Informatics PhD projects at King's College London, AY 24-25 — Cybersecurity

The PhD project proposals listed below will be considered for 2024/25 studentships available in the Department of Informatics to start 1 October 2024 or later during the 2024/25 academic year. Please note that this list is not inclusive and potential applicants can alternatively identify and contact appropriate supervisors outlining their background and research interests or proposing their own project ideas.

The PhD projects are listed in two groups. In the first group are the projects with allocated studentships: each project in this group has one allocated studentship. The remaining studentships will be considered for the projects listed in the second group. The number of those remaining studentships is smaller than the number of the projects in the second group. The allocation of studentships will be based on the merits of individual applications. Applications for PhD studies in the Department of Informatics, for all listed projects as well as for other projects agreed with supervisors, are also welcome from students applying for other funding (within other studentship schemes) and from self-funded students. See also this list of funding opportunities available at King's for post-graduate research in Computer Science.

- Scholarship Allocated
- Scholarship Not Allocated

Scholarship Allocated

(Back to Top)

- Inclusive Privacy & Security
- Mitigating Technology-Enabled Harms (Privacy & Security)
- Towards Practical Zero-Knowledge Proofs from Lattices
- Lawful Storage and Processing of Personal Data
- Exploring Alternative Permissions Mechanisms for Personal Data Sharing

Inclusive Privacy & Security

Supervisor: Kovila Coopamootoo

Areas: Cybersecurity, Human-centred computing

(Back to Scholarship Allocated)

Project Description

Privacy, security, digital safety are often designed for majority or WEIRD populations, thereby not appropriately serving all user communities (minority, with specific needs, or life situations), and risk discriminating access and appropriate engagement, while also creating concerns for them. The PhD project will focus on a (type of) privacy, security or digital safety context and seek better privacy/security design with and for at-risk communities, such as addressing:

- Privacy for women and queer communities for e.g. in online platforms
- Authentication for the visually impaired
- Digital safety for socio-economically deprived communities (where privacy/security is often traded-off for more pressing life needs)
- Privacy and digital safety programmes for adult digital starters (often from deprived or minority populations, or experiencing the digital generational divide)
- Privacy and security of refugees and migrants

Proposed PhD projects will involve a mix of qualitative, participatory research and community-centred research. The projects can involve working with stakeholders such as charities, NGOs that support vulnerable and marginalised groups and victim-survivors, as well as digital safety advocacy.

References

- Usenix Security Symposium 2023: "Un-Equal Online Safety?" A Gender Analysis of Security and Privacy Protection Advice and Behaviour Patterns. By Kovila P.L Coopamootoo & Magdalene Ng
- Usenix Security Symposium 2022: "I feel invaded, annoyed, anxious and I may protect myself": Individuals' Feelings about Online Tracking and their Protective Behaviour across Gender and Country. By Kovila P.L Coopamootoo, Maryam Merhnezhad, Ehsan Toreini.
- IEEE Euro S&PW 2023: What we do in the shadows: How does experiencing cybercrime affect response actions and protective practices. By Magdalene Ng, Maria Bada, Kovila P.L Coopamootoo.

- ACM CCS (Computer & Communications Security) 2020: Usage Patterns of Privacy-Enhancing Technologies. By Kovila P.L Coopamootoo.
- Workshop on Privacy in the Electronic Society (WPES) 2014: Sensible privacy: how we can protect domestic violence survivors without facilitating misuse. By Arief, Coopamootoo, Emma, van Moorsel.

Mitigating Technology-Enabled Harms (Privacy & Security)

Supervisor: Kovila Coopamootoo

Areas: Cybersecurity, Human-centred computing

(Back to Scholarship Allocated)

Project Description

There are growing examples of technology originally designed to enhance quality of life being mis-used to facilitate abuse, online harms and in-security. The PhD project will investigate a particular context of technology-enabled harm in intersection with the user characteristic or life event that amplifies the chances of being targeted or in experiencing online harms (e.g. gender, age, life style preference, health condition, breakdown of relationship), such as:

- Online platforms being used for hate and harassment of women, queer communities, or race minorities. This can address how platforms (are required to) mitigate harm and content moderation.
- Smart homes enabling intimate partner abuse
- Digital health and wellbeing technology involving excessive tracking and privacy invasion (e.g. Femtech, sports apps)
- Online platform's poor content moderation and lack of age-appropriate controls that threaten children's online safety
- Content sharing platform used for non-consensual intimate image distribution

