Anh V. Vu

Cambridge Computer Laboratory, William Gates Building, 15 JJ Thomson Ave, Cambridge, United Kingdom, CB3 0FD
(+44) 746 351 1312 | anhvvcs | anh.vu@cst.cam.ac.uk | www.cst.cam.ac.uk/people/vv301 | anhvvcs | anhvcs

My research offers timely empirical measurements to explore cyberspace and its social impact at scale, focusing on underground subcultures fostering online crime and harm. The resulting insights help better understand online threats and inform policy decisions for safety & security.

Education

University of Cambridge — Cambridge, UK
Doctor of Philosophy in Computer Science — Jan 2022 – Dec 2024
- Supervisor: Prof. Alice Hutchings · Advisor: Prof. Ross Anderson
- Thesis: Online Crime and Harms Following Externalities

Japan Advanced Institute of Science and Technology — Ishikawa, Japan
Master of Science in Information Science — Oct 2017 – Dec 2018
- Supervisor: Prof. Mizuhito Ogawa
- Thesis: Formal Semantics Extraction from Natural Language Specifications for ARM
- Examiners: Prof. Mizuhito Ogawa, Prof. Kazuhiro Ogata, Prof. Keita Yokoyama, Prof. Minh Le Nguyen

Vietnam National University — Hanoi, Vietnam
Bachelor of Science in Information Technology — Sep 2012 – Jun 2016
- Supervisor: Prof. Xuan Hieu Phan
- Thesis: User Behaviour Analysis and Personalisation in a Vietnamese Medical Information System
- Honours Programme · High Distinction (top 10/473) · GPA: 3.66/4.0

Work Experience

Delft University of Technology — Delft, Netherlands
Visiting PhD Student — 11 Jun 2024 - 11 Jul 2024
- **Advisor**: Dr. Rolf van Wegberg / **Skills**: Financial Cybercrime, Crypto Mixers, Anonymity Network, Python.
- Will be collaborating on a research project co-funded by the Dutch Law Enforcement on the facilitators of financial cybercrime.
- Will analyse blockchain transactions and seized ground-truth databases of crypto mixers to understand the dynamics of their ecosystem.

University of Cambridge — Cambridge, UK
Research Assistant — 21 Oct 2019 - Present
- **Advisors**: Dr. Richard Clayton, Prof. Alice Hutchings / **Skills**: Security, Cybercrime, Online Abuses and Harms, Statistics, Databases, Python.
- Collected large-scale cybercrime and extremist datasets, then analysed the longitudinal dynamics of their underground subcultures.
- Maintained a data-sharing platform, allowing over 300 researchers worldwide to access our cybercrime and extremist datasets.

National University of Singapore — Kent Ridge, Singapore
Research Intern — 6 Jan 2019 - 25 Jul 2019
- **Advisor**: Prof. Min Suk Kang / **Skills**: Blockchain, Peer-to-Peer Network Security, Partitioning Attacks, Python, C/C++.
- Conducted experiments simulating adversarial scenarios on our newly discovered partitioning attack against the Bitcoin P2P network.
- Co-authored a paper accepted at IEEE S&P'20, one of the most prestigious security conferences. The research was later featured on CoinDesk.

Vietnam National University — Hanoi, Vietnam
Teaching Assistant — 1 Oct 2016 - 30 Sep 2017
- **Advisor**: Dr. Hieu Dinh Vo / **Skills**: Software Engineering, Data Scraping, Databases, Android, Java.
- Taught undergraduate modules, including Fundamentals of Informatics and Object-Oriented Programming.
- Collected large-scale datasets for VCGate, a Vietnamese bibliographic database designed for measuring and analysing scholarly literature.

