VeTSS

UK RESEARCH INSTITUTE ON

VERIFIED TRUSTWORTHY SOFTWARE SYSTEMS

ANNUAL REPORT 2024/2025

CONTENTS

FOREWORD
(NCSC REPRESENTATIVE)

VeTSS has many remits, but pivotal to all its aims is acting as a central hub for verification research in the UK. 2025 has seem VeTSS further establish itself as a fixture in the verification community, both at home and away. Following on from a very strong 2024, we've seen VeTSS' profile rise, with increased involvement in large scale initiatives to define the future of verified software systems and software understanding. Leading figures from the institute have been invited to participate in international efforts, in what proved to be a significant year for the profile of verification and formal methods. The importance and interest in these topics is growing quickly, and VeTSS and its wider community are firmly at the heart of this movement.

This wider involvement in large initiatives has not come at the expense of VeTSS' own events and endeavours, however. The year kicked off with a brilliant workshop in Bristol, co-hosted with fellow research institute RISCS on the Value of Feminism in the Theory and Practice of Computer Science and Cyber Security. This was well attended and discussed the problems affecting diversity in these disciplines, the need for change as well as proposing gameplans and solutions for how to address them. The VeTSS Annual Conference, a highlight in many people's calendars, did not disappoint with a brilliant lineup from some of the preeminent names in the field from both academia and industry. VeTSS also hosted a specialist event exploring the potential links between AI and code vulnerability detection, and what role verification, can, should and must play in this niche. Once more, very strong representation from major players in industry, academia and government contributed to some really probing discussions on the intersection of three complicated areas.

Continuing VeTSS long standing efforts to support early career researchers, the Annual Conference also saw the launch of the VeTSS Doctoral Dissertation Awards, celebrating outstanding PhD research in the UK. As a judge on the panel, I can attest to the incredible quality of research on offer over the last few years in the community, and the recognition and prize money at the event came with a palpable sense of pride about the work being done. Later in the summer, we welcomed back the VeTSS Summer School for its 3rd consecutive year. This popular event was hosted in Glasgow this year, and provided a fantastic venue for early career researchers to meet like minded potential collaborators and open their eyes to the range of topics out there in the field. In a further show of support for early career researchers, 2025 marked the return of the VeTSS Small Grants, after significant effort and time devoted to making it happen. Aligned directly with the Problem Book published last year, with the generous funding provided by DSTL, several grants were awarded, whittled down from dozens of applications, welcoming back an invaluable source of support for researchers early in their careers.

This has felt like a year where VeTSS kicked on from the efforts of the last few years to mature as an institute, gaining greater impact and standing in the community than ever before. There have been more internal initiatives and events than in previous years and the VeTSS community is having more and more influence on the future of verified systems. The hard work invested is now paying dividends, and VeTSS is well set up for continued success and impact in the future.

Adam W1, NCSC

DIRECTORS' MESSAGE

Over the past year, VeTSS has continued to strengthen its role as the UK's central hub for verified software systems, bringing together leading expertise from academia, government, and industry. Our focus remains clear: to translate world-class research into real-world impact, support the National Cyber Security Centre (NCSC) mission, and help researchers concentrate on the core challenges that underpin trustworthy digital infrastructure. Alongside this, we are committed to supporting researchers at all career stages and promoting equality, diversity and inclusion across all our activities.

In 2025 our programme of events has grown both in scope and reach. We are proud to have supported a rich calendar of community-building activities, including: the VeTSS Annual Conference, which continues to act as our flagship national forum for verified systems research and its applications, and the VeTSS Summer School, offering an intensive programme spanning security and privacy of protocols and program synthesis to the Lean theorem prover and weak memory models.

We have also sponsored several events, attracting 100s of participants, many of whom are students and early career researchers, showcasing our continuing support of the verification and programming languages communities.

Our Doctoral Dissertation Awards remain a core part of recognising excellence in early career research. Recent winners and runners-up showcase the breadth and depth of UK leadership in verification and security. These outstanding theses illustrate the long-term investment VeTSS makes in foundational work that ultimately underpins safer and more secure systems.

A major highlight this year has been our continued partnership with DSTL, providing funding for small grants of up to £50k for projects aligned with VeTSS research themes. This scheme, designed particularly to encourage early career researchers (postdoctoral level or above), is overseen by a diverse panel of eight members from industry, government and academia. We received 39 applications and funded six projects. These projects exemplify the core VeTSS aim: combining rigorous theory with pathways to real-world security impact, from cryptographic implementations to AI-assisted secure code generation.

DIRECTORS' MESSAGE

This year we also deepened our engagement with the wider security ecosystem around AI and software assurance. Working with Stuart Murdoch from Surevine, at TechUK, we convened representatives from government (DSIT, NCSC, DSTL), industry (including Microsoft, Amazon, reveng.ai) and academia (UCL, Surrey, Imperial and others). Our focus was on the role of AI in Static and Dynamic Application Security Testing (SAST/DAST). These insights have helped us articulate a clear research and innovation agenda that aligns academic ingenuity with industrial need, and they will guide our future calls and partnerships.
We also co-organised with RISCS, The Value of Feminism in the Theory and Practice of Computer Science and Cyber Security Workshop, which brought critical, interdisciplinary perspectives into the heart of security research and practice.

Across our strategic research portfolio, we have seen notable progress and impact in several flagship projects: (1) strengthening memory management and memory safety through an open-source framework, (2) verification results and tools that uncover bugs in real-world systems like the Google Oak kernel; and (3) establishing a full-stack verification approach for RDMA-based systems, from hardware to key-value stores. These have resulted in high-impact publications at leading venues, distinguished paper awards, a HiPEAC Technology Transfer Award, etc.