Proposed PhD projects will aim to understand the lived experiences of users, make recommendations and prototype for safer/harm-mitigating experiences. Specifically, the project will look into the behavioural dynamics of interacting with such technologies (or human-computer interaction — HCI) and the design aspects, via suitable user-studies such as interviews, surveys, participatory workshops, or ethnography research methods. More information can be found here: https://kovilacoops.github.io/opportunities/

References

- Usenix Security Symposium 2023: "Un-Equal Online Safety?" A Gender Analysis of Security and Privacy Protection Advice and Behaviour Patterns. By Kovila P.L Coopamootoo & Magdalene Ng
- Usenix Security Symposium 2022: "I feel invaded, annoyed, anxious and I may protect myself": Individuals' Feelings about Online Tracking and their Protective Behaviour across Gender and Country. By Kovila P.L Coopamootoo, Maryam Merhnezhad, Ehsan Toreini.

- IEEE Euro S&PW 2023: What we do in the shadows: How does experiencing cybercrime affect response actions and protective practices. By Magdalene Ng, Maria Bada, Kovila P.L Coopamootoo.
- ACM CCS (Computer & Communications Security) 2020: Usage Patterns of Privacy-Enhancing Technologies. By Kovila P.L Coopamootoo.
- Workshop on Privacy in the Electronic Society (WPES) 2014: Sensible privacy: how we can protect domestic violence survivors without facilitating misuse. By Arief, Coopamootoo, Emma, van Moorsel.

Towards Practical Zero-Knowledge Proofs from Lattices

Supervisor: Ngoc Khanh Nguyen

Areas: Cybersecurity

(Back to Scholarship Allocated)

Project Description

Zero-knowledge proofs allow a party to convey that a given statement is true without leaking any secret information. These proofs form the foundations of many complex privacy-oriented protocols, such as electronic voting, verifiable computation, and blockchain. In such applications, it is essential to be able to prove in zero-knowledge that one knows how to open a cryptographic commitment and to prove that the committed values have particular properties or satisfy certain relations. Due to the significant progress in building quantum computers, there has been a tremendous interest in constructing such protocols from quantum-safe assumptions. This is highly evidenced by the US NIST Post-Quantum Competition [1] for standardising quantum-safe key encapsulation mechanisms and digital signatures, where the vast majority of the selected schemes rely on hardness of lattice-based assumptions. Hence, lattices are a natural candidate for building more advanced quantum-safe applications such as zero-knowledge proofs. This project will look at novel approaches to design practically efficient zero-knowledge proofs from lattices, where the particular focus will be on succinct communication size, efficient time and space complexities of the protocols. To this end, a strong mathematical background (algebra, number theory) is required. Solid cryptography background is preferred. Some knowledge and experience in coding (e.g., Python, C/C++, SageMath) is also a plus.

References

[1] https://csrc.nist.gov/projects/post-quantum-cryptography

Lawful Storage and Processing of Personal Data

Supervisor: Supreeth Shastri

Areas: Cybersecurity

(Back to Scholarship Allocated)

Project Description

Our society is in the midst of granting a new right to people: privacy and protection of personal data. For example, in 2018, GDPR accorded this right to 500+ million Europeans, and required compliance from all organizations operating on that personal data. Since then, several governments including UK, California, and China have instituted equivalent regulations covering a billion plus people. These regulations require new types of control- and data-plane operations to be executed on systems that store and process personal data. For instance, GDPR article-21 allows people to place a restriction, at any time, on their data being used for any given purpose. Yet, widely-used data analytics platforms such as Hadoop offer little or no support for the mechanisms and policies needed to comply with data rights and regulations. Indeed, they continue to be designed and implemented on the now defunct notion that data is an inert entity that could be stored and used in any way at any time by any processing system. This lack of support at the platform level not only passes the burden of compliance to applications but also makes them inefficient and error-prone. This is evident in that, on average, a GDPR penalty is issued once every 1.4 days. The goal of this proposal is to build system software that enables lawful storage and processing of personal data. We will introduce new techniques, abstractions, and tools in the Hadoop ecosystem to support the lawful computing requirements of GDPR. For more details and background, explore our research at: https://lawfulcomputing.org/

Exploring Alternative Permissions Mechanisms for Personal Data Sharing

Supervisor: William Seymour (co-supervisor TBC)

Areas: Cybersecurity, Human-centred computing

(Back to Scholarship Allocated)

