The Safety-Critical Systems Club Newsletter

Safety Systems

Vol 33 No. 1 – Feb 2025

QUANTUM LEAPS
The quantum future of Safety

TITANIC FAILURE
The fatal tale of the Titan submersible

Leading Women
Celebrating the women in engineering

A MODEL WORLD
Formalising world perspectives

SCSC

thescsc.org

SCSC Publication Number: SCSC-203

Contents

2024 Safety Systems Compendium now available on Amazon:

https://www.amazon.co.uk/dp/B0DQ4Z64RT

Editorial

Welcome to the first edition of *Safety Systems* for 2025!

As we embark on a new year and new challenges, it is timely to reflect on the year that has passed. 2024 saw the continued progression of Artificial Intelligence (AI) in many aspects of the work we do.

AI featured strongly in the annual SSS'24 symposium, with key themes being managing the risks around AI. Amongst other AI-related talks, a special panel session was held on the second day allowing delegates to pose AI-related questions to the panel of experts.

AI has also crept into the newsletter publications! There was a 60-second interview with ChatGPT and the results were, pretty remarkable, not least its predilection for extreme sports! I found myself also increasingly using AI generated images in the newsletters themselves, with the front cover of the 2024 compendium being a prime example of what can be done.

It seems unlikely that such a tailored image (through careful prompt selection) could be found from stock images. It also avoids thorny issues of rights to publish images. It is a sobering thought, but I can now no longer imagine a future where AI is not used in some capacity in all aspects of our working and domestic lives.

The SCSC Working Groups also continued to thrive with the Safe System Architecting, Safer Complex Systems and Safe AI groups being established last year. See the Working Group section for more information and if you are interested in these topics then why not get involved in the discussion!

And finally, we also saw the start of the Post Office Horizon IT Inquiry examining the Post Office's prosecution of over 700 sub-postmasters for theft, false accounting and related charges associated with technical faults in the Horizon IT system. While not usually considered a safety system, the magnitude of the impact on human life led the SCSC to extend its remit to cover any computer-based systems and services that could cause harm. Last year also saw the US Coast Guard's Enquiry into the OceanGate Titan submersible *accident*. I hesitate to call it an 'accident' –Wendy Owen's excellent leading article in this issue describing the events surrounding the implosion will hopefully show you why.

What will 2025 hold? Novel and emerging technologies like AI and Large Language Models will likely dominate, but if the Jeju accident in South Korea in December 2024 shows nothing else, system safety is never 'done' – diligence and vigilance is required even for routine, established and well-understood systems.

In this edition we also have articles discussing what Quantum computing holds in store for system safety and there's contribution on ontological modelling from the Ontology Working Group. Louise Harney talks to some of the winners of the recent IET's Young Woman Engineer Awards and it gives me great pleasure to report on the SCSC's first Seminar to be held in mainland Europe involving an epic trip to Munich, Germany!

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

In Brief






IET Young Woman Engineer of the Year 2024

The Institution of Engineering and Technology (IET) held an awards ceremony on 9th December 2024 celebrating the brilliant women showcasing their engineering excellence and truly engineering a better world. These prestigious engineering industry awards celebrate women working in modern engineering – and aim to help change perceptions of the industry.
youngwomenengineer.theiet.org

Train safety system failed in moments before fatal crash in Wales

An automated system that helps train wheels grip the tracks failed on one of the trains that crashed head-on in mid-Wales on 21 October, investigators have revealed. The Rail Accident Investigation Branch (RAIB) said the westbound train involved in the collision in Talerddig, near Llanbrynmair, in Powys, was fitted with a system to discharge sand automatically on to the rails should the wheels slide when braking. *theguardian.com*

S Korea orders air safety probe after deadly plane crash

South Korea's acting leader has ordered an emergency safety inspection of the country's entire airline operations, a day after 179 people were killed in the deadliest plane crash on its soil. The Jeju Air plane burst into flames after skidding off the runway and crashing into a wall. *bbc.co.uk*

Azerbaijan Airlines says there was 'external interference' before crash

The aircraft that crashed in Kazakhstan on Christmas Day, killing 38 people, experienced "external physical and technical interference", according to an investigation.

The plane was flying from the Azerbaijani capital, Baku, to the Russian city of Grozny in Chechnya when it crashed hundreds of miles off its planned route. The head of Russia's civil aviation agency said the aircraft tried to land in Grozny as the region was under attack by Ukrainian drones..
theguardian.com

Further disruption expected after latest NHS cyber attack

IT and security teams at Wirral University Teaching Hospitals NHS Trust worked around the clock following a major cyber incident. It is believed to have affected all clinical activity at multiple sites including Arrowe Park and Clatterbridge Hospitals.

The Trust was forced to cancel surgical procedures and turn away outpatients, although emergency care remained up and running. *computerweekly.com*

The Sinking of the Unsinkable (again)

Wendy Owen revisits a newsletter article from last year on the Titan submersible accident (June 2023), and uncovers a sorrowful tale of wayward tourism, bucket-lists of the rich, rogue designers, paid-off whistleblowers, a whole suite of independent assessors, test facilities, accreditation bodies and regulators "with no teeth", and dubious legal waivers…

As mentioned in my previous article on this topic ("A Brief Quest for Lesser-Spotted News Articles in Dependability", SCSC Newsletter Feb 2024), I feel reasonably qualified and experienced to talk about this as I've worked on similar submersibles in the distant past. The moral of the tale last time was that just because it appears reliable, dependable or a fantastic innovation, it doesn't mean it's safe, and doesn't mean you can ignore implementing a comprehensive reliability testing plan.

Over a year on, much more information has come to light as it has been through the courts – and it is worth revisiting this case. By the time of publishing this article, I'd hope that the investigation should have mostly concluded (hearings were ongoing during September and October 2024, but it then went quiet).

Firstly, I can confirm that the vessel was owned by an American tourism/expeditions company named OceanGate, and that towards the end of its dive it suffered a hull failure and imploded, with instant deaths of all five occupants. To ascertain how this occurred, we need to look back at a chain of events – in this case, the chain of events turned out to be much longer than I was expecting. Thus, this is a rather long article.

Why did Titan exist in the first place?

Tourism, for the wreck of the Titanic, sunk in 1912. The wreck lies at a depth of 12,500ft / 3,810m, 590km / 370 miles from the coast of Newfoundland. Apparently, by 2012, a century later, 140 people had visited the wreck site. One could ask why people have a need to view old shipwrecks 12,500ft below sea level, but shipwrecks feature in many maritime museums and obviously the Titanic is particularly famous. OceanGate realised that visiting shipwreck sites was a good income stream, and they could also charge a premium for the Titanic. Titan appears to have served no other purpose than for tourism, bucket-lists of the rich, and media attention – which ultimately, it received excessive amounts of, but for all the wrong reasons.

I propose a few other reasons for its continued existence during this article, namely; rogue inventors, designers who didn't whistle-blow, designers who did whistle-blow but were paid off, independent assessors who walked away from the scene at crucial points, test facilities who got worrying results but didn't halt testing, naïve businessmen, unbelievable legal waivers (presumably written by lawyers/solicitors), plus engineering industry societies, safety accreditation bodies and regulators "with no teeth". From a design and safety management perspective, this is a whole chain of very serious stuff! If you put a Swiss cheese model up against Titan's safety credentials, you would end up with grated parmesan.

What was its modus operandi?

Dives by Titan to the Titanic wreck occurred as part of multi-day excursions. The excursions were referred to as "missions", as per military language, and passengers were even referred to as "mission specialists"; they paid $250,000 each (on occasion discounted for multiple tickets). They would sail to and from the wreckage site aboard a support ship and then spend five days in the ocean above the Titanic wreckage site; eight-days in total.

Usually, two dives were attempted during each excursion, though dives were often aborted due to weather or technical malfunctions. Each dive typically had a pilot, a guide, and three paying passengers on board. Once inside the submersible, the hatch would be bolted shut and could only be reopened from the outside. The descent from the surface to the wreck took about two hours, with a full dive taking about eight hours.

According to OceanGate, the vessel contained monitoring systems to continuously monitor the strength of the hull and the vessel had life support for five people for 96 hours.

Throughout the journey, the submersible was expected to emit a safety "ping" every fifteen minutes, monitored by the above-water crew. The vessel and surface crew were also able to communicate via brief text messages. There is no GPS underwater; the support ship monitored the position of Titan relative to its target, and sent brief text messages to Titan providing distances and directions.