Across these strands, VeTSS continues to connect and support a vibrant national community in verified systems. Our events, sponsorships and partnerships reach across the UK, enabling early career researchers, deepening collaboration with government and industry, and ensuring that UK expertise in secure, verified systems remains globally influential.

We would like to thank our partners at NCSC, DSTL, industry collaborators, and the many researchers and students whose work features in this report. Their contributions are critical to realising a future where software systems are not only more capable, but demonstrably safer and more secure.

Brijesh Dongol

Azalea Raad

VeTSS ORGANISED EVENTS

THE VALUE of FEMINISM in THE THEORY and PRACTICE of COMPUTER SCIENCE and CYBER SECURITY WORKSHOP

27 March 2025, University of Bristol

A one-day interactive workshop dedicated to discussing the value of feminism in the theory and practice of Computer Science and Cyber Security was held at the Bristol Watershed on Thursday 27th of March 2025. This event is a collaboration between The Research Institute for Sociotechnical Cyber Security (RISCS) and The Research Institute on Verified Trustworthy Software Systems (VeTSS), both funded by the National Cyber Security Centre (NCSC).

The aim of the event is to discuss the challenges affecting diversity in the theory and practice of Computer Science and Cyber Security, why change is desirable, and to initiate discussion about how the theoretical computer science and cyber security communities might address these challenges.

The event was convened by Samantha Frohlich of the University of Bristol.

Highlights included:

- It featured a keynote from Felienne Hermans discussing her recent paper: 'A Case for Feminism in Programming Language Design' (2024).

- A panel discussion featuring Rebecca Jones, Sam Lindley, Wrenna Robson, Francois Dupressoir, and Samantha Frohlich

- A plenary discussion with contributions from the panel and attendees

Prof Genevieve Liveley Geneviev, the director of RISCS produced a summary report including the main headlines and key takeaways from the workshop, and personal responses from some of the panel members and from the VeTSS Director, Brijesh Dongol.

Photo credits: the Programming Languages Research Group (PLRG), University of Bristol

VeTSS ADVISORY BOARD MEETING

JULY 2025, ONLINE

Following Dstl's generous agreement to support VeTSS with £250K to fund a number of VeTSS research grants, the VeTSS Advisory Board met online on 18 July 2025. The primary purpose of the meeting was for Advisory Board members to act as the selection panel for the 2025 VeTSS Research Awards. Reviewers discussed the proposals individually, iteratively comparing scores and assessments to converge on the strongest submissions. In total, 6 proposals were selected from 39 submissions, resulting in a success rate of 15% for this year's funding call.

ANNUAL CONFERENCE

June 2025, ROYAL ACADEMY OF ENGINEERING, LONDON

The VeTSS Annual Conference 2025 brought together leading researchers working on verification, testing, and program analysis topics from across the UK for a day of presentations, panel discussions, and community networking. Key moments included:

- Opening remarks by the VeTSS Directors, highlighting achievements from the past year
- Presentations from the representatives from NCSC, DSTL, EPSRC, CRANE, and DSIT
- A keynote by Alastair Donaldson (Imperial College London): "In Search of Oracles: Testing Systems that Reason about Code"
- Seven additional talks showcasing cutting-edge work from academia and industry
- VeTSS Doctoral Dissertation Awards ceremony

Delegates praised the variety and quality of the presentations, as well as the friendly and sociable atmosphere of the event.

Opening and Panel discussion: The day began with a warm welcome and introductions from the VeTSS directors, who reflected on the achievements of the past year. This was followed by a panel discussion featuring representatives from NCSC, DSTL, EPSRC, CRANE, and DSIT, who shared their perspectives on current challenges and future directions.

Talks from academia and industry

Keynote: A highlight of the morning was the keynote lecture by Alastair Donaldson (Imperial College London), titled “In Search of Oracles: Testing Systems that Reason about Code” which provided an overview of recent and ongoing research, with various collaborators, on devising test oracles to automatically find bugs in compilers, equivalence checkers and program analysis tools, and to help map out the reasoning boundary of a given large language model.

In Search of Oracles: Testing Systems that Reason about Code
Alastair Donaldson
Imperial College London
https://youtu.be/QzqOlFcuor0

Formal methods for AI efficiency and automatic quantization
Andrey Rybalchenko
Microsoft Research
https://youtu.be/VFpQtnomgQo

Sociotechnical Futures: sociotechnical security for an age of cyber automation
Lizzie Coles-Kemp
Royal Holloway, University of London
https://youtu.be/wBU010fv5h4

Formal is Fast - Optimizing and Verifying ML-KEM
Hanno Becker
AWS
https://youtu.be/zeJ_V0NRIk0

Modelling and Verification of Quantum Systems
Rajagopal Nagarajan
QUƎNTꓤNGLE
https://youtu.be/Eq1A3GWWJ5g

Compositional automata learning for evolving systems
Mohammad Mousavi
King’s College London
https://youtu.be/rWzDU1hcVCc

Reflections on Property Based Testing
Samantha Frohlich
University of Bristol
https://youtu.be/IZjIfz_GwHc

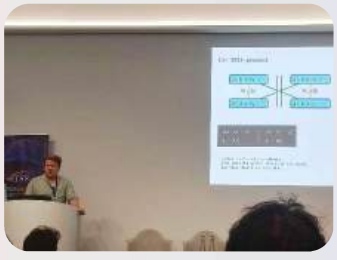

Symbolic MRD: Dynamic Memory, Undefined Behaviour, and Extrinsic Choice
Mark Batty
University of Kent
https://youtu.be/Qcf_vANZdKM

VeTSS SUMMER SCHOOL

AUGUST 2025, UNIVERSITY OF GLASGOW

The VeTSS Summer School 2025 was hosted by the Department of Computer Science at the Advanced Research Centre (ARC) of the University of Glasgow from August 11th to 14th. This year's programme featured a diverse range of lectures, tutorials, and practical sessions, covering both the theoretical foundations of verification and its real-world applications. Alongside the technical learning, participants enjoyed lively discussions, collaborative exercises, and the chance to build lasting networks with peers across academia, industry, and government.

