU trace modules. It categorizes CFA techniques by protocol type — challenge-response, continuous reporting, or local verification — and examines their applications in IoT, cloud, and embedded environments. The authors identify trends in instrumentation and runtime enforcement, outline the limitations of current CFA implementations, and call for standardized evaluation, scalability, and improved defences against advanced threats. The paper concludes with a set of open research challenges and future directions for advancing the state of CFA.

The paper 'DECML: Distributed Edge Consensus Machine Learning Framework' [2] introduces a privacy-preserving approach to collaborative machine learning for sensitive domains like healthcare and finance. Unlike federated learning, which shares model updates and risks data leakage, DECML decentralizes training: each node trains a private model and only shares derived insights. A central orchestrator issues queries and aggregates responses via a consensus mechanism, protecting both data and models. Built using Python, Scikit-learn, and Twisted, DECML runs on a peer-to-peer architecture. Evaluations on MNIST and CIFAR-10 demonstrate competitive accuracy and training performance. The framework is theoretically resilient to attacks such as inference, inversion, and data poisoning. With strong privacy guarantees and regulatory alignment, DECML presents a scalable, secure alternative to traditional collaborative learning methods.

One of the SCC's emerging research directions focuses on side-channel analysis techniques for instruction-level disassembly on RISC-V platforms, a timely and strategic area given the architecture's growing adoption across industry and security applications. The RISC-V instruction set disassembly using power analysis is an on-going area of research in the SCC. Hamza Rafi, who has recently joined the SCC as an undergraduate Computer Science student, and Amir Rafi (PhD student) are currently conducting ground-breaking work on profiling RISC-V instructions using side-channel power analysis techniques. RISC-V is an open-source and royalty free instruction set architecture (ISA) which has become popular over the recent years with currently over 2000 corporate members.

A power-based side-channel disassembler for RISC-V would enable researchers to recognise the instructions currently under execution on a RISC-V device without gaining software access to the device. This has potential applications in confirming the legitimate execution of applications and other intellectual property in constrained devices.

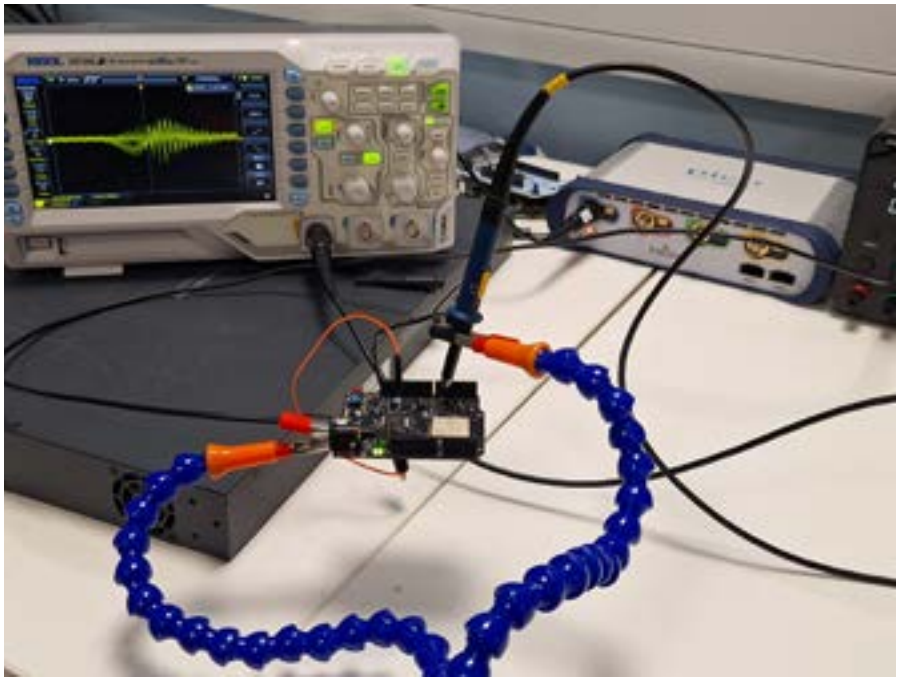

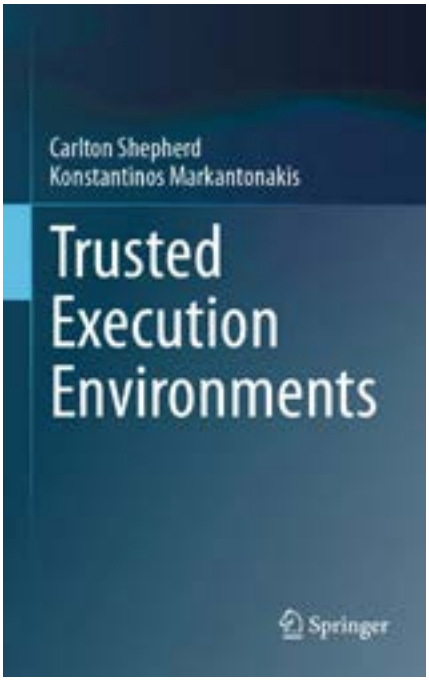

As we reflect on the evolution of the Centre, we would like to acknowledge the lasting contributions of a valued colleague. Dr. Darren Hurley-Smith, who has been an integral part of the Smart Card and IoT Security Centre (SCC), contributing significantly to our research and innovation efforts. His expertise in cyber-physical systems, network security, and autonomous systems has been invaluable. As the Technical Manager of the Omnidrome initiative at Royal Holloway, he played a pivotal role in advancing our capabilities in unmanned systems research. We extend our heartfelt gratitude for his dedication and contributions. We wish him continued success in his new role at the University of Kent.

I'm pleased to announce the publication of our book, Trusted Execution Environments, co-authored with Dr. Carlton Shepherd and published by Springer. This comprehensive work delves into the architecture, applications, and security implications of TEEs, covering technologies such as Intel SGX, ARM TrustZone, and AMD SEV. It serves as a valuable resource for practitioners, researchers, and students interested in understanding how TEEs safeguard sensitive data and code across diverse computing platforms. The book is available through Springer and major retailers.

Student engagement remains a cornerstone of the SCC's mission, in summer 2025, we are working with three teams made up of 13 Computer Science and Electronic Engineering students on distinct projects with strong research and commercial potential. Watch this space for real-world outcomes driven by their contributions.

We hope this brief insight into our recent work has captured your interest. If you see opportunities for collaboration or wish to learn more, please don't hesitate to contact us at k.markantonakis@rhul.ac.uk.

The Smart Card and Internet of Things Security Centre

REFERENCES

[1] Sha, Z., Shepherd, C., Rafi, A., & Markantonakis, K. (2025). Control-flow attestation: Concepts, solutions, and open challenges. *Computers and Security, 150,* Article 104254. https://doi.org/10.1016/j.cose.2024.104254

[2] Verdeyen, C., Shepherd, C., Markantonakis, K., Akram, R. N., Milroy, R., Abu Ghazalah, S., & Sauveron, D. (2025). DECML: Distributed Edge Consensus Machine Learning Framework. In *2024 IEEE Global Conference on Artificial Intelligence and Internet of Things (GCAIoT)* IEEE. https://doi.org/10.1109/GCAIoT63427.2024.10833588

SECURING HER HEALTH: FEMTECH AT THE CROSSROADS OF ONLINE PRIVACY AND SAFETY AND EMPOWERMENT

Dr Maryam Mehrnezhad
Reader, ISG

Since 2019, our international research team has investigated the security and privacy landscape of FemTech, a field holding immense potential for empowerment by offering tailored technological solutions for women's unique health needs. However, our findings reveal significant shortcomings in current practices, regulations, and the underlying technology, placing this promise of empowerment at a critical crossroads with online privacy and safety.

Our comprehensive analysis demonstrates widespread vulnerabilities in the handling of sensitive health information within these female-oriented health technologies, which range from intimate wellness devices to widely used fertility trackers. The very act of engaging with these tools, intended to provide greater control over one's health, well-being, and body, can inadvertently expose users to substantial risks if their data is not adequately protected. Given the rapidly expanding FemTech market, projected to exceed $103 billion by 2030, our research underscores an urgent need for technical, legal, and social reforms to better protect user data, ensure the responsible and secure development of this emerging healthcare sector, and ultimately safeguard the empowerment that FemTech aims to deliver.

USER PRIVACY IN FEMTECH APPS

In 2020, Dr Teresa Almeida and I published a paper in ACM CHI [1]. Our study titled 'Caring for Intimate Data in Fertility Technologies' was the first to investigate the privacy practices in fertility tracking apps, the most prevalent type of FemTech. We introduced the concept of 'differential vulnerabilities' in FemTech highlighting that diverse populations face varying privacy and security risks in FemTech, considering their exposure, susceptibility, and capacity to respond. We argued that fertility tracking data could be misused to identify individuals seeking reproductive healthcare, potentially leading to prosecution in some regions, or reveal infertility/pregnancy status to employers, causing discrimination and distress.

To examine this, we analyzed 30 popular fertility apps in the UK from the Google Play Store. Their initial analysis revealed that 40% of apps showed no privacy information upon opening, and 43% hid it within sign-up pages or terms and conditions. Only five apps had a dedicated privacy notice, but even these often lacked a decline option for data collection. Furthermore, using static and dynamic analysis, we found that most apps exhibited poor privacy practices, including collecting data before consent, using undisclosed trackers, and limiting user control over their data. The presence of multiple trackers within many apps was also a key finding. Several parts of our findings challenge the lack of agency in these systems when it comes to user control options.

Figure 1: Members of the Usable Security and Privacy (USP) Lab at Royal Hollowayworking on FemTech and Other Projects.

HARDWARE AND BLUETOOTH SECURITY IN FEMTECH IOT

We continued our studies over the years and published multiple papers. In one of our systems studies performed by Stephen Cook, I, and Dr Ehsan Toreini, we looked at the Bluetooth security of such systems. Our paper [2], published in 2024 in the International Journal of Information Security was titled 'Bluetooth security analysis of general and intimate health IoT devices and apps: the case of FemTech'.