Five missions occurred in the middle of each of the two previous years, 2021 and 2022. Titan imploded during the fifth mission of 2023; it was the first mission that year in which a dive came close to the Titanic, previous attempts being aborted due to poor weather.

How was it designed and by whom?

It was formerly called Cyclops 2. Apparently, there were problems with Cyclops 1 (not investigated further here). Its basic design looked like this.

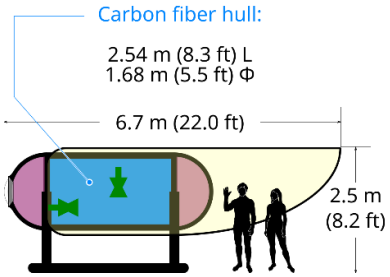

It wasn't very big. It was constructed from carbon fibre and titanium. weighed 10,432kg and could hold up to 5 people. Essentially, it was a pressure vessel made up of two titanium domes with matching titanium interface rings bonded to a carbon fibre-wound cylinder. One of the titanium hemispherical end caps could be detached to provide the hatch and was fitted with a 15-inch diameter acrylic window.

Titan could move at up to ~3 knots (5.6 km/hr, 3.5mph) using four electric thrusters, arrayed two horizontal and two vertical. Its steering controls consisted of a Logitech F710 wireless game controller with modified longer joysticks. The University of Washington Applied Physics Laboratory (UW APL) assisted with the control design on Cyclops 1 using a similar video game controller, which for Titan was substituted with the Logitech controller. Its steering gadgets were thus quite basic. The use of game controllers is nowadays not uncommon for some remote-control systems such as Uncrewed Aerial Vehicles (UAVs), bomb disposal robots and periscopes.

Titan steering control – Logitech F710

OceanGate claimed that Titan was the only crewed submersible that used an integrated real-time monitoring system (RTM) for safety. This system, patented by Rush – one of OceanGate's engineers - in 2021, used acoustic sensors and strain gauges at the pressure boundary to analyse the effects of increasing pressure as the vessel went deeper into the ocean and to monitor the hull's integrity in real time. This would supposedly give early warning of problems and allow enough time to abort the descent and return to the surface.

OceanGate claimed that Titan had several backup systems intended to return the vessel to the surface in case of emergency, including ballasts that could be dropped, a balloon, thrusters, and sandbags held by hooks that dissolved after a certain number of hours in saltwater (in theory this would release the sandbags, allowing the vessel to float to the surface). An OceanGate investor explained that if the vessel did not ascend automatically after the elapsed time, those inside could help release the ballast either by tilting the ship back and forth to dislodge it or by using a pneumatic pump to loosen the weights. Of course, this all assumed that the vessel remained intact and didn't travel beyond its bounding conditions.

In 2020, Rush said that the hull, originally designed to reach 4,000m/13,000ft below sea level, had been downgraded to a depth rating of 3,000m/9,800 ft after demonstrating signs of cyclic fatigue (recall that Titanic was at nearly 4,000m depth). A 1/3-scale model of the Cyclops 2 pressure vessel was built and tested at UW APL; the model was able to sustain a pressure of 4,285 psi (29.54 MPa), corresponding to a depth of about 3,000m. In 2020 and 2021, the hull was repaired or rebuilt (recall that dives occurred in 2021 and 2022).

In 2023, OceanGate claimed on its website that Titan was "*designed and engineered by OceanGate Inc. in collaboration [with] experts from NASA, Boeing, and UW*". After the implosion, these earlier associates disclaimed involvement with the project.

- Rush had told Travel Weekly that the carbon fibre for the hull had been sourced at discount from Boeing because it was too old for use in aeroplanes, but Boeing said they had no records of this. A Boeing spokesperson claimed Boeing "*was not a partner on Titan and did not design or build it*"
- UW claimed the APL had not been involved in the "*design, engineering, or testing of the Titan submersible*"
- A NASA spokesperson said that their Marshall Space Flight Centre had a Space Act Agreement with OceanGate, but had no connection to Titan

It subsequently became evident, however, that both UW and Boeing had put forth numerous design recommendations and rigorous testing requirements, which Rush apparently ignored despite the implosions during trials at UW APL. The partnerships dissolved as Rush refused to work within quality standards.

Safety Approvals and Use of Waivers

A note that this section is, for the most part, presented in reverse chronological order.

OceanGate had initially not sought safety certification for Titan, arguing that excessive safety protocols hindered innovation. Additionally, because Titan operated in international waters and did not carry passengers from a port, it was also not subject to (marine) safety regulations. It was also not certified by Lloyd's Register. Thus, it was not certified as seaworthy by any regulatory agency or third-party organisation.

A reporter who completed an expedition in 2022, as part of a CBS News feature, said that all passengers who enter Titan signed a waiver confirming their knowledge that it is an "*experimental*" vessel "*that has not been approved or certified by any regulatory body, and could result in physical injury, disability, emotional trauma or death*". A television producer who completed the expedition said the waiver "*mentioned death three times on page one*".

"At some point, safety just is pure waste. I mean, if you just want to be safe, don't get out of bed"

Stockton Rush

In a 2022 interview, Rush told CBS News, "*At some point, safety just is pure waste. I mean, if you just want to be safe, don't get out of bed. Don't get in your car. Don't do anything.*" In a 2021 interview he said "*I've broken some rules to make [Titan]. I think I've broken them with logic and good engineering behind me. The carbon fibre and titanium, there's a rule you don't do that. Well, I did.*" A 2019 article published in Smithsonian magazine referred to Rush as a "daredevil inventor", where he commented that the U.S. Passenger Vessel Safety Act of 1993 "*needlessly prioritized passenger safety over commercial innovation*".

Lloyd's Register refused OceanGate's request to class the vessel in 2019.

In 2018, OceanGate's director of marine operations, Lochridge, composed a report documenting safety concerns he had about Titan:

- He urged the company to have Titan assessed and certified by the American Bureau of

Shipping, but OceanGate had refused to do so, instead seeking classification from Lloyd's Register (who, as per above, refused)

- He said that the transparent viewport on its forward end, due to its nonstandard experimental design, was only certified to a depth of 1,300m/4,300ft - only a third of the depth required to reach the Titanic's wreck
- He said the RTM would "*only show when a component is about to fail – often milliseconds before an implosion*" and could not detect existing flaws in the hull before it was too late.
- He was concerned that OceanGate would not perform non-destructive testing on the vessel's hull before undertaking crewed dives and alleged that he was "*repeatedly told that no scan of the hull or Bond Line could be done to check for delaminations, porosity and voids of sufficient adhesion of the glue being used due to the thickness of the hull*"

Elsewhere, evidence pointed to the viewport being rated to only 650m/2,130ft; the viewport's engineer also prepared an analysis from an independent expert that concluded that the design would fail after only a few 4,000m dives.

OceanGate said that Lochridge – who was not an engineer – had refused to accept safety approvals from their engineering team and that their evaluation of the hull was stronger than any kind of third-party evaluation. OceanGate sued Lochridge for allegedly breaching his confidentiality contract and making fraudulent statements. He counter-sued, stating that his employment had been wrongfully terminated as a whistleblower for stating concerns about Titan's ability to operate safely. The two parties settled the case a few months later before it came to court. He (rightly) filed a whistleblower complaint with (American) Occupational Safety and Health Administration, but – crucially and unfortunately – withdrew it after the lawsuit was filed.

Late 2018, a group organized by Kohnen, the chair of the Submarine Group of the Marine Technology Society, drafted a letter to Rush expressing "*unanimous concern regarding the development of 'TITAN' and the planned Titanic Expedition*", indicating that the "*current experimental approach ... could result in negative outcomes (from minor to catastrophic) that would have serious consequences for everyone in the industry*". The letter said that OceanGate's marketing of the Titan was misleading because it claimed that the submersible would meet or exceed the safety standards of classification society DNV, even though the company had no plans to have it formally certified by DNV.

While the letter was never sent officially by the Marine Technology Society, it did result in a conversation with OceanGate that resulted in some changes, but in the end Rush "*agreed to disagree*" with the rest of the civilian submarine community. Kohnen told the New York Times that Rush had telephoned him after reading it to tell him that he believed industry standards were stifling innovation.

Another engineering signatory, Kemper, agreed to sign the letter because of OceanGate's decision not to use established engineering standards like ASME Pressure Vessels for Human Occupancy (PVHO) or design validation. He said the submersible was "*experimental, with no oversight*". Kohnen and Kemper stated OceanGate's methods were not representative of the industry; they were both members of the ASME Codes and Standards committee for PVHOs, which develops and maintains the engineering safety standards for submarines, commercial diving systems, hyperbaric systems, and related equipment. Kemper has published several technical papers on submarine windows, ironically including discourse on a need to innovate.

