Institution: The University of Edinburgh

Unit of Assessment: School of Informatics

Title of case study: Automatic detection of defects in multi-threaded enterprise Java codebases

1. Summary of the impact

A spin-out company, Contemplate Ltd, is using advanced static analysis technology in global top-ten investment banks and other clients to discover previously undetected defects in enterprise-scale business-critical multi-threaded Java codebases. The impact is in terms of the benefits delivered to Contemplate's clients by this technology and in terms of the formation and growth of Contemplate as an employer and a successful business.

2. Underpinning research

Underpinning research in the School of Informatics was undertaken during 2008–2009 by a group of researchers under the direction of Don Sannella and David Aspinall, funded by an R&D contract between the University of Edinburgh and ITI Techmedia. The researchers at the University of Edinburgh included Professor Don Sannella, Dr David Aspinall (Reader) and Dr Perdita Stevens (Reader) continuously employed through the REF period, and researchers Richard Adams, Robert Atkey, Brian Campbell, Conrad Hughes, Sam Lindley, and Jaroslav Ševčík working under contract to the University of Edinburgh between 2008 and 2009.

This research grew out of two earlier strands of work undertaken at the University of Edinburgh. The first strand of research was EC- and EPSRC-funded research on proof-carrying code for resource-bounded computation during 2002–2009 [1]. Here, the static analysis algorithms that automatically generate the proofs required for proof-carrying code, *which fail when the required property does not hold*, were seen to hold more immediate promise as an automatic bug-finding tool for software developers, in comparison with the longer-term promise of the proof-carrying approach to security. This built on University of Edinburgh research by Martin Hofmann and David Aspinall during 1999–2002 on type systems for capturing heap space usage of programs.

The second strand of research undertaken at the University of Edinburgh was the research on the semantics of memory models in the period 2005–2008 by Jaroslav Ševčík (at this time, a PhD student in the School of Informatics) and David Aspinall (his supervisor). This work showed that some important program transformations are invalid under Java's memory model and exposed some of the pathologies that make development of shared-memory multi-threaded software so challenging [2].

The underpinning research for this impact case study built on aspects of the earlier work at the University of Edinburgh and the published literature on static analysis and produced novel methods and software for performing efficient and accurate static analysis of Java codebases, and for using the information obtained to detect defects and present visual information relating to multi-threading behaviour [3-6]. The fundamental advances over existing research on static analysis were:

1. extension of ideas that had only been tried on small subsets of Java to the full Java language, including treatment of its extensive libraries;
2. combination of different ideas that had previously been studied in isolation;
3. implementation as prototype-quality software;
4. experimentation and benchmarking on a wide range of open-source codebases;
5. engineering and tuning to achieve a good balance between efficiency and quality of results, taking into account the way that Java is used in practice;
6. discovery that some experimental results reported in prominent published research appear to be incorrect;
7. development of novel visualisations of thread behaviour and interactions via shared objects; and
8. experimentation with application of the same ideas in C++.

3. References to the research

1. D. Aspinall, S. Gilmore, M. Hofmann, D. Sannella and I. Stark. *Mobile Resource Guarantees for Smart Devices.* Proc. Intl. Workshop on Construction and Analysis of Safe, Secure, and Interoperable Smart Devices, CASSIS 2004, Marseille. Springer LNCS 3362, 1–26, 2004. DOI *10.1007/978-3-540-30569-9_1*. Research funded by grant IST-2001-33149 "Mobile Resource Guarantees" under the Global Computing pro-active initiative of the Future and Emerging Technologies part of the European Commission's 5th Framework Programme, value 1.252M€, Jan 2002 – Apr 2005, partners Edinburgh and LMU Munich, coordinator Don Sannella.

2. J. Ševčík and D. Aspinall. *On Validity of Program Transformations in the Java Memory Model.* Proc. 22nd European Conf. on Object-Oriented Programming, ECOOP 2008, Paphos. Springer LNCS 5142, 27–51, 2008. DOI *10.1007/978-3-540-70592-5_3*.