Publications

Yet Another Diminishing Spark: Low-level Cyberattacks in the Israel-Gaza Conflict — Under Review
Anh V. Vu, Alice Hutchings, Ross Anderson — 2024
- Under Review · Briefing · Website — PDF
- Press coverage: Computer Weekly · Fast Company · Infosecurity · LBT

Peer-reviewed Conferences

No Easy Way Out: the Effectiveness of Deplatforming Forums to Suppress Hate and Harassment — San Francisco, USA
Anh V. Vu, Alice Hutchings, Ross Anderson — 20–23 May 2024
- **S&P'24** – IEEE Symposium on Security and Privacy · Acceptance Rate 17.8% · Briefing · CyCon — Rank A* PDF
- Press coverage: The Register · LBT

Getting Bored of Cyberwar: the Role of Low-level Cybercrime Actors in the Russia-Ukraine Conflict — Sentosa, Singapore
Anh V. Vu, Daniel R. Thomas, Ben Collier, Alice Hutchings, Richard Clayton, Ross Anderson — 13–17 May 2024
- **WWW'24** – ACM World Wide Web Conference · Acceptance Rate 20.2% · CyCon · Website — Rank A* PDF
- Press coverage: New Scientist · Associated Press · SC Magazine · The Record · LBT

Turning Up the Dial: the Evolution of a Cybercrime Market Through Set-up, Stable, and Covid-19 Eras — Pittsburgh, USA
Anh V. Vu, Jack Hughes, Ildiko Pete, Ben Collier, Yi Ting Chua, Ilia Shumailov, Alice Hutchings — 27–29 Oct 2020
- **IMC'20** – ACM Internet Measurement Conference · Acceptance Rate 24.5% · Website — Rank A PDF
- Press coverage: Hacker News · Cambridge Research · LBT

A Stealthier Partitioning Attack against Bitcoin Peer-to-Peer Network — San Francisco, USA
Muoi Tran, Inho Choi, Gi Jun Moon, Anh V. Vu, Min Suk Kang — 18–20 May 2020
- **S&P'20** – IEEE Symposium on Security and Privacy · Acceptance Rate 12.4% · Website — Rank A* PDF
- Press coverage: CoinDesk

Formal Semantics Extraction from Natural Language Specifications for ARM — Porto, Portugal
Anh V. Vu, Mizuhito Ogawa — 07–11 Oct 2019
- **FM'19** – International Symposium on Formal Methods · Acceptance Rate 30.0% · Website — Rank A PDF

Peer-reviewed Workshops

ExtremeBB: A Database for Large-Scale Research into Online Hate, Harassment, and Extremism — Toronto, Canada
Anh V. Vu, Lydia Wilson, Yi Ting Chua, Ilia Shumailov, Ross Anderson — 09–14 Jul 2023
- **WOAH@ACL'23** – ACL Workshop on Online Abuse and Harms · Website — PDF

PostCog: A Tool for Interdisciplinary Research into Underground Forums at Scale — Genoa, Italy
Ildiko Pete, Jack Hughes, Andrew Caines, Anh V. Vu, H. Gupta, Alice Hutchings, Ross Anderson, Paula Buttery — 06–10 Jun 2022
- **WACCO@EuroS&P'22** – IEEE EuroS&P Workshop on Attackers and Cyber-Crime Operations · Website — PDF

Book Chapters

Identifying and Collecting Public Domain Data for Tracking Cybercrime and Online Extremism — World Scientific
Lydia Wilson, Anh V. Vu, Ildiko Pete, Yi Ting Chua — Jun 2024
- Chapter in: Open Source Investigations in the Age of Google — PDF
- Press coverage: Center for Strategic and International Studies (CSIS)

Teaching & Supervision

Final-year Undergraduate Projects @ Cambridge: Andrei-cosmin Moroca (2023-24)

Algorithms @ Cambridge: Sutton Trust Summer School 2022 (12 students), 2023 (10 students)

Databases @ Cambridge: Michaelmas 2022-23 (15 students)

Software and Security Engineering @ Cambridge: Easter 2021-22 (14 students), Easter 2022-23 (23 students)

Object-Oriented Programming @ Cambridge: Michaelmas 2021-22 (10 students), Michaelmas 2022-23 (18 students)