We investigated technical weaknesses in 21 FemTech devices, the largest study of its kind. We analyzed the security of Bluetooth communication between these devices and their companion apps using two testing setups. The first setup employed BBC Microbits to monitor Bluetooth communication channels, revealing whether transmitted information was protected. The second setup created

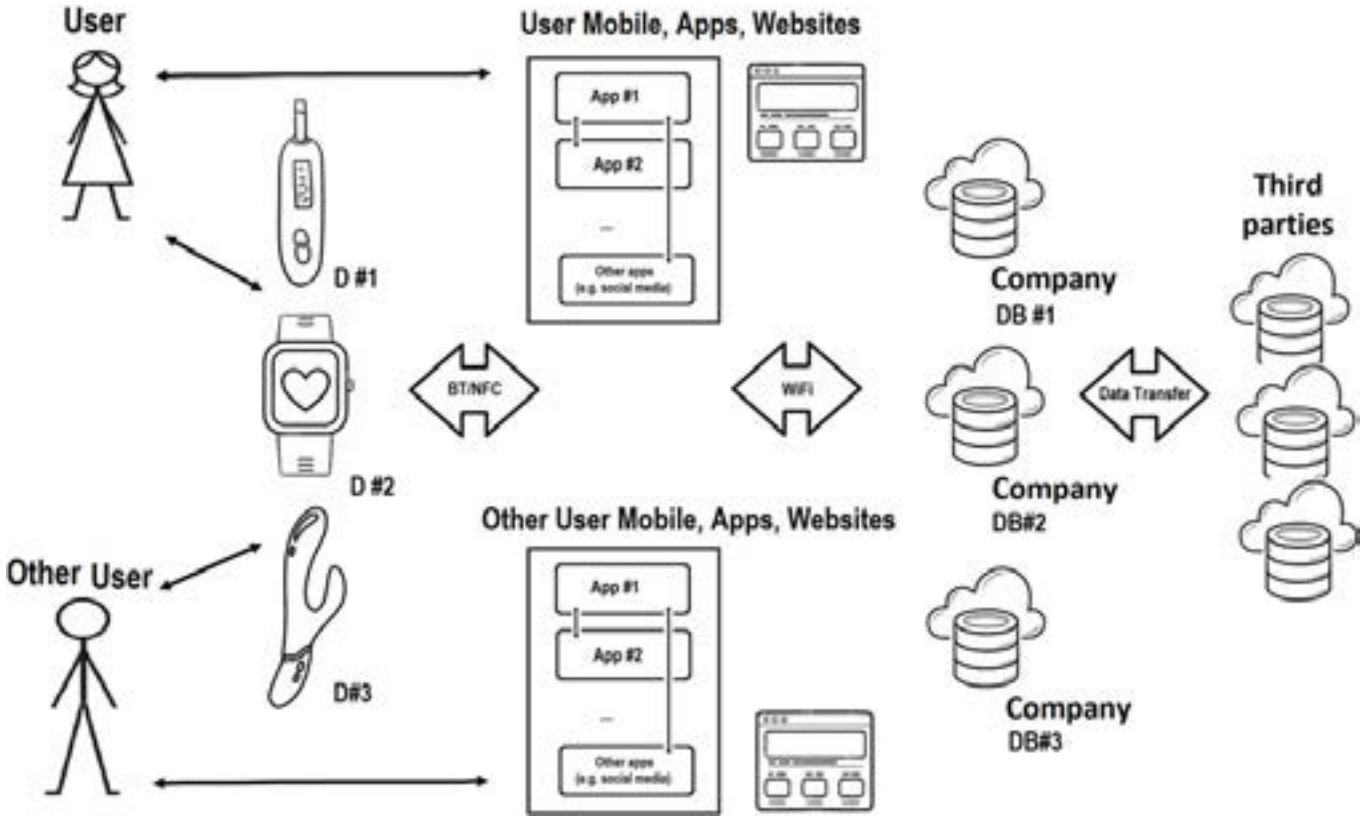

Figure 2: FemTech IoT Ecosystem; data can be entered into these systems via user putting it in the app, website, or device, or automatically collected by the smart device, app, and website. Data is transferred to other places e.g., the company's database in the cloud or sold to third parties.

a 'Man-in-the-Middle' environment to observe all data exchange, allowing us to attempt connection hijacking, denial of service, and battery draining attacks. The tests uncovered significant security flaws. For instance, one device's battery depleted rapidly during the battery draining attack. Alarmingly, most devices used outdated Bluetooth versions, and 20 out of 21 employed the insecure 'Just Works' pairing method, leaving them vulnerable to man-in-the-middle attacks, which we successfully demonstrated. These vulnerabilities have been reported to the vendors, and we have contributed to fixes where possible.

FEMTECH-RELATED REGULATIONS

Alongside multiple systems and user studies, we also looked at the related regulations. Together with Dr Thyla Van Der Merwe, and Professor Michael Catt, we published another paper [3]: 'Mind the FemTech gap: regulation failings and exploitative systems' in 2024 in the *Frontiers* journal, where we analyzed the shortcomings of current data protection regulations in safeguarding FemTech users. Examining GDPR, the Swiss Federal Act on Data Protection, and UK/EU Medical Device regulations, we identified significant gaps in addressing the specific nature of intimate health data collected by FemTech. The study categorized the diverse data collected, including personal details, lifestyle, reproductive health information, medical data, device/phone access (storage, contacts, location, sensors), and data about others (babies, partners, social media).

This complex data landscape presents unique privacy challenges. An analysis of 10 FemTech products (smart breast pump, fertility trackers, pelvic floor exercisers, sex toys, period/menopause tech, general female health devices, connected water bottle) revealed widespread GDPR violations. Only one app demonstrated valid consent mechanisms. Most bundled privacy notices with terms, used "tracking walls," or provided no privacy information. Similar to previous findings, numerous trackers were present, and websites often tracked users before cookie consent.

We found a critical disconnect between GDPR and medical device regulations, where the expected higher level of protection for personal health data in these products is not consistently enforced, even when the app isn't classified as a medical device. FemTech's combination of health solutions and user data creates intricate privacy challenges. We urge for domain-specific regulations tailored to FemTech, prioritizing marginalized users' needs. Ensuring security, privacy, and safety requires stronger enforcement and new frameworks acknowledging these technologies' unique risks.

MEDICAL VS NON-MEDICAL: MISCATEGORISING OF FEMTECH PRODUCTS

A significant issue identified across all three studies was the miscategorization of FemTech products in app stores. Many apps containing medical records or other highly sensitive health information were categorized as 'Health & Fitness' rather than 'medical devices', potentially allowing them to avoid stricter oversight. This categorization issue reflects broader challenges in regulating technologies that blur the lines between wellness tools and medical devices. We suggest that the developers be more upfront about the data they collect. They should be careful not to miscategorise the technology to bypass privacy guidelines, be clear on its usage and purposes, and state how data is shared with third parties. We also note the importance of ensuring that the most recent privacy considerations such as privacy-by-design, and minimal data collection are taken into account.

CONCLUSION

FemTech presents a powerful opportunity for empowerment by offering tailored technological solutions for unique health needs; however, this potential for enhanced agency exists at a critical juncture where online privacy and safety are significant concerns. As sensitive health data is collected, analyzed, and shared, ensuring robust security and transparent practices is paramount. As our studies prove, these issues within FemTech not only endanger women and girls but also the entire population due to its involvement in fundamental health and medical aspects like family planning. The very act of engaging with FemTech, intended for greater health control, can inadvertently expose users to vulnerabilities if these foundational aspects are not prioritized. Such a goal requires a careful and ethical approach that actively secures user data to truly empower individuals. This is only achieved via a coordinated and comprehensive effort among stakeholders (developer, policymakers, researchers, users) to reform the current practices to prevent the misuse of intimate health information.

REFERENCES

[1] 'Mind the FemTech gap: regulation failings and exploitative systems' in Frontiers, doi.org/10.3389/friot.2024.1296599

[2] 'Bluetooth security analysis of general and intimate health IoT devices and apps: the case of FemTech' in International Journal of Information Security, doi.org/10.1007/s10207-024-00883-3

[3] Caring for Intimate Data in Fertility Technologies in CHI '21: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, doi.org/10.1145/3411764.3445132

[4] Podcast: Do Security and Regulation Failures Put Women's Health Data, Their Privacy and Even Their Safety at Risk?, Link: scipod.global/do-security-and-regulation-failures-put-womens-health-data-their-privacy-and-even-their-safety-at-risk/

PORT IN A STORM: HOW IRAN'S CYBER CAMPAIGNS THREATEN GLOBAL MARITIME TRADE AND CHINA'S STRATEGIC INTERESTS

Cosimo Melella
Postdoctoral Researcher,
Centre for Maritime Cybersecurity:
Estonian Maritime Academy

Francesco Ferazza
PhD researcher, Royal Holloway,
University of London; Lieutenant,
Italian Navy

Konstantinos Mersinas (pictured)
Associate Professor, Royal Holloway,
University of London

Ricardo Lugo
Taltech, Senior Researcher,
Centre for Maritime Cybersecurity:
Estonian Maritime Academy

Over the past few years, ports in the Middle East have become more than logistical lifelines — they've become digital battlegrounds. Iran, a regional cyber powerhouse, has escalated its offensive cyber activities against maritime infrastructure, targeting ports in Israel, Egypt, and the Gulf. These state-sponsored campaigns are doing more than disrupting local operations — they're also complicating China's economic ambitions in the region.

A new wave of analysis reveals how Iran's hacking groups have aligned cyber operations with geopolitical strategy, deploying malware, phishing campaigns, and disinformation to disrupt trade, gather intelligence, and sow instability. At the same time, these actions threaten to undermine China's Belt and Road Initiative (BRI), which hinges on stable maritime trade.

PORTS UNDER CYBER SIEGE

The maritime sector has long lagged behind in cybersecurity, and Iranian actors are exploiting this vulnerability. Iranian Advanced Persistent Threat (APT) groups like **Yellow Lideric** (also known as Imperial Kitten or TA456) and **APT35** (Charming Kitten) have launched campaigns against logistics hubs and port operators using tactics such as spear-phishing, DNS spoofing, and email-based malware.

In one notable campaign, Iranian hackers compromised websites of Israeli shipping and logistics companies, embedding malicious JavaScript to spy on users and collect operational data [1]. These watering hole attacks were designed to intercept visits from port employees and maritime partners, yielding critical intelligence.

Meanwhile, in Egypt, APT35 ran a phishing campaign targeting the Port Said area. By hijacking DNS settings and mimicking trusted maritime service providers, they redirected users to attacker-controlled servers and stole credentials [2]. Their infrastructure was traced back to rogue subdomains and hosting services operating across Europe and the Middle East.

CHINA'S ECONOMIC INTERESTS AT RISK

Though Iran and China have signed a 25-year cooperation agreement, their interests in the region are far from aligned. China is focused on stability — vital for its BRI infrastructure investments, such as its 25-year contract to operate Haifa Port through the Shanghai International Port Group and its stake in Egypt's Suez Canal Container Terminal [3].

Iran's cyber aggression, however, introduces uncertainty into precisely these areas. By attacking logistics hubs and port operations in Egypt and Israel, Tehran risks disrupting infrastructure that China is actively developing. Cyber instability in these zones could compromise Beijing's long-term goal of ensuring secure and reliable trade routes through the Middle East [4].

