Institution: QUEEN MARY UNIVERSITY OF LONDON

Unit of Assessment: B11 Computer Science and Informatics

Title of case study: *Automatic memory safety verification for critical software*

1. Summary of the impact (indicative maximum 100 words)

Memory violations are a major cause of security breaches and operational flaws in today's software systems. Proving memory safety was traditionally a core challenge in program verification due to the high complexity of reasoning about pointer manipulations. Researchers at Queen Mary and Imperial jointly produced breakthrough algorithms for automatically reasoning about pointers, enabling highly-scalable automatic verification for industrial code. These techniques resulted in the industrial program analysis tool INFER developed by **Monoidics Ltd**, and used by customers across the world. The verification algorithms developed at Queen Mary and Imperial were also incorporated in Microsoft tools used to secure Windows device drivers.

2. Underpinning research (indicative maximum 500 words)

The problem of verifying that programs correctly and safely manipulate computer memory is at least 30 years old. Verification is imperative as such programs are often critical code in IT and Engineering systems: if memory errors occur there may be fatal consequences for overall system missions (e.g. safety of a plane in flight).

The development of this technology is based on research into logics for reasoning about memory use and the algorithms giving partial decision procedures. Key stages include underpinning work on substructural logic by Pym and O'Hearn at Queen Mary [P6] which was equipped with primitives to describe heap structure in joint work between researchers at Queen Mary and Carnegie Mellon (eg [P7]). This formed the first effective theory for reasoning about computer memory, with, initially, hand-coded verification of standard systems algorithms, and then the first automated verifiers.

On this base, a team of Queen Mary academics and alumni based then at Queen Mary, and Imperial designed and implemented the first separation logic-based verifier: Smallfoot [P1].

Distefano, at Queen Mary, produced the Space Invader the first program analyser based on separation logic.

Space Invader and Smallfoot opened up for the first time the possibility of automating the theoretical ideas of separation logic in mechanised proofs.

To move from academic research to automatic proofs that can be used in real-world industrial code, several steps were necessary. The first milestone was reached when Calcagno and Distefano jointly developed a technique where an analyzer can tune the precision of its analysis on-the-fly to different data structures of the analyzed program [P3].

Several routines of Windows Device Drivers were analyzed with this technique and numerous real bugs were uncovered, and confirmed by the Windows Kernel Team. The second milestone was the development of automatic proofs concerning the use of pointers in entire industrial programs (Linux and Microsoft device drivers of up to 10k LOC) [P4].

Finally, Calcagno and Distefano proposed the notion of bi-abductive inference, a breakthrough which allowed partial properties of entire open source projects in the hundreds of thousands or millions of lines of code to be shown [P5]. To put these results in context, just five years ago proofs of memory safety had been done only for toy programs in the tens or sometimes hundreds of lines

of code.

In the relatively short time since its invention, the algorithms invented by Calcagno and Distefano have been incorporated into many (>30) tool efforts in the academic sector, including in groups from Harvard, MIT, Yale, Princeton, CMU, Berkeley, Paris, Copenhagen, Brno, Tokyo, Cambridge, Sussex and elsewhere.

More importantly, the algorithms developed by Calcagno and Distefano have been developed into products that have significant industrial impact. The algorithms form the foundation for the program verifier INFER developed by the high-tech SME Monoidics Ltd. The Monoidics INFER product has attracted customers across the world, including government agencies and industrial clients like Airbus, Mitsubishi, Toyota, LLNL, and ARM. Recently Monoidics has been acquired by Facebook to improve the quality and the security of their code base. Within Facebook, the verification technology originally designed by Distefano and Calcagno will impact over one billion users.

Microsoft incorporated Calcagno and Distefano's algorithms in the Slayer Static Analyser, a product used to secure Windows device drivers.

