The Safety-Critical Systems Club Newsletter

Safety Systems

Vol 28 No. 2 - May 2020

THE SHAPE OF THINGS TO COME

Can we trust pandemic models?

NEED TO KNOW?

The criticality of organisational communication

DRIVING AMBITION

How smart are our motorways?

For everyone working in Systems Safety

SCSC

scsc.uk

SCSC Publication Number: SCSC-158

Contents

SCSC

FOR EVERYONE WORKING IN SYSTEM SAFETY

www.scsc.uk/sss

SSS'21

SCSC 29th Annual Symposium

Details of this and other events at: www.scsc.uk

www.scsc.uk

THE SAFETY-CRITICAL SYSTEMS CLUB

29th Safety-Critical Systems Symposium

9-11th February 2021, Bristol, UK

Invitation to submit an abstract

The Safety-Critical Systems Symposium in 2021 (SSS'21) returns to the Bristol Marriott Royal Hotel, Bristol, UK. The event comprises three days of presented papers, including keynote presentations, submitted papers, 3-minute 'pitches' and a poster session. Papers, pitches and posters will be selected based on abstracts.

The Symposium is for all of those in the field of systems safety, including engineers, managers, consultants, students, researchers and regulators. It offers wide-ranging coverage of current safety topics, focussed on industrial experience. It includes recent developments in the field and progress reports from the SCSC Working Groups. It covers all safety-related sectors including aerospace, defence, healthcare, highways, marine, nuclear and rail. Topics of interest are:

Accident Analysis
Agile Methods
Artificial Intelligence
Argumentation Notations
Assurance Cases
Autonomy
Change and Evolution
Cloud Systems
Data in Safety Systems
Domain Knowledge
Human Factors
Independent Review & Audit
Internet of Things
Lessons Learnt
Machine Learning
Methods and Tools
Management and Oversight
Model-Based Development
Modular Assurance
Multicore / Manycore
Ontologies / Formalisms
Regulation
Robotics and Automation
Safety Analyses
Safety Culture
Safety Management
Safety Practice and Process
Security-Informed Safety
Service-Oriented Safety
Software
Standards and Guidance
Systems-of-Systems
Training and Training Data
Validation and Verification

Authors for papers should submit a title and 200-word abstract to **mike.parsons@scsc.uk** by **30th June 2020**. Successful authors will then be asked to submit their paper by **30th October 2020**. Papers will be reviewed, and comments fed back during November 2020.

Pitch and poster authors should submit a title with abstract and/or short paper to **mike.parsons@scsc.uk** by **30th October 2020.**

The hard-copy proceedings book will be available at the event. The book, posters, videos and papers will be posted online at scsc.uk.

For **information, exhibition** and **booking** enquiries please contact: Alex King, Dept of Computer Science, University of York, Deramore Lane, York, YO10 5GH. Tel: 01904 325402; **alex.king@scsc.uk**

For **technical aspects, abstract and paper submissions**, please contact: Mike Parsons, **mike.parsons@scsc.uk**

Editorial

After over a month of lockdown, it is frightening to think how much Coronavirus has changed our world in such a short space of time. In the intervening time between two newsletter editions, we've witnessed unprecedented events that have affected everybody's lives in some way, and likely to have long lasting impacts at both a local and global level. There is also much novelty in our experiences now: who would previously have commonly used terms such as: 'self-isolation', 'furlough', 'social distancing' and 'flattening the curve', or practised those careful avoidance manoeuvres on our streets and supermarkets?

These novel circumstances have brought the art of safety assessment right into the public domain. Venturing outside is now done only after careful assessment of the potential hazards and available risk controls and mitigations. However, we find ourselves with an unfortunately tricky form of risk assessment. In practical terms, the likelihood of contracting the disease in lockdown is very low, but the criticality, death is, of course, very high. This equates to that awkward corner of the risk matrix; should some hazards therefore be eliminated entirely, or can we countenance an ALARP discussion on whether to order a take-away, open the mail or pet the next-door neighbour's cat?

Also, a critical aspect of the crisis has been the data on which everyone, including governmental policy makers are basing their safety-related decisions. There has been no shortage of data from a wide variety of sources: epidemiological data of course, but also other critical data such as the acceptable social distance to maintain and the length of time the virus can survive on different surface types. These are critical influencing factors when applied to large populations, but the data has been conflicting at times and there are often polarised 'expert' opinions on how it is to be interpreted.

To help make sense of this uncertain landscape, the SCSC has set up a Working Group that is meeting weekly to see what a systems and assurance view of the situation can bring. Visit the working group page scsc.uk/gv to find out more, and for details of how to join the group.

This edition covers a diverse range of topics, but they share a common theme: how organisations themselves approach and manage risk is equally important to managing the technical system risks. In our first article, Professor Harold Thimbleby, discusses his concerns with the reliability of the epidemiological modelling currently being undertaken to combat Coronavirus. John Ridley discusses the organisational approach to risk management adopted by the Highways Agency in developing Smart Motorways. Andrew King then discusses how tool-based GSN enhancements can help an organisation manage complex safety cases. Nicholas Hales looks at the organisational structures in place during the Chernobyl incident and assesses these from a data perspective. We conclude with insights from the Service Assurance Working Group on how organisations can assure Services.

There is a summary report of SSS'20 and I have great pleasure in including the transcript of Tim Kelly's entertaining and thought-provoking after-dinner speech.

Our 60 second interview is with Professor Erik Hollnagel.

Stay safe!

Paul Hampton
SCSC Newsletter Editor
paul.hampton@scsc.uk

In Brief

Self-driving car dataset missing labels for pedestrians and cyclists

A popular self-driving car dataset for training machine-learning systems – one that's used by thousands of students to build an open-source self-driving car – contains critical errors and omissions, including missing labels for hundreds of images of bicyclists and pedestrians.

Machine learning models are only as good as the data on which they're trained. But when researchers at Roboflow, a firm that writes boilerplate computer vision code, hand-checked the 15,000 images in Udacity Dataset 2, they found problems with 4,986 – that's 33% – of those images. *nakedsecurity.sophos.com*

'Of course it could happen again': experts say little has changed since Deepwater Horizon

A massive deepwater oil spill is nearly as likely today as it was in 2010, experts warn, 10 years after the disastrous explosion of BP's rig in the Gulf of Mexico that caused an environmental catastrophe.

The blowout killed 11 workers and spewed 4m barrels of petroleum into the ocean for 87 days before it could be capped. *theguardian.com*

Coronavirus: Can we trust the data?

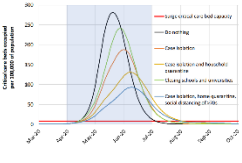

Barely a day goes by in the lifecycle of the coronavirus pandemic without a new model or analytical tool aiming to chart the spread or predict the outcome. Health authorities in different nations also release their own figures, not necessarily covering the same period. These are just a few of the sources used by politicians and journalists worldwide to talk about the number of cases and deaths, but the true totals are a mystery. *news.sky.com*

Boeing's 'culture of concealment' led to fatal 737 Max crashes, report finds

A "culture of concealment", cost cutting and "grossly insufficient" oversight led to two fatal crashes of Boeing 737 Max aircraft that claimed 346 lives, a congressional report has concluded. *theguardian.com*

Drones to deliver NHS supplies to IoW

A Windracer Ultra, will be used to carry NHS supplies to the Isle of Wight. 4 autonomous flights per day over a fixed route are expected. *bbc.co.uk*

How Reliable are Pandemic Models?

An urgent proposal to improve pandemic modelling

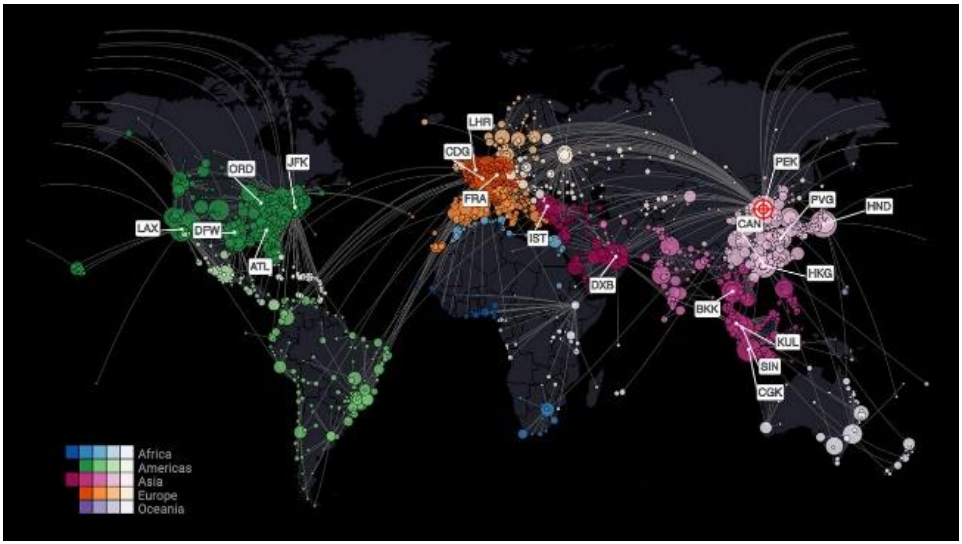

Government policies to manage the COVID-19 pandemic affect the safety, well-being and economy of the UK. Public policy decisions are informed from the use pandemic modelling. How reliable are pandemic models? How dependable are the data and the computer programs that underpin them? Professor Harold Thimbleby discusses his concerns on current practices, and presents an urgent call to action to the software safety community to help address the issues he raises.

Pandemic modelling

Government planning for responding to the COVID-19 pandemic has been based on epidemic modelling. Typically, epidemic modelling tries to make predictions of infections and death rates from mathematical models operating on medical and national data, such as demographics, social habits, and what is known about the infective agent. Of course, computers are essential to manage the data, simulate

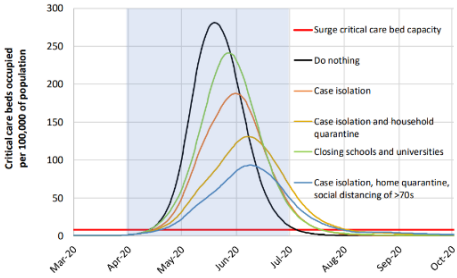

The quality of current pandemic models is not open to scientific scrutiny. The models are either not peer reviewed, or have been peer reviewed very casually.

or solve the models, and to present results. The UK's socially disruptive and economically costly isolation, is based on epidemiological models suggesting that the NHS would be overwhelmed without it.

It was a surprise to discover that the influential Imperial College model [1], which had a central role in early COVID-19 planning, was developed from an undocumented C program [2]. There is a chain of references starting from [1] showing the "same" program was in use for different diseases and different countries, in a research programme going back at least 15 years earlier [3]. A very old, undocumented program written in C cannot be reliably modified.

The epidemiological modelling literature ignores programming, and how hard good programming reliably is.

A classic epidemiology review paper [4] says, "we use the words 'computational modelling' loosely," and then, curiously, exclusively discusses mathematical modelling, as if it is unaware that programming is an essential part of modelling, and no less complex than the mathematics. Without good software engineering — including correct data management, version control, appropriate numerical methods, dependable programming, etc — the models cannot be used to reliably inform public policy.

A recent systematic review [5] of published COVID models, covering individual diagnosis and prognosis in clinical care, including apps and online tools, *completely* ignores model implementation. Apps don't work unless they have working programs! It should be noted that flowcharts and other implementations of models, which the review considered, must be designed as carefully as computer programs: they *are* programs, but intended for direct human use (which raises human factors issues).

Very few epidemiological papers provide any description of their software, let alone access to any software. Those that do, provide access to essentially undocumented code, that is, code with trivial comments.

In short: despite its critical nature, epidemiological modelling used to inform public policy has a very naïve attitude to software. The quality of the current pandemic models is not open to scientific scrutiny, and implies that the models are either not peer reviewed, or have been peer reviewed very casually.

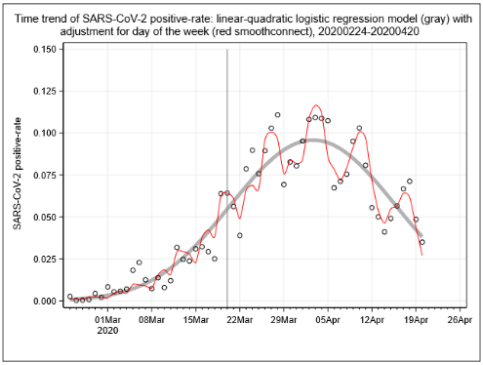

Why professional software engineering is required

Playing the piano is easy if you do not worry about quality, but playing the piano well is very difficult. We can all hear the difference when a wrong piano key is pressed, or is pressed out of time, but programming is different.

A mistake in a program may not be noticed for a very long time, if ever. It is therefore very easy for programmers to over-estimate their programming skills because they cannot see their errors — in fact, poor programmers are less able to find their own errors, and, in a vicious cycle, they therefore further over-estimate their skills.

It seems likely that epidemiology research labs give programming tasks to people who do not realise how incompetent they are at programming; indeed, probably nobody else in their lab recognises the problems with their programming. It seems to work, so they think it must be working correctly.

Software Engineering is the discipline of programming well, and is therefore the approach that epidemiologists should be taking when developing computer models.

A quick summary of Software Engineering would cover at least the following:

- Define requirements: how reliable does code have to be at what cost?
- Correct by construction (formal methods etc)
- Using dependable programming languages
- Tool use
- Defensive programming
- Version control, repositories, configuration, open source
- Rigorous testing
- Good documentation
- (at least: demonstrate independent people can use the code)
- Reusing quality-assured solutions
- Simplicity
- Explicit compliance with appropriate standards
- Effective teamwork to cover the skills required.

This list is merely an outline curriculum. Sir Tony Hoare FRS noted as far back as 1980 [6] that "in any respectable branch of engineering, failure to observe such elementary precautions would have long been against the law."

The point is: the relevant skills for turning epidemic modelling into reliable software to use those models cannot be learned quickly, even if one has the aptitude. In particular, epidemiologists alone cannot be expected to have the relevant skills.

Getting professional software engineering to where it matters

Unless professional software engineering practices are used, computer-based models, and epidemiological models in particular, and results from running them, must be considered unreliable, and unreliable in unknown ways. Unfortunately, results from poor computer models have been published and acted on without adequate scrutiny of the code in use. The accepted *laissez faire* approach is inadequate for use in any modelling informing national policies or clinical treatment.

This raises the question of how to improve.

Professional Software Engineering is essential, but how can it be applied or made available? Given the evident absence of mature software engineering awareness and skills in leading epidemiology research, what can be done?

An analogous situation arises in ethics. There are some sorts of experiments and methods that are ethically unacceptable, but few people have the ethical expertise to make mature judgements — particularly when it comes to assessing their own work. Misuse of data, exploiting vulnerable people, and not obtaining consent are typical problems. Universities, research labs, national funders, and others therefore require Ethics Boards, who review the ethical quality of proposed research. Medical journals will not publish research that has not had appropriate ethical review.

Journals should start requiring not just "reproducibility" but actual evidence of independent reproduction of results, particularly for pandemic modelling.

Quality software engineering requires an explicit proactive approach [7], including compliance monitoring. Establishing Software Engineering Boards seems the best way to make rapid improvements in software quality, without imposing onerous processes and succumbing to "the best is the enemy of the good" pitfalls.

The Software Engineering Boards would authorise as well as provide advice to guide the detailed implementation of high-quality programming. Active, senior, professors of software engineering should be on these Boards; this is not a job for people who are not qualified in the area and not actively connected with the true state of the art. There are many high-quality software companies (especially those specialising in safety-critical areas like aviation) who would be willing to help. Many university departments of computer science have qualified and experienced software engineers who can also help.

It is important to note that merely making code open source or otherwise available for scrutiny (for instance, providing it in Supplementary Information with papers) is *not* sufficient to ensure quality. Undocumented, badly-written code is inscrutable, even if its original programmers think otherwise. In fact, making poorly developed code available merely enables others to reproduce and propagate the poor work.

Software Engineering Boards must be used to ensure that critical code is of high standard, and must provide assurances for scientific papers, Government reports, and so on, that the software is of appropriate quality. Only then is it worth sharing and using more widely.

Note that an essential feature of good science is *reproducibility* — indeed, if results cannot be reproduced, what has anything contributed? When a modelling paper presents results from a model, it is very important to reproduce those results *without using the same code*. Better still, research should be reproduced without sharing libraries or APIs (for example, results from a model using, say, the language R, might be reproduced in a different language, say, Mathematica). Reproducing the same results relying on the same codebase tells you little. The more independent reproductions of results the greater the evidence for belief in the implications. Given the practical difficulties of applying professional software engineering, scientific journals should immediately start encouraging reproduction of prior results, particularly for pandemic modelling where research insights affect national and world populations.