In March 2018, one of Boeing's engineers involved in the preliminary designs, Negley, carried out an analysis of the hull and emailed Rush directly stating, "*We think you are at high risk of a significant failure at or before you reach 4,000m. We do not think you have any safety margin.*" Apparently he included an ingenious graph of the strain of the design with a skull and crossbones at a red line of 4,000m (this approach could also be adopted elsewhere!)

Also in March 2018, McCallum, a major deep sea exploration specialist, emailed Rush to warn him he was potentially risking his clients' safety and advised against the submersible's use for commercial purposes until it had been tested independently and classified: "*I implore you to take every care in your testing and sea trials and to be very, very conservative.*" Rush replied that he was "*tired of industry players who try to use a safety argument to stop innovation ... We have heard the baseless cries of 'you are going to kill someone' way too often. I take this as a serious personal insult*". McCallum then sent Rush another email in which he said: "*I think you are potentially placing yourself and your clients in a dangerous dynamic. In your race to Titanic, you are mirroring that famous catch cry: 'She is unsinkable*". This prompted OceanGate's lawyers to threaten McCallum with legal action.

> **"… in your race to Titanic, you are mirroring that famous catch cry: 'She is unsinkable'"**
>
> **McCallum**

I can envisage your hair standing on end reading these enlightening paragraphs. In a nutshell, the vessel was operating beyond several of its structural subsystem's safety margins. Yet no one stopped it. The consequences were virtually inevitable.

Let's now look at what else occurred before the inevitable happened.

Previous incidents

In 2021, as mentioned, a new hull was constructed – this was after a previous hull cracked after only fifty dives, only three of which were to 4,000m. Scale models of the hull imploded at UW APL, so a different method of curing the hull was developed and this passed a full-sized pressure test at a facility in Maryland. Rush refused to construct new domes and other components from the failed submersible and instructed the engineers to salvage and reuse parts. Anonymous former employees told Wired that damage to the components could have weakened the join with the new hull. They also added lifting rings, which was previously warned against by engineers because the submersible could not handle any tension or load.

In 2022, Pogue, a reporter, was aboard the surface ship when Titan became lost and could not locate the wreck of the Titanic during a dive. His report for CBS News Sunday Morning,

which questioned Titan's safety, subsequently went viral on social media in June 2023 after the submersible lost contact with its support ship. In the report, Pogue commented to Rush that "*it seems like this submersible has some elements of MacGyvery jerry-rigged-ness*". He mentioned the Logitech game controller and that construction pipes were used as ballast.

In a 2022 dive to the wreck, one of Titan's thrusters was accidentally installed backwards and the submersible started spinning in circles when trying to move forward near the sea floor. In the BBC documentary "*Take Me To Titanic*", this issue was bypassed by steering while holding the game controller sideways. According to court filings, OceanGate reported that, in a 2022 dive, the submersible suffered from battery problems and, as a result, had to be attached manually to a lifting platform, causing damage to external components.

On 15 July 2022 (dive 80), Titan experienced a "*loud acoustic event*" as it was ascending, which was heard by the passengers aboard and picked up by Titan's real-time monitoring system (RTM). Data from the RTM later revealed that the hull had permanently shifted following this event.

Also in 2022, British actor and television presenter Ross Kemp, who had previously participated on deep sea dives for Sky History, had planned to mark the 110th anniversary of the sinking of the Titanic by recording a documentary in which he would undertake a dive to the wreck using Titan. Kemp's agent said that the project was cancelled after checks by Atlantic Productions deemed the submersible to be unsafe and "*not fit for purpose*".

The fatal dive

Here I will only summarise what happened on the day of the accident, as to do otherwise would significantly extend this already long article, and I purposely want to draw more attention to the design and pre/post-accident discussions.

The voyage was booked in early 2023. One American businessman was offered two tickets at discount ($150,000 each) with Rush claiming it was "*safer than crossing the street*" – the businessman declined due to safety concerns. The excursion was due to happen in May but due to untoward weather (apparently the worst in Newfoundland in 40 years) it was delayed until a better window in June. The MV Polar Prince transported Titan and the expedition's crew to the dive site on 16-17th June. The dive was on 18th June.

For the first hour and a half of the descent, Titan communicated with Polar Prince via text about every fifteen minutes and received a "ping" every 5-10 seconds. Final communications were sent at 10:47 at an approximate depth of 3,341m/10,961ft. At the point of implosion, the submersible had completed, or almost completed, its dive. Comms messages indicated they had visibility of the Titanic. It was expected to resurface at 16:30. The US Coast Guard (USCG) was notified that the vessel was missing at 19:10.

Search and Rescue efforts took around 4 days; at the time, it was viewed as search & rescue rather than a wreckage recovery. The search area was informed by the US Navy's sonar detection of an acoustic signature consistent with an implosion around the time communications with the submersible ceased, suggesting the hull had imploded while Titan was descending. Search efforts involved the USCG, Canadian coast guard, US Navy (including a US Navy deep-sea salvage system that arrived but didn't get used), Royal Canadian Air Force, US Air National Guard, a Royal Canadian Navy ship, several commercial and research ships

and Remotely Operated underwater Vehicle (ROVs), and passing aircraft. The surface was searched as well as underwater, by sonar. The search covered 10,000 square miles.

The remote location, weather, darkness, sea conditions and water temperature all made the search difficult. Many submersibles have acoustic beacons, but apparently Titan did not. There were a few false alerts to noises and debris. Eventually, an ROV owned by Pelagic Research Services discovered a debris field containing parts of Titan, five hours into its search.

Remains of Titan on Ocean Floor

At a U.S. Coast Guard press conference in Boston, they said that pieces of Titan had been found in two areas on the sea floor approximately 500m northeast of the bow of the Titanic. The debris consisted of the tail cone (not part of the pressure vessel) and the forward and aft end bells (both part of the pressure vessel) with the aft-end bell lying separate from the front-end bell and the tail cone. They also said that the debris was consistent with a "catastrophic loss of the pressure chamber" i.e. an implosion.

Pelagic Research Services confirmed on 23rd June that a new mission to the Titan debris field was already underway and that it had taken the ROV one hour to reach the site to continue searching and documenting debris. The debris was too heavy for Pelagic's ROV to lift and a recovery would need to occur at a later time.

On 24th June, Polar Prince returned to St. John's harbour, where investigators boarded the ship, and also, the launch and recovery system (LARS), which Titan used.

Shortly after the disaster, a well-known deep sea Canadian filmmaker indicated that it was likely the submersible's early warning system alerted the passengers to an impending delamination of the hull, saying *"we understand from inside the community that they had dropped their ascent weights and were coming up, trying to manage an emergency*." A retired Naval officer (Ballard, the discoverer of the Titanic wreck) also said that the crew was likely "*experiencing difficulties*" and was trying to ascend at the time of the implosion. However, in September 2024, Catterson, an OceanGate contractor who was aboard the Polar Prince at the time of the disaster, testified at the United States Coast Guard's inquiry that there was no indication that the crew was aware of any problems before the implosion.

The last human-written communication by Titan indicated that they dropped two weights, amounting to 32 kg, about a tenth of the drop weights on board, but this was apparently routine to adjust the Titan's buoyancy from negative to neutral as it approached the seabed, and was an indication that the crew was not aware of any emergency situation. The last automatic ping was received by the Polar Prince approximately six seconds later, after which contact was lost.

Simulations developed in 2023 suggest the implosion of the vessel took less than one second, likely only tens of milliseconds, faster than the brain can process information; there would not have been time for the victims to experience the collapse of the hull, and they would have died immediately.

"the implosion of the vessel took less than one second, likely only tens of milliseconds, faster than the brain can process information"

On 28th June, Horizon Arctic returned to St. John's Harbour with the remains of Titan that were recovered from the debris field. Photographs and videos showed the titanium covers on both ends of Titan intact, with the single viewport missing, mangled pieces of the tail cone, electronics, the landing frame and other debris. The debris was to be transported to the U.S. as evidence for the investigation. The USCG confirmed that presumed human remains were found within the debris, and that American medical professionals would conduct an analysis. In September 2024, during the public hearing by the Marine Board of Investigation, USCG confirmed that the Armed Forces DNA Identification Laboratory positively identified DNA profiles for the five victims.

