Institution: University of York		
Unit of Assessment: 11 - Computer Science and Informatics		
Title of case study: Standardisation and Practice in Safety Critical Systems Engineering		
Period when the underpinning research was undertaken: 2000-2020		
Details of staff conducting the underpinning research from the submitting unit:		
Name(s):	**Role(s) (e.g. job title):**	**Period(s) employed by submitting HEI:**
Rob Alexander	Lecturer	Oct. 2002 to present
Radu Calinescu	Senior Lecturer	May 2012 to present
Ibrahim Habli	Senior Lecturer	Aug 2005 - Dec 2008; Jan-2012 to present
Tim Kelly	Professor	Jan 1998 to Jun 2019
Richard Hawkins	Senior Research Fellow	Nov 2001 – May 2006; Mar 2008 to present
John McDermid	Professor	Mar 1987 to present
Period when the claimed impact occurred: Aug 2013-Dec 2020		
Is this case study continued from a case study submitted in 2014? N		

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

Research at the University of York into safety-critical software development and assurance, particularly assurance principles, has influenced national and international safety standards, guidelines and regulations. These include ISO 26262, the international standard for programmable electronics and software used by all automotive manufacturers. This has led to an impact on industrial practice, in many domains and countries, influencing projects and products with a cumulative value in excess of GBP100,000,000,000. This impact has been achieved through collaborative research projects; direct engagement with system developers and regulators; contributions to standards, guidelines and other public-domain documents; and through education and training.

2. Underpinning research (indicative maximum 500 words)

The High Integrity Systems Engineering (HISE) research group at York has undertaken pioneering research into the development and assurance of safety-critical software-based systems for over 30 years. In the relevant period, HISE's research has included foundational work on assurance principles, assurance cases and safety cases (SC) and the emerging issues of autonomy, including ethics.

HISE's foundational work created and validated the '4+1' principles for safety-critical software assurance which give a sound, system- and domain independent, intellectual basis for assessing safety. The principles are articulated in [3.1] (amongst other publications) and a comparison with a standards-compliance approach is presented in [3.2]. The 4 principles relate to key aspects of the software product and its development, e.g., preservation of design intent through system decomposition; the '+1' reflects the confidence that those key aspects have been achieved and evidenced. Many standards are very complex; these principles give a simpler yet comprehensive and sound basis for analysing existing standards [3.2] and for defining new standards (e.g. 5.4A&5.4B).

HISE has improved SC argumentation with the Goal Structuring Notation (GSN) and has advanced the implementation of SC practice within regulatory structures and processes; [3.3] illustrates this research in the automotive domain. SC are essentially abductive arguments drawing the 'best' conclusions from the available evidence thus their conclusions are always open to challenge. Making explicit confidence arguments (instead of just the '+1' asserted confidence) [3.4] addresses key weaknesses of real-world arguments, such as those for the Nimrod reconnaissance aircraft SC, as identified in Sir Charles Haddon-Cave's review, and is also important in work on autonomy.

Autonomous systems, and the technology they incorporate e.g., machine learning (ML), challenge conventional certification approaches that expect and rely upon transparency of behaviour, determinism and high levels of confidence *prior to* deployment. HISE has undertaken

work to identify how safety assessment can be applied to such systems, and how it needs to be adapted to the particular characteristics of autonomous systems, since the early 2000s.

Traditionally SC (and the underpinning safety analysis) assumes that risks can be assessed statically, *prior to* system deployment. This assumption was always questionable and is clearly invalid for systems which can adapt their behaviour in operation, e.g., through use of ML. Thus, there is a need to assess risk dynamically (i.e., to understand how it changes in operation) and to reflect this in the system SC. Our work in this area, including [3.5], addresses the fundamental challenge of assuring self-adaptive software and producing dynamic assurance cases, showing how an SC can be updated as systems evolve and adapt in operation.

Finally, HISE is addressing wider concerns relating to autonomy and ML, including their consequences for accountability as well as safety [3.6] through the Assuring Autonomy International Programme (AAIP). Such issues need to be addressed in both autonomous and advisory (recommender) systems. This work [3.6] brings ideas from practical ethics to bear on a healthcare example, showing how the ethical concept of an epistemic condition for accountability can shed light on moral responsibility when decisions are transferred from clinicians to machines.