In most science it is straight forward to ensure results are "reproducible": enough details have to be provided for other researchers to reproduce the results. In software, this is not sufficient. Merely providing the code and data does not mean that others can reproduce the results. The code may not be portable, the code may require specific versions of compilers,

specific hardware, and so on. All the code might be there, but the configuration details may be missing. And so on. We need higher standards for code *if it is trying to make a scientific contribution*: it needs to *actually* be reproduced by independent parties before we can be sure it is truly reproducible. Software Engineering Boards could confirm that independent researchers have reproduced the results claimed. Conversely, the "Software statement" in a paper could just say no independent party has confirmed the results.

Just as medical papers today require conflicts of interest statements, data availability statements, and ethics board clearance, we must move to papers being required to include Software Engineering Board clearance statements. Like Ethics Boards, Software Engineering Boards might become, or be perceived as becoming, onerous and heavy handed. It is essential that Software Engineering Boards have (and perhaps are chaired by) experienced, professional software engineers to avoid this problem.

Unlike Ethics Boards, which provide hands-off oversight, Software Engineering Boards should provide professional advice where necessary, perhaps providing training or be hands-on actually helping develop appropriately reliable software. This is not free, so clearly all research, particularly medical research, should be required to include adequate funding for professional software engineering.

Can it be done?

It is striking how the NHS has very rapidly changed everyday practice to use remote consultation because COVID-19 has put pressure on it to continue work without risking cross-infecting staff and patients. Likewise, COVID-19 should have put pressure on modelling papers informing public policy to be rigorous, given the seriousness of the choices facing Government.

Therefore, as a matter of national priority, Software Engineering Boards could be rapidly established to provide direct access to mature software engineering expertise for both researchers and for journals seeking competent peer reviewers. In addition, particularly during a pandemic or other emergency, as now, Software Engineering Boards would provide direct access to their expertise for Governments and public policy organisations.

In the longer run, it is likely that a professional organisation (such as the UK's Royal Academy of Engineering) would provide a professional framework for Software Engineering Boards. In particular, until there are national qualifications, nobody — certainly nobody without professional training in software — really knows just how bad they are at software engineering.

It is hoped that this article encourages grass roots and organisational responses to the problems identified.

Conclusions

Current epidemiological practice is an alarming basis for public policy.

Mature software engineering methodologies are essential for reliable modelling, including all critical programming. We must prioritise getting appropriate professional software engineering skills and resources to bear on the COVID-19 pandemic without delay. We should also be planning to help improve modelling to better manage any future crises. (it's sobering to recall that austerity was inspired by easily avoidable bugs in an Excel spreadsheet [8]).

Since mature software engineering skills are not widely available, Software Engineering Boards, analogous to Ethics Boards, could be rapidly established to provide the necessary up-to-date knowledge and oversight for critical projects (like pandemic modelling) to succeed. Software Engineering Boards promise an effective way out of the amateurish use of computers in all research.

Indeed, epidemic modelling is not the only area affected by poor programming, and Software Engineering Boards would have far-reaching benefits across all fields. Elsewhere [9] we have shown how unreliable software undermines the safety of medical systems more generally, putting patients at risk. Current medical device regulation does not address these problems — arguably, Software Engineering Boards should have a broad remit to identify and address software problems anywhere.

References

[1] N. M. Ferguson, D. Laydon, G. Nedjati-Gilani, *et al*, "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand," http://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf, 16 March 2020, accessed 22/04/2020

[2] N. M. Ferguson, http://twitter.com/neil_ferguson/status/1241835454707699713, accessed 22/04/2020.

[3] N. M. Ferguson, D. A. T. Cummings, C. Fraser, J. C. Cajka, P. C. Cooley & D. S. Burke, "Strategies for mitigating an influenza pandemic," *Nature* 2005; 437:209–214. DOI 10.1038/nature04017

[4] H. Heesterbeek, R. M. Anderson, V. Andreasen, S. Bansal, D. De Angelis, C. Dye, K. T. D. Eames, W. J. Edmunds, S. D. W. Frost, S. Funk, T. D. Hollingsworth, T. House, V. Isham, P. Klepac, J. Lessler, J. O. Lloyd-Smith, C. J. E. Metcalf, D. Mollison, L. Pellis, J. R. C. Pulliam, M. G. Roberts, C. Viboud & Isaac Newton Institute IDD Collaboration, "Modeling infectious disease dynamics in the complex landscape of global health," *Science* 2015; 347(6227):aaa4339. DOI 10.1126/science.aaa4339

[5] L. Wynants, B. van Calster, M. M. J. Bonten, *et al*, "Prediction models for diagnosis and prognosis of covid-19 infection: Systematic review and critical appraisal," *BMJ* 2020; 369:m1328. DOI 10.1136/bmj.m1328

[6] C. A. R. Hoare, "The Emperor's Old Clothes," *Communications of the ACM* 1981; 24(2):75–83. DOI 10.1145/358549.358561

[7] B. Shneiderman, "Opinion: The dangers of faulty, biased, or malicious algorithms requires independent oversight,'" 2016; 113 (48)13538–13540. DOI 10.1073/pnas.1618211113

[8] T. Herndon, M. Ash & R. Pollin, "Does high public debt consistently stifle economic growth? A critique of Reinhart and Rogoff," Cambridge Journal of Economics, 38(2):257–279, 2014. DOI 10.1093/cje/bet075

[9] H. Thimbleby, *Fix IT: How to solve the problems of digital healthcare*, Oxford University Press, 2020.

Harold Thimbleby

Harold Thimbleby is See Change Fellow in Digital Health at Swansea University, Wales. He has recently completed *Fix IT: How to solve the problems of digital healthcare*, which will be published by Oxford University Press later in 2020. Harold won the BCS Wilkes Medal and has previously been a Royal Society Wolfson Research Merit Award Holder. He has published over 300 refereed papers, and has given over 600 presentations around the world.

How Smart Are Our Motorways?

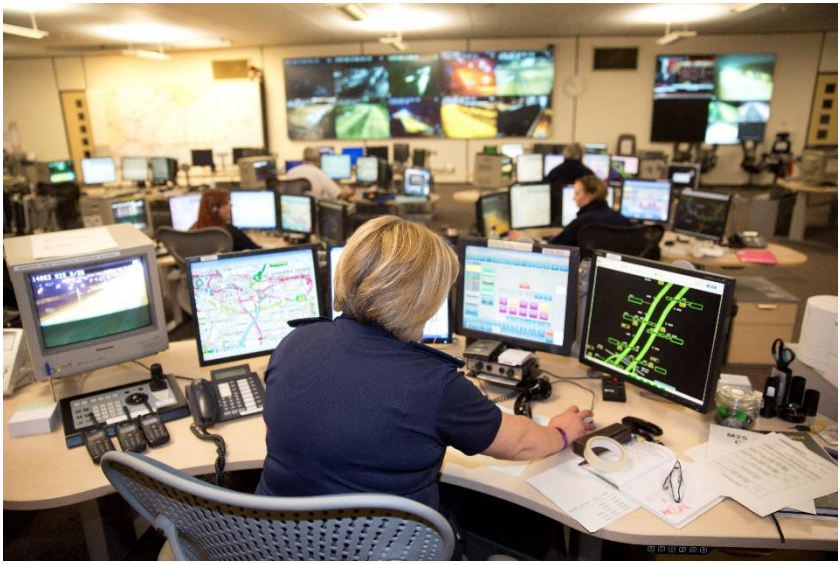

John Ridgway discusses his concerns with the safety of Smart Motorways—motorways that use the hard shoulder to increase traffic throughput during busy periods. He discusses the issues and limitations of the principles adopted by the Highways Agency in assessing Smart Motorways' safety risk, and discusses the wider implications of their use.

In view of recent adverse publicity [1] surrounding the UK government's Smart Motorways, I thought it might be useful to draw attention to two articles posted in this newsletter back in 2010. The first [2], an article written by myself, discussed ethical issues that may arise when setting safety targets. Particular reference was made to Active Traffic Management (ATM), a pilot for Smart Motorways hard-shoulder running. The second [3], an article written by David Boulton of Arthur D. Little's Risk Practice, sought to outline the approach taken by the Highways Agency (HA) in the setting of safety targets and addressed some of the issues I had raised. To summarise, I had concerns that the HA's adoption of the Globally At Least Equivalent (GALE) principle could lead to ethical difficulties, particularly if this were to entail the abandonment of the principle of reducing safety risk As Low As Reasonably Practicable (ALARP). A main purpose of [3] was to point out that the application of GALE did not preclude the aspiration to reduce risk further in a manner akin to ALARP—though I believe I would be correct in saying that any such endeavour on the HA's behalf would still not be as a result of a perceived legal imperative.

I shall not seek to review in detail the two articles referenced above. However, since I believe recent events have vindicated the concerns I had raised back in 2010, I shall expand upon those concerns below. As was the case with my initial article, I trust there will be those amongst the readership of this newsletter who will be in a position to correct any errors made and, indeed, redress any imbalance or misrepresentation of which I may be guilty.

The Safety Argument as I Understand It

In order to understand how the Highways Agency could arrive at the conclusion that Smart Motorways with live hard-shoulder running would be acceptably safe, one has to appreciate the following:

Regarding the setting of safety targets for the Smart Motorway system (and indeed for its ATM predecessor), the HA maintained that the Health and Safety At Work etc. Act 1974 (HSAWA) did not apply. Consequently, they felt under no legal obligation to reduce safety risks to a level in accordance with the ALARP principle. This position is clarified in [4].

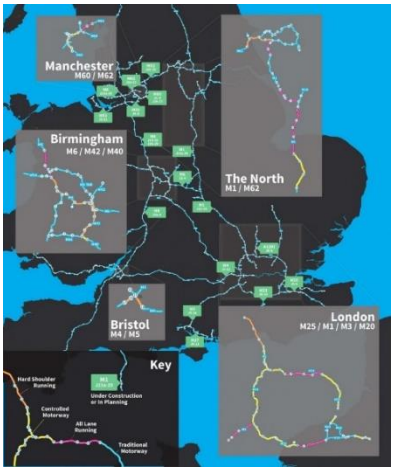

Instead of using ALARP to set safety targets, the HA looked towards the GALE principle, in which safety risks associated with a system should be Globally At Least Equivalent to the risks encountered prior to the introduction of the system (or upgrade thereof). This idea was taken from the UK rail industry (which, in turn, obtained the idea from the French railway's concept of *Globalement Au Moins Equivalent* (GAME)). When assessing whether Smart Motorways were likely to comply with the GALE principle, the HA took into account that Smart Motorway schemes require the simultaneous introduction of a previously developed technology known as Controlled Motorways, in which variable speed limits are used to achieve a 'calming' effect on traffic flow. Since previous Controlled Motorway systems had been demonstrated to reduce levels of speed-related deaths and serious injury, this could be taken into account when making a GALE calculation. Accordingly, an increase in the rate of death and serious injury resulting from the removal of the hard shoulder safety zone, could be offset by a presupposed equivalent reduction in the level of speed-related death and serious injury on the main carriageway.

The safety dividend resulting from the introduction of a Controlled Motorway scheme can be estimated based upon the accident statistics of previously operated Controlled Motorway systems (e.g. the M25's scheme, incorporating Motorway Incident Detection and Automatic Signalling System (MIDAS) queue protection). On the other hand, the estimation of the safety penalty resulting from the removal of a hard shoulder safety zone would have to be based upon a theoretical safety model. Such a model-based risk assessment was conducted, and the HA concluded that the projected increase in deaths and serious injury due to the removal of the hard shoulder safety zone, was no greater than the reduction that employing Controlled Motorway traffic calming could be expected to bestow.

The system was therefore predicted to meet the GALE target, i.e. a Smart Motorway operation (necessarily entailing Controlled Motorway traffic calming) would be at least as safe as

a stretch of motorway lacking any traffic management system (to be accurate, I should point out that an adjustment was factored into the GALE calculation, which I believe provided for a proposed 10% safety improvement).

In all of the above, it should be appreciated that a fully functional Smart Motorway scheme is not feasible without the installation of Controlled Motorway technology (traffic calming is a necessary precursor for the opening of the hard-shoulder to traffic). However, a Controlled Motorway system is itself perfectly feasible as a stand-alone traffic management scheme (e.g. as previously operated on the M25).

The Concerns

I had always thought that the Highways Agency's approach towards the setting of the Smart Motorway safety targets was problematic.

Firstly, it is ethically dubious to allow for an increase in deaths and serious injuries due to taking one action just because one can expect a compensating decrease in deaths and serious injuries due to taking another. For example, a car manufacturer that introduces air bags into the design of one of its vehicles could not use such a safety upgrade as an excuse for removing an engine fire management system from the same vehicle. By the same token, the removal of the hard shoulder as a safety zone for breakdowns cannot be justified by referencing the posited lives saved as a result of the simultaneous introduction of Controlled Motorway traffic calming. The fact that Smart Motorways require Controlled Motorway traffic calming is circumstantial and covers up the fact that key aspects of Smart Motorway operation are inherently dangerous.

GALE is supposed to apply to system installation and upgrades rather than to dynamic transitions from one mode of operation to another. Nevertheless, there is an argument that the GALE baseline for a Smart Motorway system, with live traffic running on the hard shoulder, should be a Controlled Motorway with the hard shoulder closed to traffic, since the latter provides the context in which hard-shoulder running is dynamically introduced. For its GALE baseline, the HA used a conventional stretch of motorway prior to the installation of any traffic management system. However, in those circumstances where an upgrade had actually been from a Controlled Motorway system to a full functionality Smart Motorway system, would the HA still insist on using the conventional motorway as the baseline for its GALE calculation?

Secondly, it is politically naïve to assume that the families of victims killed as a result of the removal of the hard shoulder safety zone, would be satisfied with the argument that such deaths can be condoned simply because it may be presupposed that speed-related deaths on the main carriageway of the same stretch of motorway will have been avoided as a result of separate measures taken (i.e. Controlled Motorway traffic calming). In the eyes of the public, the perception of safety for a given action will be established by considering those deaths and serious injuries that have occurred as a result of that action; there will be no allowance

Statistical arguments carry no weight with the bereaved, and the public mood is always going to be established by the factual and not the counterfactual.

made for the hypothetical deaths or serious injuries that may have been avoided by taking any separate, albeit attendant, actions. Such statistical arguments carry no weight with the bereaved, and the public mood is always going to be established by the factual and not the counterfactual.

Finally, the Highways Agency's assertion that the HSAWA does not apply when deciding upon acceptable levels of safety for safety-related systems installed on the motorway was always questionable. Certainly, such legislation would apply to the designers, suppliers and would be expected, under law, to reduce safety risks As Low As Reasonably Practicable. However, the Highways Agency (now Highways England) was the design authority for all traffic management schemes installed on England's motorway network and, as such, they stipulated the safety requirements. If a system supplier's customer (the government) is not applying ALARP when stipulating the system safety requirements, the possibility then arises that an accident resulting from system failure could place the supplier on the wrong side of the law, if the safety integrity of the system had been engineered only to meet the government's own safety targets. Furthermore, as the ultimate design authority for traffic management schemes, the HA would surely have a duty of care to ensure that whenever two GALE compliant alternatives are reasonably available, then the lower risk alternative is taken. A motorway subjected only to Controlled Motorway traffic calming provides a lower safety risk alternative to full functionality Smart Motorways, so the Highways Agency had it as a GALE and ALARP compliant alternative. Denying the legal applicability of ALARP enabled the HA to justify not taking that option. A more germane justification may be that Controlled Motorways technology on its own (whilst alleviating congestion) would not enable the HA to meet its motorway network capacity targets. To meet those targets required the adoption of the less safe system.

Some Implications

Any suggestion that the deaths and injuries caused as a direct result of the removal of the hard shoulder as a safety zone could not have been foreseen is unsupportable. They were not only foreseen by the HA, they were analysed and quantified as part of the GALE calculations.

The GALE calculation alluded to above, was based upon an assumed separation of 500m between Emergency Refuge Areas (ERAs). Having calculated that a Smart Motorways scheme employing ERAs spaced at 500m meter intervals would just meet the GALE safety target, it is inconceivable that GALE compliance could be assumed for a system employing ERAs every 1.5 miles (even working with a GALE-10% tolerance). The spacing of the ERAs was a critical factor in the risk model used by the HA's safety consultants (as, indeed was the installation of technologies to detect hard shoulder breakdowns). I have no knowledge of whether the safety case was reviewed when Smart Motorways was rolled out with increased ERA spacing and, if so, what the logic would have been for concluding that the safety case could still be made.