On 30th June, Insider published an analysis of the recovery photos by Plymouth University professor Graham-Jones. He concluded that a failure of the carbon-fibre hull was the most likely cause of the loss, given that no large pieces of carbon fibre are known to have been recovered. Another possible cause was the acrylic viewing window, since the window was absent from its bell housing when it was recovered. While the salvage team may have removed the window before salvaging its bell housing, they more likely would have left it in place. He said that if the window had failed before the hull rather than after, he would have expected larger pieces of carbon fibre to be recovered.

In early October, engineers recovered the endcap and the rest of the debris.

Financials and Taxpayers

A defence budget expert estimated the costs of U.S. Coast Guard operations alone at about $1.2 million of taxpayers' money, with the additional operations to recover the debris not included. A Canadian taxpayer's estimate was at least $3 million. The US National Association for Search and Rescue said the search for Titan was likely to cost millions of dollars of public funds; however, the USCG refused to give an estimate, saying they "*do not associate cost with saving a life*". According to a U.S. attorney, the USCG is generally prohibited by federal law from collecting reimbursement related to any search or rescue service.

The incident renewed past debates about whether taxpayers should bear the cost of search and rescue missions involving wealthy people engaged in high-risk adventuring. The costs and scale of the search and rescue efforts for Titan sparked criticism in the media when compared to those for the Messenia migrant boat disaster. This occurred only days earlier and involved a fishing boat that sank while carrying an estimated 400 to 750 migrants,

resulting in around 100 deaths, 100 rescued, and hundreds more missing and presumed dead. This event also received significantly less media attention (noted by former US president Barack Obama).

Investigations and Lawsuits

On 23rd June, both the Canadian and the United States federal governments announced that they were beginning investigations of the incident. They were subsequently joined by authorities from France (Bureau d'Enquêtes sur les Événements de Mer, BEAmer) and the UK's Marine Accident Investigation Branch (MAIB); the final report is expected to be issued to the International Maritime Organization (IMO). Whether any lasting reforms will result from the investigation is uncertain, because the IMO may not have appropriate regulatory authority. Either way, parties are advising that the investigation is a complex and ongoing effort.

Even the Royal Canadian Mounted Police (RCMP) announced that it was performing a preliminary examination of the incident in order to determine whether to begin a full investigation, which will occur if it determines that criminal, federal or provincial laws were broken.

From 2021 to 2023, the Titan conducted seven dives in Canadian waters and three dives in Canada's Exclusive Economic Zone (EEZ). During this same timeframe it also conducted nineteen dives outside Canadian waters and Canada's EEZ, which included its dives to the Titanic. For each of these dives, the Titan was transported to the dive site from a Canadian port and returned to a Canadian port, using a Canadian-flagged vessel. During these operations, the Titan and its launch platform were not registered or certified in Canada or any other country. The investigation also identified other submersibles operating within Canadian waters and Canada's EEZ, both before and after June 2023. Some are registered in Canada; some are registered outside of Canada; and some are not registered. As a result of this information, the TSB issued a marine safety information letter (MSI 01/24) to Transport Canada advising of the risk posed by submersibles operating in Canadian waters.

The investigation is now in the report phase. On 6th August 2024, the family of Nargeolet (one of the deceased) sued OceanGate for wrongful death.

In the Media

Scott-Beddard, the CEO of White Star Memories Ltd, a Titanic exhibition company, observed that the likelihood of performing future research at the Titanic wreck decreased due to the incident.

James Cameron, who directed the 1997 movie Titanic, visited the Titanic wreck 33 times, and piloted Deepsea Challenger to the bottom of the Mariana Trench, said he was *"struck by the similarity*" between the submersible's implosion and the events that resulted in the Titanic disaster. He noted that both disasters seemed preventable and were caused indirectly by someone deliberately ignoring safety warnings from others. He criticised the choice of carbon-fibre composite construction of the pressure vessel, saying it has "*no strength in compression*" when subject to the immense pressures at depth; the wound carbon fibre of Titan's hull had seemed like a bad idea to him from the beginning. He stated that it was long known that composite hulls were vulnerable to microscopic water ingress, delamination, and progressive failure over time.

He also criticised Rush's real-time monitoring of the hull as an inadequate solution that would do little to prevent an implosion. Cameron expressed regret for not being more outspoken

about these concerns before the accident and criticised the "*false hopes*" presented to the victims' families; he and his colleagues realised early on that for communication and tracking to be lost simultaneously, the cause was almost certainly a catastrophic implosion.

Also in the media:

- Apparently the Logitech F710 game controller used to steer Titan sold out on Amazon soon after the incident, described as "*a more benign form of disaster tourism*" by the New York weblog The Cut
- Titan became widely discussed on social media as the story developed and was the subject of "*public schadenfreude*" inspiring grimly humorous Internet memes, namely interactive video game recreations and image macros that ridiculed the submersible's deficient construction, OceanGate's perceived poor safety record, and the individuals who died. The memes were criticised as insensitive, and some have felt the negative reaction to the victims may be a response to past news coverage of other expeditions by billionaires. The incident was widely treated on social media as entertainment. Major elements include the allure of disasters, fascination with the wealthy, conspiracy theories, uncertainty, and the mythology of the Titanic, as well as the romance of rescue operations
- In September 2023, it was announced that a new movie about the Titan submersible incident, named *Salvaged*, was in development
- The 2024 American Broadcasting Company (ABC) special *Truth and Lies: Fatal Dive to the Titanic* examined the implosion of Titan
- In February 2024, a movie inspired by the events of the Titan submersible incident, entitled *Locker*, was announced
- In March 2024, a two-part documentary by ITN Productions*, Minute by Minute: The Titan Sub Disaster*, was broadcast by Channel 5 (UK). The documentary included interviews with the Canadian air crew that searched the surface, the Pelagic ROV team that found the wreckage, and members of the Marine Technology Society (Kohnen and Kemper, mentioned above). Analysis of the mysterious "*banging*" sounds that seemed to indicate the occupants were still alive was a main feature of the first part

End Notes. What is Tolerability?

The investigation discovered the company had not sought certification for the vessel, arguing that "*excessive safety protocols and regulations*" hindered innovation. I commented previously that in this case it became obvious that 'successful' operations and the purported reliability of the missions to date (not many) had been prioritised over real design safety margins, by a long chalk. Issues were raised by the design team but ignored. It hadn't been rigorously tested before going into operation. These assertions still hold.

Yet, the court case has also revealed a plethora of other issues; **it was clearly being operated beyond its design safety margins but the passengers signed legal waivers essentially alerting them of their impending death. This gives a whole different perspective to the concept of 'tolerability' and 'acceptability' beyond the norms of what one would usually come across** in a civil or military ship safety case. Safety in the maritime sector covers concept, feasibility, design, build, test, operation and maintenance, plus intentional, inadvertent or accidental disposal, and now, evidently, disaster tourism.

The lack of regulation – or, more specifically, regulation with teeth – for such vessels is an obvious loophole in the current system. Maybe regulators didn't think such a dangerous vessel operating well beyond its safety margins would ever be encountered, so it was not on anyone's radar. The rogue inventor scenario was not considered. The wealthy who would take the risk in the rogue invention were not considered. There were too many 'Yes Men' in the chain of events, and one too many out-of-court settlements for whistleblowers. The fact that there had been whistleblowing on technical matters should have raised red flags somewhere. It seems lawyers were happier to support settlements than safety.

Moreover, safety and reputation don't matter to some people when their end goal is publicity, bravado and extreme tourism, regardless of what anyone else thinks. Not 'following the rules' and ignoring red is no excuse to you and me, but you must get inside the mind of those who see it as a big blocker to understand why they evade them – it's a blocker to their own personal pursuits. But it is to their own blinkered peril.

This is a very different situation to the Titanic; a well-built, large, purportedly unsinkable ship with many safety features that happened to be unlucky enough to crash into a ship-sinking iceberg. No one who designed or built Titanic was seen as a rogue inventor, although there are reports of the captain being a bit of a rogue operator ignoring warnings of icebergs. The fares were also a lot cheaper – apparently £1,000-£3,000 in today's money.

Wendy Owen

Wendy has 35 years' experience in systems safety and assurance across a wide range of highly regulated industries covering energy, defence, transport and process sectors. Her career has largely been in engineering consultancy, also with time in product development, research and regulatory. She is a Fellow of the Safety & Reliability Society and WES, and within SCSC a task lead in SiSWG and a member of the Steering Group. Early in her career she was involved in safety cases for various frigates, S, T & V Class submarines and other naval submersibles, and dockyards. One of her great-grandfathers was a 'marker-off' at Barrow shipyard, one grandfather was a sea-based merchant who rode on similar vessels to the Titanic, the other was a printer for a national newspaper covering such stories, her father was a Royal Marines Commando, and a former manager was an advisor to the Kursk submarine accident. She lives by the sea and does not have a boat.