SOCIAL EVENT HIGHLIGHTS

Beyond the lecture sessions, the social events were a real highlight, providing excellent opportunities to network and make new friends. In addition to our traditional Pizza Night, all participants were invited to attend a guided tour and drinks reception at the historic Glasgow City Chambers as part of the celebrations for Glasgow's 850th birthday. We received a Civic Welcome from officials of Glasgow City Chambers, and the VeTSS Directors also gave a speech outlining the history of VeTSS and its importance in the development of UK cyber security research at the reception.

PROGRAMME HIGHLIGHTS

This year's Summer School covered many interesting topics like program synthesis, model checking, and functional programming.

THE SIX MAIN LECTURES

How can we avoid conflicts of interest?
Ross Horne
University of Glasgow
https://youtu.be/rboNhfpY4lM

The Lean Programming Language and Theorem Prover, *Sebastian Ullrich and Joachim Breitner, Lean Focused Research Organization*
https://youtu.be/FVHxu4h3BKA

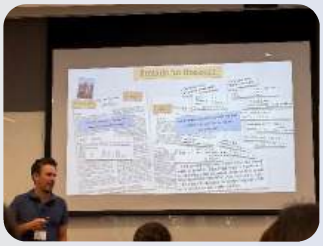

Algebraic Effects
Nicolas Wu
Imperial College London
https://youtu.be/wgfYkQVpC-M

Introduction to Program Synthesis
Elizabeth Polgreen
University of Edinburgh
https://youtu.be/ii9Zj2YFP_s

Session Types, Asynchronous Communication, and Subtyping
Laura Bocchi
University of Kent
https://youtu.be/a8fiUbWjAbo

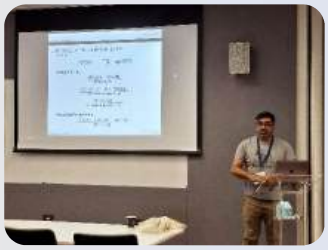

Owicki-Gries Reasoning for Weak Memory Models
Brijesh Dongol
University of Surrey
https://youtu.be/9jibN7HWbZM

IMPACT AND ACKNOWLEDGMENTS

We are delighted to congratulate all participants on completing this year's Summer School. Over the past 4 days, students, researchers, and practitioners from across the UK and beyond came together to explore the cutting edge of formal verification, program synthesis, and functional programming. We are especially proud of the enthusiasm and dedication shown by this year's attendees. Completing an intensive week of deep technical study is no small achievement, and we hope participants leave inspired to continue developing and applying the skills they've gained.

Finally, a huge thank you to the local organisers Jeremy Singer and Kai Feng, and all speakers who made this year's Summer School such a success. Your hard work ensures that VeTSS continues to play a central role in building the next generation of experts in verification and security.

AI CODE VULNERABILITY ANALYSIS WORKSHOP

OCTOBER 2025, TECHUK, LONDON

BACKGROUND

During the latter part of 2025, Surevine were engaged by the UK National Cyber Security Centre (NCSC) to conduct a research project into the state of the art of AI code analysis tools. Whilst it was recognised that there are opportunities for AI tools to support every point for the Software Development Lifecycle (SDLC), this research was focussed on AI used to support Static Application Security Testing (SAST). The research looked at commercial vendors, open source projects and academic research. A systematic approach was used to compare the candidates, with an agreed set of common criteria. The AI enhancements to existing SAST tools discovered were mostly in triage (reducing false positives) or suggesting code remediations. Almost all of the tools required networked access to the vendor's model to enable the AI enhancements, or to a frontier model in the case of many of the open source projects. Hardly any of the tools reviewed used AI to discover new vulnerabilities which could not be discovered using existing SAST techniques. There were some very promising academic approaches, which did seem to report improved performance over traditional SAST. None of the academic approaches had been pulled-through into a production-quality tool suitable for use by industry. The approach to the research included a literature review, direct engagement with vendors and technical testing of some of the most promising tools. As part of the engagement, the VeTSS Research Institute organised a one-day workshop with the leading researchers in this field in the UK, the findings from which were incorporated into the research.

EVENT SUMMARY

Primary Focus of the Workshop:

Static Analysis and Testing: Applying LLMs for vulnerability detection and mitigation suggestions (e.g., during SAST)

Secondary Focus Areas:

- Dynamic Testing (DAST) - ML/LLM support for runtime vulnerability detection
- Code Creation - Security-aware IDE plugins preventing common pitfalls
- Code Reviews - LLM assistance for independent code review processes
- Testing - Automated generation of security-focused unit/integration tests
- Production Monitoring - Vulnerability detection in live systems
- Impact Metrics - Measuring AI-enabled detection effectiveness vs. conventional methods

NEXT STEP

Twenty-two delegates with diverse expertise from academia, start-ups, and industry attended the workshop. The day comprised a series of technical talks and roundtable discussions on the topics outlined above. Meeting minutes, including a summary of the main themes, topics discussed and future actions, were prepared and circulated. Three main future actions are:

- Present the findings to relevant stakeholders, including the NCSC, and explore potential funding opportunities to assess the feasibility, scalability, and impact of further work.
- Schedule a follow-up workshop or technical discussion to review progress, refine priorities, and identify emerging challenges or opportunities.
- Integrate a concise, technically focused summary of the discussion into the VeTSS Problem Book, explicitly articulating its relationship to the ongoing Problem Book review and highlighting concrete research challenges.