Current statements from politicians, along the lines of 'one death is one too many' or 'we need the motorways to be as safe as possible' are inconsistent with the safety policies that had been adopted by the HA throughout the specification, development and operational phases of Smart Motorways. These are political sound-bites aimed at an electorate; they have not previously informed the safety management process and it would be naïve to believe that they will do so in the future.

It is ethically dubious to allow for an increase in deaths and serious injuries due to taking one action just because one can expect a compensating decrease in deaths and serious injuries due to taking another.

The HA's insistence that they had no legal requirement to apply the ALARP principle when deciding upon safety targets is now likely to be challenged in the courts. I understand that at least one bereaved family plans to charge Highways England with corporate manslaughter for failing in its duty of care [5]. This duty is premised in law upon the application of ALARP, therefore I assume that adherence to the GALE principle would not be a sufficient defence (notwithstanding a 10% margin). Even if my concerns regarding the HA's approach to GALE calculations should be ruled as immaterial, it is difficult to see how extending the spacing between ERAs could be seen as an ALARP approach to safety risk management. Highways England may still be able to make a case based upon an improvement in risk efficiency (i.e. the number of deaths for a given traffic volume) but I have no idea where this would stand legally.

References

[1] See, for example, https://www.bbc.co.uk/news/uk-51236375

[2] J. Ridgway, "The Bogeyman – Fact or Fiction" Safety Systems, Vol 19, Number 2, 2010, https://scsc.uk/r111.2:1

[3] D. Boulton, "GALE or ALARP: Which to Choose" Safety Systems, Vol 20, Number 1, 2010, https://scsc.uk/r114.1:1

[4] M. Halbert, S. Tucker, "Risk Assessment for M42 Active Traffic Management, Developments in Risk-based Approaches to Safety", Springer, London, 2006, https://scsc.uk/r6/2:1

[5] https://www.transport-network.co.uk/Widow-to-sue-Highways-England-over-smart-motorway-safety/16123

John Ridgway, Retired

John Ridgway is a retired traffic systems consultant who led the software development team responsible for the M25 MIDAS computer control system. Subsequently, he undertook a functional safety management role whilst working for a contractor commissioned to provide a computer control system for the Highways Agency's M42 Active Traffic Management system. All opinions expressed in this article are entirely his own and do not reflect the views of any previous employers.

SCSC
FOR EVERYONE WORKING IN SYSTEM SAFETY

THE SAFETY-CRITICAL SYSTEMS CLUB, 155th Seminar:

Safe Use of Multi-Core and Manycore Processors

Thursday 24 September, 2020 - London, UK

This seminar will consider how to use processors with multiple cores in a way that safety can be assured, and such that the resulting system can be certified against industry standards and guidelines.

Critical systems, such as those used in avionics, are moving from single core processor to multiple core (multi-core) processor architectures. This enables a reduction in size, weight and power and the use of common processing platforms, reducing costs and allowing common spares. Software certification policies and guidance are currently evolving as experience is gained with creating certification evidence for multi-core processor architectures.

There are some unique challenges for using multi-core processors in certified platforms and these will be highlighted and discussed, including the investigation of multi-core interference channels.

This seminar will be held in central London at the Radisson Blu Edwardian Grafton, 130 Tottenham Court Rd, Bloomsbury, London W1T 5AY

(A multi-core processor is typically made of several independent processor cores on the same chip, connected through an on-chip bus. Manycore processors are specialist multi-core processors designed for a high degree of parallel processing, containing a large number of simpler, independent cores.)

This seminar is relevant to all those involved with Multi-core and Manycore processing systems, including software and hardware developers, certification bodies and sector regulators.

Safe Use of Multiple Core Systems in Critical Applications

https://SCSC.uk/e638

WWW.SCSC.UK

Complex Safety Cases – Enhancing The Goal Structuring Notation

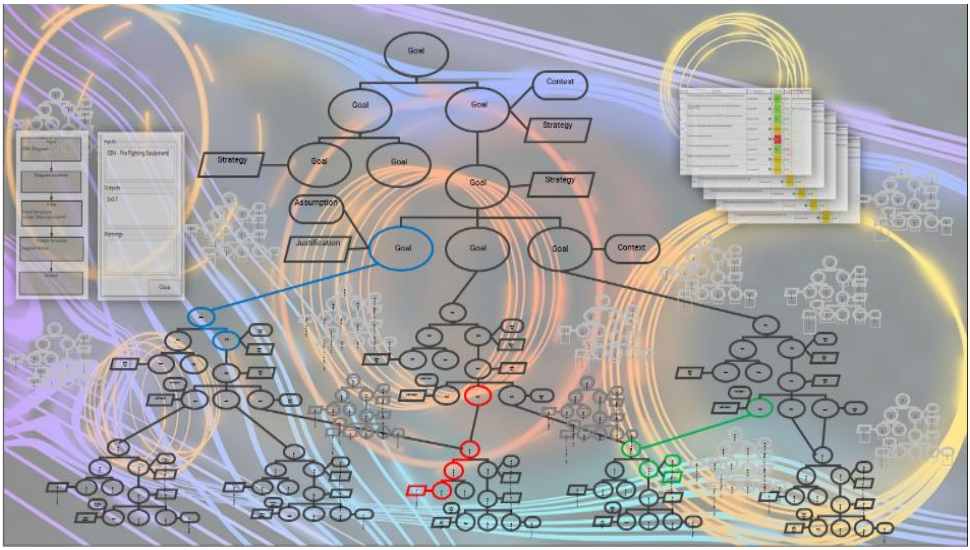

Andrew King discusses how the Goal Structuring Notation (GSN), when encapsulated into a safety modelling tool, can enhance the ability of the GSN to manage large and complex safety cases. He shows how a software architecture can be built around GSN and how, using reference and data fields, that architecture can be used to manage modular safety cases.

Goal Structuring Notation

A safety argument is a representation of a number of safety claims and comprises goals, assumptions, justifications, strategies and solutions (evidence). The GSN captures these elements in a graphical notation and provides a clear representation of complex arguments with their supporting evidence. The GSN makes explicit the reasoning behind a safety case and thus makes it easier for stakeholders to understand. For complex arguments, the GSN breaks down the argument into manageable sections and shows how the safety argument has been constructed to meet the top-level claim.

A large GSN can, however, consist of thousands of elements with complex interdependencies making tracking progress and managing changes a demanding and time-consuming task.

When a challenge to a complex safety case is made, the owners must first identify all the areas that are affected by the challenge, assess if the challenge is valid, and then make changes to rectify the problem and restore the validity of a safety case. These new changes must be readily apparent to the regulator if the safety case is to be reviewed and approved.

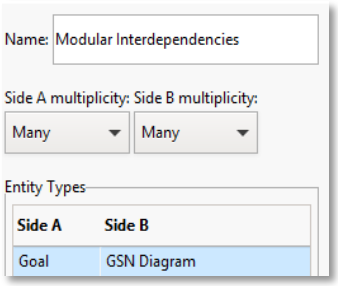

the duty holder found it almost impossible to keep track of developments and identify areas of weakness using the GSN

When working with the Senior Engineering Officer at RAF HQ 22 Group some years ago, it quickly became apparent that the scope and depth of the Organisation's safety case was so extensive, that it was very difficult to assess the GSN for progress visually – platforms, equipment and operating bases stretched the GSN structure into the distance – and the duty holder found it almost impossible to keep track of developments and identify areas of weakness using the GSN. This has been an experience that I have seen duplicated in many organisations and has been one of the main tenets behind the development of a new safety modelling tool.

Building a safety modelling tool where GSN elements have embedded reference and data fields, will deliver rapid search, analysis and reporting for large and complex safety cases.

GSN Model Architecture

Large and complex safety cases are easier to manage when they are broken down into subsystem safety cases or modules. This helps in the identification and isolation of areas where the change applies and also allows the development of the safety case by different teams. To break down a safety case into sections, it is necessary to build a safety case architecture. This now common term is defined as:

"the high-level organization of the safety case into components of argument and evidence, the externally visible properties of these components, and the interdependencies that exit between them". [1]

For example, a safety case covering the operation of a maritime oil tanker will be divided into separate safety cases for separate systems, such as the navigation equipment, ship construction, firefighting, cargo handling etc, with a high level safety case that covers the general operation of the vessel. The individual safety case modules will have their own clear boundaries but will also have interdependencies, such as ship construction and the ability to fight fires (thermal insulation between accommodation and cargo). In the architecture definition, equal weight is placed on the dependencies between safety case modules as on the safety components themselves. So dependencies must also be recorded as part of any interface definition, including arguments requiring support from other modules and reliance on objectives, evidence and context presented elsewhere.

The need to identify objectives, evidence and context for each module or component and the interdependencies between them is crucial to a successful modular safety case. Kelly outlines extensions to the GSN to support the concept of modular safety case construction [1].

In a safety modelling tool, this can be achieved through the use of a reference field sitting behind the 'Goal' entity. If a goal is supported by a GSN diagram elsewhere in the safety case, this interdependency can be imbedded in the reference field for that goal, and subsequently, can be shown to exist through analysis of the safety case (see analysing the GSN structure). In a safety modelling tool, if a goal is used to support an argument in several GSN structures, it is simply dragged onto the GSN diagram with the knowledge that if a change is made to that goal in one section, it will automatically be reflected wherever else that goal is used.

Analysing the GSN Structure

To effectively analyse a GSN Structure it is necessary to identify not only the goals, strategies, assumptions, justifications and solutions, but also the links (support and context arrows) down through sub-goals to solutions and the evidence. Thus GSN elements and GSN arrows need to be created as entities in the software model. Each entity can then have its own set of properties, either inherent or added by the user and which contain key data for analysis. References will hold the links to other entities and will appear in the entity and the entity being referenced; an important factor when tracing up the GSN from evidence, or down from goals.

In order to analyse a GSN structure effectively, you need to be able to examine each element in the GSN and follow the links up from a piece of evidence or down from a high-level goal. A query can achieve this and other analytical functions by identifying a start point (input) and through a number of steps to arrive at an analytical result (output). The following functions are required to analyse a GSN safety argument:

- Identifying GSN structures that contain certain GSN entities – Where is the evidence?
- Identifying certain entities shown in a GSN structure – What are the strategies?
- Finding all the children of a GSN entity – Where are the GSN structures?
- Finding all the parents of a GSN entity – In what section is this GSN structure?
- Following arrows down from an entity - What are the sub-goals and solutions?
- Following arrows up from a GSN entity – What are the high-level goals for this evidence?
- Filtering by GSN entity – Only show the contexts for this diagram.
- Filtering GSN entities by selected criteria – date, time, key word etc.
- Select the type of GSN entity – Only look at these entities.

With this capability, the query can really start to address the business questions many owners of large safety cases will want answered:

- What areas of the safety case are affected by this challenge?
- What changes have been made over a certain period?
- Who is responsible for those changes?
- What other areas are affected by this change?
- What needs to be completed in the safety case?

An example is given here. A safety case manager wants to know what has changed in the last 24 hrs.

Using queries, he can select the GSN diagram or diagrams he is interested in, highlight the contents, and then set the period – in this case 1 day.

The query output will then show the elements that have been changed during the last 24 hrs.

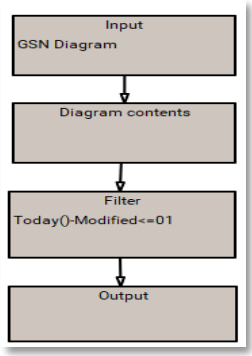

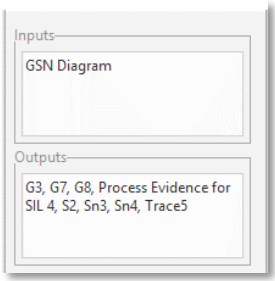

The query could be targeted at key words if assessing a challenge to the safety case, or at a status level or section owner. Queries are a very powerful tool for conducting GSN analysis of a safety case, and they can be exported easily for use in other safety cases.

Presenting the information from analysis is equally important for assessment and understanding. Matrices provide a tabular output for the analysis that can show references and data fields as columns in a table. A matrix uses the analysis queries to build the table and can provide simple (single query) output, or an in-depth (multiple queries) analysis.

Enhancing The GSN Structure Visually

GSN is not just about making the reasoning behind safety cases explicit, it is also about a visual representation of the individual elements of the safety case. Assessing large and complex safety cases is challenging, even using GSN, but this can be made simpler when the GSN structures are broken down into more manageable sections. The HQ 22 Group GSN structure is an example of a complex and extended safety case that proved difficult to navigate and manage. By breaking down the high-level goals into a number of smaller groups, and with hyperlinks from one group to another, navigation and visual assessment becomes much easier. In addition, the ability to tab through a number of selected GSN structures and view two or more structures alongside each other, provides more clarity and aids assessment.

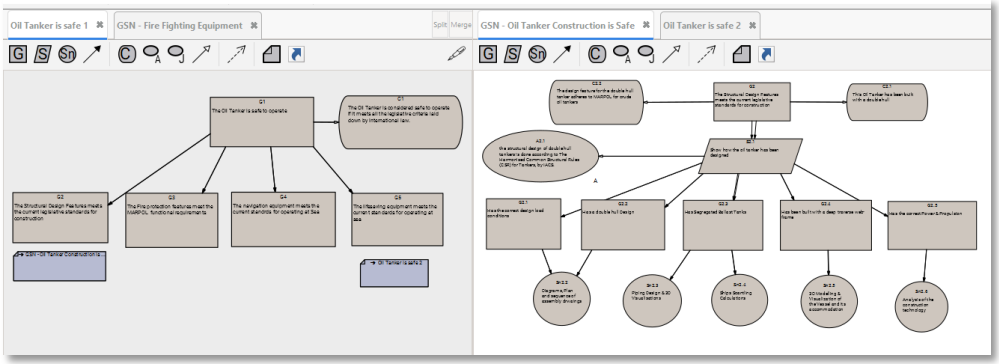

Safety cases are no longer just an "into service" item. Throughout the life cycle of the system numerous challenges and changes will occur and when a challenge is made against a safety case, the applicable areas must be identified before rectification can take place.

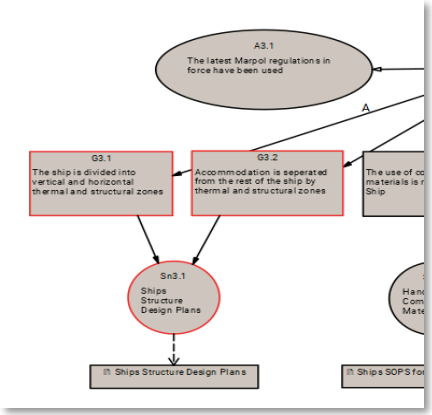

Where goals and evidence are used in a number of GSN structures, they must all be identified if the situation is to be resolved satisfactorily. Queries can also help in this area and can be used to first identify those areas that are affected by the challenge and then colour code them so they stand out from the rest of the structure. This makes identification much easier and additional colour coding can be applied to new GSN structures that have been put in place to show what changes have been made to address the challenge.

Summary

Today's safety cases are becoming ever larger and more complicated. The GSN captures safety case arguments in a graphical notation that provides a clear representation of complex arguments with their supporting evidence. A large GSN can, however, consist of thousands of GSN elements with many interdependencies that make tracking progress and managing change exceedingly difficult. Moreover, as safety cases embrace through-life management, many organisations are finding it more and more difficult to keep track of the required changes.

By building the GSN in a safety modelling tool where elements are created as entities within that model, and adding reference and data fields to those entities, it becomes possible to handle modular safety cases more easily and efficiently. Queries and matrices provide a powerful analytical function that can track development, identify change requirements and report on the actions taken to meet any challenge. By enhancing the visual representations of GSN structures through easier navigation and colour coding, managers and system engineers can quickly identify key elements without the need for lengthy and time-consuming reviews.

References

[1] Dr T Kelly. "Managing Complex Safety Cases", https://www.users.cs.york.ac.uk/tpk/sss03.pdf accessed, 21/04/2020

Andrew King – Diametric Software

Andrew worked in the RAF as the Safety Manager Executive for Air Traffic Control and Air Surveillance & Control specialisations before becoming a client director for a leading information management & computer modelling company in the private sector. He now works as the Operations Director for Diametric Software.

SCSC
FOR EVERYONE WORKING IN SYSTEM SAFETY

This seminar is relevant to safety engineers and safety consultants who have to perform analysis of systems. It will also be useful for safety auditors and assessors who may have to interpret or review analyses in these new methods.

Safety Analysis Techniques of the Future

https://SCSC.uk/e654

WWW.SCSC.UK

THE SAFETY-CRITICAL SYSTEMS CLUB, 156th Seminar:

POSTPONED: New Safety Analysis Techniques

TBD - London, UK

This seminar will be postponed until the autumn.