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

3.1. **Hawkins, R.**, **Clegg, K.**, **Alexander, R.** and **Kelly, T.**, 2011, September. Using a software safety argument pattern catalogue: Two case studies. In International Conference on Computer Safety, Reliability, and Security (pp. 185-198). Springer, Berlin, Heidelberg. [Paper in the leading international conference on computer system safety.] DOI: doi.org/10.1007/978-3-642-24270-0_14

3.2. **Hawkins, R.**, **Habli, I.**, **Kelly, T.** and **McDermid, J.,** 2013. Assurance cases and prescriptive software safety certification: A comparative study. Safety Science, 59, pp.55-71. [Paper in one of the two leading international journals on system safety.] DOI: doi.org/10.1016/j.ssci.2013.04.007

3.3. Birch, J., Rivett, R., **Habli, I.**, Bradshaw, B., Botham, J., Higham, D., Jesty, P., Monkhouse, H. and Palin, R., 2013, September. Safety cases and their role in ISO 26262 functional safety assessment. In International Conference on Computer Safety, Reliability, and Security (pp. 154-165). Springer, Berlin, Heidelberg. [Paper in the leading international conference on computer system safety.] DOI: doi.org/10.1007/978-3-642-40793-2_15

3.4. **Hawkins, R.**, **Kelly, T.**, Knight, J. and Graydon, P., 2011. A new approach to creating clear safety arguments. In Advances in Systems Safety (pp. 3-23). Springer, London. [Paper in an annual, international conference on computer system safety.] DOI: doi.org/10.1007/978-0-85729-133-2_1

3.5. Calinescu, R., Weyns, D., **Gerasimou, S.,** Iftikhar, M.U., **Habli, I.** and **Kelly, T.,** 2017. Engineering trustworthy self-adaptive software with dynamic assurance cases. IEEE Transactions on Software Engineering, 44(11), pp.1039-1069. [Paper in one of the leading software engineering journals.] DOI: https://ieeexplore.ieee.org/document/8008800

3.6. **Habli, I.**, Lawton, T., **Porter, Z.** 2020, 'Artificial intelligence in health care: accountability and safety', Bulletin of the World Health Organization, vol. 98, pp. 251-256, DOI: http://dx.doi.org/10.2471/BLT.19.237487 [A high impact journal in the health sector.]

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

Sectoral Impact

1) Automotive

York's research has had an impact on automotive standards [5.1] and on major automotive suppliers JLR (an Original Equipment Manufacturer (OEM)), Bosch (a tier 1 supplier) and HORIBA-MIRA (an independent assessment house) [5.2]. York (Kelly) was invited onto the British Standards Institution (BSI) committee developing the automotive standard ISO 26262 specifically to draft the Part 10 material on SC development [5.1B]. ISO 26262 is the international standard for programmable electronics and software used by all automotive

manufacturers. York and JLR were the founding members of the Motor Industry Software Reliability Association (MISRA) SC Working Group. York's work, with the other working group partners [3.3], has led to our recognition as a significant contributor to the resulting MISRA SC Guidance [5.1A, 5.2D]. York has also contributed to the work on "Safety of the Intended Function" which addresses some of the issues arising from autonomy [5.1C, 5.2D]. The global automotive market is circa USD2,000,000,000,000 with the UK alone having a turnover of GBP82,000,000,000 and employing 168,000 people directly. ISO 26262 and the SC guidance apply mainly to the electronic elements of vehicles; these represent more than one third of the vehicle value. As an example, the SC work now used by Bosch "established release conditions for such systems … spread across worldwide locations" [5.2A]. Ongoing work is also influencing German standards [5.2A], British and International standards and guidelines [5.2B]. The work has had impact beyond York's direct collaborators [5.3]; for example, the German Pegasus project has adopted York's SC approach for Safety Argumentation [5.3A] and Uber uses GSN for their autonomous vehicle SC [5.3B].