EXPLOITING MARITIME VULNERABILITIES

Despite the critical role ports play in global logistics, they remain highly vulnerable to cyberattacks. Poor segmentation, outdated software, abandoned subdomains, and exposed control systems leave them easy targets for advanced adversaries.

A 2023 incident involving Israel's largest oil refinery, the **BAZAN Group**, illustrates this risk. The pro-Iranian hacktivist group **Cyber Avengers** claimed responsibility for a Distributed Denial of Service (DDoS) attack that temporarily disabled public access to BAZAN's online systems [5]. Although the refinery denied any serious breach, screenshots shared by attackers hinted at potential access to sensitive SCADA (industrial control) systems — whether real or exaggerated, the psychological impact was profound.

Even more alarming was the group's claim to have exfiltrated 1TB of sensitive power infrastructure data from Ashdod Port, offering it for sale on Telegram. Analysts later concluded that this was likely part of an influence campaign rather than a genuine breach, but the narrative alone sparked global concerns [6].

MEET THE THREAT ACTORS

The Iranian cyber threat landscape in the maritime space is largely dominated by three groups:

- **Yellow Lideric (Imperial Kitten):** Active since at least 2018, this group has used phishing campaigns and custom malware such as IMAPLoader to collect intelligence from port logistics companies [1]. Its operations are aligned with the Islamic Revolutionary Guard Corps (IRGC) and are highly tailored to maritime sectors.

- **APT35 (Charming Kitten):** This prolific actor uses social engineering, DNS manipulation, and typosquatting to steal credentials from maritime stakeholders, particularly in Egypt. Their tools include Hyperscrape, a credential-harvesting utility [2].

- **MuddyWater:** Believed to be tied to Iran's Ministry of Intelligence, this group focuses on spear-phishing campaigns targeting logistics companies in the Gulf. In a 2024 case, it targeted **AllTrans Freight & Logistics LLC**, using fake subdomains and remote access tools like BugSleep and MuddyRot [7].

These groups often reuse infrastructure and evolve their malware to evade detection. Their goal is not just disruption but long-term espionage and strategic advantage.

CYBERSECURITY AS STATECRAFT

These attacks are not isolated incidents. They are part of a broader geopolitical strategy in which Iran leverages its cyber capabilities to weaken adversaries, project regional power, and complicate foreign interests — particularly those of the U.S. and its allies.

As digital and physical infrastructure converge, so do the risks. Cyberattacks on ports affect not only the flow of goods but also the security of energy supplies, military logistics, and regional diplomacy. And with increasing Chinese investment in these same regions, Tehran's actions may trigger a quiet recalibration of Sino-Iranian cooperation.

WHAT NEEDS TO CHANGE

Our research concludes with a clear warning: maritime cybersecurity is a global blind spot. Most port operators still lack basic protections such as vulnerability monitoring, incident response plans, and access controls on critical systems. As a result, cybercriminals and state actors alike can exploit weaknesses at scale.

To address this, the report recommends:

- **Sector-specific cybersecurity standards** tailored for port operations and maritime logistics.

- **Improved threat intelligence sharing** between governments, port authorities, and private security firms.

- **International cooperation** to align legal frameworks and create accountability mechanisms for cyberattacks on critical infrastructure.

Ultimately, protecting the Middle East's ports is about more than keeping ships moving. It's about shielding global supply chains, regional stability, and the future of international economic cooperation.

SELECTED REFERENCES

1. ClearSky Cyber Security. Fata Morgana: *Watering Hole Attack on Shipping and Logistics Websites.* May 2023. Link

2. SecurityScorecard. *Charming Kitten Campaign: Maritime Targets in Egypt and Israel.* 2022. Link

3. Shanghai International Port Group (SIPG). *Haifa Port Operations and Strategic Importance.* 2015. Link

4. Miao, Qiang. *China's Strategic Junctions: The Role of Haifa in the Belt and Road Initiative.* China Daily, July 2023. Link

5. BleepingComputer. *Israel's Largest Oil Refinery Website Offline After DDoS Attack.* July 2023. Link

6. SecureWorks CTU. *Analysis of Cyber Avengers' Influence Campaigns.* 2024. Link

7. Sekoia. *MuddyWater Replaces Atera by Custom MuddyRot Implant in a Recent Campaign.* June 2024. Link

8. International Maritime Organization (IMO). *Cybersecurity in the Maritime Sector: Addressing Global Challenges.* October 2024. Link

9. CrowdStrike. *IMPERIAL KITTEN Deploys Novel Malware Families.* Nov. 2023. Link

10. ATS Advanced Technical Solutions. *The Rise of Cyber Security Threats to Critical Infrastructures in the Middle East.* 2024. Link

11. Melella, C. Ferazza, F., Mersinas, K., and Lugo, R. Port in a Storm: Iranian Cyber Operations and Chinese Strategic Interests in Middle Eastern Maritime Infrastructure. In *2025 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW).* [under publication]

CRYPTOGRAPHY GROUP SUCCESSFUL AT EUROCRYPT 2025

Dr Christian Weinert
Lecturer, ISG and Director of the Cryptography Group

The Cryptography Group has had three papers accepted at EUROCRYPT 2025, a CORE A*-ranked conference and one of the premier venues for cryptographic research. This year, the conference will be held in Madrid, Spain, from May 4 to 8. Special congratulations go to our doctoral students from the Centre for Doctoral Training (CDT) of the Everyday who co-authored the accepted papers: Benjamin Benčina, Daniel Jones, and Simon Pohmann.

The acceptance of these papers at EUROCRYPT 2025 not only reflects the high calibre of research conducted by the Cryptography Group but also shows the breadth of research expertise, addressing both theoretical and practical challenges in cryptography – from enhancing privacy-preserving protocols to analysing real-world communication systems. Below, we briefly highlight the contributions made in each paper.

1. HOLLOW LWE: A NEW SPIN, UNBOUNDED UPDATABLE ENCRYPTION FROM LWE AND PCE

Authors: Martin R. Albrecht (King's College), Benjamin Benčina (Royal Holloway), and Russell W. F. Lai (Aalto University)

Public-Key Encryption (PKE) is central to securing digital communications. However, traditional PKE systems become inefficient when keys need frequent updating – an essential requirement in dynamic environments such as secure messaging. Updatable Public-Key Encryption (UPKE) addresses this issue, but lattice-based constructions, while being post-quantum secure, suffer from noise accumulation during key updates, which restricts the number of possible updates. In this work a novel variant of the prominent Learning with Errors (LWE) problem is proposed that – coupled with a rotation-based key update mechanism – allows to eliminate noise growth completely. As a result, unlimited secure key updates are possible, significantly enhancing the practicality of encryption systems in dynamic environments.

2. FORMAL ANALYSIS OF MULTI-DEVICE GROUP MESSAGING IN WHATSAPP

Authors: Martin R. Albrecht, Benjamin Dowling (both King's College), and Daniel Jones (Royal Holloway)

Secure messaging platforms like WhatsApp are integral to modern communication, often promising strong privacy guarantees via end-to-end encryption. Yet, extending these guarantees seamlessly to multi-device scenarios is challenging due to complexities around synchronizing and securing multiple devices. Unfortunately, the implementation of multi-device group messaging for WhatsApp, one of the largest secure messaging platforms worldwide, remains largely undocumented publicly. This work attempts to provide a rigorous formal security analysis by reverse-engineering WhatsApp's protocol and modelling it within an extended 'Device-Oriented Group Messaging' framework. The study systematically assesses WhatsApp's security properties, uncovering both its strengths and limitations.

3. ON ALGEBRAIC HOMOMORPHIC ENCRYPTION AND ITS APPLICATIONS TO DOUBLY-EFFICIENT PIR

Authors: Hiroki Okada (KDDI Research and University of Tokyo), Rachel Player, Simon Pohmann, and Christian Weinert (all Royal Holloway)

Private Information Retrieval (PIR) protocols enable users to retrieve items from a server's database without revealing which items are being retrieved. Unfortunately, this often results in severe computational and communication overhead. For example, many traditional PIR protocols require the server to process the entire database for each query. Doubly-efficient PIR (DEPIR) protocols address this inefficiency by introducing a client-independent preprocessing phase that allows subsequent queries to be answered with sublinear communication overhead and in sublinear time relative to the database size. Unfortunately, while asymptotically optimal, the only two known DEPIR constructions based on algebraic homomorphic encryption (HE) are concretely inefficient.

To address this issue, this work presents optimizations for DEPIR that fundamentally rely on a better understanding of algebraic HE. This includes deriving new lower bounds and defining a new relaxed notion of algebraic HE that is sufficient for constructing DEPIR. While the resulting performance might still not practical for large-scale applications, significant improvements in both computational speed and memory usage compared to the previously best DEPIR construction are evident.

For more information on the Cryptography Group at Royal Holloway and our ongoing research, visit cryptography.isg.rhul.ac.uk

NEW THIRD EDITION OF EVERYDAY CRYPTOGRAPHY

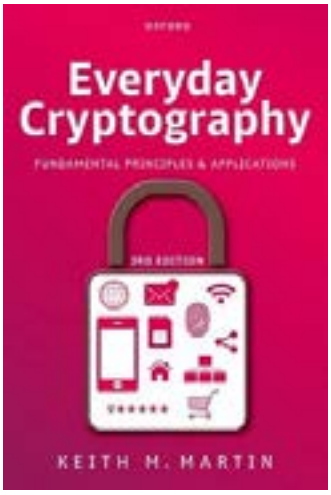

Back in 2012, when Oxford University Press (OUP) published the first edition of Everyday Cryptography, I regarded it as a project complete. The book was written to support the Introduction to Cryptography module component of Royal Holloway's MSc Information Security programme and was focused on fundamental principles and applications. By avoiding the mathematical details of the latest algorithms and treating cryptographic tools as black boxes, I did not anticipate that the book would date significantly. Unlike some areas of cyber security, cryptography evolves slowly and the fundamentals don't change much. Right?

Well, mostly right! In 2017, OUP requested a second edition and suggested that they expected around 20% new content. This seemed a manageable request. Cryptographic principles had indeed not changed, but the latest applications of them had, so I added new case studies on emergent technologies such as Bitcoin and WhatsApp, updated existing case studies to incorporate new versions such as TLS 1.3 and, following the Snowden revelations, added a new chapter on control of use of cryptography.