This seminar will look at emerging, novel and recently established techniques for analysing aspects of safety systems: their overall properties, their architecture and interactions, their software, hardware and their data.

Safety systems require analysis for potential failures that can lead to hazards. Traditional techniques tend to have limited applicability in today's world of highly complex, interconnected, continually updated systems. Learning systems bring new analysis problems as the faults may be contained in the training data rather than the system itself.

Techniques such as STAMP/STPA will be covered as well as emerging methods for analysing hazards in context (Environmental Hazard Analysis). The uses and abuses of Bow-Ties will be covered. The Functional Resonance Analysis Method (FRAM) will be explained. Model-based safety assurance will be considered. Tools and techniques for analysing service aspects (e.g. based on Swimlane Diagrams and Business Process Model and Notation, BPMN) and data safety will also be covered (e.g. use of data FMEA).

There will be a range of speakers covering different techniques. A wrap up session at the end of the day will discuss the most promising contenders.

Speakers include:

Waleed Chaudhry, EDF - "MBSE for Safety Assurance of COTS devices with embedded software"

Chris Harper, Bristol Robotics Laboratory - "Environmental Survey Hazard Analysis"

Mike Parsons, SCSC - "Data FMECAs"

Mark Sujan, Human Reliability - "FRAM TBC"

Simon Whiteley, Whiteley Aerospace - "STAMP/STPA" TBC

TBC, "Service Bow-Ties"

Formalising Communication On Potentially Catastrophic Safety Projects

Nicholas Hales examines the retrospective and thereby hypothetical use of portions of a seven layer data safety model in the design, implementation and operation of the Chernobyl RBMK reactor, and how that would have made avoidance of the disaster far more probable. Nicholas proposes a model for the safety issues in all potentially catastrophic projects, stretching from Board Room to safety experts and operators.

On Saturday 26 April 1986, the No. 4 nuclear reactor in the Chernobyl Nuclear Power Plant exploded. The reactor had entered a meltdown situation which was entirely preventable. Like many catastrophic accidents, a whole series of oversights, mistaken behaviours, design faults and communication failures led to the accident.

The main data related reasons are as follows:

1. Operational problems in the earlier small prototype were not advised to the Chernobyl project manager. (Data missing)
2. Disagreements among experts in the reactor design agency about the safety of the design were not resolved and the fact that disagreements existed was suppressed. (Data suppressed)
3. The Chernobyl reactor was so large that what was going on in one part, about which there was sensor-supplied understanding, did not necessarily reflect what was happening in another area of the reactor. (Data not available)
4. Tests to be conducted went ahead late at night after management with safety responsibilities, that should have been consulted, had gone home. (Data maintaining safety processes ignored)
5. The shutdown process was much slower than necessary, so that the material on the tips of the control rods actually allowed acceleration of the reactions for a few critical seconds when emergency shutdown of the runaway reaction was absolutely essential. (Reported design and operation data ignored)

This article examines the retrospective, and thereby, hypothetical use of portions of a seven layer data safety model in the design, implementation and operation of the Chernobyl RBMK reactor, and how that would have made avoidance of the disaster far more probable. In doing so, a model for the safety issues in all potentially catastrophic projects is proposed, stretching from Board Room to safety experts and operators, and which is compatible with actual engineered projects. The comprehensive source document for much historical fact comes from "Midnight in Chernobyl", [1]. The recommendation is for openness about safety data and principles, both between members of an enterprise and stakeholders in associated enterprises, notwithstanding the need for secrecy in cyber-attack capability etc. In short, communication is what it is all about, but safe communication.

Layers previously not considered formally safety-critical are engaged in the on-going safety process, namely layers 7-6, running from Enterprise to Organisational Unit as illustrated in Figure 1. Layer 5, known as the 'Optimising' layer in the original Data Safety Guidance version 1.3 document [2], is also reviewed. The safety issues are isolated by considering the idea of a vertical tranche in an organisation, like an annex, such that the ordinary working of the Enterprise is unaffected by the necessity for the higher two data safety layers 7 and 6 to be involved in safety-critical data issues without all working becoming absurdly cumbersome. The use of it brings a quasi-military approach to command structures with feedback, as in modern military command structures. It is not dissimilar to how machines operate, where we humans input commands at a high level and expect that our intent will not alter, no matter how deep into the machine the command goes, via various translations to

machine code and binary numbers, to communication with another machine and the sending of the data back up to the Human Computer Interface of another machine. Thus, the data model provides a template for process and procedure dissemination, command, maintenance process approval and design, whenever critical safety is involved.

The question of how to engage responsible management is answered. The top layers 7-6 are considered part of the safety related aspects. This then assigns responsibility ultimately, and in all cases, to the Enterprise. This is akin to how a Field Marshall has responsibility for the behaviour of sub-ordinates, such as a Colonel who gets his regiment captured. It thereby becomes a recommended model for financial Enterprises such as banks, where a lack of concern can lead to catastrophic consequences. An example of this is a Gaussian related theory refinement developed by Merton and Scholes, which won them the Nobel Economics prize in 1997. One year later, in 1998, their bank, the ironically named Long Term Capital Management,(LTCM), collapsed, taking a lot of Wall Street with them, because they had taken massive risks based on the infallibility of their model [4]. In this latter case the ultimate Enterprise is the financial regulating authority, which should have taken evidence from the doubters that enabled it to take a view on how much risk LTCM could take.

To recall, the seven layer data model that appeared in the first version of the Data Safety Guidance document, [2] is shown in Figure 1:

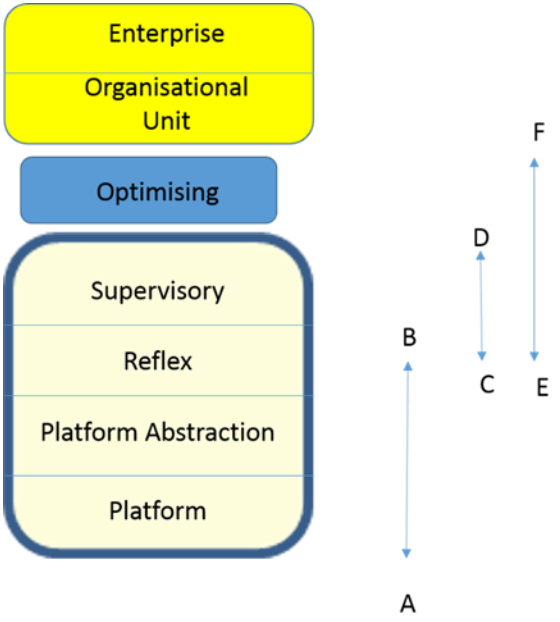

Figure 1

In the figure, as given in the original Data Safety Guidance document, line A-B represents the maximum jurisdiction of SILs 3 and 4, line C-D represents the maximum jurisdiction of SILs 1 and 2 and line E-F represents the combined SIL in a Computer Based Technology and/or human system. The proposition is that, in the context of data, that can simply be a piece of paper passed down through a hierarchy, some level of safety criticality must be assigned. Thus, in some future enquiry into a disaster, the degree of effort made to maintain safety integrity even up to the managerial layer can be asserted through the provision of evidence. It is not easy to prove 'best efforts' by all those involved in each of the Enterprises that makes up a project like Chernobyl, but the complex paths safe data needs to travel from initiation of a project at Government level, to the operational level many years later, are far more visible and thereby better assured than under some current business practices where management, having delegated, wishes not to be troubled by voiced concerns. The Boeing 737 MAX programme being one such where airline pilots and test pilots raised issues [5].

In starting a project, to minimise investment that may result in no return, some Enterprises involved may be linked to others in a minimal form or only using one or not all layers, until the full commitment is developed from established interactions. It should always be the case though that, initially, an Enterprise at layer 7 is created. People can be layers; the important principle being that, as in the 7 layer electronic model, they must be able to supply guaranteed data and receive dependent data safely. The principles can be applied at any abstraction level in safety engineering. It should be seen as a method of maintaining rigour in all aspects of safety related system production.

Due regard to data safety using this model should, as well as defining management responsibility, also ensure management awareness of the safety commitment required. The intention is that arguments over responsibility will not occur as the evidence will be available. That evidence should show that the full enterprise from layer 1 to layer 7 was cognisant of the risks and who was responsible so that all personnel would know exactly who to communicate with whenever issues arose.

It is possible to envisage situations where the Enterprise Layer has not been created. Then, the Enterprise considering setting up an Enterprise at a sub-ordinate layer (the Soviet Government in the Chernobyl case), would have communicated with experts, (from all-encompassing nuclear issue scientists at Sredmash, discussed later), who in time would remain consultants to an, as yet, not established Enterprise, Chernobyl, dedicated to the construction of the reactors. Thus, communication and the tasks of assuring data exchange is safe, begins with the Soviet Enterprise layer 1 communicating with Sredmash at layer 3 Optimising layer and above, without finalised evidence of the need for an Enterprise. Just as with staff turnover in safety cells in industry, there would need to be assurance that data safety expertise from this Enterprise was passed safely to the new Optimising, Organisational Unit and Enterprise layers in the new Enterprise that was to construct and operate Chernobyl.

None of what is written here in any way detracts from the principles of the latest edition of the Data Safety Guidance [3] document. It is intended to broaden the scope out to encompass the wider risks, where a narrow engineering view of software and electronics may lead to inadvertent exclusion of those risks.

The definitions of the model and their relation to the nuclear command structure in 1980s Russia are examined next. These are base starting point definitions, but as this theory progresses, the definitions may alter:

Layer 7 – The 'Enterprise' layer

The enterprise layer is the corporate entity, which in the case of the Soviet Union under the Communist Party, can be taken as the entire country as it saw itself as having a common purpose. The layer is defined as being responsible for the planning and execution of large scale changes to the infrastructure, responding to changes in legislation, setting and maintaining standards, procedures and competency requirements.

In terms of a communist monolith, the changes to the infrastructure and responding to changes in legislation equate to influences outside the enterprise forcing competition to produce cheap electricity, in other words, other countries. Setting and maintaining standards, procedures and competency requirements, more than educational and experience requirements at this level, implied membership of the communist party. This of course meant some talent would never achieve greatness for those who were considered less than committed to the party. The cracks then begin to show because data has a source and a sink and if the source is less than competent, the sink, whether human or machine, will not operate to best intent, though one may get away with it for a while.

Layer 6 – The 'Organisational Unit' layer

This layer is described originally as being responsible for delivery of the planned service. In the communist world this naturally describes the "Government", as described in the Cast of Characters in [1]. The layer is described as not playing any part in the day to day running of the system, whatever it may be. Of course, in authoritarian societies that is exactly what happens and unfortunately the second level of cracks start to show as authoritarian means an understanding of the enterprise layer above but very little hands on understanding of the problems those delivering these dreams, in the lower layers of the system, have to deal with. The Data Safety Guidance further describes this layer as likely to become "involved in the short term operation of the system in response to a serious incident that causes substantial impact on the delivery of the system". In the case of the explosion in reactor 4 at Chernobyl, this definition precisely describes the substantial impact on the Soviet Enterprise intent, to create machines that a perfect society could rely on to run smoothly and be able to outcompete the West. Hence what was required, was much closer safety data definition and communication, but that did not happen.

Layer 5 - The 'Optimising' layer

The optimisation layer is described as the most sophisticated control layer. The optimisation layer respects the performance and safety constraints of the underlying system and the information contingency plans. Crucially, this layer is described by the words "information demands on the optimising layer are high requiring a full understanding of the underlying system, the planned service and contingency plans." This, in Chernobyl terms, represents "The Nuclear Experts" as also given in the Cast of Characters in [1]. They are spread over a number of organisations illustrated in Figure 2.

The Chernobyl Disaster

Below are some of the involved characters and Enterprises as defined by this article:

- Anatoly Aleksandrov – Chairman of the Soviet Academy of Sciences, responsible for nuclear technology development. – **Classed as an Enterprise** and also Director of the Kurchatov Institute – **Classed as an Enterprise**
- Nikolai Dollezhall - director of The Scientific Research and Design Institute of Energy Technology (Russian acronym NIKIET) – **Classed as an Enterprise**
- Victor Brukhanov – Director of Chernobyl plant – **Classed as an Enterprise**
- Efim Slavsky - Ministry of Medium Machine Building (Russian acronym 'Sredmash') – **Classed as an Enterprise**
- Leningrad Nuclear Plant (1st RBMK design) – **Classed as an Enterprise**
- Leonid Brezhnev - The Soviet Union – **Classed as an Enterprise**

So here we have one engineer**ed** Enterprise, (The Leningrad Prototype Reactor), four engineer**ing** Enterprises, an engineering project barely underway, (Chernobyl itself), and a political Enterprise that all the others serve, though there would be more in time, such as the Enterprises providing components and cement for the construction of the Chernobyl reactors. All of whom could have functioned more efficiently, had data safety level Enterprise analysis been a design criteria throughout the industries involved.

Early in the development of the Chernobyl Plant, when only the director had been appointed, the 7 layer data model may have looked something like figure 2. As can be seen, Victor Brukhanov has been appointed director but without sub-ordinate levels at this stage. This then dates to early 1970. Problems with the first RBMK, the Leningrad Reactor, needed to be reported by the operator to whomsoever had taken the role of optimising at Layer 5. To recap, that person must understand how it should operate, the planned service and contingency plans, how the operations are to be achieved and what the safety plans were. Where that information should have gone is shown on the diagram. What actually happened and the relation to data types is given in the discussion below.

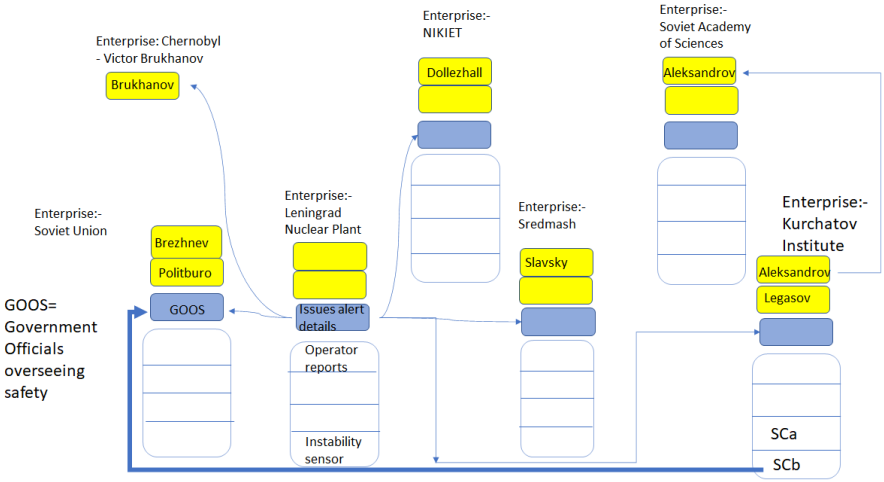

Figure 2 - Chernobyl Related Enterprises in 1970 and Some Illustrative Links

Here are some examples to illustrate how 7 layer diagrams simplify the visualisation of what should have happened and if used from the design stage, could help prevent disaster in the future. It is proposed that this type of diagram will enable the construction of safety cultures by the simplicity that pictures afford, (the actual story of Chernobyl, in written words only, is quite complex to grasp). The simple idea is to observe where data should have gone and where it either did not leave, was not received or was suppressed on the journey.

Note:

a) The Politburo is used as an embracing term in the Soviet Union Enterprise to include the Soviet Council of Ministers and the Central Committee of the Communist Party.
b) Aleksandrov had a position as director of both the Academy and the Kurchatov Institute so a link is given there
c) The Leningrad reactor lower layers, (1-3) are actual machines producing data for the layer 4 operator.

The Leningrad Reactor problem

By 1973, the first ever RBMK reactor had started up. Its construction had been ordered by Sredmash. Aleksandrov and Dollezhall were there witnessing it. Soon after start up, faults became apparent. Earlier at the Kuchatov Institute a scientist, (SCa in Fig 2), warned that the design was too dangerous and another, (SCb), warned that the positive void coefficient made the reactor susceptible to explosions. This latter scientist has been positioned at the bottom of the Enterprise to bear some resemblance to the original 7 layer data model intent in Figure 1, i.e. dealing with the highest Data Safety Assurance Levels or DSALs in the lower orders. The positive void coefficient made the Leningrad reactor harder and harder to control as it reached the latter part of its three year maintenance cycle. In terms of the Data Safety Guidance document there are issues for 'training', 'design' and other data types, which did not change despite the known problems emerging at the Leningrad plant, thus nothing was added to the manuals for the reactors. The alarm bells rung by SCb were ignored and although attempts were made to sack him, his letters eventually reached the Soviet Council of Ministers (the thick line on the diagram). Too late though, as the decision to build the 4 Chernobyl reactors had already been taken. Sredmash made the fateful decision after redesigning the RBMK as best they could but rejected a prototype and went straight into building the Leningrad reactor. The point of using diagrams like these, as proposed by this article, is that those who read SCb's earliest letters, on looking at these diagrams, would be more likely to realise that the alarm bell was not a case of crying wolf but had been issued by the trusted data integrity layer.