Project Description

Installing a new app or unpacking a new smart home device almost always involves granting a range of permissions about how those products can use your personal data. Browsing the web similarly requires making a vast number of decisions about cookies and other tracking technologies. But giving consent in this way often doesn't really feel like consenting at all, and prior work in this area has highlighted the implausibility that the mechanisms we use fulfil the requirements of what is understood as informed consent. This is particularly apparent in contexts where multiple people are using the same device. "Bystander" partners, children, family, and housemates are often left out of the installation process, even though information about them is also being collected. Instead, one member of the household unilaterally makes decisions for everybody else. This echoes the model adopted by data protection regulation, which (broadly) controls the use of data on those within the home by organisations outside of it, avoiding the topic of privacy between cohabitants. This PhD project will build on my and other existing work in this area to explore alternatives to the status quo described above, including the potential for novel automated and group privacy decision making mechanisms. It will do so across a range of different interaction modalities (such as graphical, voice, and touch interfaces) and social contexts. You'll be able to shape the direction of the project to fit your own interests. Key skills and research methodologies will include some of the following, tailored to your preferences: -Qualitative interviews and/or focus groups with users and designers -Quantitative surveys with users -Lightweight prototyping and speculative design (potentially including programming prototypes) -More creative methods such as home deployments and creating interactive experiences While you don't need to have experience with any of these techniques before starting, you do need to have a demonstrable interest/experience in online privacy and human computer interaction.

References

William Seymour, Mark Cote, and Jose Such. 2023. Legal Obligation and Ethical Best Practice: Towards Meaningful Verbal Consent for Voice Assistants. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 166, 1—16.

https://doi.org/10.1145/3544548.3580967
William Seymour, Reuben Binns, Petr Slovak, Max Van Kleek, and Nigel Shadbolt. 2020. Strangers in the Room: Unpacking Perceptions of 'Smartness' and Related Ethical Concerns in the Home. In Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20). Association for Computing Machinery, New York, NY, USA, 841—854. https://doi.org/10.1145/3357236.3395501
Bettina Nissen, Victoria Neumann, Mateusz Mikusz, Rory Gianni, Sarah Clinch, Chris Speed, and Nigel Davies. 2019. Should I Agree? Delegating Consent Decisions Beyond the Individual. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19). Association for Computing Machinery, New York, NY, USA, Paper 515, 1—13. https://doi.org/10.1145/3290605.3300745

Scholarship Not Allocated

(Back to Top)

- Privacy in the Internet of Things
- Designing with and for at-risk populations
- Female mobile health apps: understanding user needs and concerns
- Privacy, safety, and trust in emerging AI technologies
- Understanding Software Security: Unveiling Vulnerabilities through Binary-based Testing Strategies
- Enhancing Cybersecurity Education through Gamification: using AI and Immersive Technology
- String Sanitization with Applications to Internet of Things Data

Privacy in the Internet of Things

Supervisor: Maribel Fernandez

Areas: Cybersecurity

(Back to Scholarship Not Allocated)

Project Description

Data Collection policies are used to restrict the kind of data transmitted by devices in the Internet of Things (e.g., health trackers, smart electricity meters, etc.) according to the privacy preferences of the user. The goal of this project is to develop cloud/IoT architectures with integrated data collection and data sharing models, to allow users to specify their own policies and trade data for services. For this, new data collection and data sharing models will have to be developed, with appropriate user interfaces, policy languages, and policy enforcement mechanisms. An important aspect of the project is the development of policy recommendation systems that can suggest/create policies based on user profiles, making privacy an integral part of the system (according to the privacy-by-design" IoT paradigm).

References

A Privacy-Preserving Architecture and Data-Sharing Model for Cloud-IoT Applications, Maribel Fernandez; Jenjira Jaimunk; Bhavani Thuraisingham
IEEE Transactions on Dependable and Secure Computing, vol. 20, no. 4, pp. 3495-3507, 1 July-Aug. 2023, doi: 10.1109/TDSC.2022.3204720.

Designing with and for at-risk populations

Supervisor: Ruba Abu-Salma

Areas: Cybersecurity, Human-centred computing, Social computing

(Back to Scholarship Not Allocated)

Project Description

Recently, researchers have started to realize that designing digital technologies for one population in mind risks ignoring the security, privacy, and safety needs of, as well as creating concerns for, other populations. While research on understanding the needs and concerns of at-risk populations like older adults or migrant domestic workers, is evolving with the US dominating the field, the research is still in its infancy without a clear understanding of the scale and impact of technologies on such populations. The objective of this project is to empirically study the security, privacy, and safety needs — based on lived experiences — of at-risk populations like the poor, the young, and the disabled using qualitative (eg, interviews, focus groups, participatory design workshops) and quantitative (eg, surveys) methods. The empirical evidence gleaned from these studies will inform the design of current and future technologies.