When OUP came calling again in 2024, I foresaw a similar exercise. The fundamental principles had probably still not changed, so it looked like a bit of updating of applications would do the job. However, there was a catch. The second edition of Everyday Cryptography (20% larger than the first) had reached the maximum acceptable book length. Therefore, the 20% new content required for the third edition would need to come at the expense of existing material. In other words, writing the third edition would require a complete rewrite of the book.

Over a decade since my first attempt, it was now time to rethink whether the fundamentals of cryptography had changed in any way. On reflection, there were four primary issues that I felt a new third edition needed to tackle.

Firstly, there has been evolution in best practices in cryptography. For example, previous editions of the book skipped lightly around the concept of formal provable security, but this has now become so essential to design of new cryptographic tools that I felt compelled to provide some intuition as to what lies behind this idea. There has also been a shift towards demanding properties such as perfect forward secrecy, which merited a more detailed look at different means of using asymmetric cryptography to support symmetric key exchange. Most significantly, there is much greater awareness of the pitfalls that can arise during implementation of cryptography. Previous editions dismissed these as beyond the book's scope, but the third edition devotes a new chapter to this topic.

Secondly, there is a revolution underway regarding the core cryptographic tools supporting everyday uses of cryptography. This is driven by the potential future threat presented by quantum computers. The third edition considers some of the new post-quantum algorithms that have emerged and considers issues such as migration processes. In 2017, we knew these issues were coming but they seemed a long way off. Today, while the manifestation of the quantum threat may well be distant, the need to prepare has definitely arrived.

Thirdly, there has been an emergence of interest in the different ways that cryptography can support privacy. This is partially demand driven following the exposure of national surveillance programmes, ongoing data breaches and increased storage of data by third party (cloud) service providers. However, it is also supply driven through the improved diversity and performance of cryptographic privacy-enhancing technologies. While the likes of zero-knowledge proofs, secure multiparty computation and homomorphic encryption all existed in 2012, there was very little demand for them and they were impractical to deploy. The situation is quite different today. The third edition thus features a new chapter that reviews these various tools and explains the ways they support different aspects of privacy.

Finally, as in 2017, updates to the various application case studies considered in the book are necessary. The third edition refreshes these, introducing cryptographic material on the likes of WPA3 security for Wi-Fi, 5G mobile telecommunications and contactless payments. In addition, I added some material on ransomware since this has risen to become a cyber-scourge, alas one created by cryptography.

What material was thrown out? Amazingly, not too much. As every writer knows, careful rewriting achieves a fair degree of content pruning! However, some cryptographic concepts are less important today, some restructuring of previous dedicated chapters on digital signature schemes and cryptographic protocols created space, and older case study material has become redundant.

There was just one more thing to consider. The cover! Anyone who has authored a book will appreciate just how contentious this issue can be. Back in 2012, the OUP design team offered up a set of covers that were either clichéd (keys spinning in a sea of ones and zeros) or simply inappropriate (a slow-shutter night image of car lights on a motorway). The editor and I eventually won a battle to have a cover created by graphic designer Allyson Waller, who produced the ISG newsletter for many years. Ally designed an artistic abstraction of applications of cryptography presented on a bold black background. I loved it! Come the second edition in 2017, OUP simply suggested replacing the black background with white. For the third edition… we're going red!

The third edition of Everyday Cryptography was published by OUP on 30th June 2025.

THE MODERATION GAP: MEASURING ONLINE HATE ON 4CHAN USING AI

Dr Adrian Bermudez-Villalva
Visiting Researcher, ISG

Dr Maryam Mehrnezhad (shown right)
Reader, ISG

The internet has revolutionised communication, enabling connection and collaboration across the globe. However, it has also fostered a darker side, online hate speech. While mainstream platforms have taken steps to moderate harmful content, less regulated spaces such as 4chan remain breeding grounds for extreme ideologies. Our latest research [1-2] delves into the troubling prevalence of hate speech on 4chan's politically incorrect board (/pol/), using cutting-edge deep learning techniques.

Figure 1: An example of a 4chan post.

EXPLORING ONLINE HATE: THE 4CHAN /POL/ STUDY

4chan is infamous for its anonymous nature and minimal moderation, making it a unique case study for analysing hate speech. Our research employed state-of-the-art Natural Language Processing (NLP) models, including RoBERTa and Detoxify, to measure the extent and types of hate speech found on /pol/. With a dataset of 500,000 posts collected over time, we aimed to quantify the scale of harmful discourse and uncover its hidden patterns.

In this project, our key objectives were:

- To measure the prevalence of different forms of hate speech on /pol/.
- To assess the toxicity of discussions and their impact on digital safety.
- To identify recurring topics within hate speech to better understand its context.

DISTURBING FINDINGS: A BREAKDOWN OF ONLINE HATE

Our analysis revealed that **11.2% of posts contained hate speech**, targeting various communities based on race, religion, gender, and sexual orientation.

Among the different categories:

- **Racism** was the most prevalent, making up 35.9% of hateful posts.
- **Religious hate** followed closely, accounting for 23.3%.
- **Sexual orientation hate** comprised 16.5%.
- **Sexism** was found in 12% of hateful discussions.

Beyond hate speech classification, our research also evaluated toxic content, uncovering high levels of obscenity, identity attacks, and threats. Discussions involving sexual orientation and racism were among the most toxic, with nearly 99% of flagged posts classified as highly offensive.

BEYOND NUMBERS: THE LANGUAGE OF HATE

Using topic modelling techniques, we identified recurring themes within these hateful discussions. Racial hate speech often incorporated conspiracy theories and dehumanising rhetoric. Religious hate was largely directed at Jewish and Muslim communities, filled with stereotypes and aggressive language. Meanwhile, sexism on /pol/ displayed strong misogynistic tendencies, reducing women to derogatory labels and objectification.

The findings highlight the complexity of online hate, where discussions are not limited to overt slurs but extend into coded language, political rhetoric, and ideological indoctrination.

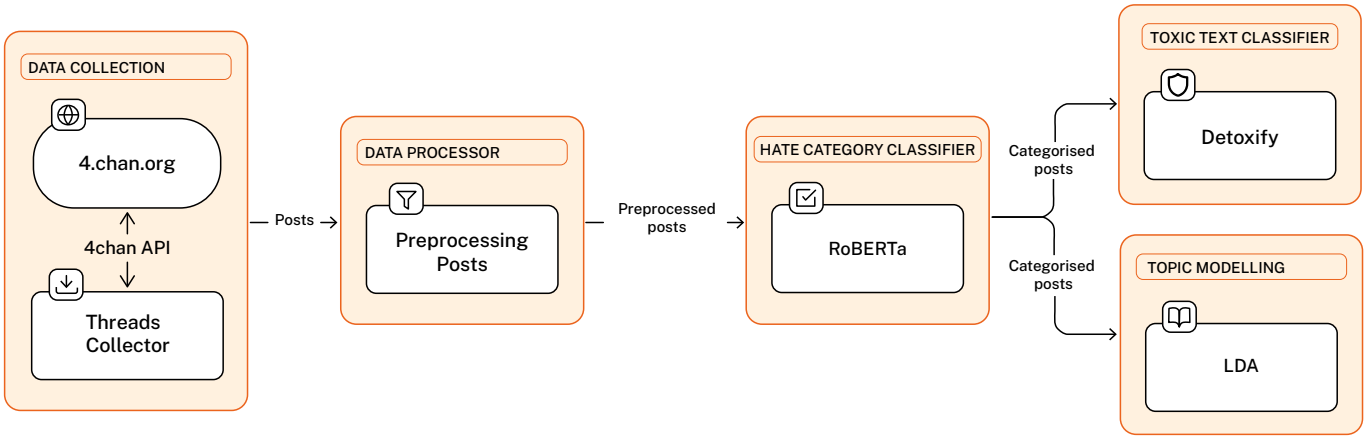

Figure 2: Overview of the methodology used to measure online hate on 4chan.

IMPLICATIONS FOR MODERATION AND POLICY

This study has profound implications for policymakers, platform developers, and researchers:

- **For policymakers**, understanding the nature and prevalence of online hate can help craft more effective regulations, such as the UK Online Safety Act and EU content moderation policies.
- **For tech developers**, our findings stress the importance of advanced AI-driven moderation tools that can detect not only explicit hate speech but also nuanced and coded language.
- **For society**, these results underscore the need for digital literacy initiatives to educate users on identifying and countering online hate speech.

CONCLUSION: ADDRESSING ONLINE HATE IN UNMODERATED SPACES

The findings from this study expose the dangers of unmoderated digital spaces and their role in spreading hate speech. While freedom of expression is a fundamental right, it must be balanced against the risks of online harm. In the future, we aim to develop actionable solutions that empower individuals to navigate online spaces safely while holding platforms accountable for fostering inclusive digital environments.

ACKNOWLEDGEMENT:

This work is supported by the UK Research and Innovation (UKRI), through the Strategic Priority Fund as part of the Protecting Citizens Online programme (AGENCY: Assuring Citizen Agency in a World with Complex Online Harms, EP/W032481/2).

REFERENCES

[1] Adrian Bermudez-Villalva, Maryam Mehrnezhad, Ehsan Toreini, Measuring Online Hate on 4chan using Pre-trained Deep Learning Models, IEEE Transactions on Technology and Society, 2025, link: arxiv.org/abs/2504.00045

[2] Dataset: Measuring Online Hate on 4chan Using Pre-Trained Deep Learning Models A Bermudez-Villalva, M Mehrnezhad, E Toreini, link: zenodo.org/records/14219048

THEN WHAT? RECIPROCITY IN INFORMATION SECURITY RESEARCH

Jessica McClearn (picture block, top left)
PhD Researcher, CDT in Cyber Security for the Everyday

Rikke Bjerg Jensen (picture block, bottom left)
Professor ISG;

Reem Talhouk
Associate Professor, Northumbria University

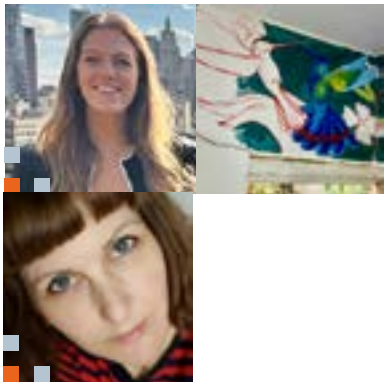

The three of us – Jess, Rikke and Reem – were recently asked to contribute a piece to the IEEE Security and Privacy Magazine, to draw greater attention (and more research) to the intersection of conflict and information security. In discussing what we would want to cover in such a piece, we quickly landed on how the prefix 'cyber' often masks the lived experiences of communities that live through and flee from war. We titled our piece 'The Everyday Security of Living through Conflict'. But it also led us to discuss how we do reciprocity in information security research and how our approach does not necessarily fit into existing metrics and frameworks in UK higher education. We reflect on this here.