The Technology Development and Design Authority Problem

Aleksandrov believed that big was beautiful and ignored the problems of scale that occur when handling dangerous materials, especially nuclear materials. That size meant reactivity in the core in one measured area may have little to do with the reactivity elsewhere in the core. The problems were not explored further by either Aleksandrov or Dollezhall. Brukhanov had no idea of the problems he was to face at Chernobyl because the suppression of the information by Sredmash meant all design data was out of date or just plain wrong. That compounded the problem as the scale of the machines increased. Whatever the data type, suppression of important information on the poor performance of a prototype is never going to help.

NIKIET's lack of comprehensive training initiatives

No reasons were given for new procedures by NIKIET, so operators were less likely to stick to those rules as they were not explicit safety instructions. Although in nuclear power stations perhaps one should have assumed that operators would take care, trust in shutdown systems was high since NIKIET had approved the designs. Of course, any data on accidents would never reflect anything other than the false information that the reactors were considered safe. [Data Safety 'Instructional' type defined as "data used to warn, train, or instruct users about the system"]

The exclusion of communication to the Chernobyl Enterprise

In the typical manner of Soviet paranoia and secrecy, there is clearly a line of communication that could have gone straight to Brukhanov from the Leningrad plant, as shown in Figure 2. In fact, the information did get to Sredmash but they suppressed it. Thus, vital design and operation data became data destined never to be used again.

Using extracts from Table 4 of the Data Safety Guidance Document Version 3.2, [3], the following becomes clear as shown in Table 1 here below:

Table 1: Data Types

Nos	Title	The Problems Caused
1	Predictive	Lost because lessons were not learnt from the Leningrad reactor.
6	Design and Development	Lost because scientists at the Kurchatov Institute were ignored.
11	Behavioural	The larger RBMK's at Chernobyl could never be controlled accurately. Operators could only guess and rely on the trusted shutdown systems, which in the event could not be trusted.
13	Staffing and Training	Training was totally inadequate as the reactor's behaviour was not predictable.
16	Release	The data to ensure safe operations was more appropriate for the smaller Leningrad reactor.
17	Instructional	Warnings were not in any way clear.
+1	Trustworthiness	This meta-data could never be delivered because the Chernobyl reactor could not be trusted. How can one trust something one knows nothing about?

Conclusion

By preparedness of necessary communication, any project can be made much safer. Greater care over communication will be necessary as systems that humans depend on become more integrated, aka the internet of things. Increasing complexity also demands better communication. In a world of increasingly limited resources, cost miscalculations through a lack of understanding of safety may become unbearable and thereby wasteful.

Even with our current epidemic of Covid-19, better communication to co-ordinate the fight would have helped in the initial stages. The military analogy at the beginning of this article holds, in that, recognition of the need for coordination and accuracy of data drives the demand for trustworthy command and control systems. Though there is no reason to doubt the wet market in Wuhan was responsible for the coronavirus epidemic initiation, the same effect could occur from an accident in a laboratory. Much of the near exponential growth has been thought to be due to a lack of co-ordination between stakeholder health agencies. Therefore, there has to be absolute certainty in the future that all programmes with safety related aspects must make all parties aware of all issues and all current understanding of risks involved. This coronavirus and Covid-19, has caused a global economic meltdown, which in severity, matches the effect of Chernobyl on the Soviet Union, and the only way to avoid a repeat is through data safety.

It is proposed that the 7 layer data model, used to identify multi-agency communication and protocols for the safe, timely dissemination of data, could assist greatly in the planning of those necessary command, coordination and control actions.

References

[1] Adam Higginbottom, "Midnight in Chernobyl", Bantam Press, 2019, ISBN: 9780593076835

[2] The Data Safety Initiative Working Group, "Data Safety Guidance", The Safety Critical Systems Club, 2016, Version 1.3, ISBN-13: 9781519533579

[3] The Data Safety Initiative Working Group, "Data Safety Guidance", The Safety Critical Systems Club (SCSC-127E), 2020, Version 3.2, ISBN-13: 9798601577359, scsc.uk/scsc-127E

[4] Nicholas Taleb, 'The Black Swan', Penguin Books, 2007, 9780141034591

[5] https://www.nbcnews.com/news/us-news/u-s-pilots-complained-about-boeing-737-max-8-months-n982651 (last checked 16/04/20)

Author: Nicholas Hales IET Chartered Engineer (retired)

Nicholas Hales is the former Senior Engineer for Critical Avionics in the Defence Engineering and Support agency. He has been involved in Safety Critical Systems Club (SCSC), activities since 1991. He has written a number of articles for the SCSC Newsletter and the IET members Newsletter, and has presented at various associated events.

As a postscript I would like to thank Paul Hampton, the Editor of the SCSC Newsletter, for his patience in the production of this article and the numerous suggestions that have made it more readable.

THE SAFETY-CRITICAL SYSTEMS CLUB, Seminar:

Management and Oversight of Complex Systems

Thursday 3 December, 2020 - Reading, UK

This seminar is aimed at those who have to manage, approve, regulate and operate complex systems which have a safety aspect, e.g. Air Traffic Control systems, Nuclear Plant, Aircraft Carriers or National Power Transmission systems.

It will cover aspects such as design, operation and manning of control centres, and the use of dashboards and metrics used for monitoring. It will consider the key performance indicators that can be used to measure safety performance.

It will also cover the tools and skills that are required for management and understanding of such systems.

This seminar will be held at Thales offices, Reading, UK: 350 Longwater Avenue, Green Park, Reading, Berkshire, RG2 6GF

Further details TBC.

SCSC
FOR EVERYONE WORKING IN SYSTEM SAFETY

This 1-day seminar will be useful for all those involved in running a complex operation that involves safety. It is aimed at Managers, Operators, Regulators and Assurance staff. If you operate a management or operations centre for your organisation then this seminar is for you.

How to gain an overview of complex, safety-related Systems

https://SCSC.uk/e661

WWW.SCSC.UK

Service Assurance Guidance

By the SCSC Service Assurance Working Group (SAWG)

The SCSC Service Assurance Working Group (SAWG) gives insights into the Service Assurance Guidance document, published in February 2020, which provides guidance for the assurance of services when working in a safety context.

The Service Assurance Guidance document [1] was produced because many safety-related systems are now implemented and used in a service context, specifically:

- There is a significant shift to a service-based approach for delivery, especially in information technology with the use of cloud computing
- It is now recognised that collaborative working of systems, organisations, people and processes all contribute to safety
- A service-based approach to assuring safety provides a different perspective that extends beyond system analysis

The guidance document applies to the assurance of services when there are safety implications associated with the use of those services. These services are '*Safety-Related Services*'. Examples might be an ambulance dispatch service or an air traffic control service.

Introduction

Many safety-critical systems utilise services that are designed, developed, operated and maintained outside the immediate boundaries of the system. Future developments in business and technology are likely to mean that this *service paradigm* will become increasingly prevalent in the next generation of safety-critical technologies.

Example: Passengers struck by a flying cable at a station

As a brief example, here is an excerpt from a service-related accident in the rail sector. This example is discussed in more detail in the paper (King et al 2020) [5].

"At about 18:05 hrs on 28 July 2017, as a northbound passenger train entered Abergavenny (Y Fenni) station, a cable drooping from the station footbridge became caught on the train's roof. The train dragged the cable and caused it to be pulled from the footbridge until its end broke free from a distribution cabinet. Once free, the end of the cable struck a group of passengers on the footbridge stairs and caused minor injuries to three of them. A member of station staff who was on the platform, close to the footbridge, was nearly struck by the cable. The accident also caused damage to cabling running over the footbridge, the station buildings, and a signal at the end of the platform.

The cable, which provided the signal box at Abergavenny with its electrical power supply, had become detached from the cable tray running over the footbridge and was drooping down to the extent that it was foul of the train. It then caught on an antenna fixed to the roof of the rear vehicle. The cable was drooping because the nylon cable ties used to attach it to the cable tray had broken. The RAIB found that the cable had not been inspected periodically as required for electrical installations and the drooping cable was not identified during footbridge inspections. It was not reported during routine station safety checks, or after it was drooping below the bottom of the footbridge. An underlying cause was that Network Rail had no controls in place for the management of low voltage electrical supply cables that cross operational railway lines via its overline structures." (RAIB, 2018) [6].

This incident highlights the importance of services that should be utilised at regular intervals. In this case the 'cable inspection' service and 'station safety checks' service both of which were probably carried out by the same staff some time prior to the incident. Both services

involve people, process and equipment, and should have been carried out with a level of assurance. A case of over-familiarity was likely a flaw in the service execution.

Future developments in business and technology are likely to mean that this *service paradigm* will become increasingly prevalent in the next generation of safety-critical technologies

A Service View of Safety

A service-oriented view of safety may be able to identify and manage safety risks more effectively since it highlights the collaboration of various elements of the socio-technical situation (people, organisations, processes, maintenance, change, automation, through-life aspects etc.) and their contributions to the overall safety of the operation.

This service paradigm presents considerable challenges for safety engineering and assurance for a variety of technical and non-technical reasons, often extending beyond the traditional concerns of system safety engineering (e.g. into the realm of commercial contracts, service-level agreements and cross-organisational concerns).

What is a Service?

The term "Service" is not clearly defined across standards and has many uses. The way that a Service is normally described or defined is very different from the specifications and descriptions more commonly employed in safety-related systems. An individual Service (or *Service Component*) is typically described by a Serv*ice Provider* via an entry in a *Service Catalogue*. The Service Catalogue describes the externally visible capabilities/functionality offered to a *Service Consumer* without providing implementation detail, in fact, it is unusual for any aspects of the design and implementation of the Service to be visible to the Consumer.

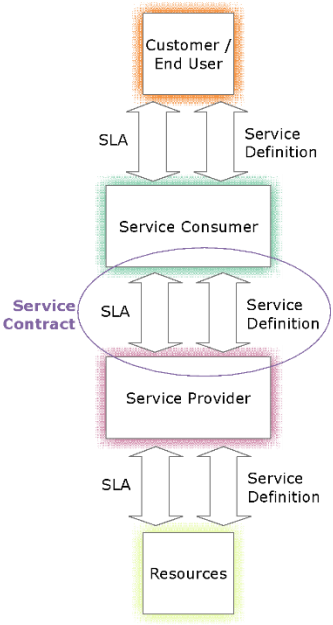

A summary of key terms is given below:

- A **Service Provider** provides one or more Services.
- A **Service Consumer** consumes one or more Services.
- A **Service Definition** describes the Services available for consumption, which may include technical and/or commercial aspects. It may include deliverables, prices, contact points, availability, ordering and processes to request Services. It often includes a service catalogue.
- A **Service Level Agreement (SLA)** is the agreement between the Service Provider and Consumer that defines the level of service that the Consumer will receive. It usually specifies responsibilities of both parties and defines the penalties in the event the specific targets in the SLA are not met.
- The **Service Contract** is the contractual agreement between Service Provider and Service Consumer.
- A **Service Based Solution (SBS)** comprises the systems, organisations, processes and resources to deliver and manage the services through life. It may consume other services. An **SBS** delivers capabilities to its customers via a set of collaborating services.

Characteristics of Services

It is important to understand the key characteristics of a service and establish the distinct nature of that service (as opposed to a system or product), especially with respect to assurance for safety. Six key service characteristics are given in the following table:

ID	Service Characteristic	Context	Compared to Systems / Products
C1	Services are provided for the duration of the service contract	By definition, services provide features with a given performance for the duration of the contract between provider and consumer. So e.g. the consumer does not have to be concerned with the design, maintenance or disposal of the service components.	Products usually provided with warranties but there may be no further involvement from the provider after product acquisition. The consumer is responsible e.g. product maintenance and disposal.
C2	Services often designed to meet the needs of a broad range of consumer needs	Providers want to attract a broad range of customers and may wish to avoid too many bespoke solutions	Similar, but products often designed to exacting specifications. Can be easier to remove features of products for particular customers.
C3	Services likely to be used by more than one consumer	Providers usually desire to sell similar service offerings (catalogues) to a wide range of customers, so multiple consumers can be using the same service offering at the same time (or sharing resources of other services)	Not so easy to segregate services or have sufficient visibility to understand and control potential forms of interference and disruption
C4	Services are implemented through a combination of people, procedures, products and other services	Service implementation requires the collaborative working of people, processes, products and other services to deliver a set of features with the desired performance	Similar to systems that consider people procedures and equipment and all aspects of in-service operation. However, systems and product suppliers are not usually responsible for live system operation.
C5	Services may be designed without recognition of the full context of use	Providers may desire a quick route to market, so may release their service before the full context is understood. The service may be designed to adapt to a context of use or may evolve over time to meet the emerging context of use. Services may be developed for specific purposes that other users decide to exploit (unexpected uses).	Similar to off-the-shelf products and systems, however, once established and demonstrated within the context of use, products and systems are usually only changed under the direct control of the consumer. Changes in the underlying service provision may increase the likelihood of undesired emergent properties.

ID	Service Characteristic	Context	Compared to Systems / Products
C6	Service implementation details may not be visible to the consumer	Whilst the performance and features of a service should be clear to the consumer, the details of how the service is delivered may be kept confidential or be hidden within other lower-tier services	Very similar to COTS products/software. However, once established and demonstrated for a given installation, products and systems are usually only changed under the direct control of the consumer.

Service Assurance Principles

The following table presents the Service Assurance Principles as a way of structuring the service assurance activity together with brief rationale:

1	**Service assurance requirements shall be defined to address the service-based solution's contribution to both desirable and undesirable behaviours**
	There must be an overall definition of what the service is trying to achieve (formulated as requirements) and this must be within an expected usage scenario (e.g. concept of operations). There must be requirements addressing known behaviours that are unwanted or unsafe.
2	**The intent of the service assurance requirements shall be maintained through the service definitions, service levels, the service architecture and the agreements made at service interfaces**
	This relates to the way the service hierarchy and service decomposition is constructed. It is saying that the intent of the assurance requirements must be shown to be met by the service elements comprising the service, and that the overall service architecture or hierarchy supports this flow down (i.e. that all service elements together meet the overall intent, and nothing is missing). Service elements can be of various types, including other services, systems, subcontracts, and agreements.
3	**Service assurance requirements shall be satisfied**
	Service requirements must be satisfied, i.e. verified as-is or decomposed into further requirements which are subsequently verified in some way. The methods by which service requirements are verified are wider than traditional systems, often including extensive use of proven-in-use (service history) and commodity-usage arguments, and also some specific contractual mechanisms. This principle (together with (4) below) creates the need for assurance "wrappers". (A *wrapper* is an assurance augmentation which addresses the assurance deficit inherent in the consumed service in some way).
4	**Unintended behaviours of the service-based solution shall be identified, assessed and managed**

	All undesired or unintended behaviours which may impact safety properties or safe behaviour of the overall system must be identified and assessed within the usage context. They must be appropriately managed (e.g. mitigated, avoided or accepted in some way). This is not always possible to the extent desired, especially when commercial "commoditised" services are involved. Hence this may create the need for additional wrappers to make up the assurance gaps (see also principle (3) above).
5	**The confidence established in addressing these principles shall be commensurate with the level of risk posed by the service-based solution**
	This is the proportionality principle: the level of (safety) risk must be used to determine the amount of effort (resources, time, etc.) put into assurance and mitigation activities. This principle can be used to underpin a set of levels of service assurance, where applicable activities are defined in bands derived from the risk level.
6	**These principles shall be established and maintained throughout the lifetime of the service-based solution, resilient to all changes and re-purposing**
	Services may have a long lifetime and the service offering may evolve significantly over this time. These principles must be established and maintained throughout life: through e.g. usage change, technical change, subcontractor change, supplier or process and personnel change. This principle must also hold in service failure scenarios (contingency situations) where the service might temporarily employ manual or procedural activities to achieve its aims. It might be thought that this principle is implied by the others, but continuous evolution and change is a key property of services; in this they are different to (largely) static systems.

The principles are similar to those used for software safety assurance, with the addition of #6 which is concerned with change and evolution. This set has been developed over the last few years [2], [4]. This set of principles are considered necessary and sufficient for assurance of services in a safety context; how they are achieved is through objectives, see below.