References

- https://arxiv.org/pdf/2112.07047.pdf.
- https://arxiv.org/pdf/2206.15037.pdf.
- https://www.usenix.org/conference/soups2022/presentation/bernd.
- https://www.usenix.org/conference/foci20/presentation/bernd.
- https://www.usenix.org/conference/usenixsecurity22/presentation/slupska-vulnerability.
- https://petsymposium.org/popets/2024/popets-2024-0009.pdf.

Female mobile health apps: understanding user needs and concerns

Supervisor: Ruba Abu-Salma

Areas: Cybersecurity, Social computing, Human-centred computing

(Back to Scholarship Not Allocated)

Project Description

Mobile apps which support women's health have developed rapidly. However, the ubiquity of these apps has advanced the practice of intimate surveillance and sensitive data collection. While the overturning of Roe v. Wade has prompted reflection on the privacy and safety implications of female mobile health apps, the needs and concerns of the users of these apps are yet to be explored. The project aims to generate empirical evidence of these needs and concerns, as well as improve the design of female health apps, with a focus on user privacy and safety.

References

- https://petsymposium.org/popets/2020/popets-2020-0083.pdf.
- https://arxiv.org/pdf/2306.05956.pdf.
- https://dl.acm.org/doi/pdf/10.1145/3411764.3445132.

Privacy, safety, and trust in emerging AI technologies

Supervisor: Ruba Abu-Salma

Areas: Cybersecurity, Social computing, Human-centred computing

(Back to Scholarship Not Allocated)

Project Description

Research examining the impact of AI on society is concerned with the need to develop AI technologies that are ethical, human-centered, responsible, safe, and trustworthy. However, the body of human-subjects research investigating how people may define those goals is not large, and little of it compares how views vary across countries and populations; eg, with respect to demographics, social norms, technological experiences, etc. The proposed research aims to expand our empirical understanding of people's experiences with and views on AI technologies, including GenAI tools, across different countries and populations. In particular, we are interested in how privacy and safety concerns affect people's trust in these technologies.

References

- https://arxiv.org/pdf/2302.05284.pdf.
- https://arxiv.org/pdf/2305.06415.pdf.
- https://arxiv.org/pdf/2302.08157.pdf.
- https://arxiv.org/pdf/2310.06778.pdf.
- https://dl.acm.org/doi/pdf/10.1145/3613249.
- https://www.usenix.org/system/files/soups2023-kelley.pdf.

Understanding Software Security: Unveiling Vulnerabilities through Binary-based Testing Strategies

Supervisor: Dr Karine Even-Mendoza, Dr Hector Menendez Benito

Areas: Machine Learning (ML), Foundations of computing, Cybersecurity, Systems (SE, programming, autonomous systems, robotics, ...), Computing Applications

(Back to Scholarship Not Allocated)

Project Description

Software ecosystems rely on the way operating systems distribute resources. By creating the address space, the process space, the threads and the security tokens of the running program, the system provides an execution context that changes depending on the kernel version or even the compiler used to execute the programs. The integration of a program into a systematic environment that evolves depending on kernel and compilation version might not imply security vulnerabilities, but in the presence of crashes, the exploitability of the system will directly depend on how it deals with resources, as obfuscations proved [2]. Under these conditions, there are a few strategies that can provide some light on ways to identify these vulnerabilities. The first one is to employ semantic equivalent transformations to the software and study the behavioural changes in the system. The second is to study the decompilation of the final PE or EFL file and investigate how it changes under different compilation options. The third is to employ various testing strategies, such as differential testing, to analyse how the environment is changing the execution, often tracked through profiling strategies. These three strategies will define the three parts of the thesis. Part 1: Process Resources. The student will work by extending the previous work on the security of obfuscations [2]. The extension will focus on the way the heap and the stack are affected in terms of the address space and the managed resources. With this information, the student will better understand the exploitability of specific parts of the system and work on potential mitigations that can support the system's security. Part 2: Compiler's configurations. Compilers optimise code by adding transformations that reduce the way the process collects and manipulates resources. It is also affected by the scheduler. The student will catalogue the effect of optimisations in software, especially bugs, and how they change their nature and exploitability when the system is more vulnerable. Based on these principles, the student will extrapolate the previous knowledge on exploitability to the compilers' context. Part 3: Testing improvements. Based on the previous discoveries about how the system interacts with processes and how compilers and contexts change this, this last part of the thesis focuses on changing the ways testing is applied in the

context of vulnerabilities with the aim of making it more focused to unmask the risk that the context and the compiler can associate with the execution of the files.