THEN WHAT?

As ethnographers who engage a diversity of communities in our research, often in marginalised and/or higher-risk contexts and often over longer periods of time, to understand their information security needs, practices and perspectives, we are repeatedly faced with, and ask ourselves, the question: *then what*? How will our research benefit these communities? Why should they (want to) engage in this research?

These are not unique questions to the study of information security. However, as ethnographers we spend extended periods of time, sometimes several months, with the communities we work with, hearing and observing the threats they face and the insecurities they feel. We are welcomed into people's homes and social networks, introduced to distinct practices and ways of life, and entrusted with people's personal – often intimate and sensitive – stories and experiences. Doing ethnographic research with a focus on information security means that we, quite literally, position ourselves in the midst of the daily lives of people and everything that this entails; we observe and aim to understand their fears and worries, the threats they face from often strong and powerful adversaries, and the practices they try to incorporate into their daily lives to protect themselves and each other. In our experiences, this leads to a feeling of wanting to do something, an almost heightened sense of responsibility to make our work meaningful to the communities we engage with. But, *then what*?

It is important to note that what we consider here is not driven by a want to demonstrate impact to some external party in line with existing impact or knowledge exchange frameworks in UK higher education. In our vocabulary, reciprocity is realised through much more subtle ways of combining our knowledge of information security and our ethnographic positioning within the communities we study to benefit them in diverse ways. Academia also trades in the currency of research

publications, preferably at respectable or 'top-tier' venues. Indeed, when submitting ethnographically founded papers, we are of course required to demonstrate that our work is ethically sound, while being asked to provide recommendations for information security interventions, future research directions, technology design and/or policy-making. But, *then what?*

RECIPROCITY

Then what might reciprocity look like for those we involve in the kind of research we do? The answer to this question will likely look different to each researcher and feature differently across research designs; in many ways, it *ought* to differ since communities differ and will likely have distinct (information security) needs. Fundamentally, however, reciprocity is often experienced as a vouch of gratitude for the time and trust invested in the relationship with ethnographers during longer-term fieldwork. Indeed, reciprocity has become increasingly foundational to ethnographic work in response to the historic critique of extractivist approaches which clouds the discipline of social anthropology.

Ethnographically speaking, reciprocity might mean multiple visits and extended engagements where research insights are shared and/or co-designed with communities, where findings are brought to communities through different forms of dialogue and mutual meaning-making, where research outputs are translated into local languages and where collaborative workshops, training sessions and discussion forums are held on the topic of information security as experienced and understood from the perspective of the communities themselves. We exemplify how we have approached reciprocity in ethnographic fieldwork carried out by Jess in Colombia.

I established connections with communities in Colombia in 2018-2019 through a research aim seeded in Belfast, with the ambition to share learnings between two post-conflict societies. After completing ethnographic research in Cauca, Colombia, for my Master's thesis, I continued to engage in differing capacities (as a researcher, volunteer and friend) both in-person and virtually, until the opportunity arose to pursue my PhD in information security where I could further expand this work and build on existing relations.

During my current PhD research, a one-month pilot field study in 2023 and four months of fieldwork in 2024, reciprocation has changed several times in response to changing community needs. For example, one participant shared the need for tools to announce and denounce violence in their community, especially in relation to youth. They asked me: "Are there people in our extended research and practitioner networks who we can connect to develop this idea further?" Multiple participants wanted to be connected to other local organisations, which I could facilitate through my research. I have reviewed people's CVs, painted murals in a local coffee shop, provided advice on international scholarships, amongst many other things. I spoke at the National Day for Victims to support and raise awareness about the communities that I work with and I held local workshops on security futures. Reciprocation is ongoing; as I continue my research I also continue the dialogue with those who I know in Colombia. Indeed, sometimes as ethnographers we quite simply listen to people's stories, putting aside our research goals, closing our notebooks and switching off our recorders.

When researching in Colombia one interlocutor shared how they "do not come from a culture of written word", when speaking on behalf of other female conflict victims who had been forcibly displaced to the urban centre. Indeed, an academic article, even if translated to the local language, may be of limited or no relevance to the communities we work with, reinforcing that reciprocity will take place in different forms, including the co-development of blogs, toolkits, talks, art pieces and other artefacts that may be more 'useful' to such communities.

As this example illustrates, reciprocity is not a fixed idea or approach, but depends on the specific setting as well as the level and length of engagement. Often, however, what communities need is not something we as researchers can provide. We therefore also manage expectations, recognising that our research might have little impact on the communities we engage in the research. While some field relations continue long after the fieldwork has completed, other relations do not. Thus, reciprocity is often uneven across research sites and contexts. It often requires time and sometimes resources. It may lead to long-term and continuous engagements with communities through shared learning and/or follow-up research, with many ethnographers returning to the same sites over several years. Regardless of what might constitute reciprocity in a given context, we consider the question *then what* well worth asking.

[1] The piece is now out and can be found here: computer.org/csdl/magazine/sp/2025/02/10942495/25p31ql8zTy; open access version here: researchportal.northumbria.ac.uk/en/publications/the-everyday-security-of-living-with-conflict

WHEN SECURITY BECOMES EXPOSURE: A SECURITY AND PRIVACY EVALUATION OF IP CAMERAS ON SHODAN

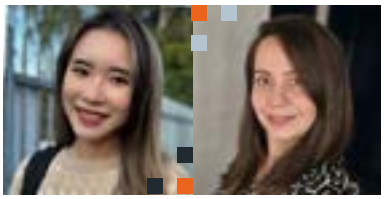

Cheok Ieng Ng (picture block, left)
Former MSc Student, ISG
Dr Maryam Mehrnezhad (picture block, right)
Reader, ISG

In today's world, where smart devices are ubiquitous in our homes and workplaces, IP cameras stand out for their practicality and ease of use. They provide real-time video feeds for home security, pet monitoring, and more. However, like many IoT devices, their rapid adoption brings significant security and privacy concerns. In our research, we investigated UK publicly accessible IP cameras available on Shodan, an online search engine for internet-connected devices, highlights these risks.

WHAT IS AN IP CAMERA?

An IP camera is a digital camera that uses the internet to send and receive video. Unlike traditional CCTV cameras, IP cameras don't need a local recorder. They can be accessed directly over a network, making them ideal for surveillance. Additionally, IP cameras support advanced features such as motion detection, two-way audio, and cloud-based storage, further enhancing their capability for both residential and commercial security systems.

IP CAMERAS: A DOUBLE-EDGED SWORD

IP cameras have transformed the surveillance landscape, replacing traditional CCTV systems with internet-enabled devices that offer features like motion detection and cloud storage. However, these attractive features also introduce vulnerabilities. Misconfigurations during setup or manufacturing can leave these cameras exposed to unauthorized access.

Figure 1: An AI-generated visualisation, based on real IP camera footage, highlighting the privacy risks of exposed residential spaces captured by unsecured IP cameras. Generated by ChatGPT

This research examined 281 publicly accessible UK-based IP cameras indexed on Shodan, uncovering alarming vulnerabilities, such as unsecured RTSP (Real-Time Streaming Protocol) connections and footage accessible without authentication. Following our initial findings, we then focused on a subset of 7 captured footage and qualitatively analysed the privacy risk exposure.

KEY FINDINGS

1. Exposure of Residential Spaces
   Over 76% of the IP cameras analysed were found to be installed in residential areas. These devices were often placed outdoors to monitor entry points and surrounding areas. The footage captured by these cameras frequently included sensitive details, such as the layout of homes, daily routines of residents, and even identifiable individuals, posing significant privacy concerns.

2. Shodan as a Risk Multiplier
   Shodan, a search engine for connected devices, amplifies the risks associated with unsecured IP cameras. By indexing these devices, Shodan provides malicious actors with easy access to locate and exploit vulnerable cameras. Additionally, cached footage and metadata provide insights into a camera's location and usage patterns, further compromising security.

3. AI Enhancing Risk
   The use of advanced AI tools like ChatGPT demonstrates both the capabilities and potential dangers of modern technology. Analysis of camera footage using AI can extract detailed information, such as property type, the presence of individuals, and environmental context. While valuable for legitimate purposes, such tools could also be misused to automate malicious activities, including stalking or targeting individuals.

4. Legislative Gaps
   Existing regulations, such as the GDPR and the UK's PSTI Act, fail to adequately address the security challenges posed by IP cameras. These legislative frameworks lack specificity regarding IP camera security requirements, resulting in inconsistent levels of protection. Manufacturers also face significant challenges due to the diversity of global standards, which increase costs and complicate compliance efforts.

5. The Broader Implications
   The study highlights a troubling paradox: while IP cameras are designed to enhance physical security, their vulnerabilities often lead to online privacy breaches and security risks. For example, an outdoor camera designed to deter intruders might unintentionally broadcast sensitive footage accessible via Shodan.

   These risks are not limited to individual users. Compromised cameras can serve as entry points for larger cyberattacks, such as Distributed Denial-of-Service (DDoS) campaigns, or act as nodes in botnets like Mirai.

MITIGATION STRATEGIES

To address these challenges, the study suggests a comprehensive approach:

- **For Users:** Users should prioritise changing default passwords on their IP cameras and enabling encryption to protect their data. It is equally important to regularly update the firmware on these devices and avoid using outdated models that may lack critical security updates. When selecting a camera, users should opt for devices with robust built-in security features to minimize vulnerabilities.

- **For Manufacturers:** Manufacturers should adopt security-by-design principles, ensuring that features like encryption and authentication are enabled by default in their devices. Providing clear and accessible user guides is essential to help users understand how to secure their cameras effectively. Additionally, manufacturers must commit to post-sale support by offering timely firmware updates to address emerging security threats.

- **For Regulators:** Regulators should enforce uniform global standards for IoT security to eliminate regional loopholes that may weaken protections. Mandating basic security measures, such as encryption, for all IoT devices is critical. Public awareness campaigns should also be promoted to educate users about the importance of securing their IoT devices and the steps they can take to safeguard their privacy.

CONCLUSION

The increasing popularity of IP cameras highlights the critical need for stronger security measures. This study shows that failing to act can lead to serious consequences, from privacy breaches to large-scale cyber threats. By fostering user vigilance, ensuring manufacturer accountability, and implementing robust regulatory frameworks, we can ensure that these devices fulfil their promise of safety without compromising security.

Are your IoT devices secure? Check their settings today and ensure they are configured to protect your privacy. Together, we can make the digital world a safer place.

REFERENCES

[1] Ng, Cheok leng, and Maryam Mehrnezhad. "Security and Privacy Evaluation of IP Cameras on Shodan." *International Conference on Research in Security Standardisation*. Cham: Springer Nature Switzerland, 2024.