Objectives for Principles

The principles form the high-level "mission statement" of service assurance. These are mapped to a set of Service Objectives forming the next level of specification. If the objectives for a principle are met, then the principle is considered satisfied. The following table shows how the principles may be mapped to lower-level objectives:

1	**Service assurance requirements shall be defined to address the service-based solution's (SBS) contribution to both desirable and undesirable behaviours**
	a. Context and intended use of the SBS SHALL be established
	b. States of the SBS SHALL be defined including normal, abnormal and degraded modes, as well as transitions between the states
	c. Key stakeholders of the SBS SHALL be identified
	d. Service assurance requirements for desirable behaviours, including service and performance levels of the SBS, SHALL be defined
	e. Service assurance requirements to mitigate undesirable behaviours of the SBS SHALL be defined
	f. A high-level service architecture SHALL be defined

	g. Historical accidents and incidents related to the service offering SHOULD be assessed and any relevant recommendations considered.
2	**The intent of the service assurance requirements shall be maintained through the service definitions, service levels, the service architecture and the agreements made at service interfaces**
	a. Service assurance requirements SHALL be decomposed and assigned to service elements within the service architecture of the SBS b. The service architecture including sub-services SHALL be defined c. Service assurance requirements SHALL be defined for each sub-service d. The agreements made at service interfaces SHALL be defined e. Service assurance requirements tracing through the service architecture SHALL be established f. Methods and techniques used to provide service assurance within each level of the service architecture SHALL be defined and implemented g. Assurance wrappers SHALL be identified and defined for service elements to make good any known assurance shortfalls
3	**Service assurance requirements shall be satisfied**
	a. Verification evidence SHALL be produced to show that service assurance requirements are met by the architecture and the elements of the SBS b. Assurance wrappers SHALL be implemented and verified c. Evidence SHOULD include proven in use and service history evidence
4	**Unintended behaviours of the service-based solution shall be identified, assessed and managed**
	a. Residual risks SHALL be identified and linked to service artefacts and service properties b. The residual risk of the SBS SHALL be reduced to an acceptable level c. Unintended behaviours resulting from the service architecture and service elements SHALL be identified, assessed and managed d. Unintended behaviours resulting from fault-free cases SHALL be identified, assessed and managed e. Service-service interactions SHALL be considered f. Service assurance artefacts SHALL be identified and produced
5	**The confidence established in addressing these principles shall be commensurate with the level of risk posed by the service-based solution**
	a. Levels of Service Assurance (LSAs) SHALL be established based on the level of risk that the service presents to the service users b. LSAs SHALL be decomposed and assigned to service elements within the service architecture of the SBS c. Service assurance artefacts SHALL be produced according to the LSA d. Activities, methods, analyses and tools used to provide service assurance SHALL be appropriate for the LSA
6	**These principles shall be established and maintained throughout the lifetime of the service-based solution, resilient to all changes and re-purposing**
	a. All changes to the SBS that impact these objectives SHALL be assessed and managed b. Service assurance artefacts SHALL be maintained

	c. Use of the SBS SHALL be monitored for change and a safety impact analysis shall be undertaken d. Use of the SBS for a new purpose, or changed scope SHALL cause a re-evaluation of the compliance with the objectives e. Degraded and contingency modes of the SBS SHALL maintain the defined set of these objectives f. Lessons learnt SHALL be incorporated in the SBS

The objectives can be treated as requirements (with further elaboration), and therefore give a means of compliance to the principles. The guidance document introduces selection of objectives depending on the level of service assurance required.

Who is Responsible?

The overall responsibility for specifying and showing overall achievement of objectives sits with the organisation **consuming** the service, however the objectives will typically be met by other organisations within the service hierarchy (which can be the service provider or a third party contracted to provide service assurance).

Level of Service Assurance

The concept of Level of Service Assurance (LSA) is an important one. Differing amounts of assurance are required for different types and instances of service usage. The LSA is defined by the level of risk in using the service (defined by the consumer of the service):

LSA	Definition (Service Consumer View)	Additional Clarification
LSA 0	No safety aspects present in service	There is no obvious route to harm to humans or the environment from use of the service
LSA 1	Minor safety aspects with little impact of failures (minor injury possible but unlikely)	Harm is unlikely as there are many mitigations in place and plenty of time to recover the situation
LSA 2	Safety aspects with some impact of failures (several injuries possible)	Some harm is possible (to one or more people), but there are mitigations in place and some time to recover the situation
LSA 3	Significant safety aspects with service with major impact (could indirectly lead to single death or multiple injuries)	Significant harm is possible to several people; there are a few mitigations in place; there may be little time to recover
LSA 4	Service is safety-critical: service failures could have catastrophic impact (could directly lead to multiple deaths)	Major harm is immediately possible to many people and there are limited or no effective mitigations if the service fails. There is almost no time to recover.

NB: Here harm refers to people or the environment. This could be extended to cover other aspects if required, e.g. assets or platforms.

Five levels are considered appropriate and give enough granularity to be distinct. The key aspect to remember about these levels of service assurance, is that they are based on the service consumer's view. Services are often generic (e.g. a wireless communication system), used by many consumers and it all depends on how the service is used by this particular consumer (i.e. for what safety-related applications it is intended).

Objectives and Applicability

The guidance document has a table that links the objectives to a set of methods and techniques tables, giving suggested approaches to satisfying the objectives. Each objective is considered applicable (i.e. required) at a certain LSA or higher. Hence the LSA defines the quantity, breadth and rigour of the service assurance required from the assurance provider.

Evidence to meet Objectives

The guidance gives tables of example techniques for the production of evidence against each objective. These are:

SP - Service Scope

SD - Service Design

SA - Service Analysis

SC - Service Contracting

SY - Service Delivery

SS - Service Assurance

SV - Service Verification

SH - Service Change

SR - Service Regulation

SF - Service Staffing

Service Analyses

It is recognised that different (or modified) safety analysis techniques may be required to analyse a safety-related service, including systems, people and process aspects. Some of these are:

- Service Functional Failure Analysis [SFFA]
- Service Failure Modes and Effects Analysis [SFMEA]
- Service Hazard Analysis [SHA]
- Service Business Process Failure Analysis [SBPFA]
- Service Interaction Analysis [SIA]
- Failure Analysis of Agreements [FAA]
- Service Structuring Analysis [SSA]

Further details are given in the guidance.

Case Studies

Examples are discussed in the guidance document, including incidents and accidents from across industry where service failures are considered to be significant contributory factors. The Deepwater Horizon accident is considered in some detail [3].

Conclusion

Services are increasingly being used to provide safety-related functionality to end users. There is currently no standard or guidance that addresses the general assurance of services in a safety context.

The service assurance guidance document is important because, for the first time, it recognises this delivery of safety functionality via services and gives a framework for assuring those services. The framework includes principles, objectives and means of achieving those objectives and assurance levels thus giving a method for achieving an assurance position in a particular situation.

The guidance has now been issued for comment and any feedback on its application is appreciated.

An updated version of the guidance will be issued at SSS'21 in February 2021.

References

[1] Service Assurance Guidance, The SCSC Service Assurance Working Group (SAWG), Feb 2020, https://scsc.uk/scsc-156

[2] Durston N, Scott A, Parsons M, Simpson A (2019), The Principles of Service Assurance, in Kelly T and Parsons M, "Engineering Safe Autonomy", SCSC-150, 2019, https://scsc.uk/rp150.6:1, accessed April 2020

[3] DHSG Deepwater Horizon Study Group. 2011, Final Report on the Investigation of the Macondo Well Blowout. Deepwater Horizon Study Group. March 1, 2011

[4] Harris C, Parsons M and Simpson A (2018) Service-Based Safety Assurance in Kelly T and Parsons M, "Evolution of System Safety", SCSC-140, 2018, https://scsc.uk/r140/8:1, accessed October 2018

[5] King, K, Parsons M, Sujan, M (2020) A Service Perspective on Accidents, in Nicholson M and Parsons M, "Assuring Safe Autonomy", SCSC-154, 2020, https://scsc.uk/rp154.6:1, accessed April 2020

[6] RAIB, Report 06/2018, Passengers struck by a flying cable at Abergavenny (Y Fenni) station https://www.gov.uk/raib-reports/report-06-2018-passengers-struck-by-a-flying-cable-at-abergavenny-y-fenni-station, accessed April 2020

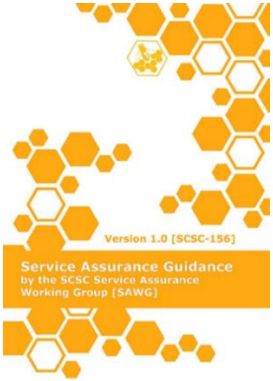

scsc.uk/SCSC-156

SCSC Service Assurance Working Group (SAWG)

The SCSC Service Assurance Working Group (SAWG), led by Mike Parsons, has been set up to produce clear and practical guidance on how services should be managed in a safety related context, to reflect emerging best practice.

Comments or suggestions for the Service Assurance Working Group (SAWG) and the Guidance Document are welcome at:

sawg-comments@scsc.uk

SCSC Symposium 2020

In February 2020, the SCSC held its 28th annual 3-day Symposium at the Principal Hotel in York. The event was well-subscribed, with over 150 delegates attending from all over the world, including the UK, Europe, Asia, the US and as far as Australia. Key themes covered during the event were: Autonomy, Machine Learning and AI, Assurance, Security informed Safety, Human Factors and Data Safety. Paul Hampton reports on the key note speeches given during the event.

The Boeing 737 MAX Accidents

Dewi Daniels, from Software Safety Ltd, opened the event with a key note speech on the Boeing 737 Max crashes and the role the Manoeuvring Characteristics Augmentation System (MACS) played in the accidents. In the talk, he presented the history and evolution of the aircraft design and provided a detailed and informative description of the events surrounding both accidents, including extracts from the actual flight logs.

Dewi gave his personal perspective on many of the myths surrounding the accidents that have been reported in the news, such as causes being due to poor regulatory oversight and insufficient pilot training. In countering these myths, he illustrated, through a cockpit simulator video, how difficult it would have been for *any* pilot to save the aircraft, even when following the FAA published guidelines for dealing with uncommanded horizontal stabiliser rotation. Dewi concluded that one of the key contributors to the accidents was the misclassification of the safety dependency of the MCAS system.

Satellite Navigation

John Spriggs, from NATS, presented the history of Global Positioning Systems (GPS) and the challenges that surround the implementation of accurate GPS systems. For example, he explained how difficult it is to determine the position of the actual satellites and the errors that can be introduced from the Ionosphere and 'space weather' in the Troposphere.

Safety in Space

Emma Taylor from the Safety and Reliability Society (SaRS), described her experiences with the dangers of space debris and how the rapid expansion of commercial space based operations is accelerating the risks to space users. She explained that there are now millions of pieces of debris in orbit, each travelling at high velocities. She showed from experimental results, the damage even the smallest size of debris can do at a speed of 7kps, and the challenges posed with trying to recover debris at these velocities. She noted how increased space use is bounded, not by technical challenges, but by political and spectrum constraints.

Safety of the High Speed Rail Network (HS2)

Reuben McDonald, of High Speed Two (HS2) Ltd, described the safety challenges of the planned implementation of the High Speed Rail Network (HS2) linking (amongst others) London and Birmingham and providing improved services to 25 stations around the UK. He noted that the 15 year HS2 programme is an unprecedented engineering undertaking, requiring of the order of 300 under/over bridge works to deliver a service that will see 18 trains every hour travelling at speeds of up to 360kph. Reuben gave some examples of the safety challenges with HS2, such as assessing the risks of running the rail track in close proximity to an airport runway, and the cost benefit analysis when developing an ALARP case for metro-style platform barriers.

Safety of New Nuclear Builds

Alastair Crawford of EDF Energy, described the progress and plans for the build of the new multi-billion-pound nuclear reactors being built at Hinkley Point in Somerset, UK. He described the architecture of the EPR reactors highlighting the key safety concerns of: core, cooling and containment. He described many of the safety features being employed, such as backup generators and the lessons that have been learned from the Fukushima Daichi accident, such as, introducing flood resident doors.

Responsibility Sensitive Safety

The final key note speech was from Jack Weast, of Intel and MobileEye. Jack described the challenges of developing Autonomous Vehicle (AV) algorithms that are intended to keep the vehicle free from conflict with other road users and pedestrians. He noted that making the algorithms overly defensive would significantly limit the utility of the vehicle's operations. For example, it may stop so frequently to avoid perceived hazards that it causes intolerable congestion.

He described work that he and his team have been doing in the development of a more flexible algorithm that tries to emulate human driving behaviours, such as cautious but assertive driving. He noted that although the algorithms improve the utility performance of AVs, they are not sufficient to prevent collisions in cases where other road users act outside the boundaries of anticipated behaviours, such as, travelling at 70 mph in a 35 mph zone.

Jack stressed that the intention is to share the work his team have done and to engage with industry, academia and government bodies in an open, transparent and collaborative way to arrive at AV behaviours that are societally acceptable.

Delegates meet for refreshments and the tools and services exhibition

Interactive Voting

The 'question time' panel session was replaced this year with an interactive voting session chaired by Catherine Menon.

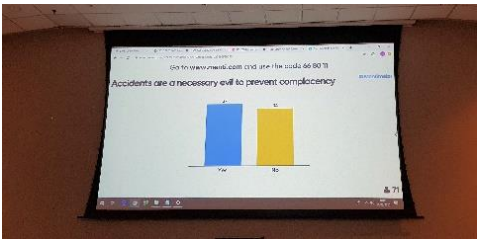

In this session, delegates voted electronically via a smart phone app on questions in the world of system safety, such as, "Would you be happy for an AI GP to treat your family?" and "Accidents are a necessary evil to prevent complacency". Full results can be found on the SCSC website: scsc.uk/re619.4:1

Social Events

During the symposium, delegates were invited on one of three social events. These comprised a tour of York Minster, a guided ghost-walk around York, and gin tasting including samples of 'Old Tom' and 'Grey Lady' provided by the local brewery York Gin.

Evening Banquet

Delegates attending the evening banquet were entertained to not one, but two, after-dinner speakers. Peter Ladkin gave some interesting insights into the emerging risks of Coronavirus and concluded with an amusing rendition of 'There was an old lady'!

Tim Kelly, the previous SCSC Director, provided an entertaining and thought provoking speech, which drew parallels between system safety and his new career, following his ordination in the Church of England. A full transcript of his speech is provided on page 46.

Tim Kelly's SSS'20 After-dinner Speech

First of all, I must say how delighted I was to be invited back to do the after-dinner speech at this year's symposium. I've always enjoyed this event, my first being 26 years ago.

It's now approaching eight months since leaving my friends and colleagues at the University of York, becoming ordained as a Church of England minister – a Reverend! (some people have mistakenly suggested that I just like collecting titles) and upping sticks with my wife and family to the East Riding of Yorkshire to become part of the clergy team at Beverley Minster. So, it's lovely to have this reason to be here and to catch up with old friends.

Despite speaking at the symposium many times, I've never before given the *after-dinner speech*. So this remains for me a new and daunting prospect. However, this may be one area where advice from my new vocation could readily apply. George Burns, the American comedian once said that:

> *The secret of a good sermon is to have a good beginning and a good ending; and to have the two as close together as possible.*

I think this advice probably applies quite well to a good after-dinner speech too, so don't worry, I won't be keeping you all too long from the bar.

One of my most recent talks (last week) in my role as curate at Beverley Minster has been to primary school aged children. People have asked me which is the most challenging crowd. Is it delivering a keynote at a conference or speaking to 6 and 7 year olds? Perhaps you can already guess the answer! All I will say is that 6 and 7 year olds ask some *really* good questions if you like being put on the spot! And they often have a refreshing perspective on longstanding and challenging problems. Such as the story of the teacher who was walking around observing her classroom of children while they were drawing pictures. As she got to one girl who was working diligently, she asked what the drawing was.

> *The girl replied, "I'm drawing God."*
>
> *The teacher paused and said, "But no one knows what God looks like."*
>
> *Without looking up from her drawing, the girl replied, "Well they will in a minute!"*

You might wonder, having made the leap from the world of safety critical systems engineering to the world of the Church of England, what could possibly read-across from my former career to my new vocation (and vice-versa). Well I'm going to take you on a journey of a few thoughts that I've had that connect these, perhaps at first glance, extremely diverse domains. Firstly, where better to start than a discussion of *risk*. Now, as I'm sure this audience appreciates well, risk is everywhere, risk unites us all.

As humans we are physically born, we live (variously) and then we die.