PROGRAM of THE WORKSHOP

Exploring Semi-automatic Agentic Performance Engineering on the GitHub Platform
Don Syme, (Microsoft)

Documentation-guided Taint Analysis for the Android Platform
Santanu Dash, (University of Surrey)

AI code analysis tools for vulnerability detection
Stuart Murdoch & Matthew Vivian, (Surevine)

Source code vulnerability detection with deep learning and static analysis
Sergio Maffeis, (Imperial College London)

Trustworthy AI... for Systems Security: LLMs for Malicious Code Analysis
Lorenzo Cavallaro, (UCL)

Tool Assurance: Some Considerations from Standards
Rob Ashmore, (DSTL)

Building foundational LLMs for automated vulnerability analysis in compiled software binaries
James Patrick Evans, (Reveng.ai)

VeTSS DOCTORAL DISSERTATION AWARDS

One of the standout moments of the Annual Conference 2025 was the presentation of the VeTSS Doctoral Dissertation Awards, which celebrate outstanding PhD research in the UK. This award recognises outstanding research by UK PhD students working on topics aligned with the VeTSS agenda, as outlined in the VeTSS Problem Book. Three prizes were awarded for dissertations completed in 2022, 2023, and 2024, following evaluation by a panel of four judges. Each winner received a £1,000 prize and was invited to present their research at the VeTSS Annual Conference. Please find the winners below:

- **2022 Winner**

Simon Cooksey: *Automating C++ Execution Exploration to Solve the Out-of-thin-air Problem*

- **2023 Winner**

Kayvan Memarian: *The Cerberus C Semantics*

- **2024 Winner**

Yann Herklotz: *Formal Verification of High-Level Synthesis*

- **2024 Runner up**

Zhixuan Yang: *Structure and Language of Higher-Order Algebraic effects*

FUNDING OPPORTUNITIES

Dstl FUNDING

VeTSS was pleased to announce that it has secured £250,000 in funding from the Defence Science and Technology Laboratory (Dstl) in early 2025. Building on this support, the VeTSS Research Award 2025 funding call was launched in July 2025, inviting proposals for research projects to be conducted between 1 September 2025 and 31 August 2026. This call aimed to support projects aligned with VeTSS research on topics including, but not limited to, those outlined in the VeTSS Problem Book. Problems not listed in the Problem Book were also eligible for funding, provided that the proposal clearly links the project with VeTSS objectives.

In total, 39 applications were received and evaluated by a panel of eight experts. From these, six projects were selected for funding, corresponding to a success rate of approximately 15%.

Congratulations to the awardees, especially the proposals led by early career researchers (ECRs) and by researchers currently underrepresented in computer science and / or belonging to marginalised communities. The details of awardees are:

- **Unlocking Hardware Verification via Software Algorithm Specifications**, PI: Dr Jianyi Cheng, University of Edinburgh, Co-I: Dr Yann Herklotz, EPFL.
- **Probabilistic Precision Tuning**, PI: Fredrik Dahlqvist, Queen Mary University of London.
- **Relational Dualities for Program Logics**, PI: Alex Kavvos, University of Bristol.
- **Secure Code Generation with Large Language Models**, PI: Arindam Sharma, Co-I: Cristina David, University of Bristol.
- **Formalised Software Requirements for Human-Robot Teamwork**, PI: Hazel M. Taylor, Co-I: Marie Farrell, University of Manchester.
- **Translation Validation for SuperOptimized Cryptographic Software**, PI: François Dupressoir, Co-I: Daniel Page, University of Bristol.

SPONSORED EVENTS

26th INTERNATIONAL SYMPOSIUM on TRENDS in FUNCTIONAL PROGRAMMING

13th - 16th JANUARY 2025, UNIVERSITY OF OXFORD

From 13th to 16th January 2025, the 26th International Symposium on Trends in Functional Programming took place in the Department of Computer Science at the University of Oxford. Generous sponsorship from VeTSS and Epic Games supported over 50 participants, about half of whom were students.

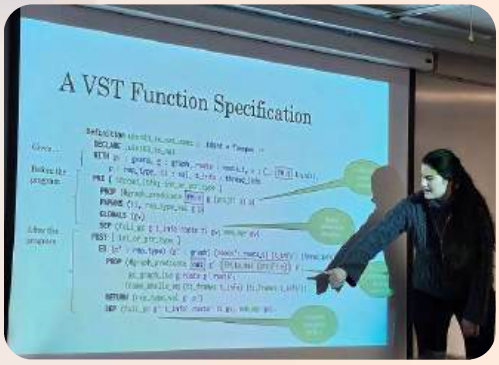

The main conference featured keynote talks from Graham Nelson from the University of Oxford on ‘Literate Programming and Cultural Practice’, Mike Sperber from Active Group in Tübingen on ‘Things We Never Told Anyone About Functional Programming’, and Kathrin Stark from Heriot-Watt University on ‘A Verified Foreign Function Interface Between Coq and C’, in addition to 28 contributed talks.

FSCD 2025 CONFERENCE and WOMEN in LOGIC

14-20 July, University of Birmingham

The FSCD 2025 conference, hosted at the University of Birmingham, was a highly successful event, attracting 155 registered participants from over 20 countries. Attendees included 53 students, 18 postdocs, and 67 faculty members, with a strong international presence and a gender breakdown of 131 male, 29 female, and 3 other participants. Women in Logic (WiL 2025) was an associate event to FSCD. A free virtual participation option welcomed 126 registrants, with 10-20 joining online at any given time. Thanks to the generous support of VeTSS, the event maintained competitive pricing (e.g., early student registration at £430 for the full event), ensuring accessibility while delivering an enriching academic experience.