2) Defence and Aerospace

York's work has influenced UK Defence Standards [5.4] and the UK Defence and Aerospace sector [5.5]. York's 4+1 principles [3.1] are the basis for Defence Standards 00-056 and 00-055 [5.4A, 5.4B]; York (McDermid) was the lead author on 00-056 [5.4A] and a contributor to 00-055 [5.4B]. York's research on risk, confidence, and compliance arguments directly shaped the UK Military Aviation Authority's (MAA) Manual of Air System Safety Cases [5.4C] to support MAA Regulatory Article RA1205. The Defence Science and Technology Laboratory (Dstl) confirm the scale of impact of this work on projects of value circa GBP20,000,000,000 per annum [5.5C] and attests to its wider impact: *"I doubt that there is a safety case worldwide that has not been influenced by this work"* [5.5C]. Also, referring to the AAIP, the international impact is clear: *"this initiative provided me support at a NATO research meeting on managing and assuring autonomy"* [5.5C]. The industrial impact is reflected in [5.5A, 5.5B] (including influences on software development beyond the work referenced above [3.1-3.6]). For example, [5.5B] identifies influences on a range of certification problems including data-intensive systems (see also [5.6A]); it also identifies two programmes "over which York's research has provided influence and direction" with a value of over GBP2,000,000,000.

3) Other/Non-Domain-Specific

York has played a leading role in running the Safety Critical Systems Club (SCSC) a UK-based, but international, professional community. The SCSC has produced a range of guidelines which are used widely in the UK and internationally [5.6]. For example, the Data Safety Guidance [5.6A] addresses the problems of assuring safety of data intensive systems, adopting the '4+1' principles (albeit citing another York paper, rather than [3.1]); there is a similar influence on the Service Assurance Guidance [5.6B]. York (Alexander) led the development of work on autonomous systems [5.6C] which builds on the work of the AAIP and is already beginning to be used.

The Global Mining Guidelines Group (GMG) is a worldwide consortium of all the major players in the mining and quarrying industry, including manufacturers, e.g., Caterpillar, operators, e.g., Rio Tinto, and government agencies, e.g., the Government of Alberta. The sector increasingly uses autonomy and is developing guidelines on the design, assurance and use of autonomous systems in mines and quarries (there are significant safety risks, with individual machines carrying circa 500t of material). New guidelines have been developed with input from York (McDermid) on safety of autonomous systems [5.7], building on the principles developed by the AAIP and "not only helped GMG complete and publish the guideline, but it also helped the mining industry align on the topic" [5.8B] since the presence of operators, manufacturers and regulators in the GMG inevitably brings tensions. Further, although recent, the "guideline is being used in a college curriculum and in an application for regulatory approval" [5.8B].

Impact on Industry 4.0 by York (through the AAIP) is further illustrated through work in support of the Global Manufacturing and Industrialisation Summit in the area of autonomy, where "the University of York will continue to bring a unique and impactful understanding and approach to

the challenges we face in assuring the safety of complex manufacturing systems in the future" [5.8C].

In healthcare the CONSORT-AI (Consolidated Standards of Reporting Trials–Artificial Intelligence) Guidelines [5.9] provide "a new reporting guideline for clinical trials evaluating interventions with an AI component" which directly references our work on accountability and safety [3.6].

The impact of the HISE group's work internationally (through one of the Fraunhofer Institutes) can also be seen [5.8A]; for example, this shows "GSN … implemented into Fraunhofer's safety engineering tool (safeTbox) … successfully applied assurance cases" and cites influences on other standards in Germany and internationally, e.g. the Structured Assurance Case Meta-Model (SACM), and highlights the influence of the AAIP on projects in Germany.

Pathways to Impact

A) Direct Engagement with solution developers and regulators
HISE has worked with a wide range of developers of safety-critical systems, e.g., BAE Systems including funding for PhD students [5.5B], Bosch (through the AAIP) [5.2A], and JLR including EngD projects [5.2C]. Work with regulators includes the MAA [5.4C] and ongoing work through the AAIP embraces the Health and Safety Executive (HSE), and the Medicines and Healthcare products Regulatory Agency (MHRA) and the Care Quality Commission (CQC) in healthcare.