In a way, I feel much more aware of this now that part of my job involves the very beginning of life – baptism – and the very end of our earthly life – funerals. Risk is a function of living. We can't escape it. In my new role I've certainly discovered a fair few new occupational hazards: Clergy Robes (Cassocks and Surpluses – the white frilly bit with the long sleeves) mixed with candles, many candles, is definitely a hazard. And then there's the hazard of a large 12th century font, baby in one hand, with a towel and an order of service in the other and making sure that if you're going to drop one thing it's not the baby. I've also been grateful for the hazard mitigation advice of funeral directors of where and exactly how far back I am to stand to avoid falling in when conducting a grave side funeral. But back to the fundamental question of *risk aversion*. I'm going to question for a moment; should we ever really want to totally escape risk when it comes to our own experience?

I can remember a discussion with John Knight when talking about a concept that he was trying out with his PhD student at the time (Patrick Graydon), which they had called *Assurance Based Development*. The basic premise of the idea was that they wanted to have the development of a safety critical system being totally driven by the incremental development of an assurance case (or safety case to most of us). However, there was a problem. When they tried out the idea on the development of an artificial heart device, *it simply didn't work.* The problem was that the development of the safety case has as its goal (and rightly so) the continuous minimisation of risk – how could we reduce risk. And, of course, that can never work as the sole, or primary driver. It has a tendency to promote *nothing* … doing nothing … risk nothing. If in doubt, say no! Some of you may have heard me say before that whenever a company says "Safety is our top priority" or "Safety is our number one goal", they are clearly, in my opinion, *lying*. How can it be? If safety was genuinely the number one priority, in many cases the answer as to what do is simple – do nothing! Safety always has to be addressed in juxtaposition to other competing objectives. We manage the risks of high speed train travel in relation to our goal of wanting to move large numbers of people from city to city in as short a time as possible, or the risks associated with space travel in order that we can explore the wonders of the universe. Putting it this way, it can perhaps start to be seen as important that we shouldn't let safety win. We shouldn't want safety always to dominate. Safety needs to be put in its place, and perhaps we still don't have the tools to manage this. The ALARP framework focuses on risk versus cost and is not adept at managing multi-attribute tradeoffs – something once explored by my PhD student George Despotou.

This balancing of risk and reward is something that I can recognise in my move from safety to church and clergy life. I understand when some even described my move itself as an excessively risky thing to do. Why would I want to leave my secure job and career to receive less pay for potentially more hours? Why would I risk uprooting my family – new house, new schools for my boys – what could possibly be worth it? If my sole interest was the minimisation of risk, this was clearly an unnecessary move. However, the risk needed to be balanced with other competing forces – and in my case, I felt one pretty large competing force. This was something that I felt I was being clearly prompted and nudged to do. The risk of doing nothing, although the apparently safe option, was now having to be weighed against the lost opportunity of doing something that I was meant to do. Thinking about the end of that bathtub curve again, the end of life, was I going to reach the end and look back and see there was so much more to be done? Safety critical systems engineering, safety cases, and dare I say it even GSN, are not the be all and end all!

And of course, having made the move, there's the question of the appetite for risk in the life of a Christian. Jesus himself wasn't averse to a little risk. He challenged injustice and oppression where he saw it, didn't insulate himself in fancy temples and palaces but instead lived lightly on the road, and of course, he wasn't afraid to upset the authorities and the established order of things. Someone could pipe up at this point and rightly say, "and you can see where that landed him", and you'd be right to point that out. Risk was balanced against reward – just not necessarily the kind of rewards everybody associated with success – money, military and political power. You could say that a life of love, a life driven by love, looks very different to a life driven by safety and risk management. Love isn't safe – it drives us to do things that sometimes look *far* from safe.

Last autumn I was delighted to be asked to be involved in the marriage ceremony of a friend and safety engineering colleague from BAE Systems, and I was determined to see if I could get a mention of some safety and security terms into his wedding sermon. In the end it wasn't that difficult at all. You see there's perhaps few greater examples than a wedding of trust and of course, it's flip side of risk. Where the instinct of the world is so often to try to stay as safe and secure as possible, where we try to reduce risks ALARP, and to not trust anyone unless you absolutely have to, when two people get married together love and trust clearly wins over risk.

So where does this discussion of risk leave me now professionally? Well, I think the church is having to challenge itself on this very topic of risk. How should the church move forwards? Play it safe or live dangerously? Well you might guess where my thoughts start to land. To play it safe would be to change nothing, risk nothing, maintain the status quo and avoid trying or exposing ourselves to anything new – better not, it might backfire. Surely, we should minimise the risk surface not extend it? But you see the Christian church is called to a life of love, and some might say, as a consequence, called to life set with risk. It was never meant to be about power, influence, or wealth (as some might suggest from the outside). It was meant to be about love, including sometimes wilful sacrifice. (Now, don't get me started on what can be considered a reasonable risk!) I'm not about to suggest to the Church of England, following the fashion of John Knight, a new process called "Love Based Development" – but it wouldn't necessarily be a bad place to start!

And where does this discussion of risk leave the world of safety engineering? I think it shows we have to be careful about sometimes playing the safety hand too strongly, and to the detriment of other reasonable societal goals. The field is still open for the development of a new multi-attribute ALARP framework!

Now I want to move on to discussing a man that, I suspect, many of you may not have heard of before – Richard Hooker. He was a priest and influential theologian in the life of the Church of England and lived between 1554 and 1600.

Amongst other things, he was famed for, what is now known as, his *three-legged stool*. Not an actual stool, but theological framework that said *scripture*, *reason* and *tradition* all had a role in determining the thought and practice of the life of the Anglican Church of England. If you permit me to explain a little of what he thought, I think there's some merit in reading across some of his principles.

Martin Luther, the great theologian of the protestant reformation age, was a big fan of *sola scriptura* (scripture alone). This is the theological principle that Christian scriptures are the sole infallible source of authority for Christian faith and practice. Luther railed against some of, in his opinion, the accreted practices and ceremony of the Roman Catholic church. This protestant position was the one adopted and endorsed by Henry VIII. However, adding to this position, Richard Hooker introduced the other legs of reason and tradition. First, by introducing reason, he suggested that where the bible was plain and clear, it should simply be followed, but where it is not, we should be using all of our critical reasoning skills and faculties to help us determine our position. Secondly, he said (if I summarise simplistically for a moment) that the church would be unwise to completely ignore the tradition handed down to it, given that in many cases, it was the result of much discussion, debate and experience.

So, *scripture*, *reason* and *tradition*. Does this model have any merit for safety critical systems and software development and assurance? Let's tackle them one by one. Scripture … *Ah!* This is perhaps where we first become unstuck. Do we have anything approaching an analogue of scripture in the field of safety critical systems development? I'm sure many of you remember John McDermid's article – *Software Safety, Where's the Consensus?* If you haven't read it, it's worth a read. Within that paper, John highlights the many safety standards (or dare I say '*holy books*') that exist within the safety domain. Sometimes we almost seem to divide ourselves into different churches – The church of DO-178C, or the church of IEC 61508, or the church of ISO 26262. *Which one do you belong to?* Of course, unfortunately, as John highlights, they don't all agree. It's hard to adopt a 'Lutheran' position that these safety standards present an infallible source of authority when such inconsistency is present. It's also easy to provide counter examples to their scripture and verse (e.g. the ineffectiveness of using MCDC testing to reveal functional failures in testing Bayesian Networks and Artificial Neural Networks, as discussed by my PhD student Mark Douthwaite some conferences back).

When discussing the idea of this talk with Mike Parsons, he suggested that it would be great if there were some equivalent of the ten commandments for system safety. Indeed it would, but alas, no such 10 commandments exist. I've seen a few attempts online (it's worth a Google) but it's hard to have faith that these are any kind of infallible set.

Some of my work over the last ten years or so with the 4+1 Software Safety Assurance principles (that can now be found in Def Stan 00-55 and 00-56) was an attempt to provide some immutable principles that sit over and above the variations in detail of the different safety standards. However, I'd be one of the first to spot another flaw in providing such lofty principles, good as they may be. The more abstract the principles the harder they can sometimes be to apply. They need interpretation and implementation (cue discussion of reason and tradition). But before I leave the topic of standards, the safety scriptures if you will, I just wanted to flag one other connection I've made between my two worlds. Whilst I do believe in the holy and inspired position of the Bible for Christians, I also acknowledge that ultimately, the words were written down and arranged by human authors. As such we need a *hermeneutical* approach when reading the bible. Hermeneutics is the theory and methodology of interpretation, especially the interpretation of biblical texts. Alongside a literal reading of the words, we should consider the context of writing, the significance of the chosen words of the text, and the challenges of application of the text to a contemporary situation.

It strikes me that such an approach has merit in adopting a thinking approach to the use of safety standards. Safety standards, despite what some may think, do not simply appear on tablets of stone. They are written by humans, in many cases committees of humans, trying their best to capture their thinking and advice. In many cases, they are a product of their time, context and authorship. Sensibly therefore, when reading standards for contemporary application, we would be wise to consider the context in which they were originally written. We also should think carefully about their choice of words. I forget where it was, but I was particularly aggrieved when I saw a rewriting of Principle No. 2 (From the 4+1 principles) which I originally stated as being that *the intent of safety requirements should be preserved throughout system development.* The rewrite simply paraphrased this as: *requirements traceability should be maintained*. No … no … no! My original wording was carefully chosen and meant something. Alter this wording at your peril. Finally, a critical mind is required when thinking about how best to apply the requirements of a safety standard to a current project. To apply unthinkingly the letter of the law, rather than think how best to understand and apply the intent of the law, would be a mistake, and so easily could miss the point – e.g. congratulating ourselves on a SIL 4 Artificial Neural Network inference engine that perfectly executes faulty inferences learnt through poorly selected data.

And so to *reason*. This is an easy one for me in terms of its merits for safety engineering. If anyone knows anything about what I have attempted to promote through the development and application of the Goal Structuring Notation (GSN), it was the clear application of reason and hopefully (when done well) critical thinking skills. In the domain of theology, Hooker said that we should never be afraid to test what we read and inherit against the whole armoury of our critical thinking and reasoning skill. Similarly, it is essential to the development and assurance of safety critical systems that reasoning is explicated, and tested. This should be the main aim of safety case development. Two of my SCSC newsletter articles come to mind – *Are Safety Cases Working?* and *There's No Substitute for Thinking*. The first of these captured my pre-Haddon Cave concerns of situations where safety cases somehow were missing the mark through becoming mundane, routine, or simply an end in themselves. The later article, *There's No Substitute for Thinking* was written in response to suggestions from some at the time, that all we really needed was safety evidence and what really was the point of this argument stuff anyway, and suggestions that all would be well if we could just persuade people to reduce their safety justifications to a few sides of A4. Both positions terrified me at the time because they served to diminish the position of exposing and criticising the otherwise implicit reasoning as to why people believed their position of acceptable safety. So, I'm with Richard Hooker – *reason* is a key pillar!

And then to *tradition*. In Hooker's case, he wanted to encourage respect for the handed down traditions of the church. They were often there for good reason. They had been found to be useful practices. As with reason, I don't believe it's difficult to see the utility of this concept when read across to the safety domain. Notwithstanding my earlier comments of sometimes needing a hermeneutical approach to reading safety standards, safety standards *do* capture much collective wisdom. Another of my biggest fears of people unthinkingly adopting a safety case approach, was that the safety case author would simply eschew and reject all existing safety guidance and standards "Don't worry I'll do it my way" (indeed this was one of Nancy Leveson's erroneous strawman objections to the safety case approach). I would say instead, that it's a foolhardy safety engineer that says they can ignore the body of knowledge contained in existing safety standards.

I'm a fan of Hooker's three-legged stool, and not just because I'm a good Anglican. Although I've talked about how each leg is useful and has read-across when applied to the safety domain, it is perhaps the concept of the stool itself that is perhaps the most useful thing – being willing to hold these three things in tension. I've already mentioned how in the safety domain, false oppositions have sometimes been manufactured between different legs. Such as, it's safety standards or safety cases, but not both – a hopefully obviously fallacious position. We are required to hold the positions in tension and check them against one another. Reason should not be allowed to overrule fundamentally understood principles. Tradition shouldn't be blindly followed if it runs contrary to reason or defies critical investigation.

This talk has given me only a chance to scratch the surface of the links that I find myself making as I settle into my new vocation. My thoughts on the importance of robots that are capable of following Jesus' commandment to us all to love one another, will have to wait for another occasion. Watch this space – although I have little desire to gain another academic qualification – this is the area I'm studying for my Theology master's degree. Incidentally, let me just say for now, that you have to be very careful when you put the phrase "love robots" into Google. And of course, to that million dollar question – have I found myself being able to deploy my beloved GSN in my new context? Well so far, I've yet to explicitly inflict GSN on the congregations of Beverley Minster and associated churches. However, I *have* received some positive feedback on my sermons along the lines of "Your sermons always seem very logical and easy to follow". There may just be a reason for that.

Finally, given that I've been talking about the relationship between my former and present careers, I wanted to end with one particular joke that struck me as being appropriate:

> *An engineer dies and reports to the Pearly Gates. Saint Peter checks his dossier and not seeing his name there, accidentally sends him to Hell. It doesn't take long before the engineer becomes rather dissatisfied with the level of comfort in Hell.*
>
> *He soon begins to design and build improvements. Shortly thereafter, Hell has air conditioning, flush toilets and escalators. Needless to say, the engineer is a pretty popular guy.*
>
> *One day, God calls Satan and says: "So, how are things in Hell?"*
>
> *Satan replies: "Hey, things are going great. We've got air conditioning, flush toilets, and escalators. And there's no telling what this engineer is going to come up with next."*
>
> *"What!" God exclaims: "You've got an engineer? That's a mistake - he should never have been sent to Hell. Send him to me."*
>
> *"Not a chance," Satan replies: "I like having an engineer on the staff, and I'm keeping him!"*
>
> *God insists: "Send him back or I'll sue."*
>
> *Satan laughs uproariously and answers: "Yeah, right. And where are you going to get a lawyer?"*

Connect

The Newsletter

The newsletter is published three times annually, in February, May and October and sent to paid-up members of the Safety-Critical Systems Club.

Do you have a topic you'd like to share with the systems safety community? Perhaps an interesting area of research or project work you've been involved in, some new developments you'd like to share, or perhaps you would simply like to express your views and opinions of current issues and events. If you are interested in submitting an article, then get in touch with the Newsletter Editor to discuss ideas: paul.hampton@scsc.uk

The SCSC Website

Visit the Club's website scsc.uk for more details of the Safety-Critical Systems Club including past newsletters, details of how to get involved in working groups and joining information for the various forthcoming events.

Twitter

Follow the Safety-Critical Systems Club's Twitter feed for brief updates on the club and events: @SafetyClubUK

LinkedIn

You can find the club on LinkedIn. Search for the Safety-Critical Systems Club or use the following link:

www.linkedin.com/groups/3752227

Advertising

Do you have a product, service, event or job vacancy you would like to advertise in the Newsletter? The SCSC Newsletter can reach out to over 1,000 members involved in Systems Safety and so is the perfect medium for engaging with the community. For prices and further details, please get in touch with the Newsletter Editor.

SCSC Working Groups

The Safety-Critical Systems Club is committed to supporting the activities of specialist working groups for areas of special interest to club members. The purpose of these groups is to share industry best practice, establish suitable work and research programmes, develop industry guidance documents and influence the development of standards.

Assurance Cases

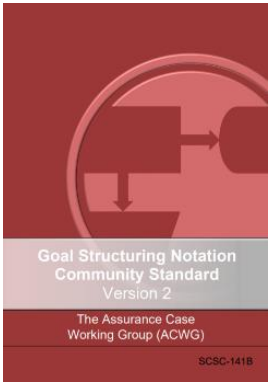

The Assurance Cases Working Group (ACWG) has been established to provide guidance on all aspects of assurance cases including construction, review and maintenance. The ACWG will:

- Be broader than safety, and will address interaction and conflict between related topics
- Address aspects such as proportionality, rationale behind the guidance, focus on risk, confidence and conformance
- Consider the role of the counter-argument and evidence and the treatment of potential bias in arguments

One of the working group's initial activities is to take on board the maintenance of the Goal Structuring Notation (GSN) Community standard.

Lead Phil Williams phil.williams@scsc.uk

Security Informed Safety

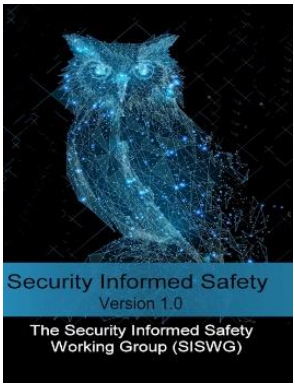

The Security Informed Safety Working Group (SISWG) aims to capture cross-domain best practice to help engineers find the 'wood through the trees' with all the different security standards, their implication and integration with safety design principles to aid the design and protection of secure safety-critical systems and systems with a safety implication.