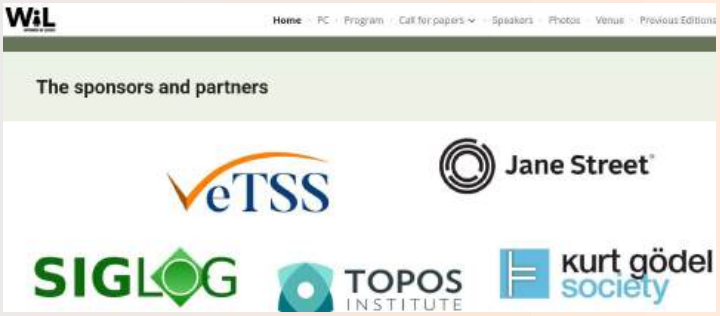

S-REPLS 16

1 April 2025, Imperial College London

VeTSS supported the 16th edition of S-REPLS, which took place at Imperial College London on 1st April 2025. With more than 100 attendees from all around the country, the programme opened with a keynote from Sophia Drossopoulou (Imperial) and Matthew Parkinson (Azure Research, Microsoft) about behaviour-oriented concurrency, its adoption in Project Verona, and the use of behaviour-oriented concurrency ideas in bringing concurrency to Python in a principled manner.

MGS 2025

7-11 April 2025, University of Sheffield

On 7-11th April, the School of Computer Science at the University of Sheffield hosted the 25th edition of the Midlands Graduate School in the Foundations of Computing Science (MGS 25). MGS is a major yearly event in the national and international landscape of PhD summer/spring schools, and a long-standing collaboration between the Universities of Birmingham, Leicester, Nottingham and Sheffield. MGS 25 had 66 student participants and featured eight intensive courses on a variety of foundational topics (category theory, logic, proof theory, verification and quantum computing) from both world-leading and rising-star academics. VeTSS was one of the sponsors.

VeTSS-ALIGNED PROJECTS (2024/25)

In 2022, the EPSRC, NCSC and the DSTL jointly agreed to fund a number of research projects aligned with the NCSC-funded Research Institutes ukri.org/opportunity/research-aligned-with-cybersecurity-research-institutes

Within this call, three VeTSS-aligned projects were funded, which together received £3.4M of investment. You will find the progress and impact they have achieved in the first two years of their research.

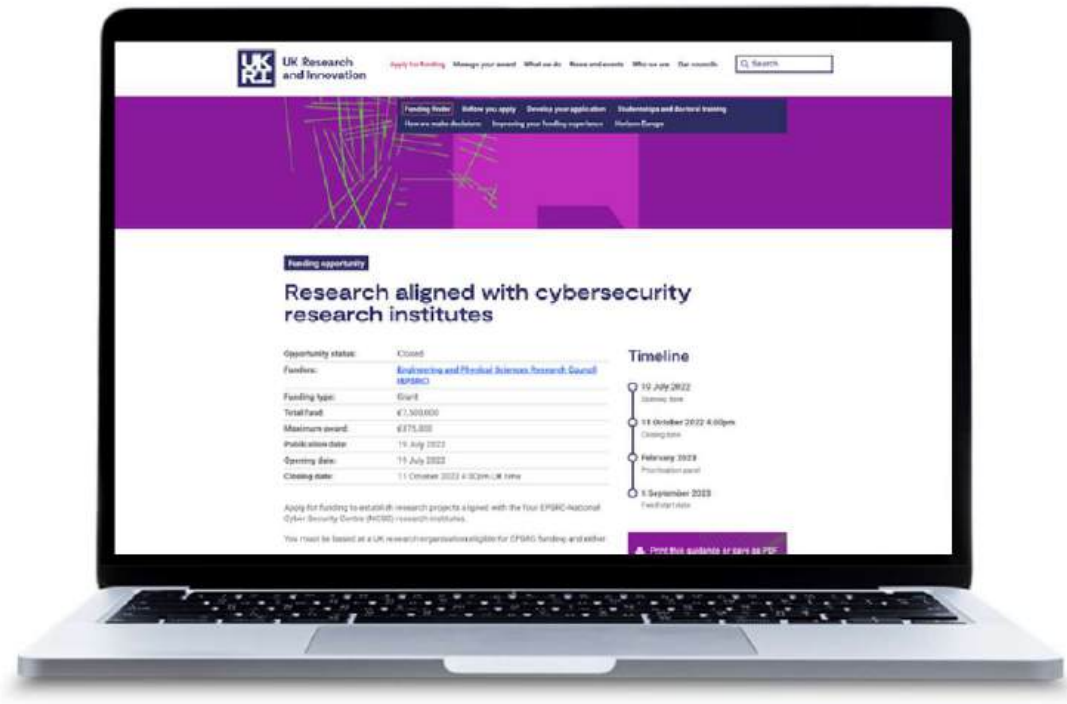

MAKING MEMORY MANAGEMENT MORE SECURE (M4SECURE)

OBJECTIVES

M4Secure aims to develop an automatic framework to create high-performance and secure memory allocators to take advantage of hardware security features. This is achieved by combining machine-learning-based code optimisation for memory allocator code synthesis and formal methods for verifying memory allocator security properties.