Fundamental of Informatics, Object-Oriented Programming @ Vietnam National University: 2016-17

Professional Activities

Reviewers: IMC'22 (shadow), WWW'24

External Reviewers: EuroS&PW'22, USENIX Security'23, USENIX Security'24

Research Talks and Presentations:

- Yet Another Diminishing Spark: Low-level Cyberattacks in the Israel-Gaza Conflict
  @ **CCC'24** · Cambridge Cybercrime Conference · University of Cambridge, UK · **in-person** (upcoming) — 10 Jun 2024
- No Easy Way Out: the Efficacy of Deplatforming an Extremist Forum to Suppress Hate and Harassment
  @ **SHB'24** · Workshop on Security and Human Behaviour · Harvard University, USA · **in-person** (upcoming) — 05 Jun 2024
  @ **S&P'24** · IEEE Symposium on Security and Privacy · San Francisco, USA · **in-person** (upcoming) — 22 May 2024
  @ **CCC'23** · Cambridge Cybercrime Conference · University of Cambridge, UK · **in-person** — 22 Jun 2023
- Bored of Cyberwar: the Role of Low-level Cybercrime Actors in the Russia-Ukraine Conflict
  @ **DSO** · DSO National Laboratories of Singapore · Singapore · **in-person** (upcoming) — 16 May 2024
  @ **WWW'24** · ACM World Wide Web Conference · Sentosa, Singapore · **in-person** (upcoming) — 15 May 2024
  @ **CCC'22** · Cambridge Cybercrime Conference · University of Cambridge, UK · **in-person** — 05 Sep 2022
- ExtremeBB: A Database for Research into Online Hate, Harassment, the Manosphere and Extremism
  @ **WOAH@ACL'23** · ACL Workshop on Online Abuse and Harms · Toronto, Canada · **online** — 13 Jul 2023
- PostCog: A 'Search Engine' Enabling Interdisciplinary Research into Underground Forums at Scale
  @ **WACCO@EuroS&P'22** · IEEE EuroS&P Workshop on Attackers and Cyber-Crime Operations · Genoa, Italy · **in-person** — 06 Jun 2022
- Turning Up the Dial: the Evolution of a Cybercrime Market Through Set-Up, Stable, and Covid-19 Eras
  @ **IMC'20** · ACM Internet Measurement Conference · Pittsburgh, USA · **online** — 29 Oct 2020
- Formal Semantics Extraction from Natural Language Specifications for ARM
  @ **FM'19** · International Symposium on Formal Methods · Porto, Portugal · **in-person** — 11 Oct 2019

Professional Skills

Security · Cybercrime · Data Science · Statistics · Malware Analysis · Reverse Engineering · OSINT · ML · Software Engineering · Databases

Programming: Python · Java · C/C++ · JavaScript · **Languages**: English (full professional proficiency), Vietnamese (mother tongue)

Honours & Awards

Selected as a Cambridge representative at the Global Young Scientists Summit	Singapore, Jan 2023
Awarded the Monbukagakusho Honours Scholarship in financial support for my master's study	Japan, 2017 – 2018
Awarded Outstanding Undergraduate Student at Vietnam National University	Vietnam, Jun 2016
Awarded the Shinnyo-en Japan Scholarship in financial support for my undergraduate study	Vietnam, 2012 – 2016

References

Prof. Alice Hutchings

PhD Supervisor

- Computer Laboratory, University of Cambridge
- Email: alice.hutchings@cl.cam.ac.uk
- Office: (+44) 1223 763 660

Prof. Ross Anderson

PhD Advisor

- University of Cambridge and University of Edinburgh
- Email: ross.anderson@cl.cam.ac.uk
- Office: (+44) 1223 334 733

Dr. Richard Clayton

Former Line Manager

- Former director, Cambridge Cybercrime Centre
- Email: richard.clayton@cl.cam.ac.uk
- Office: (+44) 1223 763 570

Prof. Mizuhito Ogawa

MSc Supervisor

- School of Information Science, JAIST
- Email: mizuhito@jaist.ac.jp
- Office: (+81) 761 511 247

Research Statement

My research offers timely empirical measurements to explore cyberspace and its societal impact at scale, focusing on underground subcultures fostering online crime and harms. The resulting insights help us better understand online threats and inform policy decisions for online safety and security. My primary approach is data-driven, with research questions addressed by rigorous quantitative and qualitative measurements of large-scale real-world evidence.