The working group aims to produce clear and current guidance on methods to design and protect safety-related and safety-critical systems in a way that reflects prevailing and emerging best practice. The guidance will allow safety, security and other stakeholders to navigate the different security standards, understand their applicability and their integration with safety principles, and ultimately aid the design and protection of secure safety-related and safety-critical systems.

Lead Tom Turner tom.turner@scsc.uk

SCSC Working Groups

Data Safety Initiative

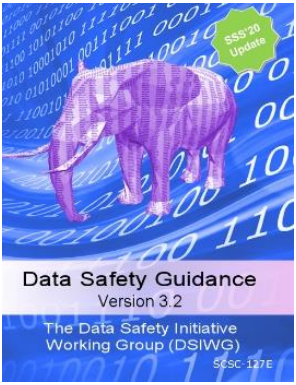

Data in safety related systems is not currently sufficiently addressed in current safety management practices and standards.

It is acknowledged that data has been a contributing factor in several incidents to date. There are clear business and societal benefits, in terms of reduced harm, reduced commercial liabilities and improved business efficiencies, in investigating and addressing outstanding challenges related to safety of data.

The Data Safety Initiative Working Group (DSIWG) aims to have clear guidance on how data (as distinct from the software and hardware) should be managed in a safety related context, which will reflect emerging best practice.

Lead **Mike Parsons** mike.parsons@scsc.uk

Safety of Autonomous Systems

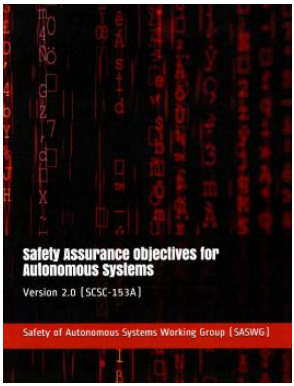

The specific safety challenges of autonomous systems and the technologies that enable autonomy are not adequately addressed by current safety management practices and standards.

It is clear that autonomous systems can introduce many new paths to accidents, and that autonomous system technologies may not be practical to analyse adequately using accepted current practice. Whilst there are differences in detail, and standards, between domains many of the underlying challenges appear similar and it is likely that common approaches to core problems will prove possible.

The Safety of Autonomous Systems Working Group (SASWG) aims to produce clear guidance on how autonomous systems and autonomy technologies should be managed in a safety related context, in a way that reflects emerging best practice.

Lead **Rob Alexander** rob.alexander@scsc.uk

SCSC Working Groups

Multi- and Manycore Safety

It is becoming harder and harder to source single-core devices and there is a growing need for increased processing capability with a smaller physical footprint in all applications. Devices with multiple cores can perform many processes at once, meaning it is difficult to establish (with sufficient evidence) whether or not these processes can be relied upon for safety-related purposes.

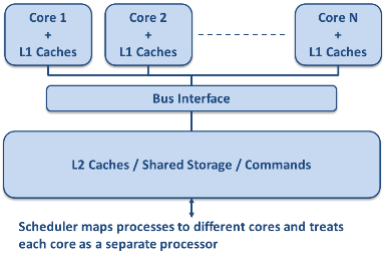

Parallel processes need to access the same shared resources, including memory, cache and external interfaces, so they may contend for the same resources. Resource contention is a source of interference which can prevent or disrupt completion of the processes, meaning it is difficult to know with a defined uncertainty the maximum time each process will take to complete (Worst Case Execution Time, WCET) or whether the data stored in shared memory has been altered by other processes.

The Multi- and Manycore Safety Working Group (MCWG) has been established to explore the future ways of assuring the safety of multi- and manycore implementations.

Lead Louise Harney louise.harney@scsc.uk

Ontology

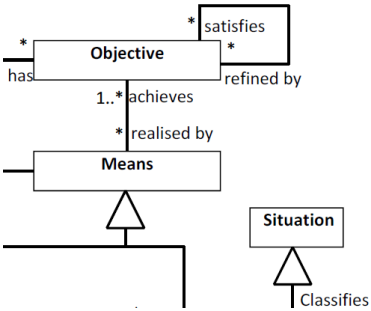

The Ontology Working Group (OWG) develops ontologies that will form the basis of SCSC guidance, as well as having wider industrial and academic applications.

The OWG is currently working on the definition of an ontology of risk for application in guidance for risk-based decision making - notably safety and security - and for which ISO 31000 Risk Management principles are to be applied.

The Data Safety Working Group (DSIWG) developed the core aspects of the Risk Ontology, which has been migrated to this working group. The Risk Ontology will form the upper ontology to the Data Safety Ontology that the DSIWG will continue to develop.

Lead Dave Banham ontology@scsc.uk

SCSC Working Groups

Covid-19

The Covid-19 Working Group is involved with discussion, analysis and assistance related to the Coronavirus. The group meets remotely to see what a systems and assurance view of the situation brings. The group has also created a discussion group on the SCSC website, for thoughts and ideas related to the work of the group.

Members are all experienced engineers, used to making reasoned arguments about safety. The aim is to apply the groups considerable technical expertise to the problem and find and assure appropriate solutions.

Lead Mike Parsons mike.parsons@scsc.uk

Service Assurance

Risks presented by safety-related services are rarely explicitly recognised or addressed in current safety management practices, guidelines and standards. It is likely that service (as distinct from system) failures have led to safety incidents and accidents, but this has not always been recognised. The Service Assurance Working Group (SAWG) has been set up to produce clear and practical guidance on how services should be managed in a safety related context, to reflect emerging best practice.

Lead Mike Parsons mike.parsons@scsc.uk

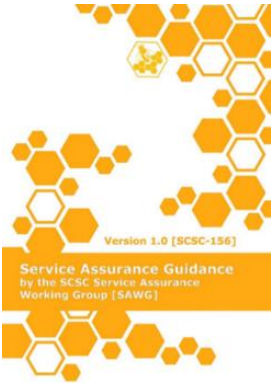

SCSC Safety Culture Working Group (SCWG)

A new working group is being established to provide guidance on creating and maintaining an effective safety culture. The group seeks to improve safety culture within teams working with safety-critical systems. This is a well-researched area with a lot of guidance already in place, so the inaugural meeting will consider what is working well, what is not effective, and where guidance is required. The meeting will also provide an opportunity to discuss any particular aspects attendees are interested in taking forward, and to agree a firm purpose for the SCWG.

Please contact waleed.chaudhry@scsc.uk or michael.wright@greenstreet.co.uk if you are interested in joining this group.

60 Seconds with… Professor Erik Hollnagel

Erik is an internationally recognised specialist in the fields of resilience engineering, system safety, human reliability analysis, cognitive systems engineering, and intelligent man-machine systems. He is the author of more than 500 publications including twenty-seven books, articles from recognised journals, conference papers, and reports.

What first attracted you to working in the field of System Safety?

It just happened

What aspect of your career are you most proud of?

That I am still learning

What advice would you give to yourself age 12?

You can do it!

What worries you the most about the future of System Safety?

The inertia of industries

What's your most favourite quote or motto?

"The difference between what we can imagine and what can happen is larger than we can imagine"

"The difference between what we can imagine and what can happen is larger than we can imagine"

If you could learn to do anything, what would it be?

To make a perfect Salzburger Nockerl

If you could be any fictional character, who would you choose?

Alice (in Wonderland)

The SCSC Steering Group

Tom Anderson

Honorary member

Stephen Bull

stephen.bull@scsc.uk

Jane Fenn

jane.fenn@scsc.uk

Louise Harney

louise.harney@scsc.uk

Brian Jepson

brian.jepson@scsc.uk

Graham Jolliffe

Honorary member

Alex King

alex.king@scsc.uk

Mike Parsons

mike.parsons@scsc.uk

Roger Rivett

roger.rivett@scsc.uk

Phil Williams

phil.williams@scsc.uk

Robin Bloomfield

Honorary member

Dewi Daniels

dewi.daniels@scsc.uk

Paul Hampton

paul.hampton@scsc.uk

James Inge

james.inge@scsc.uk

Nikita Johnson

nikita.johnson@scsc.uk

Tim Kelly

Honorary member

Mark Nicholson

mark.nicholson@scsc.uk

Felix Redmill

Honorary member

Emma Taylor

emma.taylor@scsc.uk

Sean White

sean.white@scsc.uk

Club Positions

The current and previous (marked in italics) holders of club positions are as follows:

Director

Mike Parsons 2019-

Tim Kelly 2016-2019

Tom Anderson 1991-2016

Steering Group Chair

Roger Rivett 2019-

Graham Jolliffe 2014-2019

Brian Jepson 2007-2014

Bob Malcolm 1991-2007

Programme & Events Coordinator

Mike Parsons 2014-

Chris Dale 2008-2014

Felix Redmill 1991-2008

Manager

Alex King 2019-

Newsletter Editor

Paul Hampton 2019-

Katrina Attwood 2016-2019

Felix Redmill 1991-2016

University of York Coordinator

Mark Nicholson 2019-

Website Editor

Brian Jepson 2004-

Administrator

Alex King 2016-

Joan Atkinson 1991-2016

Safety Future Initiative Lead

Nikita Johnson 2019-

Calendar

January

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31		

February

M	T	W	T	F	S	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	

March

M	T	W	T	F	S	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29

April

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30			

May

M	T	W	T	F	S	S
				1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

June

M	T	W	T	F	S	S
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30					

July

M	T	W	T	F	S	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31		

August

M	T	W	T	F	S	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30

September

M	T	W	T	F	S	S
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30				

October

M	T	W	T	F	S	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

November

M	T	W	T	F	S	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29

December

M	T	W	T	F	S	S
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

Events Diary

27 Apr-21 May 2020
Online Course

Online Course: Security for Safety Critical Systems (SESA)

scsc.uk/e692

18 May 2020
SCSC Group Meeting

Assurance Cases Working Group Meeting #12

London, UK

scsc.uk/e672

2 June 2020
SCSC Group Meeting

Service Assurance Working Group meeting #23

Farnborough, UK

scsc.uk/e681

POSTPONED
SCSC Seminar

New Safety Analysis Techniques

London, UK

scsc.uk/e654

7-10 Sep 2020
Conference

The 16th European Dependable Computing Conference - EDCC 2020

Munich, Germany

edcc.dependability.org

15-18 Sep 2020
Virtual Conference

39th International Conference on Computer Safety, Reliability and Security - SafeComp 2020

safecomp2020.di.fc.ul.pt

CANCELLED
SCSC Tutorial

Combining Safety with Security

London, UK

scsc.uk/e666

24 September 2020
SCSC Seminar

Safe Use of Multi-Core and Manycore Processors

London, UK

scsc.uk/e638

3 December 2020
SCSC Seminar

Management and Oversight of Complex Systems

Reading, UK

scsc.uk/e661

9-11 February 2021
SCSC Symposium

29th Safety-Critical Systems Symposium (SSS'21)

Bristol, UK

scsc.uk/e683

NB: all events are subject to postponement and cancellation due to the Covid-19 situation. Please check the SCSC website for up-to-date information: scsc.uk/events

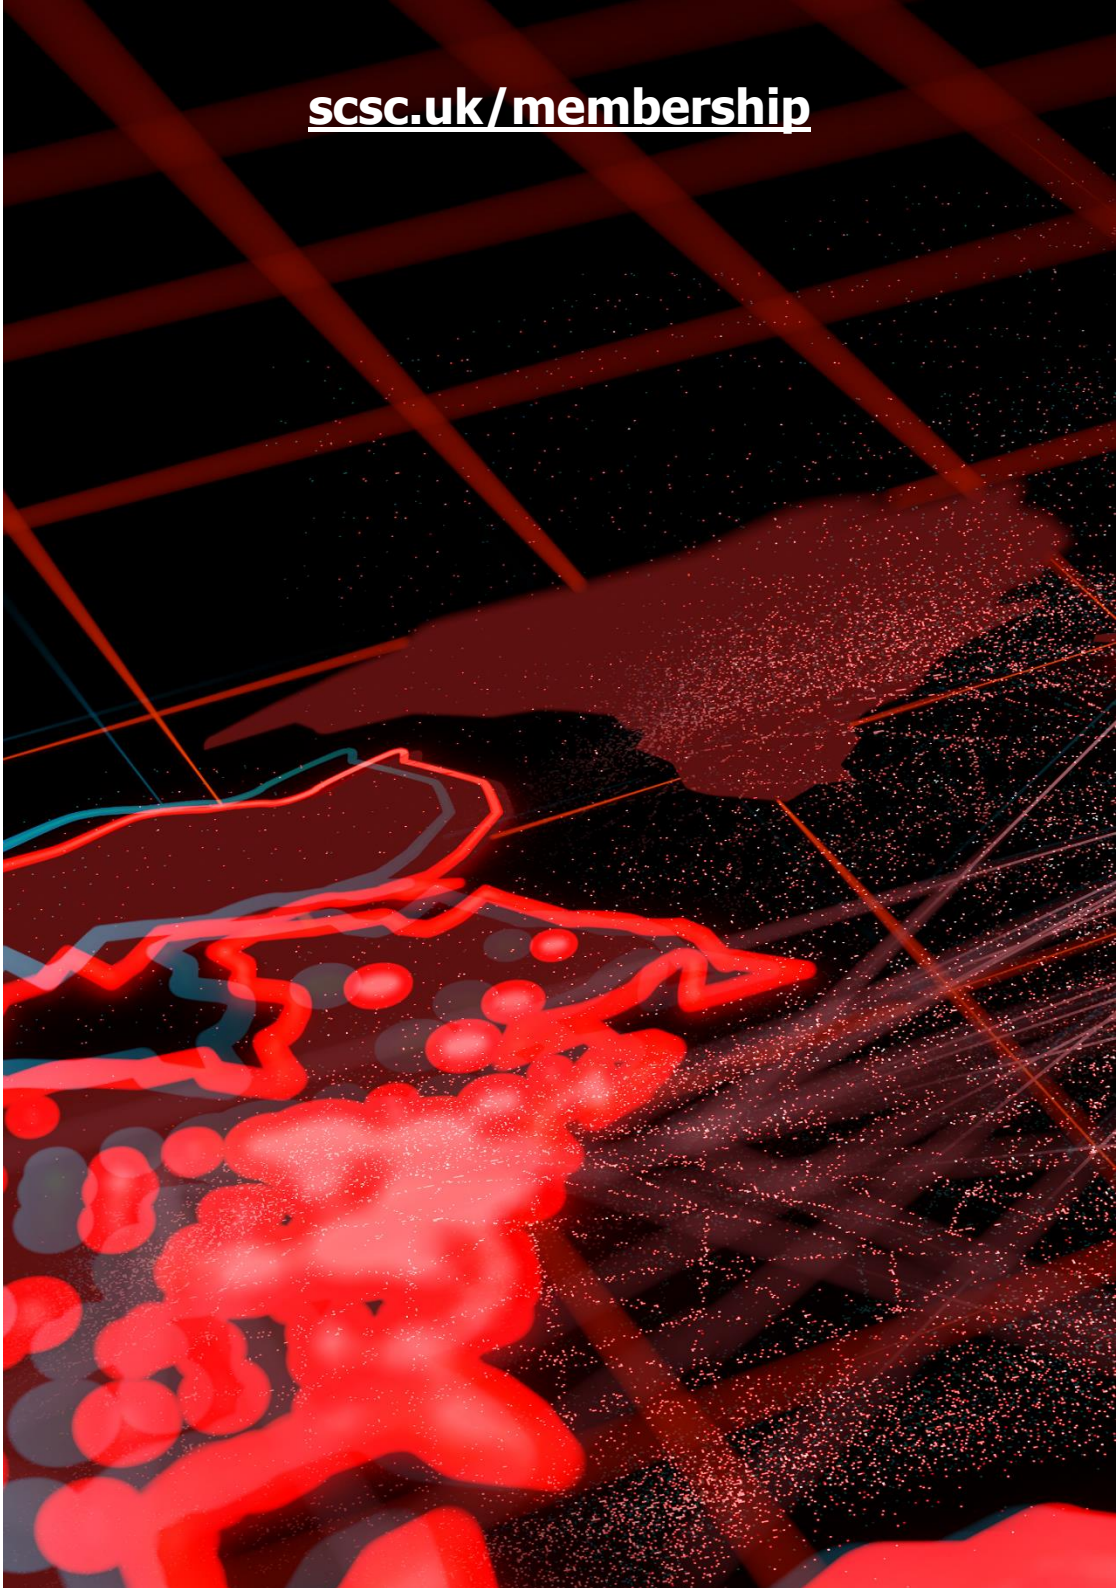
scsc.uk/membership