References

At time of writing, the Wikipedia entry for the Titan accident has nearly 200 references, such is the appetite for people writing about this case. In writing this article, we have distilled data/information from various websites, to make all of it more readable for a safety-oriented audience. A references list is therefore not included on this occasion as it would take up the whole newsletter, although ironically perhaps this article may end up on that list one day.

Key websites referred to are those of the 'regulators' and other industry bodies discussing the court case and incident reports, as mentioned in the text.

Photographic credits to:

schematic: Mliu92, CC BY-SA 4.0 <https://creativecommons.org/licenses/by-sa/4.0>, via Wikimedia Commons
logitech control: Wstomsk, Creative Commons Attribution-Share Alike 3.0 Unported license.
Titan under water: https://en.wikipedia.org/wiki/File:Titan_submersible.jpg, reproduced under US fair use policy
Madelgarius, CC BY-SA 4.0 <https://creativecommons.org/licenses/by-sa/4.0>, via Wikimedia Commons
Wreck images: United States Coast Guard, Public domain, via Wikimedia Commons

Tangling with the Entangled – The impacts of quantum computing on safety

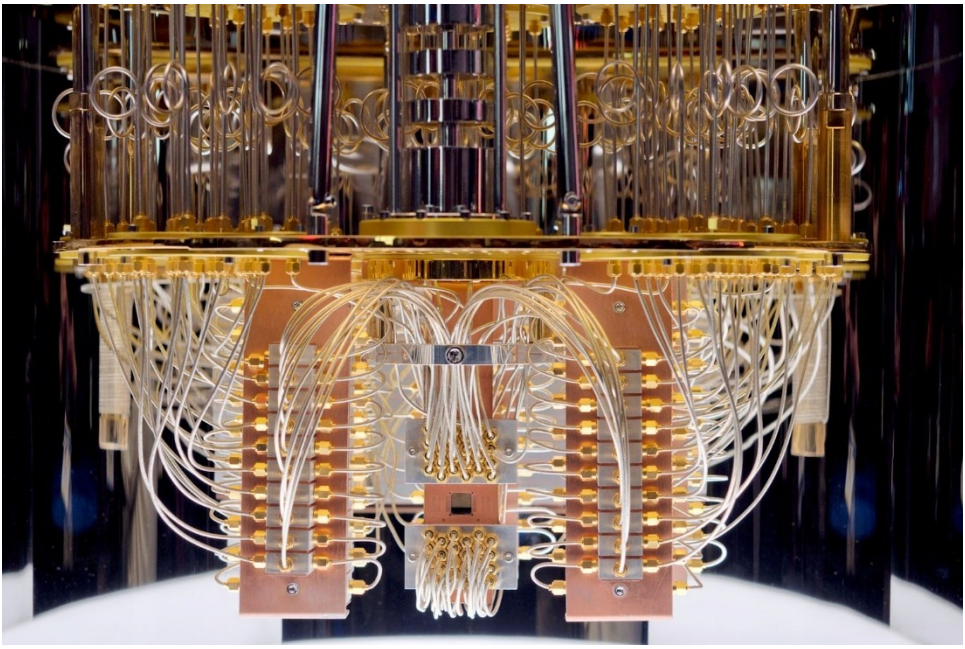

Quantum computers are a new way of processing information that promise to solve a host of problems facing humanity, but they also carry risks. James Cruise and Bob Oates discuss the quantum future of safety.

The safety community has long known that when new technologies emerge, they bring both new risks and new opportunities. On the horizon is quantum computing: complex, difficult to understand, but incredibly powerful.

While widespread access to this technology is likely to be at least a decade away, the rapid advances in the size and quality of quantum computers mean that now is the time to be preparing for how this new technology will interact with safety.

But if we are to prepare, we must first understand what the impact of quantum computers will be, in the world of safety. What is quantum computing? What opportunities do quantum computers present to enhance safety? What new risks do quantum computers introduce? Most importantly, what challenges must be addressed before quantum computers can be trusted to make safety critical design decisions?

We'll address these questions in turn, providing a high-level introduction to the technology, and going on to present the good (safety enhancements), the bad (risks that they undermine safety), and the ugly (problems that have not yet been well characterised.)

Quantum Computing

The proponents of quantum computing view the emergence of quantum computers as a revolutionary step for society akin to the way steam power drove the industrial revolution and our ability to manipulate electrons drove the computing revolution. The ability to control and interact with the quantum mechanical world promises an alternative computing paradigm that natively processes information by representing data and transformations using structures that are more analogous to mathematical concepts such as matrices and linear algebra than the Boolean algebra used by classical computers. This new way of representing and reasoning about information makes some problems, thought to be intractable for even the most powerful supercomputers, solvable for the first time in human history.

A common challenge for quantum computing is that in this early stage of development, many of the discussions about its potential are driven by quantum physicists, who have their own terminology and concepts, such as the Bloch sphere pictured as a representation of a quantum bit (or "qubit.") But just as modern programmers do not need a working knowledge of the semiconductor junctions that enable transistors, it is useful for many to think about what these devices can do, rather than focus on how they do it.

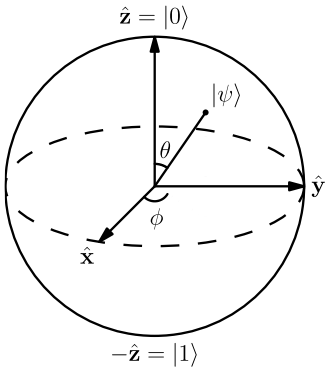

There are a wide variety of fantastic references for interested readers to look at if you want to explore the "how" of quantum computers. We suggest Nielson and Chuang's, "Quantum Computation and Quantum Information" [1] as a good starting point.

For our discussion it is sufficient to know that the output of a quantum computer will still be a sequence of numbers, the output of a chain of operations. As for classical computers, this output will just be a binary string. However, unlike classical computing while carrying out the computation, quantum computing considers all binary strings by associating an 'amplitude' with each one, this describes the probability for each string that this will be returned as the output string. Computation is then carried out by manipulating these amplitudes.

For example, consider a computation on 3 qubits. There are 8 possible binary strings that they could represent. For argument's sake we will assume that each string is equally likely at the beginning (not necessarily true in all cases!)

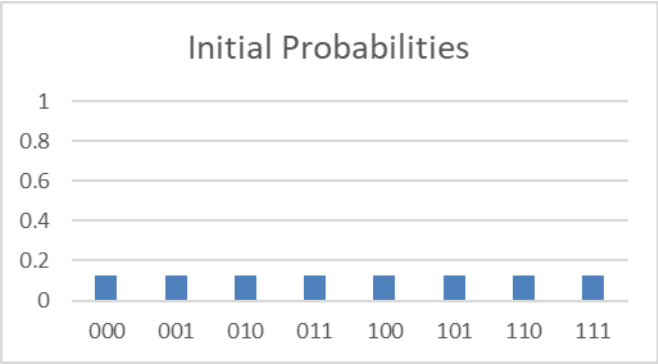

When a quantum computer processes those bits, it uses gates to manipulate those probabilities.

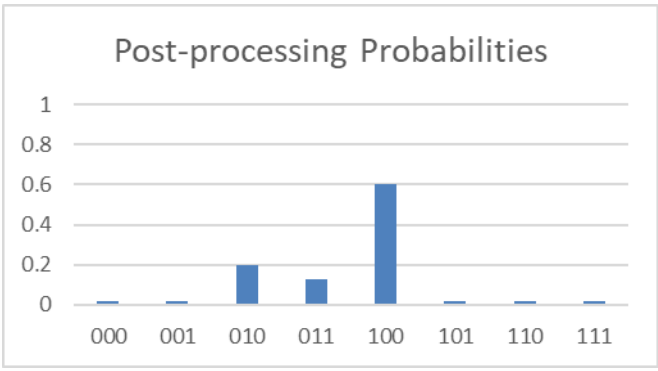

At the end of the calculation we "measure" the qubits which produces a random binary string using the assigned distribution. In our example we are more likely to get the answer "100" but "010" and "011" are also reasonably likely to occur, and all the strings have a chance.