UPDATES ON STANDARDISATION OF FULLY HOMOMORPHIC ENCRYPTION

Dr Rachel Player
Senior Lecturer, ISG

Fully homomorphic encryption (FHE) enables the evaluation of arbitrary functions on encrypted data. This is a powerful primitive that could enable applications in a variety of sectors, including healthcare and finance. As well as the usual key generation, encryption, and decryption algorithms, fully homomorphic encryption schemes are equipped with an evaluation algorithm, which operates on ciphertexts to provide the additional functionality.

In recent years, large technology companies as well as a growing number of start-ups have been driving the commercialisation of this technology. This has been accompanied by a large-scale community effort aimed at the standardisation of homomorphic encryption, which was initiated in 2017. The community effort is known as HomomorphicEncryption.org and is an open consortium of participants representing industry, government, and academia. Royal Holloway academics have been actively participating in this initiative since 2018, with former CDT researcher Ben Curtis and me having been invited speakers at the 4th and 5th meetings.

Since July 2023, I have been serving on the steering committee for HomomorphicEncryption.org and in this context I have co-organised the 7th and 8th HomomorphicEncryption.org Standards Meetings that took place respectively in Salt Lake City, USA (October 2024) and Istanbul, Turkey (March 2025). Alongside updates from related standardisation initiatives and keynote talks, both meetings devoted significant time to breakout sessions that focus on advancing key directions for the community. There were sessions on software tools for developing and benchmarking homomorphic encryption solutions, hardware acceleration approaches, and industry applications of homomorphic encryption. Breakout sessions on security in homomorphic encryption were organised by Yuriy Polyakov (Duality Technologies), Ro Cammarota (Intel), and me. In the Salt Lake City meeting, I was glad to co-lead the session on the day. In Istanbul, the security breakout session was led by CDT researcher Erin Hales, Yuriy Polyakov, and Ilaria Chillotti (Desilo).

Alongside the activities of HomomorphicEncryption.org, there is a formal effort to standardise FHE through ISO/IEC JTC1 SC27 WG2, a body within which Royal Holloway academics have been actively involved for over 30 years. The formal standards for FHE have been in development for several years but are in a late stage of the process, with publication of the standards expected in 2026. Royal Holloway academics have been among the UK experts providing comments on the draft standards as they have moved through the development process.

To support the development of the ISO/IEC FHE standards, an unofficial working group on security in FHE was established in late 2021. The group included academics and industry practitioners based in South Korea, USA, France, Spain, China, and the UK, with members both internal and external to ISO. The UK-based members of the group included participants from Royal Holloway.

The main goal of the working group was to produce a white paper on parameter selection. Selecting parameters that balance security, correctness, and performance considerations is one of the key challenges in practical FHE. The paper provided parameter sets that target particular security levels, as well as providing example parameter sets that also include the parameters relevant for correctness and performance. The paper was recently peer-reviewed and published in IACR Communications in Cryptology [1], with Erin Hales and me from Royal Holloway being among the 18 co-authors. To evaluate security of the parameters underlying FHE schemes, the paper crucially relies on the Lattice Estimator [2,3], which was developed at Royal Holloway. The white paper is now cited in the draft ISO/IEC standards, demonstrating that research from Royal Holloway is having direct impact on industry standards.

Following the publication of [1], and to build on activities at the Salt Lake City meeting, in January 2025 a HomomorphicEncryption.org Security Working Group was formed by Yuriy Polyakov, Ro Cammarota, and me. The group currently has over 20 participants, including authors of [1], participants at recent HomomorphicEncryption.org events, and others interested in homomorphic encryption security. The Security Working Group has initiated projects that will further extend the work of [1] and support the future development of practical FHE. Participants are collaborating mainly online, with the opportunity to meet in person at future HomomorphicEncryption.org meetings. I am looking forward to the next steps as FHE becomes practical and more widely deployed!

[1] J.-P. Bossuat et al. Security Guidelines for Implementing Homomorphic Encryption. *IACR Communications in Cryptology*, vol. 1, no. 4, 2025.

[2] github.com/malb/lattice-estimator

[3] M. R. Albrecht, R. Player, S. Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology.* 9 (3), pp 169–203, 2015.

STAFF PROFILE

Dr Olga Angelopoulou
Senior Lecturer, ISG and Programme Director,
MSc Information and Cyber Security

Q1 HOW DID YOU BECOME INTERESTED IN COMPUTER SCIENCE?

A I was first introduced to computers during my secondary school years in Athens, Greece. It was then that I discovered the fascinating power of programming — the ability to solve problems, build solutions, and bring ideas to life through code captivated me. Despite this early interest, I had not considered computing as a field of study. It was until a motivating conversation with my father near the end of my school years that I began to seriously reflect on my future. His encouragement during that conversation ultimately shaped my path and led me to the UK to begin my undergraduate studies in Computer Science.

Q2 HOW DID YOU BECOME INTERESTED IN INFORMATION SECURITY?

A Towards the final stages of my undergraduate studies, I found myself increasingly interested in web development and databases. I decided to build an e-commerce website for my undergraduate project and combine everything together. I started looking into secure transactions as a result of that. When the time came to decide on an area I would like to 'specialise' in postgraduate studies, information security came as a natural choice. At the time, however, academic programmes in this field were still relatively limited in the UK. I studied at the University of Glamorgan – now South Wales – specifically because their curriculum included modules on computer crime and computer forensics. This journey inspired me to continue my studies with a doctorate focusing on Identity Theft Investigations immediately after.

Q3 TELL US ABOUT YOUR RESEARCH

A My research has always been driven by my interest in digital investigations, cybercrime, and the security challenges posed by emerging technologies. Over the years, my focus has evolved in response to the increasingly complex nature of cyber threats. Currently, much of my work focuses on digital forensics in the context of incident response, as well as the forensic challenges presented by the expanding Internet of Things (IoT).

Earlier in my career, I conducted extensive research into data remanence in second-hand digital devices — a topic that remains highly relevant today. The aim was to raise awareness about the disposal and reuse of devices in the second-hand market. Through empirical analysis of hard disks, mobile phones and memory cards, my studies highlighted the significant privacy risks posed by residual data left behind by previous users.

In addition to my academic and research work, I have served as the Editor-in-Chief of Information Security Journal: A Global Perspective (Taylor & Francis) since 2018. This role has given me broad exposure to a wide spectrum of developments within the field of information security. It allows me to engage with emerging trends, and a diverse community of scholars, further enriching my own work and contributing to the advancement of the discipline as a whole.

Q4 WHAT OTHER ACTIVITIES ARE YOU INVOLVED IN?

A I am one of the founding members of Security BSides Athens, part of a global framework of community-driven annual events focused on information security. My primary role involves coordinating the call for presentations and organising the consequent reviews, while still featuring current topics from across the field. This year marks the tenth anniversary of the event, and reflecting on this milestone, I feel grateful to be part of this organising committee. Being involved has given me a meaningful professional connection to my hometown, while also offering opportunities to engage with new people, build links with the industry, and create valuable experiences for security enthusiasts to participate and learn within the cybersecurity community.

Q5 HOW DO YOU FEEL ABOUT BEING THE NEW MSC INFORMATION SECURITY COURSE LEADER?

A I am genuinely honoured to be in this role. As a student and throughout the early stages of my career, this course stood out as a programme in information security to look up to and having this role now is deeply meaningful for me. I am equally excited I joined the Information Security Group (ISG) and have the opportunity to work alongside such a respected and dedicated team. I believe the timing of my appointment is particularly significant, as the course enters an era of transformation. It is a privilege to be involved at such a significant moment and to contribute, in whatever way I can, to shaping its future.

STAFF PROFILE

Waleed Shahid
Lecturer, ISG

Q1 HOW DID YOU BECOME INTERESTED IN COMPUTER SCIENCE?

A My interest in computer science began during my school days when I got my first computer. I started learning Visual Basic, and the process of creating programs fascinated me. After a brief internal battle between chemistry and computer science, the latter clearly emerged as the winner of my interest.

Q2 HOW DID YOU BECOME INTERESTED IN INFORMATION SECURITY?

A My interest in information security began during my bachelor's in Information and Communication Systems Engineering at NUST (National University of Sciences and Technology) in Islamabad, Pakistan. While working on my final year project on frequency planning and optimization, we were introduced to a module on Data Communications and Security, taught by a Royal Holloway graduate. That course resonated deeply with me. Around the same time, my personal computer was infected with a virus, and dealing with that experience only intensified my curiosity. I was fortunate to be mentored by the same instructor. This led me to pursue an MSc in Information Security (2010–2013), where I focused on malware analysis and proposed a detection and analysis framework for Android. After working in the industry for four years, I eventually transitioned into teaching information security and completed a PhD, where I researched cyber deception systems using AI. To further deepen my expertise, I also earned several professional certifications, including CEH (Certified Ethical Hacker), CHFI (Computer Hacking Forensic Investigator), ECIH (EC-Council Certified Incident Handler), CSCU (Certified Secure Computer User), and CISSP (Certified Information Systems Security Professional).

Q3 TELL US ABOUT YOUR RESEARCH

A My research focuses on the application of artificial intelligence (AI) in cybersecurity, particularly in developing intelligent systems for threat detection, web security, and cyber deception. I use machine learning to create adaptive and proactive defence mechanisms capable of identifying and responding to evolving cyber threats in real time.

During my MSc I published my research about detection and analysis of Android malware. During my PhD, I worked on thwarting web attacks using AI-driven cyber deception techniques, developing innovative methods to mislead attackers and protect web applications. I published my research across diverse areas including web security, cyber deception, digital forensics, and threat intelligence. I have strong expertise in conducting thorough research surveys and have published in distinguished venues such as *ACM Computing Surveys*. Additionally, I have a keen interest in ethical hacking and its role in strengthening security defences.

More recently, I have developed a focus on the security of AI systems themselves, investigating vulnerabilities in machine learning models and designing robust defences against adversarial attacks. My broader interests also include cybersecurity governance and the responsible integration of AI into digital infrastructure.

STAFF PROFILE

Yiannis Tselekounis
Lecturer. ISG

Q1 HOW DID YOU BECOME INTERESTED IN COMPUTER SCIENCE?

A I've been fascinated by computers since a very early age. I remember being around eight years old, trying to read encyclopedias about the ENIAC and other early machines. When I got my first computer at the age of ten, I quickly became immersed in building and upgrading systems, experimenting with hardware, and diving into operating systems.

That naturally led me to programming -starting with languages like BASIC and PASCAL. Around the same time, I developed a curiosity for information security. I began reading tutorials and white papers, trying to understand how attacks worked and what made systems vulnerable. Even then, I was drawn not just to using technology, but to deeply understanding how it all worked under the hood.