References

[1] A. Dakhama, K. Even-Mendoza, W.B. Langdon, H. Menendez, and J. Petke (2023). SearchGEM5: Towards Reliable gem5 with Search Based Software Testing and Large Language Models. Symposium on Search Based Software Engineering (SSBSE). https://tinyurl.com/2u2aeb4r

[2] H. D. Menendez and G. Suarez-Tangil. 2022. ObfSec: Measuring the security of obfuscations from a testing perspective. Expert Syst. Appl. 210, C (Dec 2022). https://doi.org/10.1016/j.eswa.2022.118298

[3] K. Even-Mendoza, C. Cadar, and A. F. Donaldson. 2022. Csmithedge: more effective compiler testing by handling undefined behaviour less conservatively. Empirical Softw. Engg., 27, 6, (Nov. 2022), 35 pages. doi: 10.1007/s10664-022-10146-1.

[4] X. Hou, Y. Zhao, Y. Liu, Z. Yang, K. Wang, L. Li, X. Luo, D. Lo, J. Grundy, H. Wang. Large language models for software engineering: A systematic literature review. arXiv preprint arXiv:2308.10620. 2023 Aug 21

[5] V. Le, M. Afshari, and Z. Su. 2014. Compiler validation via equivalence modulo inputs. In PLDI '14. ACM, New York, NY, USA, 216—226. https://doi.org/10.1145/2594291.2594334

[6] E. T. Barr, M. Harman, P. McMinn, M. Shahbaz and S. Yoo, "The Oracle Problem in Software Testing: A Survey," in IEEE Transactions on Software Engineering, vol. 41, no. 5, pp. 507-525, 1 May 2015, doi: 10.1109/TSE.2014.2372785.

[7] X. Yang, Y. Chen, E. Eide, and J. Regehr. 2011. Finding and understanding bugs in c compilers. In Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '11). Association for Computing Machinery, San Jose, California, USA, 283—294. isbn: 9781450306638. doi: 10.1145/199349 8.1993532.

[8] J. Regehr. Finding Bugs in C and C++ Compilers using YARPGen. SIGPLAN. PL Perspectives. https://blog.sigplan.org/2021/01/14/finding-bugs-in-c-and-c-compilers-using-yarpgen/

[9] AFL Michal Zalewski, "Technical "whitepaper" for afl-fuzz," https://lcamtuf.coredump.cx/afl/technical_details.txt

Enhancing Cybersecurity Education through Gamification: using AI and Immersive Technology

Supervisor: Tasmina Islam

Areas: Cybersecurity, Education, Computing Applications

(Back to Scholarship Not Allocated)

Project Description

People are spending more time online due to the increasing digitisation of society. This also means the security measures need to be stronger and awareness of the cybersecurity risks, and implications is crucial. However, traditional educational approaches often struggle to engage and effectively educate learners especially young learners in cybersecurity practices, often due to lack of personalisation. This project aims to explore the integration of Artificial Intelligence (AI) and Virtual Reality (VR) technologies into educational games as a means to enhance cybersecurity awareness and skills among children, catering different types of learners. This project aims to explore the design and development of immersive, interactive, and age-appropriate AI-driven VR/XR games in an engaging and accessible manner for educating children about cybersecurity fundamentals, risks and consequences of the risks being materialised.

References

Islam, T & Zou, Y 2023, ChildSecurity: A Web-based Game to Raise Awareness of Cybersecurity and Privacy in Children. in Cybersecurity Challenges in the Age of AI, Space Communications and Cyborgs - Proceedings of the 15th International Conference on Global Security, Safety and Sustainability, London, October 2023. Springer.

String Sanitization with Applications to Internet of Things Data

Supervisor: Dr Grigorios Loukides, Professor Luca Vigano

Areas: Cybersecurity, Data science, Foundations of computing, Computing Applications

(Back to Scholarship Not Allocated)

Project Description

The overall aim of the project is to develop and evaluate a robust and efficient approach that allows organisations and businesses to protect the privacy of data represented as strings. The project will consider the protection of aggregated data (event sequences), as well as string databases, and it will also address the interrelated issues of usefulness, security, and scalability. It aims to develop a methodology (model, algorithms, protocols) for sanitising (i.e., transforming) data that is: (I) privacy- preserving, by designing and applying a privacy model along with algorithms for sanitising string data. (II) Utility-preserving, by designing measures and tools for quantifying the level of usefulness of data that must be traded-off for achieving privacy. (III) Secure and scalable, by designing efficient protocols that allow multiple parties to protect their data securely and jointly. The methodology will be evaluated on data from the Internet of Things (IoT) domain.

References

Bernardini et al. Hide and Mine in Strings: Hardness, Algorithms, and Experiments. IEEE TKDE 2023. https://ieeexplore.ieee.org/document/9732522