This departure from traditional binary representation allows some computational problems to be solved more efficiently. In some cases, this means a quantum computer offers solutions that are polynomial faster and even exponentially faster for specific examples. Problems that have already been shown to benefit from this approach include:

- Finding the best entry in an unstructured list – Grover's algorithm
- Simulation of quantum mechanical systems – Hamiltonian evolution
- Factoring large numbers – Shor's algorithm

These abstract problems have huge real-world ramifications. Grover's algorithm can be applied to combinatorial optimisation or satisfiability problems ultimately offering more environmentally friendly logistics, efficiency savings for organisations, and better utilisation of resources. Hamiltonian evolution can be used to design new materials with beneficial properties for sectors such as aerospace, classical computing, and communications. The most explored application of Shor's algorithm is in the field of cryptanalysis (more on that later!)

A small number of subroutines have also been identified with the potential of creating other beneficial algorithms. These include "the quantum Fourier transform", and amplitude and phase estimation algorithms.

It's important to temper the enthusiasm for quantum computers with a dose of reality. They will not be useful for all problems, and they have a several inherent drawbacks that means that rather than replacing traditional computing, quantum computers will likely augment traditional computing, allowing us to take advantage of the best paradigm for the job on a case-by-case basis.

"It's important to temper the enthusiasm for quantum computers with a dose of reality ... the primary challenges facing quantum computers can be summed up as them being error prone, slow, and complex."

The primary challenges facing quantum computers can be summed up as them being error prone, slow, and complex.

Error-prone because unlike classical computing where the data representation using voltages has huge redundancy and hence natural error resilience, in qubit representation there is no such redundancy. This means any slight interference or misalignment can lead to an error that needs to be dealt with. Further, the probabilistic nature of quantum computers means that each algorithm really needs to be run several times before you can be confident in an answer. In short, if you want to add two numbers together, you're always going to be better off with a classical computer!

Slow because there are hard physical limits to the speed of quantum computers that are dictated by the physical properties of the systems used to implement them. For trapped-ion quantum computers this limit is the rotational speed of an ion, which fixes quantum computers to be in the region of 1000 times slower than a modern GPU. Now consider that you have to run error correction and programs multiple times on top of this, all compounding the slowdown in comparison to classical computers.

Complex because the laws of quantum physics introduce counter-intuitive properties into the system. For example, a property called "the no cloning theorem" means that you cannot just copy data from one place to another preventing branching calculations. Entanglement means that changing the copy also changes the original. This means that for computational tasks that rely on iterating around a single data set (for example, training AI) you are forced to reload the data into the computer every iteration and you cannot use checkpointing to protect your calculation.

These limitations ultimately prevent quantum computers from solving the data deluge that we're experiencing with modern computing. Quantum computers will be limited to high computation, low data applications, with classical, and maybe even newer paradigms, picking up the slack everywhere else.

The Good – Enhancing safety with quantum computers

With these limitations in mind, we can explore the areas where quantum computers could potentially benefit the safety community. Here we present two scenarios, one rooted in a well-understood and researched application domain for quantum computers, and a second, in a more speculative space that is still being characterised by researchers.

Our first scenario builds on one of the most promising commercial uses for quantum computing: chemistry and materials science. Consider a world where chemical experiments can be analysed entirely virtually, without the need for human operators to go anywhere near volatile or harmful byproducts. Automated systems, capable of exploring a previously inconceivable range of possible permutations in order to develop exotic materials with properties that enable lighter, stronger, more conductive building blocks for the systems of tomorrow. As this technology becomes more common-place, bespoke materials, with exactly the right properties for their application can be specified, requested, and formulated to resist environmental wear, avoid the creation of stress fractures, or prevent catastrophic collapse.

The virtualisation of chemistry not only makes experiments more efficient, but also enables properties to be explored that can't be directly observed during live experiments, such as reaction rate. These insights will allow for the creation of high-fidelity digital twins that can pre-empt maintenance needs and schedule corrective actions for potential failures before they happen.

Our second scenario targets one of the most challenging problems in software safety today: software verification for complex systems. As software becomes more complex there is an explosion in the space of input parameters and possible outputs, that makes testing every scenario impossible. This is made more complex still when software systems interact, generating even more permutations. One approach to give confidence that software will act as intended is formal verification. Formal verification can be viewed as a logical satisfiability problem whereby a number of logical clauses are created, to represent undesirable states within the software. For example, one clause may represent the activation of a function that should not be reachable without certain safeguards being active (such as allowing a dangerous piece of machinery to operate whilst its protective inspection cage is open.) We can now search the software for input parameters that cause undesirable conditions. If the resulting clauses produce an unsatisfiable problem, then there exists no combination of events in the software's execution that can result in the undesirable condition. If a solution is found, then that solution represents a way for the software to behave in an unsafe manner. Currently, a key problem with this approach is that solving large-scale satisfiability problems is computationally expensive. However, quantum computers have the potential to tackle much larger satisfiability problems, and thus formally verify much more complex software.

"Every discipline, safety included, is likely to be challenged by quantum computers to change the way they operate in response to revolutionary new possibilities."

Of course, both these scenarios raise new questions about how we assure the output of our new quantum algorithms. How much testing needs to be done to gain confidence that our highly specified super-material doesn't have a new failure mode that wasn't tested for? How much tool validation is required of a quantum verification to provide the confidence to stop running traditional tests? The former question starts to blend the problems of software validation with the problems of mechanical engineering, whereas the latter can be viewed as an extreme case of tool qualification.

Every discipline, safety included, is likely to be challenged by quantum computers to change the way they operate in response to revolutionary new possibilities. In fact, quantum computing is in such an exciting period of discovery that the most important contributions to safety have probably not been found yet. We need to be ready to embrace those new opportunities as they arise. It is noteworthy that we have constrained ourselves in this article to only explore the benefits of quantum *computers*, not the wider benefits from the broader family of quantum technologies such as quantum sensors (like the quantum magnetometer pictured), which promise new opportunities for non-invasive medical monitoring, and industrial system monitoring.

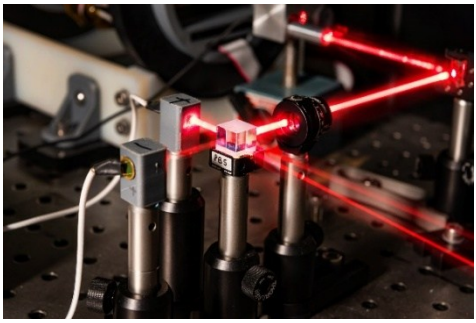

The Bad – The risks to safety posed by quantum computers

The most talked about quantum computing risk is the threat that they pose to cryptography.

Modern cryptography relies on a family of mathematical problems being "computationally infeasible" to solve with traditional computers. Computational infeasibility implies that a traditional computer would take in the order of millions of years to find a solution. Sadly, several of these cryptographic mathematical problems could be solved in a matter of hours by quantum computers with the right specification. This raises the spectre that an adversary with access to a quantum computer could launch sophisticated cyber-attacks against systems whose safety and security relies on cryptographic techniques to prevent malicious harm.

The attacks that quantum computers could enable include reading confidential information in transit, sending fake messages/commands that look like they come from trusted devices or people, and signing malicious software so that it looks like it was produced by a trusted source. In terms of the SCSC's Data Safety Guidance [2] HAZOP guidewords: integrity, intended destination/usage, traceability, and fidelity can all be compromised by an attacker with a quantum computer, if vulnerable cryptography was the sole protection in place.

The complexity of quantum computing has made it difficult for many people to tell vulnerable algorithms and quantum resistant algorithms apart. Judging if an algorithm is vulnerable requires an understanding of the underlying mathematical problem that a cryptographic control relies on and, where appropriate, the size of the key that it uses. Looking at encryption algorithms specifically, an important variable to consider is if the algorithm of interest is asymmetric or symmetric. For asymmetric encryption, the key used to encrypt the message (the public key) is different to the one needed to decrypt the message (the private key.) For symmetric encryption the same key both encrypts and decrypts messages.

Quantum computers pose a much more pressing danger to asymmetric algorithms than symmetric algorithms. Asymmetric algorithms are vulnerable to quantum computers calculating the private key from the public key, as the mathematical relationship between the two is commonly a prime number factorisation or a discrete logarithm problem (both of which can be solved by variants of Shor's algorithm.) In comparison, quantum computers only offer a modest speed up for attacks against symmetric algorithms, by optimising the search process for identifying the key used (using Grover's algorithm.)

The limited scope of the attacks enabled by quantum computers offers little comfort for the millions of safety systems that rely on asymmetric encryption, including https connections for data transfer, digital signatures for validating firmware, and most remote network access software for allowing operators to interact with difficult-to-access systems.