Q2 HOW DID YOU BECOME INTERESTED IN INFORMATION SECURITY?

A I studied Mathematics and Computer Science in Athens, but my interest in information security had always been present. As an undergraduate, I took courses in Cryptography which planted the seeds early on. A pivotal moment came in 2010, during my MSc studies, when I attended a summer school in Information Security in Crete. There, I had the opportunity to attend inspiring talks on cryptography and security by leading figures in the field — including Prof. Adi Shamir, Prof. Steven Bellovin, Prof. Kenny Paterson (who was at Royal Holloway at the time), and Prof. Bart Preneel. Their talks were incredibly motivating and solidified my decision to focus on Cryptography and Information security. Following that experience, I chose to write my Master's thesis on the topic, under the supervision of Prof. Aggelos Kiayias, who later became my PhD advisor in Cryptography.

Q3 3. TELL US ABOUT YOUR RESEARCH

A My research interests lie in the fields of Cryptography and Information security, both theoretical and applied, and include themes such as the security of cryptographic protocols and primitives, hardware security, and blockchain technologies. As a researcher, my philosophy is to always tackle important theoretical problems that are well-motivated by practical applications, and to construct adversarial models that are strong enough to capture realistic and complex adversaries. Selected highlights of my contributions as a researcher include the following results: (i) the formal introduction and study of the cryptographic primitive so-called continuous group key agreement, which is a core component of the rapidly evolving area of end-to-end encrypted group messaging, (ii) the full security analysis and improvements of the highly complex messaging layer security (MLS) protocol, which is the Internet Engineering Task Force (IETF) standard for secure group messaging that will be used by billions of users, the construction of efficient private information retrieval, which is a core cryptographic protocol that has been extensively studied by the cryptographic community for almost 30 years, and (iv) the game-theoretic analysis of Bitcoins backbone protocol, that motivated the introduction of the emerging eld of distributed ledger technologies.

ALUMNUS PROFILE

Lyan Moe Kyaw

Royal Holloway's MSc in Information Security was launched in 1992 and, since that date, over 5000 students have graduated from the programme. We've lost count of the number of countries represented by our alumni, who are now scattered throughout the globe. In 2024, Lyan Moe Kyaw graduated from the programme and, we think he may have added one new country to this list. We caught up with Lyan to discuss his experience on the programme and his new career.

Q1 LYAN, TELL US A LITTLE BIT ABOUT YOUR BACKGROUND

A I'm originally from Myanmar and have a technical background, having completed a BSc in Computing. While my undergraduate degree gave me foundational knowledge, it was only after graduating that I began focusing seriously on cybersecurity.

Q2 WHY DID YOU CHOOSE TO STUDY AT ROYAL HOLLOWAY?

A I decided to pursue an MSc because I wanted to go beyond just knowing how to do things — I wanted to understand the why behind them. Certifications and hands-on experience had taught me practical skills, but I felt it was important to gain a deeper understanding of the broader context: the underlying principles, the risk landscape, the human factors, and the legal and ethical dimensions of cybersecurity.

I chose Royal Holloway specifically because of its strong reputation in the field and its long-established MSc Information Security programme. The course offered the academic depth I was looking for and the opportunity to learn from experts who are not only teaching but actively contributing to cybersecurity research. It felt like the right place to bridge the gap between technical execution and strategic thinking.

Q3 WHAT DO YOU FEEL YOU GAINED MOST FROM YOUR MSC DEGREE?

A The biggest thing I gained from my MSc at Royal Holloway was a broader perspective on what cybersecurity really involves. Before the degree, my focus was heavily on the technical side — attack and defence. But through the programme, I came to understand that security isn't just about breaking into systems or securing them — it also involves areas like compliance, privacy, and human behaviour and, more importantly, why those areas matter.

One course that particularly stood out to me was Usable Privacy and Security. It made me realise how often usability and human factors are overlooked in security solutions, even though they can make or break the effectiveness of those solutions in the real world. It was an area I had previously underestimated and it completely changed the way I think about security design.

I felt like Royal Holloway gave me the perspective that it might have taken me years to gain on my own, but I was able to see it within just one year.

Q4 YOU HAVE CHOSEN TO UNDERTAKE SEVERAL CERTIFICATIONS. WHAT ADVICE DO YOU HAVE FOR OTHERS ABOUT UNDERTAKING BOTH DEGREE PROGRAMMES AND CERTIFICATIONS?

A Apart from the bragging rights that let you talk about your cert all day and add a shiny little badge to your LinkedIn profile, a good certification (and I mean one that truly challenges you) forces you to develop strong, hands-on, practical skills. These certifications are usually tailored to specific job paths, making them a great investment if you're aiming for a particular role after graduating. Of course, whether a certification feels "hard" or "easy" is subjective — it really depends on how much you already know and your prior experience in the field.

From my own perspective, my growing interest in offensive security led me to pursue the OSCP (Offensive Security Certified Professional) certification before starting my MSc in Information Security at Royal Holloway. After completing the MSc,

I have since earned the CRTO (Certified Red Team Operator) to deepen my skills in penetration testing and red teaming.

An MSc, on the other hand, dives deeper, but at a higher level. While certifications focus on providing the skills required for specific roles, an MSc takes a broader, industry-wide approach. It helps you understand how different areas of cybersecurity — such as risk management, policy, compliance, and privacy — fit together. The industry isn't just about hackers and stopping hackers; it's about the larger, more complex issues that affect the entire ecosystem.

As I started my career, I really began to appreciate the MSc course more. There were a lot of "Aha!" moments for me, especially during conversations and meetings. Concepts and ideas that I learned during the MSc became much clearer and more relevant when I could apply them to real-world situations.

Q5 WHAT DOES YOUR CURRENT JOB INVOLVE?

A I am an Offensive Security Analyst at ThreatSpike Labs, which offers managed security services to companies of all sizes. I focus on performing penetration testing and red team engagements. Not limited to the offensive side of things, I also get exposure to the Security Operations Center (SOC).

Q6 YOU BADGE YOURSELF AS AN OFFENSIVE CYBER ENTHUSIAST: WHY DO YOU THINK OFFENSIVE CYBER IS SO IMPORTANT?

A Personally, I enjoy Capture The Flag (CTF) challenges — they're like a game where you learn something new with each round. I prefer hands-on learning, and CTFs offer the opportunity to practice various techniques without breaking the law. Over time, I've also developed a deeper appreciation for the methods used by real-world threat actors — that's why I would say I am so enthusiastic about offensive security. It's not their motives that interest me, but the technical elegance behind their techniques. Offensive security isn't just about "hacking" the client; it's about uncovering the gaps they may have missed. Ultimately, its value lies in identifying overlooked vulnerabilities, which helps organizations strengthen their defences and improve their overall security posture.

Q7 DO YOU COME ACROSS ETHICAL CHALLENGES AS PART OF YOUR WORK AND INTEREST IN OFFENSIVE CYBER?

A Definitely. As you dig around there are moments when you realize that you could really harm a website or system. Even for things that are within the scope of the engagement, I always pause before exploiting to consider whether it could break the system or cause any harm. Clients want insights, not destruction. The goal of offensive security is to identify gaps and weaknesses, not to cause disruptions.

Q8 YOU ARE POSSIBLY THE FIRST RHUL MSC INFORMATION SECURITY GRADUATE FROM MYANMAR. WHAT IS THE STATE OF CYBER SECURITY IN MYANMAR?

A Since COVID, things in Myanmar have changed drastically. Before the pandemic, the cybersecurity community was tight-knit and collaborative, where people openly shared resources and knowledge, creating a welcoming environment for anyone interested in the field. It was a unique space where you didn't need a degree to work, as long as you had the skills and contributed to the community.

This was also the time when Myanmar began developing its own digital infrastructure, and opportunities in the cybersecurity space were starting to grow.

However, with the onset of political instability, many people were forced to focus on their own survival. The situation worsened with unreliable electricity, currency devaluation, and other hardships. These factors led many to leave Myanmar, which has had a significant impact on the cybersecurity industry.

Additionally, recent events such as natural disasters, including the earthquake, have further strained the country's infrastructure, adding to the already difficult situation. Despite these challenges, there are still individuals and groups working hard to maintain Myanmar's digital assets. However, there's a decrease in new talent entering the field, as career prospects in Myanmar have become more limited compared to other countries. Furthermore, with the government tightening its controls, using VPNs has become illegal, and various social media platforms are censored, making it even harder for citizens to access information and communicate freely.

Q9 WHAT DO YOU THINK ARE THE UPCOMING CHALLENGES IN CYBER SECURITY THAT MOST WORRY YOU?

A One of the challenges that worries me in the UK cybersecurity industry is the high entry-level requirements. As the job market becomes more competitive, it has become increasingly difficult for people to break into cybersecurity, especially for international students who are passionate about the field. For many, the biggest hurdle isn't the technical skills or knowledge, but the requirement for security clearances and a visa. While it's understandable from a security perspective, this requirement can be a significant barrier for those just starting out in their careers.

Q10 WHAT ARE THE MOST EXCITING CYBERSECURITY DEVELOPMENTS THAT YOU ARE MOST EAGERLY ANTICIPATING?

A One of the most exciting developments in cybersecurity right now is the increasing accessibility of resources and platforms like Hack The Box (HTB) and TryHackMe. Ethical hacking and offensive security used to be behind paywalls or accessible only to those with significant resources or connection. Now, these platforms are making it easier than ever for people to break into the field. With affordable learning materials and hands-on challenges, cybersecurity has become more inclusive, allowing a broader range of individuals to get involved and gain valuable skills.