3. D. Aspinall, R. Atkey, C. Hughes, D. Sannella and P. Stevens. *Design Level Technology Study: Reverse Engineering and Refactoring Concurrency Final Deliverable WP2.3.1 D2, Version 1.1.* ITI Techmedia Software Integrity Engineering programme, Dec 2008. Can be supplied on request. This gave rise to 8 "Innovation Disclosures" for ITI Techmedia, as input to potential patent applications, of which 5 relate to the impact in question. An independent written assessment of this work was produced by Roke Manor Research as SIE deliverable WP2.5 D4.2a, which can be supplied on request. Research funded by R&D contract "Design Level Technology Study: Concurrency" from ITI Techmedia's Software Integrity Engineering programme, value £97.9k, Jul–Nov 2008, coordinator Don Sannella.

4. R. Adams, D. Aspinall, R. Atkey, B. Campbell, C. Hughes, D. Sannella, J. Ševčík and P. Stevens. *Program Analysis Demonstrator: Early Alpha Demonstrator WP 2.1.6 D3*. ITI Techmedia Software Integrity Engineering programme, Jan 2009, report and software. Can be supplied on request. An independent written assessment of this work was produced by Roke Manor Research as SIE deliverable WP2.5 D4.2b, which can be supplied on request. Research funded by R&D contract "Program Analysis Demonstrator" from ITI Techmedia's Software Integrity Engineering programme, value £110.8k, Nov 2008 – Jan 2009, coordinator Don Sannella.

5. D. Aspinall, R. Atkey, B. Campbell, S. Lindley, D. Sannella and J. Ševčík. *Annotations for Concurrency: Final Report WP 2.1.7 D2.* ITI Techmedia Software Integrity Engineering programme, Mar 2009, report and software. Can be supplied on request. This gave rise to an "Innovation Disclosure" for ITI Techmedia, as input to a potential patent application. Research funded by R&D contract "Extensions of Program Analysis Demonstrator" from ITI Techmedia's Software Integrity Engineering programme, value £102k, Jan–Mar 2009, coordinator Don Sannella.

6. D. Aspinall, R. Atkey, S. Lindley, D. Sannella and J. Ševčík. *Further Research on Technologies for Concurrency Analysis, Final Report WP4.3.1 D2.* ITI Techmedia Software Integrity Engineering programme, Jul 2009, report and software. Can be supplied on request. Research funded by R&D contract "Further Research on Technologies for Concurrency Analysis" from ITI Techmedia's Software Integrity Engineering programme, value £33.6k, May–Jul 2009, coordinator Don Sannella.

Items 3–6 above were funded by contracts for individual workpackages within ITI Techmedia's £4.3m Software Integrity Engineering programme [E], 2008–2011. These items contain commercially sensitive IP and, if reviewed, should be treated by REF subpanel 11 panel B in the same way as confidential research outputs. The University of Edinburgh held contracts for 15 workpackages for a total of £778.8k, all awarded following a competitive tender process.

Among the above-listed research outputs, references [1], [2] and [6] are most indicative of the quality of the underpinning research.

4. Details of the impact

4.1. Formation of the company

Contemplate Ltd was formed in March 2009 to commercialise the research described above [A,B], and has obtained an exclusive licence to the corresponding IP for this purpose from ITI Techmedia (which owns the IP under the terms of its R&D contract with the University of Edinburgh). The static analysis technology in Contemplate's first tool, ThreadSafe, builds directly on the prototype software developed in the course of the underpinning research. The substantial performance improvements over competing static analysis products offered by this technology has been key to Contemplate's success with clients.

4.2. Reach and significance of the technology

ThreadSafe provides mechanisms for developers of highly sophisticated multi-threaded enterprise Java applications to understand their code and ensure that it is free from errors. Contemplate's initial market focus is in Investment Banking where in-house developed low-latency trading platforms and their front-ends provide banks with competitive advantage. Bugs in these applications can be extremely expensive, whether they cause crashes or deadlocks or unintended or missed transactions. Estimates of the cost of downtime range from US$25k to US$250k per minute. Even more costly is the reputational damage caused by client-visible bugs. Recent high-profile examples in financial services are:

- the cancelled IPO of BATS Global Markets, caused by bugs in its software;
- glitches in NASDAQ's software (reportedly caused by a concurrency bug) which affected the Facebook IPO and led to losses at leading investment banks of more than $500m; and
- bugs in Knight Capital's automated trading software leading to a loss of $440m in 40 minutes.