In addition to the confusion about what algorithms are vulnerable, nobody knows precisely *when* quantum computers will have the power to target cryptographic algorithms. The specification of a quantum computer that is "cryptographically relevant" is not, as is commonly reported, simply a case of how many qubits the computer is built from. The connectivity between the qubits, their error rate, and the underlying materials that the qubits are built from, all have an effect. The specification required is also a function of the underlying mathematical problem. Predicting the future is always a dangerous thing to do in technology, but the general trend in the literature is that asymmetric encryption that relies on "the elliptic curve discrete logarithm problem" (for example TLS 1.3 which is used by modern https connections) is likely to be the first to be challenged by quantum computers. Followed by encryption that relies on the "prime number factorisation problem" (for example RSA, which is used by older https connections.) This is in part because of the smaller key sizes used by elliptic curve algorithms. The uncertainty around when algorithms will become vulnerable makes prioritising which systems to protect extremely difficult, especially for industries with long-lifetime assets (e.g. digital control systems in industries such as energy and aerospace) and sectors with long-term confidentiality goals such as healthcare and law enforcement.

But engineers should not be complacent about when to address these risks. Whilst most quantum-enabled attacks are several years away, nation states are already reportedly harvesting data encrypted by vulnerable algorithms, in the hope that it will still be useful when they acquire the power to read it. This means that from an intended destination/usage perspective, attacks have already started.

Lengthening the keys used, and increasing the frequency of key rotation may offer short-term protection against some quantum-enabled attacks, depending on context. Architectural changes may address other vulnerable systems, by no longer relying on asymmetric cryptography to perform key exchanges. But wholesale conversion of systems from asymmetric to symmetric cryptography simply isn't practical, as that transformation introduces new, often intractable, key management challenges. But the future is not entirely bleak. Significant efforts have been made around the world to identify new cryptographic techniques that are resistant to both traditional and quantum computers. Post-quantum cryptography has made great strides, with three specifications being officially approved by the US National Institute of Standards and Technology (1 key encapsulation mechanism and 2 digital signatures.)

Many large companies such as Microsoft, Amazon, and Google have already begun moving towards those algorithms under the hood, but for many organisations upgrading will require managed change, working with their software supply chain, and ensuring that they are ready for the emergence of quantum computers. As with any large-scale digital transformation there are going to be challenges along the way. The new algorithms have different computational and communication characteristics which have already exposed latent errors in networking infrastructure for early adopters. The extra bandwidth consumed by the new algorithms, and the much longer packets and signatures are going to have far-reaching consequences for digital infrastructure. In addition, despite the rigorous testing, there is no guarantee that an algorithm, or the implementation of that algorithm, will remain secure in the

face of future developments. In this regard, the fact that NIST has only approved a single key encapsulation mechanism so far means there is a single point of failure in the security for early adopters who have decided to follow their advice. Though it should be noted that there are several other algorithms that are already available but are still in various stages of being assessed.

The Ugly – Challenges yet to be faced

We've made the argument that quantum computers bring both opportunities to make the world a safer place, and a more dangerous one. But there are several challenges that need to be addressed before quantum computers are ready to be applied. Here we present two of those challenges that the safety community may be able to offer insights into.

As we've frequently alluded to, quantum computers are inherently error prone. There is a very active research space looking for approaches to perform automated error correction on quantum computers. Large organisations such as Google have made progress [3] but the overheads from quantum error correction are large. Some people predict that we will need redundancy in the order of 100-1000 times, i.e. up to 1000 physical devices implementing qubits to realise a single, reliable logical qubit. If we wait for quantum computers to reach the sizes where this level of redundancy is possible it is quite possible that we will miss opportunities to make the world a safer place. Some researchers are now investigating what we can achieve with devices that have inherent error rates. There's potentially a lot to learn from both communications and safety science in this space. Communications specialists build reliable systems from components with well-characterised noise properties every day. Similarly, safety engineers build trustworthy systems from components that have known failure conditions.

Our second challenge is that of quantum verification. Classical software needs to be analysed and tested to make sure that it has been specified and implemented correctly. Likewise, quantum algorithms, represented by 'quantum circuits' also need verification and validation. The image below is an example of a quantum circuit representation for an error correction function. In this example, five physical qubits are used to represent the same "logical" qubit. Errors that form in individual physical qubits can be detected and corrected to make the overall system more robust.

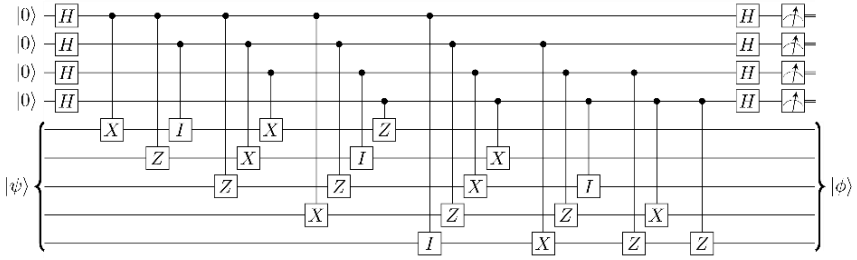

From a hardware perspective the current architecture for quantum computers makes this a much more challenging problem. Quantum gates, the implementation of operations on qubits, are realised by applying physical processes to qubits, for example irradiating a qubit with a microwave pulse. Therefore, each quantum gate is implemented separately for each qubit. In a classical computing architecture, we can be confident that if the adder within the ALU of a computer works for one pair of registers, that it will work for any other pairing. But

for quantum computers, there can be latent errors in quantum gates that will lie undetected until that specific hardware qubit has that specific operation performed on it.

From a software perspective there is little knowledge about how simple test cases can be used as building blocks to test the validity or otherwise of more complex systems, a problem that is not dissimilar to modular safety cases.

We suggest that there may be lessons to be learned from the way AI components will be ultimately integrated into safety systems. Design patterns that build trustworthy systems when one component is known to have a significant error rate but provides useful input nonetheless may be the key to integrating quantum computers into safety critical workflows.

Conclusion / Call to arms

It's fair to say that quantum computers are an emerging technology. In fact, we believe that they're roughly where classical computing was in the 1940s. The uses of this technology in a few decades time are likely to surprise all of us!

Quantum computers are not a competitor to classical computers, but another device to be integrated by the dynamic and flexible silicon computer, alongside other quantum technologies such as sensors and timers.

We are likely to see impacts within the safety community as we seek to enhance safety with this exciting new technology and seek to protect systems from the harms that it makes possible. The safety community should be ready for a quantum computing revolution and be prepared to contribute their knowledge and expertise to this entangled web.

Bob Oates, Cambridge Consultants
Dr Bob Oates is a specialist in the cyber security of safety critical systems and the assurance of AI-enabled systems. He has worked for over ten years on a wide variety of problems in the defence, maritime, aerospace and energy sectors. He advises critical national infrastructure organisations on how to prepare for threat actors equipped with quantum computers.

James Cruise, Cambridge Consultants
Dr James Cruise is an expert in quantum computing and quantum algorithms. He is a mathematician by training but has worked in quantum computing for over ten years with range of organisations including government, startups, and consulting. He has particular interest and history in helping organisations understand the specifics and practicalities of how quantum computing will change their business including from a security perspective.

References

[1] Michael A. Nielson and Isaac L. Chuang, "Quantum Computation and Quantum Information", Cambridge University Press, 2010, 2nd Edition, ISBN: 9781107002173

[2] Data Safety Guidance (version 3.7), SCSC-127J, Jan 2025, https://scsc.uk/scsc-127J

[3] Google Quantum AI and Collaborators, "Quantum error correction below the surface code threshold", Nature, 2024, DOI: https://doi.org/10.1038/s41586-024-08449-y

Seminar: Safe Agile Developments (TBC)

THE SAFETY-CRITICAL SYSTEMS CLUB, Seminar:

Safe Agile Developments (TBC)

Thursday 1 May, 2025 - London, TBA

This 1-day seminar will look at techniques and approaches for managing Agile developments for safety-related systems. It will include development, maintenance and support activities and will look at example projects and standards and guidance in the area.