Jeremy Singer
University of Glasgow

Alice Miller
University of Glasgow

Zheng Wang
University of Leeds

PROGRESS

The project focuses on optimising application performance and enhancing memory security using automated and AI-driven techniques. They have made multiple key advances that support the goals of M4Secure:

- **Developed a CHERI-aware superoptimiser:** The first framework to model capability semantics in both its search space and correctness checks.
- **Carried out the most comprehensive performance study of CHERI:** Using on-chip performance counters on the Arm Morello platform, we analysed 20 C/C++ applications, including SPEC CPU2017, a SQL database engine, a JavaScript engine, and an LLM inference workload, across three CHERI ABIs.
- **Developed SecureMind:** SecureMind provides a simple Python interface and automates data preparation and benchmarking.
- **Enhanced the open-source MicroPython interpreter:** The new MicroPython port supports hardware-enabled spatial memory safety, mitigating a large set of common runtime memory attacks.
- **Developed Prom:** An open-source library to improve the robustness of machine-learning models for code analysis and optimisation. Prom uses statistical tests to flag likely mispredictions and uses feedback to update deployed models.

IMPACT

This work shows that strong memory safety can be achieved with practical performance through CHERI-aware optimisation, while providing clear evidence to guide future hardware improvements. Key impacts of the study:

- Delivered an open, **reusable framework** to evaluate and improve how AI models detect and repair memory-safety vulnerabilities, supporting more secure software development, and improving the reliability of machine learning based code optimisation
- **Prom**, which was published in ACM CGO 2025, a premier compiler conference, won a distinguished paper award.
- A University of Leeds spin-out company, **ACE3 Ltd**, co-founded by two M4Secure members, has received a 2025 HiPEAC Technology Transfer Award, recognising the successful transfer of research outcomes to market.

SECURITY SUMMARIES FOR SECURE SOFTWARE DEVELOPMENT (TRUSTED)

OBJECTIVES

TRUSTED aims to enhance program security by designing efficient algorithms to generate and maintain security summaries for Java and Rust programs. It seeks to develop theories that refine and compose these summaries, improving precision and enabling cross-language information flow analysis. Additionally, it focuses on creating validation and lightweight verification techniques to assess the security of third-party libraries. Developers will receive meaningful feedback to help eliminate illegal information flow, ensuring more secure software. The effectiveness of these techniques will be evaluated through real-life case studies.

Narges Khakpour
University of Newcastle

Sven Schewe
University of Liverpool

Dominik Wojtczak
University of Liverpool

PROGRESS

- Focused on algorithms and techniques for learning the behaviour of systems represented of formal acceptors (like security summaries or automata) in foundational and application facing automata learning approaches for non-terminating systems and protocols.
- Considered the impact of changes from computationally simple reachability and (generalised) repeated reachability to generalised ordinary reachability, which is closer to security guarantees as it reflects conjunctive protections
- Considered efficient treatment of properties, e.g. in the treatment of LTLf specifications and in obtaining unambiguous automata efficiently from a fairly general class of specifications (alternating weak automata, which include behaviours expressed by security summaries as well as specifications in LTL or LDL) to unambiguous automata.
- Developed a fast method to calculate fixed points in quantitative systems with contraction properties, like discounted games (stochastic or not).
- Verified the bitmap allocator of the Google Oak kernel in RoCQ and detected some bugs.
- Used zero-knowledge proof systems for sharing SBOMs (Software Bills of Materials) in a secure way. This method will be used for propagating security summaries.
- Defined formal semantics for Rust+C programs to build the foundation for security analysis and summary generation (in progress).

IMPACT

- The immediate impact is tool support, and we have started to complement foundational results with tools, in particular for the efficiency gains in model checking.
- Finding bugs in Google Oak kernel and collaboration with industry

SAFE AND SECURE REMOTE DIRECT MEMORY ACCESS (SACRED-MA)

OBJECTIVES

Remote Direct Memory Access (RDMA) is a modern technology enabling networked machines to exchange information without involving the operating system of either side, and thus significantly speeding up data transfer in computer clusters. This project aims to build up the understanding of RDMA from hardware models to distributed applications. We use formal approaches to develop and adapt reasoning techniques and proof methods to the unique design challenges of RDMA. Our goal is to develop tools to empower programmers and future developers of RDMA-based systems.

Brijesh Dongol
University of Surrey

Azalea Raad
Imperial College London

Gregory Chockler
University of Surrey

PROGRESS

- Dartagnan and RDMA Integration: Dartagnan encodes program logic and memory model semantics (CAT), uses relation analysis and compact SMT encodings, and supports a wide range of architectures. RDMA support was integrated to enable analysis of distributed algorithms.
- Fault Handling with RDMA: To issue parallel RDMA operations to all servers and wait for acknowledgements from a quorum (majority). RDMA's reliable transport and direct access help clients detect failed servers via missing completions or timeouts, ensuring operations complete if a quorum responds.
- Created new reasoning techniques to modularly specify library behaviours and verify implementation correctness. These methods enable the construction of abstraction layers, providing efficient, verified data structures while abstracting away network complexities.
- Collaborating with NVIDIA, we extended the formal semantics of RDMA to encompass atomic operations.
- Looking into the semantics of combining RDMA with other architectures, such as GPUs and ARM CPUs, aiming to broaden its applications beyond high-performance computing and data centres.

IMPACT

- Enables developers and researchers to verify correctness and portability of communication libraries and concurrent algorithms in distributed settings.
- Adding support for atomic compare-and-swap (CAS) operations, which will allow Dartagnan to verify a broader class of synchronization mechanisms and concurrent algorithms under weak and RDMA memory models.
- Dartagnan now supports distributed programs by encoding node identities and RDMA-specific primitives, allowing formal analysis of RDMA_TSO semantics.
- Provided a formal basis for state-of-the-art network algorithms.
- We verified basic libraries that already outperform established network protocols such as MPI.
- This opens the door to manipulating network memory similarly to shared memory models.