Controlling and quantifying the risk introduced by the possible presence of such bugs is also growing in importance, with ever-increasing regulation of risk in banking including the US Securities and Exchange Commission's proposed new rules ("Regulation SCI") that regulate the quality of software used by key market participants.

4.3. Details of the problem

Multi-core hardware is becoming mainstream, yet most mainstream software development is currently unable to exploit it. Although Java has been designed to support concurrency from the ground up, most Java programmers are insufficiently versed in the subtleties of concurrent programming, since it is technically demanding and they have not needed to use it. Concurrent programming introduces new classes of highly unpredictable failures, for which current testing methodologies are inadequate. This leads to increasing numbers of bugs, which are difficult to resolve without supporting technology.

4.4. Advantages of the static analysis technology

Contemplate's products, beginning with ThreadSafe, allow companies to reduce costs and minimise operational risk by producing better quality concurrent code. The earlier that problems are identified and fixed, the greater the saving or risk reduction. The tools help software developers eliminate bugs at an early stage, by providing:

- interactive assistance to support good concurrency idioms
- warnings by analysing concurrent interactions in programs as they are written
- visual feedback on the analysis to help the developer understand the tool's findings.

In 2011-2012, Contemplate worked with two of the world's top-ten global investment banks in an Early Adopter Programme in which pre-release versions of ThreadSafe were applied to sample production codebases. These trials were successful, with ThreadSafe automatically revealing a range of previously undetected important defects in these business-critical applications.

4.5. Release of the product

ThreadSafe was first released for sale in October 2012 [F,G,H]. It has been purchased for use in a software project in one of the Early Adopter banks, with negotiations underway with other projects in the same bank, and is undergoing trials in other investment banks and several other companies. Contemplate has agreed a partnership with GrammaTech, a leading US-based static analysis supplier, to make ThreadSafe's analysis engine available as an added-cost plug-in to its CodeSonar product, which targets mainly the embedded systems and aerospace sectors. A reseller arrangement has also been agreed with ArchitectGroup in South Korea for distribution in East Asia.

Contemplate currently employs 5 FTE staff. Strong growth in revenue and staff is predicted [C,D].

5. Sources to corroborate the impact

A. CEO, Contemplate Ltd can provide information about all aspects of Contemplate's business, including copies of purchase orders and invoices to confirm sales. All information about customers and sales is commercially sensitive.
B. Contemplate's business plan, including financial projections, can be supplied on request; this contains commercially sensitive information. This plan has been subject to repeated financial and technical due diligence by Scottish Enterprise in connection with negotiation of Contemplate's IP licence agreement, applications for funding from SMART:SCOTLAND, and investment from the Scottish Investment Bank.
C. Copies of IP due diligence reports commissioned by SMART:SCOTLAND can be supplied on request; these contain commercially sensitive information.
D. A senior member of the High-Growth Start-Up Unit at Scottish Enterprise can provide information about Scottish Enterprise's support for Contemplate. See also *http://www.scottish-enterprise.com/~/media/SE/Resources/Documents/ABC/Contemplate.pdf*
E. A commercialisation advisor to Software Integrity Engineering programme and former chair of the ICT Advisory Group for Scottish Enterprise can provide information about the role of the University of Edinburgh and Contemplate in the ITI Software Integrity Engineering programme.
F. A software developer at Attachmate in Seattle can comment on ThreadSafe from a user's point of view.
G. Press articles on Contemplate and ThreadSafe:
    a. *Scotland on Sunday*, Dec 2012, *http://www.scotsman.com/business/management/software-firm-unveils-program-to-beat-costly-computer-failures-1-2682298*
    b. *CFO Insight*, 1 May 2013, *http://www.cfo-insight.com/risk-management-it/it/software-errors-new-technology-briefing-for-cfos/*
    c. *InfoQ*, Aug 2013, *http://www.infoq.com/news/2013/08/ThreadSafe-Public-Release*
    d. *DevX*, Sep 2013, *http://www.devx.com/Java/contemplate-delivers-threadsafe-java-concurrency-static-analysis-tool.html*
H. Contemplate's website *http://www.contemplateltd.com*: contains press releases concerning product releases and partnership agreements.

Copies of these web page sources are available at *http://ref2014.inf.ed.ac.uk/impact*