I. Measuring Security, Cybercrime, and Online Wickednesses

Externalities, such as Covid-19, may cause significant shifts and enduring changes in human behaviour, both offline and online. My recent work explores the effects of major incidents that *intensify* or *disrupt* online crime and harms.

▸ Intensifying Online Wickednesses: The Pandemic and Armed Conflicts

Turning Up the Dial @ IMC'20 [1] evidences the significant influence of Covid-19 on illicit trading activity on the largest underground hacking forum. This empirical observation can be attributed to the increasing time spent online by individuals during lockdowns. The paper reveals that various forum users overcame the 'cold-start problem' – where new traders face difficulty settling transactions due to a lack of reputation, yet they cannot gain reputation without trading – by engaging in low-value exchanges to build their trustworthiness before developing larger trading volumes.

Getting Bored of Cyberwar @ WWW'24 [2] explores the involvement of volunteer hacktivists and low-level cybercrime actors in the Russia-Ukraine conflict. Although they promptly participated in targeting digital assets of both countries after the invasion using DDoS and defacement attacks, this intensification was short-lived, with a clear loss of interest after a few weeks. Their activity may cause immediately noticeable effects, but the impact was mainly propaganda dissemination instead of contributing to the 'hard' digital frontline. While popular narratives tend to overhype and conflate these actors with persistent state-sponsored hacktivists, we believe they should be considered separately.

Yet Another Diminishing Spark @ Under Review [3] compares defacement attacks seen in the Russia-Ukraine conflict to those in the Israel-Hamas war, discovering similar patterns peaking shortly after the war started. While attacks were two-sided in the case of Russia-Ukraine, they have been mostly one-sided in the Israel-Hamas war: most targeted Israel while no significant waves have hit Palestine, presumably as Palestine has far fewer sites, many of which are hosted overseas. The scale of attacks on Israel and Palestine has been much less than those on Russia or Ukraine.

▸ Disrupting Online Wickednesses: Industry and Police Interventions

No Easy Way Out @ S&P'24 [4] examines a concerted effort to dismantle Kiwi Farms, the largest forum for online hate and harassment. We show that solely relying on deplatforming, even by swift actions of several competent tech firms, can be insufficient. The forum traffic and activity were quickly disrupted, but gradually recovered after a few months. Many users temporarily decamped to Telegram, but returned when the forum was back and became even more connected. The industry often does better than government actions, but this extraordinary event suggests that shutting down a dispersed community is unlikely to be effective if the censor cannot incapacitate or deter the key operators.

Assessing the Aftermath @ Under Review studies a global takedown of DDoS-for-hire services (or *booters*) involving the FBI, the NCA, and the Dutch Police. The first wave on 14 December 2022 seized 49 domains, and 13 more were seized in the second wave on 5 May 2023. These domains were redirected to a police-deployed page hosted by us to collect access flows across booters. We found that 26 first-wave seized domains quickly returned under new domains, while all 13 second-wave seized ones reappeared. These emergence, however, failed to recover their visit traffic, with over 80% being lost. The global DDoS attack volume quickly declined by half, with the effect lasting for only 8 weeks.

▸ Data Licensing: Reproducibility and Extensibility

Reproducibility and extensibility are crucial to me. All data used in my work can be shared with researchers, enabling them to immediately pursue ideas to address real-world problems without spending months or years collecting data.