UNIVERSITY OF SURREY IMPERIAL

NVIDIA arm TEL AVIV UNIVERSITY Cornell University MAX PLANCK INSTITUTE FOR SOFTWARE SYSTEMS University of Colorado Boulder

VeTSS ALIGNED PROJECT PUBLICATIONS

Year 2024/25

Xiaoyang Sun, Jeremy Singer, and Zheng Wang. Sweet or Sour CHERI: Performance Characterisation of the Arm Morello Platform. In Proceedings of the 2025 IEEE International Symposium on Workload Characterization (IISWC), 2025. https://eprints.whiterose.ac.uk/id/eprint/231424/

Huanting Wang, Dejice Jacob, David Kelly, Yehia Elkhatib, Jeremy Singer, and Zheng Wang. SecureMind: A Framework for Benchmarking Large Language Models in Memory Bug Detection and Repair. In Proceedings of the 2025 ACM SIGPLAN International Symposium on Memory Management (ISMM), 2025. https://doi.org/10.1145/3735950.3735954

Huanting Wang, Patrick Lenihan, and Zheng Wang. Enhancing Deployment-Time Predictive Model Robustness for Code Analysis and Optimization. In Proceedings of the 21st ACM/IEEE International Symposium on Code Generation and Optimization (CGO), 2025. Distinguished Paper Award. https://doi.org/10.48550/arXiv.2501.00298

Duncan Lowther, Dejice Jacob, Jacob Trevor, and Jeremy Singer. Secure Scripting with CHERIoT MicroPython. In Proceedings of the 34th ACM SIGPLAN International Conference on Compiler Construction (CC), 2025. https://doi.org/10.1145/3708493.3712694

Jeremy Singer and Steve Draper. Let's Take Esoteric Programming Languages Seriously. In Proceedings of the 2025 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!), 2025. https://doi.org/10.48550/arXiv.2505.15327

Gaojie Jin, Xinping Yi, Wei Huang, Sven Schewe, and Xiaowei Huang. S2O: Enhancing Adversarial Training with Second-Order Statistics of Weights. Transactions on Pattern Analysis and Machine Intelligence 47(10):8630-8641, TPAMI25. 2025. https://doi.org/10.48550/arXiv.2203.06020

Yong Li, Sven Schewe, and Moshe Y. Vardi. Singly exponential translation of alternating weak Büchi automata to unambiguous Büchi automata. Theoretical Computer Science 1006:114650, TCS24, 2024. https://doi.org/10.48550/arXiv.2305.09966

Daniele Dell'Erba, Sven Schewe, and Ashutosh Trivedi. Objective Improvement: Algorithm for Controller Synthesis in Uncertain Environments. In Proceedings of the 64th IEEE Conference on Decision and Control (CDC 2025), 2025. https://doi.org/10.1109/CDC57313.2025.11312600

Mona Alluwaym, Yong Li, Sven Schewe, and Qiyi Tang. Efficient Learning of Weak Deterministic Büchi Automata. In Proceedings of the 28th European Conference on Artificial Intelligence (ECAI 2025), 2025. https://doi.org/10.48550/arXiv.2508.14274

Wanrong Yang, Manhui Wang, and Dominik Wojtczak. Abstract Attack Intention Inference Using Low-Rank Gated Arithmetic Interactive Attention. In Proceedings of the 2025 IEEE International Conference on Cyber Security and Resilience (CSR25), 2025. https://doi.org/10.1109/CSR64739.2025.11129983

Wanrong Yang and Dominik Wojtczak. Efficient Inference of Sources and Targets in a Graph with Limited Observations. Volume 413: ECAI 2025. https://doi.org/10.3233/faia251389

Soumyajit Paul, David Purser, Sven Schewe, Qiyi Tang, Patrick Totzke, and Di-De Yen. Resolving Nondeterminism by Chance. In Proceedings of the 36th International Conference on Concurrency Theory (CONCUR 2025), 2025. https://doi.org/10.48550/arXiv.2504.10234

León Bohn, Yong Li, Christof Löding, and Sven Schewe. Saturation Problems for Families of Automata. In Proceedings of the 52nd EATCS International Colloquium on Automata, Languages, and Programming (ICALP 2025), pp. 146:1-19, 2025. https://doi.org/10.48550/arXiv.2506.13197

Sougata Bose, Daniel Hausmann, Soumyajit Paul, Sven Schewe, and Tansholpan Zhanabekova. Generalised Reachability Games Revisited. In Proceedings of the Sixteenth International Symposium on Games, Automata, Logics, and Formal Verification (GandALF 2025), 2025. https://doi.org/10.48550/arXiv.2509.14091

Yong Li, Soumyajit Paul, Sven Schewe, and Qiyi Tang. Accelerating Markov Chain Model Checking: Good-for-Games Meets Unambiguous Automata. In Proceedings of the 37th International Conference on Computer Aided Verification (CAV 2025), 2025. https://doi.org/10.1007/978-3-031-98679-6_13

Giuseppe De Giacomo, Yong Li, Sven Schewe, Christoph Weinhuber, and Pian Yu. Solving MDPs with LTLf+ and PPLTL+ Temporal Objectives. In Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2025), 2025. https://doi.org/10.48550/arXiv.2505.17264

Hugo Gimbert, Soumyajit Paul, and B. Srivathsan. Simplifying Imperfect Recall Games. In Proceeding of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), 2025. https://doi.org/10.48550/arXiv.2502.13933

Guillaume Ambal, Ori Lahav, and Azalea Raad. Sufficient Conditions for Robustness of RDMA Programs, European Symposium on Programming (ESOP), 2025. https://doi.org/10.1007/978-3-031-91118-7_3