This research has attracted significant academic and industrial funding [G1-7]. Two EPSRC projects happened at QM in the period 2003-2011 (Local Reasoning about State [G7], Adaptive Heap Analysis [G3]). Calcagno received an EPSRC Advanced Fellowship [G2], and Distefano was the recipient of an RAEng fellowship [G1]. The outcome of this fellowship was used as a "Case Study" by the Royal Academy of Engineering in 2013 [P8].

Distefano is the 2012 recipient of the BCS Needham Award, made annually for a distinguished research contribution in computer science by a UK based researcher who has completed up to 10 years of post-doctoral research.

Monoidics Ltd is a participant in the Correct and Efficient Accelerator Programming (CARP) Consortium, along with project partners ARM, Realeyes, Rightware, École Normale Supérieure, University of Aachen, University of Twente, and Imperial College. In 2011 CARP has been awarded a three year Strep European grant worth 4 million euro.

3. References to the research (indicative maximum of six references)

This section refers to the Microsoft Academic Search citation analyser (MAS): http://academic.research.microsoft.com/

Citation Number from Google Scholar.

[P1] J. Berdine, C. Calcagno, P. O'Hearn, Symbolic Execution with Separation Logic. In Proc. Asian Symposium on Programming Languages and Systems, APLAS 2005, 2005, LNCS Vol 3780/2005, pp. 52-68, http://www.springerlink.com/content/p6x5787r63560l88/fulltext.pdf **(184 citations)**

[P2]. D Distefano, PW O'Hearn, H Yang: A Local Shape Analysis Based on Separation Logic. TACAS 2006: 287-302. **(According to MAS, 2nd most highly- cited paper of TACAS'06.) (253 citations)**

[P3] Josh Berdine, Cristiano Calcagno, Byron Cook, Dino Distefano, Peter W. O'Hearn, Thomas Wies, Hongseok Yang: Shape Analysis for Composite Data Structures. CAV 2007, Lecture Notes in Computer Science, volume 4590, pages 178-192. Springer 2007. **(168 citations)**

[P4]. H Yang, O Lee, J Berdine, C Calcagno, B Cook, D Distefano, PW O'Hearn: Scalable Shape Analysis for Systems Code. CAV2008:385-398. **(Most highly- cited paper of CAV'08.)**

(166 citations)

[P5]. C Calcagno, D Distefano, PW O'Hearn and H Yang: Compositional Shape Analysis by means of Bi-abduction.Journal of the ACM, 2011. 73 pages. **(JACM is usually considered the top journal in Computer Science.)**
(183 citations, in conjunction with the conference version of the paper)

[P6] O'Hearn, Peter W., and David J. Pym. "The logic of bunched implications." *Bulletin of Symbolic Logic* (1999): 215-244.
(375 citations)

[P7] Ishtiaq, Samin S., and Peter W. O'Hearn. "BI as an assertion language for mutable data structures." *ACM SIGPLAN Notices*. Vol. 36. No. 3. ACM, 2001.
(501 citations)

[P8] http://www.raeng.org.uk/research/researcher/postdoc/current.htm

This work has been carried out with the support of the following grants:

[G1] Distefano's RAEng Research Fellowship, 2007 to 2012; see [P8].

[G2] Calcagno's EPSRC Advanced Fellowship; EP/C544757/1, 2005 to 2011, £274k, "A Unified Theory of Structured Update".

[G3] EPSRC grant EP/G006245/1, 2009 to 2012, £253k, "Adaptive Heap Analysis", PI: Distefano.

[G4] EPSRC grant EP/H011749/1, 2010 to 2013, £219k, "jStar: making java verification practical", PI: Distefano.

[G5] EPSRC grant EP/E002536/1,2006 to 2009, £357k, "Smallfoot: Static Assertion Checking for C programs"; PI: Calcagno.

[G6] EPSRC grant EP/H008373/1, 2010 to 2012, £3.2m, "Resource Reasoning", PI: O'Hearn.

[G7] EPSRC grant GR/R17034/01,2001 to 2004, £150k, "Local Reasoning about State", PI: O'Hearn.