ExtremeBB @ ACL WOAH'23 [5] is a dataset for large-scale research into online hate, harassment, and extremism. Outside of Cambridge, it has been licensed to 67 scholars in 27 groups across 21 institutions in 7 countries. I co-authored PostCog @ EuroS&P WACCO'22 [6], an interactive 'search engine' enabling researchers, especial non-technical scientists, to analyse our data visually and straightforwardly. I contributed to a chapter outlining how we ethically collect and share cybercrime and extremist datasets, as part of the book Open Source Investigations in the Age of Google [7].

II. Attacks and Defenses

I believe the best way to safeguard things is to figure out how to attack them. I am particularly fascinated by fundamental research uncovering novel practical attack vectors, and research that develops tools to counter such threats.

Erebus @ S&P'20 [8] is a new attack capable of partitioning the Bitcoin P2P network in a stealthy manner, without the victims realising they are being isolated. It leverages the advantages of big Internet entities such as ASes and ISPs to intercept and monitor thousands of 'shadow' malicious Bitcoin nodes. These can subsequently establish connections with legitimate nodes, gradually populating a large number of IP addresses into the victims' peering tables, and ultimately isolating them from the rest of the network. This vector also affects many other Bitcoin-based cryptocurrencies, such as Litecoin, Bitcoin Cash, and Dogecoin, necessitating a few protocol tweaks in the Bitcoin codebase.

Corana @ FM'19 [9] is a dynamic symbolic execution engine designed for multiple ARM Cortex variants. Complicated obfuscations such as packers, indirect jumps, and dead conditional branches (often referred to as opaque predicates) pose challenges to traditional approaches for malware analysis, including both static and dynamic methods. Corana is capable of effectively tracing IoT malware in the presence of these obfuscations. It was partly generated from natural language specifications of ARM instructions, and our paper demonstrates that this method can be systematically generalised to other variants, opening a new research direction in applying formal methods to malware analysis.

References

[1] A. V. Vu, J. Hughes, I. Pete, B. Collier, Y. T. Chua, I. Shumailov, and A. Hutchings, "Turning Up the Dial: the Evolution of a Cybercrime Market Through Set-up, Stable, and Covid-19 Eras," in *Proceedings of the ACM Internet Measurement Conference (IMC)*, 2020.

[2] A. V. Vu, D. R. Thomas, B. Collier, A. Hutchings, R. Clayton, and R. Anderson, "Getting Bored of Cyberwar: Exploring the Role of Low-level Cybercrime Actors in the Russia-Ukraine Conflict," in *Proceedings of the ACM Web Conference (WWW)*, 2024.

[3] A. V. Vu, A. Hutchings, and R. Anderson, "Yet Another Diminishing Spark: Low-level Cyberattacks in the Israel-Gaza Conflict," *Under Submission*, 2024.

[4] A. V. Vu, A. Hutchings, and R. Anderson, "No Easy Way Out: the Effectiveness of Deplatforming an Extremist Forum to Suppress Online Hate and Harassment," in *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2024.

[5] A. V. Vu, L. Wilson, Y. T. Chua, I. Shumailov, and R. Anderson, "ExtremeBB: A Database for Large-Scale Research into Online Hate, Harassment, the Manosphere and Extremism," in *ACL Workshop on Online Abuse and Harms (WOAH@ACL)*, 2023.

[6] I. Pete, J. Hughes, A. Caines, A. V. Vu, H. Gupta, A. Hutchings, R. Anderson, and P. Buttery, "PostCog: A Tool for Interdisciplinary Research into Underground Forums at Scale," in *Proceedings of the IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2022.

[7] L. Wilson, A. V. Vu, I. Pete, and Y. T. Chua, "Identifying and Collecting Public Domain Data for Tracking Cybercrime and Online Extremism," in *Open Source Verification in the Age of Google*, 2024.

[8] M. Tran, I. Choi, G. J. Moon, A. V. Vu, and M. S. Kang, "A Stealthier Partitioning Attack against Bitcoin Peer-to-Peer Network," in *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2020.

[9] A. V. Vu and M. Ogawa, "Formal Semantics Extraction from Natural Language Specifications for ARM," in *Proceedings of the International Symposium on Formal Methods (FM)*, 2019.