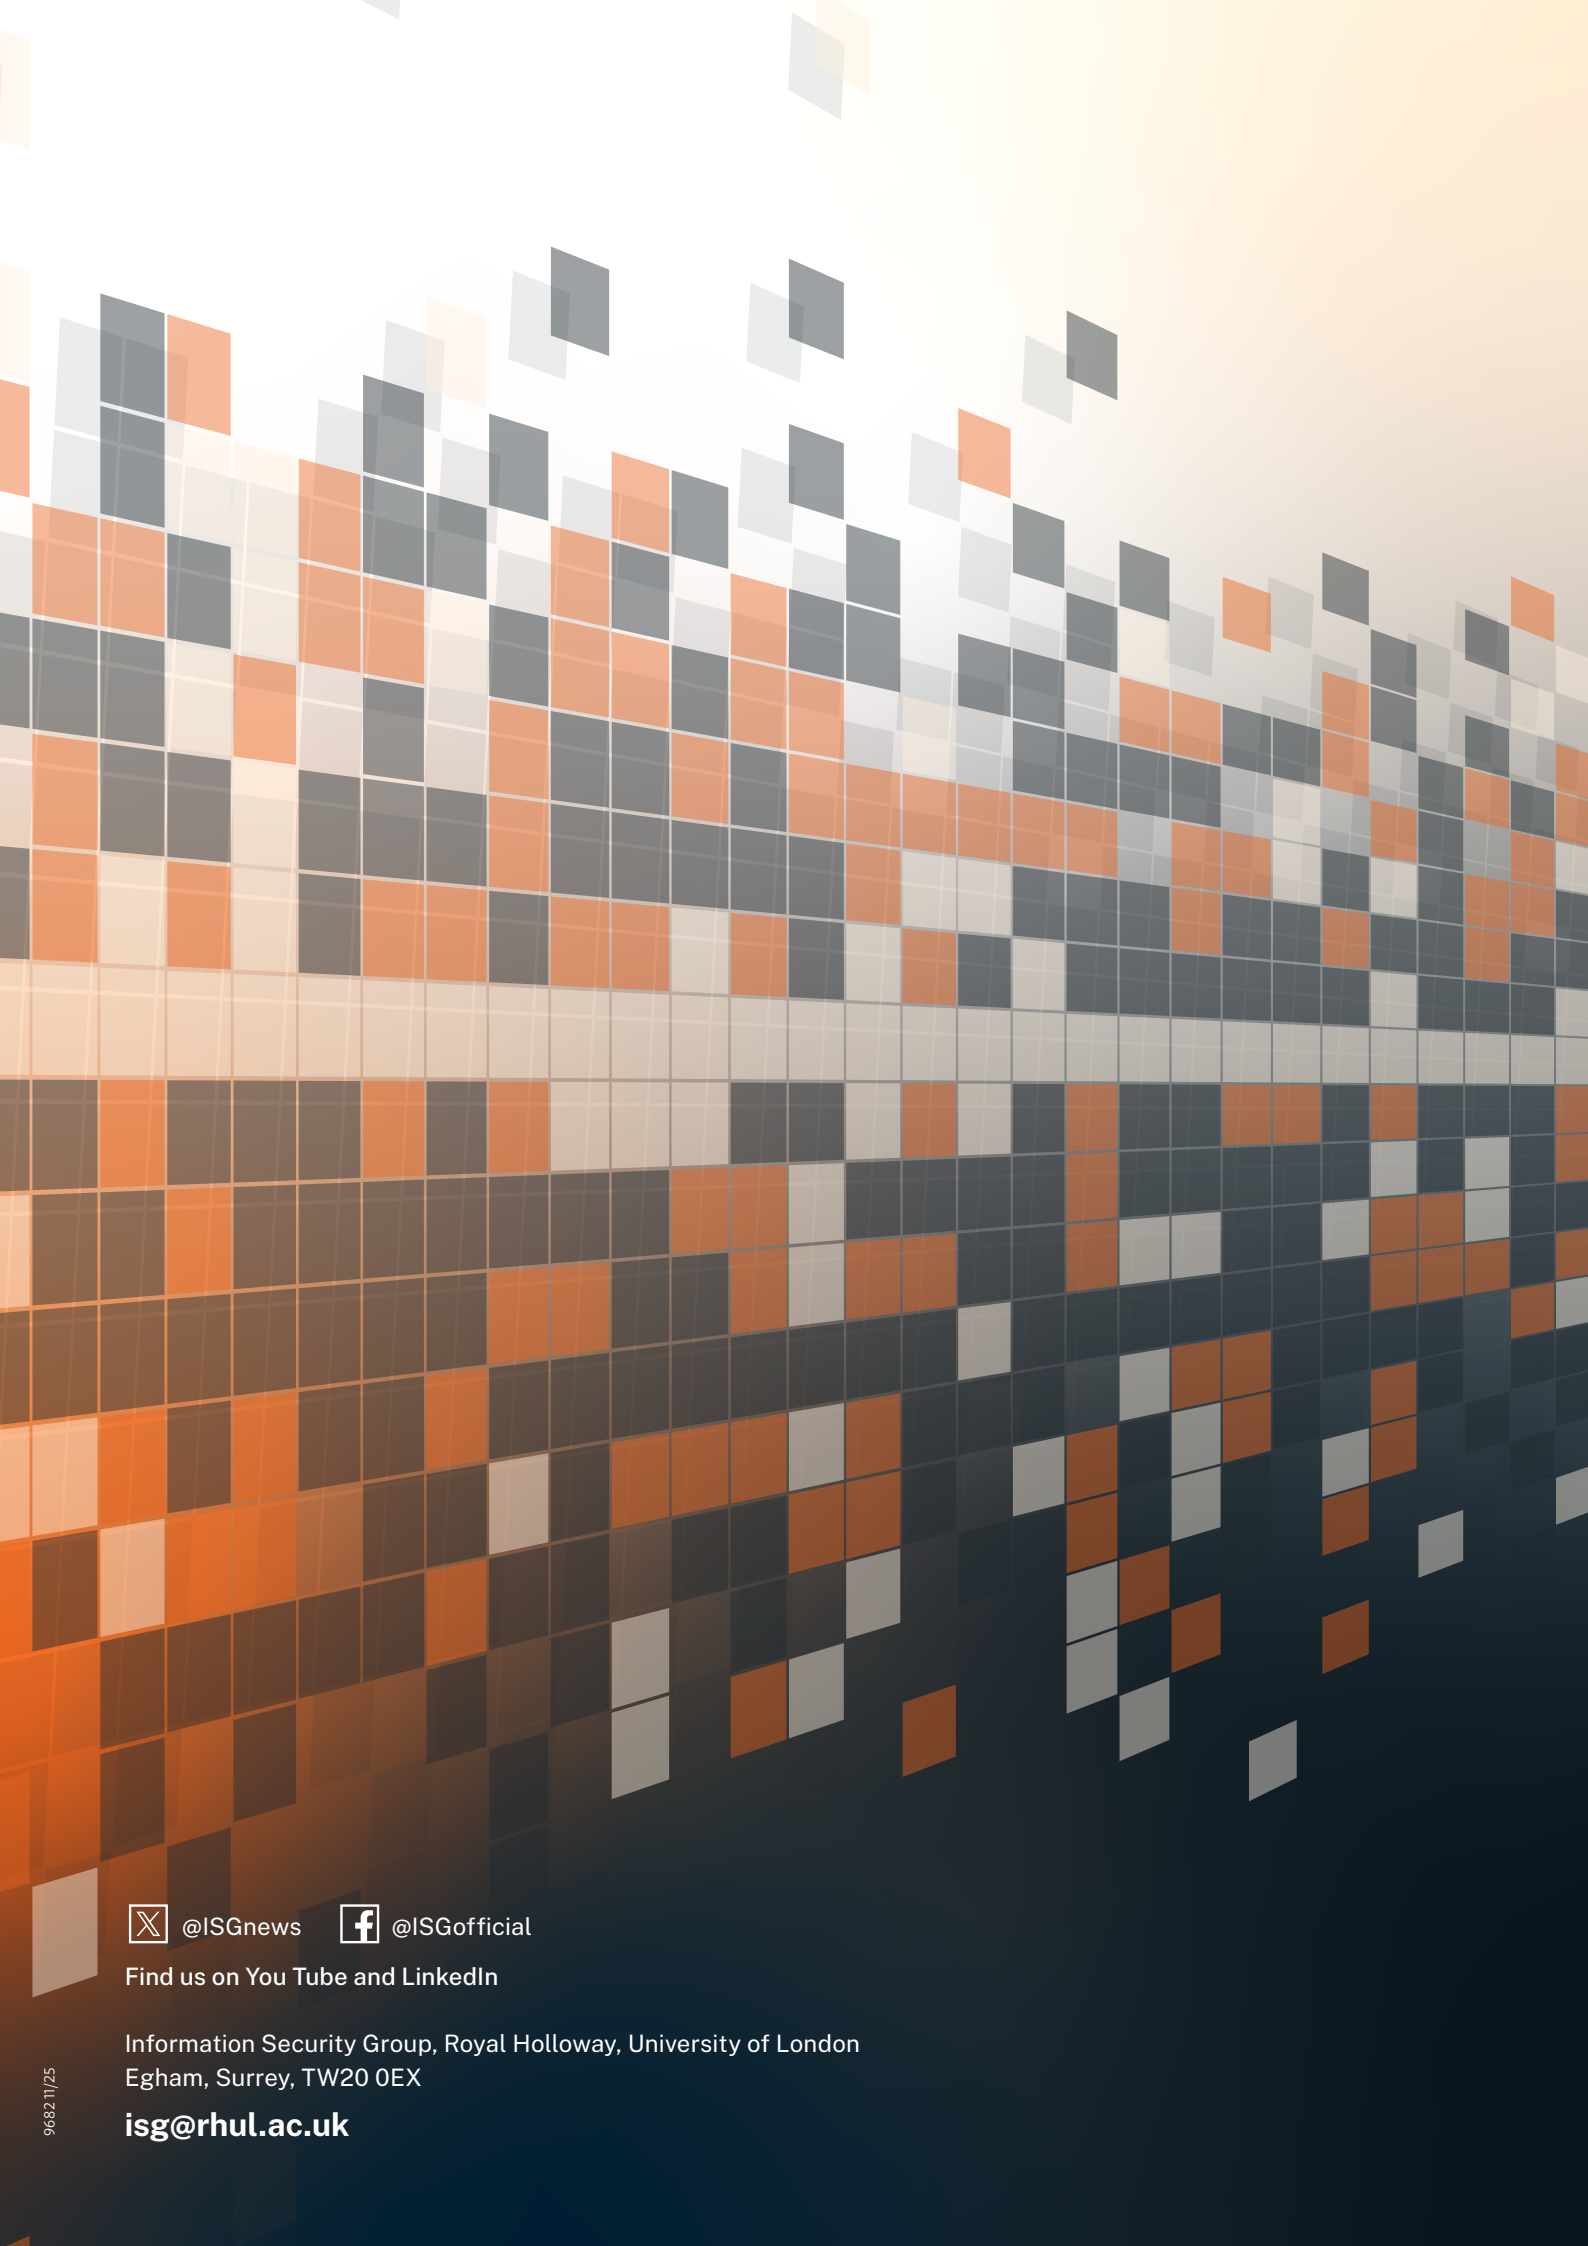

@ISGnews @ISGofficial

Find us on You Tube and LinkedIn

Information Security Group, Royal Holloway, University of London
Egham, Surrey, TW20 0EX

isg@rhul.ac.uk

9682 11/25