B) Standards and Guidelines
The references (section 5) contain 12 standards or guidelines which build on York's work; the letters cite a further 5; York's research is (and has been) influential in a number of standards activities, e.g. ISO 21448 and ISO/TR 4804 [5.2A], VDE-AR-E 2842-61, ISO/IEC AWI TR 5469, and the Structured Assurance Case Metamodel (SACM) [5.8A].

C) Collaborative Research
In the period HISE has been involved in 11 major projects of value circa GBP24,200,000 (spend in the period of circa GBP17,200,000) with 27 industrial partners. This has led to direct impact in 8 countries. Ongoing research projects such as the ICON project with two Fraunhofers, Fraunhofer IESE (Institute for Experimental Software Engineering) and Fraunhofer IKS (Fraunhofer Institute for Cognitive Systems), [5.8A] is building on the work on dynamic risk [3.5].

D) Education and Training
York has taught an MSc programme in Safety Critical Systems Engineering (SCSE) since 1995 and has also delivered continuing professional development (CPD) courses direct to industry; the teaching is research-led. The York SCSE CPD team has taught 666 individuals from 111 organisations in 12 countries over the REF period; in addition, there has been in-house training for specific organisations, e.g. to 300 staff at NHS Digital [5.10A] and 100 at JLR [5.2C]. The SCSE MSc is increasingly introducing material on autonomy, based on the work of the AAIP, and specialist CPD courses are also being presented on the challenges of safety engineering for autonomy. Whilst only recently started, this material has already been presented to more than 70 engineers, from 45 organisations and 3 countries, and to 80 within the NHS [5.10A].

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

5.1. A) Guidelines for Automotive Safety Arguments, ISBN 978-1-906400-23-1, September 2019; B) Automotive Standards and Guidelines: ISO 26262-10:2018(en) Road vehicles — Functional safety — Part 10: Guideline on ISO 26262. International Standardization Organization; C) ISO/PAS 21448:2019. Road vehicles — Safety of the intended functionality, International Standardization Organization.

5.2. Letters from the automotive sector: A) Director Vehicle Systems Safety, Bosch GmbH; B) Director of Sectors, BSI; C) Functional Safety Technical Specialist, Jaguar Land Rover (retired); D) Chief Engineer, Functional Safety, HORIBA MIRA Ltd.

5.3. Other automotive use: A) Pegasus Safety Argumentation (see: https://www.pegasus projekt.de/files/tmpl/pdf/PEGASUS%20Safety%20Argumentation.pdf); B) Uber ATG Safety Case (see: https://uberatgresources.com/safetycase/gsn);

5.4. Defence Standards: A) DEF STAN 00-056: PARTS 1 and 2 Revision 5, 2017; B) DEF STAN 00-055: PART 1 Revision 4, April 29, 2016; C) UK Military Aviation Authority, Manual of Air System Safety Cases, 2019.

5.5. Letters from defence and aerospace sector: A) Technology Manager, BAE Systems (Rochester); B) Engineering Capability Director, BAE Systems (Warton); C) Senior Fellow, Dstl.

5.6. SCSC Reports: A) Data Safety Guidance (Version 3.0), SCSC, Jan 2018, ISN 978-1981662463; B) Service Assurance Guidance (Version 1.0), SCSC, Jan 2020; C) Safety Assurance Objectives for Autonomous Systems V2, Ref: ISBN: 978-1654029050.

5.7. Functional Safety for Autonomous Equipment Sub-committee (GMG), Guidelines for applying functional safety to autonomous mining systems, 2020

5.8. Letters/emails showing wider impact: A) Executive Director, Fraunhofer IESE; B) Technical Editor, GMG; C) Principal Policy Consultant, IfM Education and Consultancy Services.

5.9. Liu X, Rivera SC, Moher D, Calvert MJ, Denniston AK. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI Extension. BMJ. 2020 Sep 9;370.

5.10. Letter from Principal Safety Engineer, NHS Digital.