4. Details of the impact (indicative maximum 750 words)

The research has had an economic impact and an impact on public services. Academics at Queen Mary and imperial looked at a series of fundamental problems in computer science. Their work turned into research that focussed on automatic memory safety verification for critical software. That innovative and breakthrough research had a significant impact on the academic world, and also proved to have enormous industrial potential. Partners across academia and industry collaborated to maximize the commercial and real-world impact of this game changing research, including:

Monoidics: This London-based start-up company was founded in 2009 by Distefano & Calcagno, and has grown to include offices in the UK, USA and Japan. Monoidics markets INFER, an automatic program verification tool which builds on ideas in P1-P5 above. Monoidics has ten employees (8 in UK), seven of whom hold advanced degrees, and a revenue of £588,000 in 2012. Customers include government agencies and industrial clients like Airbus, ARM, Lawrence Livermore National Laboratory, Mitsubishi and Toyota.

Recently Monoidics has been acquired by Facebook to improve the quality and the security of their code base. Within Facebook, the verification technology originally designed by Distefano and Calcagno will impact over one billion users of this very public social networking service. *Head of Facebook London office said to The Telegraph [l9]: "This asset acquisition represents our investment in the quality of our mobile applications platform and also our people, as members of*

Monoidics talented engineering team will join us to work at Facebook's London office once the deal closes," [I2]

"One example of deep computer science being put into practice in the Old Street area is Monoidics, a software verification provider whose technology is based on research at Queen Mary, University of London and Imperial College London." – says InformationAge[I3].

Microsoft: There are separation logic-based tools being implemented in Microsoft's labs in Redmond and Cambridge. One of these, SLAyer, was released in 2011, and has been applied to numerous Microsoft device drivers [I1], where it has found bugs missed by Microsoft's powerful collection of analysis tools. SLAyer builds on the ideas of Calcagno and Distefano found in publications P1 and P2 above. Microsoft's Cambridge lab has at least two full-time researchers and one software engineer working on Separation Logic based Analysis Tools (based on Calcagno and Distefano's work) from PhD's in the UK (one from Cambridge, two from QM). Microsoft has hired upwards of 10 post-docs and interns to work on Automatic Analysis Tools based on Separation Logic.

Several routines of Windows device drivers were analysed with this technique and numerous real bugs were uncovered, and confirmed by the Windows Kernel Team.[I1]. The impact of this work on improved Microsoft Windows reliability is worldwide via release of improved drivers.

Press Coverage: The work of Monoidics Ltd has received extensive press cover: in Japan, a primary market for the tool, see for example [I4]; in UK in the context of the Tech City [I3,I9]; In other international press in context of the CARP project [I5,I6,I7, I8]

5. Sources to corroborate the impact (indicative maximum of 10 references)

[I1] Microsoft, Researcher. http://research.microsoft.com/en-us/um/cambridge/projects/slayer/ Corroborate about impact of separation logic based tools inside Microsoft.
[I2] Head of Facebook London office. Corroborate impact on the use of Monoidics Technology inside Facebook.
[I3] Monoidics in Tech City: http://www.information-age.com/industry/uk-industry/2137348/monoidics-applies-mathematical-analysis-to-software-verification
[I4]Monoidics in Japan: http://itpro.nikkeibp.co.jp/article/Interview/20120926/425324/?utm_source=ITPro総合&utm_medium=twitter
[I5] ARM supports the CARP project 2012: http://www.hpcwire.com/hpcwire/2012-05-29/arm_gets_behind_accelerator_programming_project.html
[I6] ARM supports the CARP project 2012: http://www.eetimes.com/electronics-news/4373938/ARM-European-research-on-GPU-programming-language
[I7] The CARP project: http://www.design-reuse.com/news/29342/carp-correct-and-efficient-accelerator-programming.html
[I8] The CARP project: http://www.eetimes.com/document.asp?doc_id=1266565
[I9] Monoidics and Facebook: http://www.telegraph.co.uk/technology/facebook/10188628/Facebook-buys-UK-startup-Monoidics.html