Year 2023/24

Guillaume Ambal, Brijesh Dongol, Haggai Eran, Vasileios Klimis, Ori Lahav, Azalea Raad. 2024. Semantics of Remote Direct Memory Access: Operational and Declarative Models of RDMA on TSO Architectures. In: Proceedings of the ACM on Programming Languages, Volume 8, Issue OOPSLA2.Article No.: 341, Pages 1982 – 2009. doi.org/10.1145/3689781

Narges Khakpour, Charilaos Skandylas. Compositional Security Analysis of Dynamic Component-based Systems. ASE '24: Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering, Pages 1232 – 1244. doi.org/10.1145/3691620.3695499

Narges Khakpour, David Parker. Partially-Observable Security Games for Attack-Defence Analysis in Software Systems, 22nd International Conference on Software Engineering and Formal Methods (SEFM 2024). doi.org/10.1007/978-3-031-77382-2_9

Daniele Dell'Erba, Yong Li, Sven Schewe. DFAMiner: Mining minimal separating DFAs from labelled samples. In Proceedings of the 37th International Symposium on Formal Methods (FM 2024), pp. 48–66, 2024. doi.org/10.48550/arXiv.2405.18871

Ernst Moritz Hahn, Mateo Perez, Sven Schewe, Fabio Somenzi, Ashutosh Trivedi, Dominik Wojtczak. Omega-Regular Decision Processes. In Proc. Thirty-Eighth Conference on Artificial Intelligence (AAAI 2024), pp. 21125–21133, 2024. doi.org/10.1609/aaai.v38i19.30105

Ernst Moritz Hahn, Mateo Perez, Sven Schewe, Fabio Somenzi, Ashutosh Trivedi, and Dominik Wojtczak. Omega-Regular Reward Machines. In Proc. 26th European Conference on Artificial Intelligence (ECAI 2023), pp. 972–979, 2023. doi.org/10.48550/arXiv.2308.07469

Sougata Bose, Thomas A. Henzinger, Karoliina Lehtinen, Sven Schewe, Patrick Totzke. History-deterministic Timed Automata. 2024. In Logical Methods in Computer Science, 20(4:1). doi.org/10.46298/lmcs-20(4:1)2024

Nicolas Berthier and Narges Khakpour. Symbolic Abstract Heaps for Polymorphic Information-Flow Guard Inference. In the 23rd International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI 2023), 66-90, 2023. doi.org/10.48550/arXiv.2211.03450

Armando Castañeda, Gregory V. Chockler, Brijesh Dongol, Ori Lahav. What Cannot Be Implemented on Weak Memory? DISC 2024: 11:1-11:22. arxiv.org/abs/2405.16611

Eleni Vafeiadi Bila, Brijesh Dongol. A verified durable transactional mutex lock for persistent x86-TSO. Formal Methods Syst. Des. 64(1): 237-282 (2024). doi.org/10.1007/s10703-024-00462-1

Maurice H. ter Beek, Manfred Broy, Brijesh Dongol. The Role of Formal Methods in Computer Science Education. Inroads 15(4): 58-66 (2024). dl.acm.org/doi/10.1145/3702231

Sharar Ahmadi, Brijesh Dongol, Matt Griffin. Operationally proving memory access violations in Isabelle/HOL. Sci. Comput. Program. 234: 103088 (2024). doi.org/10.1016/j.scico.2024.103088

Brijesh Dongol, Matt Griffin, Andrei Popescu, Jamie Wright. Relative Security: Formally Modeling and (Dis)Proving Resilience Against Semantic Optimization Vulnerabilities. CSF 2024: 403-418. ieeexplore.ieee.org/abstract/document/10664336

Azalea Raad, Ori Lahav, John Wickerson, Piotr Balcer, Brijesh Dongol. Intel PMDK Transactions: Specification, Validation and Concurrency. ESOP (2) 2024: 150-179. doi.org/10.48550/arXiv.2312.13828

Azalea Raad, Ori Lahav, John Wickerson, Piotr Balcer, Brijesh Dongol. Artifact Report: Intel PMDK Transactions: Specification, Validation and Concurrency. ESOP (2) 2024: 180-184. doi.org/10.1007/978-3-031-57267-8_7

Lara Bargmann, Brijesh Dongol, Heike Wehrheim. Unifying Weak Memory Verification Using Potentials. FM (1) 2024: 519-537. doi.org/10.1007/978-3-031-71162-6_27

Sergey Egorov, Gregory V. Chockler, Brijesh Dongol, Dan O'Keeffe, Sadegh Keshavarzi. Mangosteen: Fast Transparent Durability for Linearizable Applications using NVM. USENIX ATC 2024: 799-815. dl.acm.org/doi/10.5555/3691992.3692041

Stefan Bodenmöller, John Derrick, Brijesh Dongol, Gerhard Schellhorn, Heike Wehrheim. A Fully Verified Persistency Library. VMCAI (2) 2024: 26-47. link.springer.com/chapter/10.1007/978-3-031-50521-8_2

We look towards developing the new version of Problem Book and are keen to know the users' opinion on the VeTSS Problem Book 2024. A user survey on the VeTSS Problem Book 2024 has been created to gather feedback on how the community engages with it and suggestions for improvement. Insights from this survey will help the development of the 2025 edition.

THE RESEARCH INSTITUTE ON VERIFIED TRUSTWORTHY SOFTWARE SYSTEMS

CONTACT US

vetss.org.uk

contact@vetss.org.uk

linkedin.com/in/VeTSS/

National Cyber Security Centre
a part of GCHQ

UNIVERSITY OF SURREY

IMPERIAL