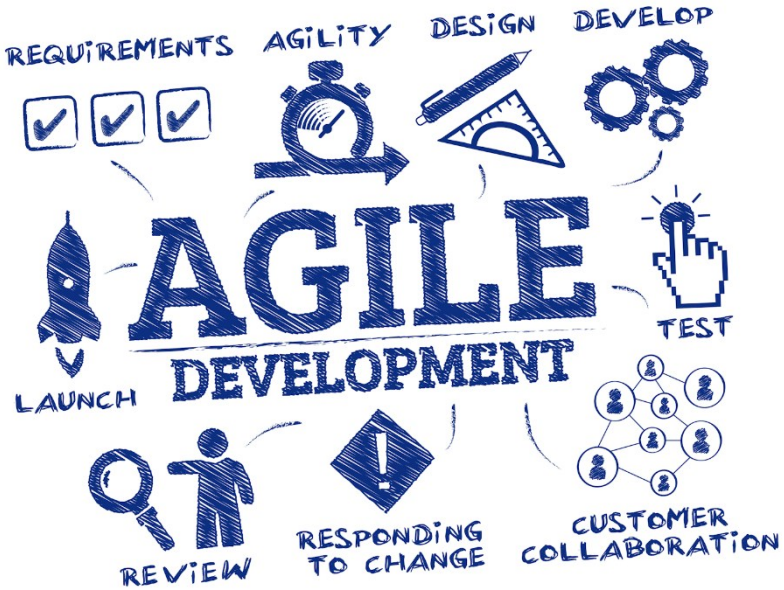

Modelling Our World

With origins in antiquity and ancient philosophers speculating on the nature of being, "ontology" and "ontological modelling" have perhaps a reputation of being esoteric, impenetrable and only for the purist of academics.

In the last couple of decades however, there has been great work done to make ontological concepts and models readily accessible and to build the essential toolkits required to apply them at a practical level.

The Ontology Working Group (OWG) provides an introduction to ontology, explains how relevant and important it is and describes the OWG's current objectives in developing new ontological models that will be of benefit to everyone working in domains that involve the management of risk, such as safety and security.

What Is an ontology?

Isn't an ontology some esoteric formalisation, which is nice to have, but not really relevant or accessible to my day-to-day work on real systems?

Ontology is **everything** to do with your day-to-day work! When we build a system, we are modelling our world and this involves developing concepts, terms and their relationships – this is called an ontology, or more formally: "**a specification of a conceptualization**" [1].

Why do we need them?

Even if not being done explicitly, an ontology will always be developed – it's therefore never a case of whether or not an ontology is needed. If they are not explicitly developed, and developed well, you only end up with bad ontologies.

For the system to operate correctly, the model needs to be correct and consistent. Hopefully it is written down, but at worst, it can just be held in the minds of the designers.

Poor ontologies can lead to poor design, misunderstanding, ambiguities and inconsistencies. The system may be difficult to explain or may diverge from the real-world intended behaviour. Many of these issues may actually be difficult to spot or test and may lie latent for years until the right set of circumstances occurs (c.f. the NATS outage [2]). Consider, for example, two people with the same name and date of birth, which can occur in real life, but might not be handled well by a system.

"It's never a case of whether or not an ontology is needed. If they are not explicitly developed, and developed well, you only end up with bad ontologies"

How do we capture models?

The idea of modelling relationships is not new and there are notations such as the **Unified Modelling Language (UML)** [3] to help us to do just that. However, UML can be too general for modelling ontologies and without a framework to impose some constraints, it is very easy to model things that are inconsistent or ambiguous.

This is why the OWG have chosen to use the **Unified Foundation Ontology (UFO)** [4]. This is a specialised framework for modelling in UML tailored for ontological modelling and implemented in **OntoUML** [5] (this is a customisation of UML with associated software support via application plugins).

The following diagram shows an example of a UFO model using OntoUML:

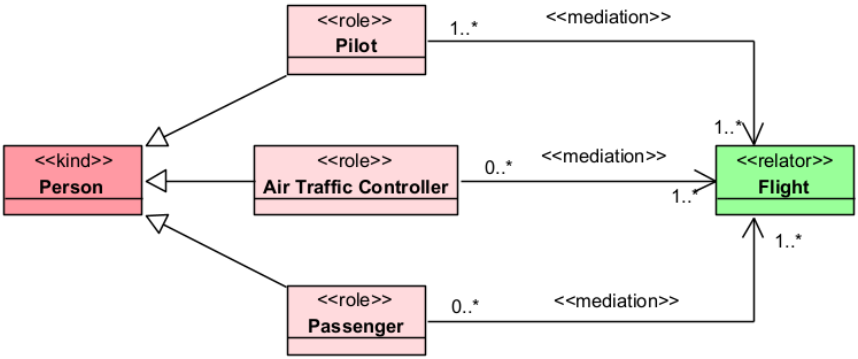

In the diagram, boxes (or more formally *classes)* represent 'things' in our world view. The ontological aspects are expressed as stereotypes (that's the part of the class definition with the Guillemets like this: «role»). Classes of the same stereotype share the same colour in

the diagram. Lines between the boxes show their relationships with other things. Arrows with a filled white ▽ triangle are *specialisations*, or in other words subtypes of the thing that is being pointed to (for example, a **Passenger** is a subset of all **Persons**). An open arrow → relationship shows a looser *association* between things and these also have stereotypes. Expression such as **0..*** and **1..*** show *cardinality* – that is the number of each thing that can be related. These particular ones are read as *zero or more*, and *one or more* respectively so for example, a **Flight** will have at least one **Pilot** but may have more, and those **Pilot**s may be involved in one or more **Flight**s. A particular **Flight** may have one or more **Passenger**s but may have none if the **Pilot** is the only occupant. A **Person** may also be a **Passenger** or more than one **Flight**.

Different Perspectives

This first model shows the concept of a **Flight**, which is hopefully well-known to everyone. In this model, the flight only exists as a relationship between a **Pilot**, and usually the **Air Traffic Controller** and **Passenger**s. In other words, there is no flight without a **Pilot**. **Passenger**s are often also on the flight (but not always!) and the **Air Traffic Controller** usually helps with the navigation (but not always!). Note that this model is likely under-constrained because it does not preclude, for example, **Pilot**s from being **Passenger**s on the same flight. Dealing with this type of issue require additional constraint logic, such as described by the OntoUML anti-pattern catalogue [6].

This perspective might provide a useful way of modelling flights at a small private airfield, where flights are created on an ad hoc basis when a **Pilot** and **Passenger**s are in the plane and want clearance to take off. There might be no prior concept of an actual flight until that moment. In this way of modelling, if the flight never took place, it never existed.

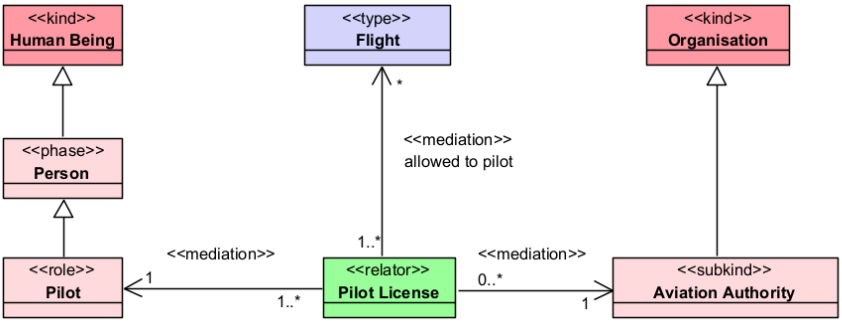

In this second model, **Flight** is considered a type of thing that has its own independent existence, not one that relies on relationships to exist. The **Pilot**'s relationship with the **Flight** is expressed in their ownership of a **Pilot License** to give them authority to pilot the aircraft.

This second perspective could be used to represent flights by a commercial transport airline. In this situation, **Flight**s are planned months or years in advance and might take place, be delayed or even be cancelled. The purely relational approach of the first diagram does not support the permanency of a flight even if it never took place. For example, the cancelled

flight still has relevancy for operational needs (charges and payments to airlines for airport services) or to manage passenger financial compensation.

Importantly, **Flight** is modelled completely differently from before, but neither model is incorrect! Also note the differences in how a **Person** is modelled. In the first, a **Person** has identity and is a **Kind**, in the second, the **Person** is seen as a **Phase** of a human being.

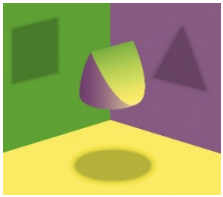

These examples illustrate the idea that there can be several perspectives of something in the world, which are different but are not incorrect. Consider the various shadows cast by the object – which is the correct shadow for this object? They all are!

The Unified Foundation Ontology (UFO)

UFO comes in three parts: a top-level model for "things", (UFO-A), a model for "events" (UFO-B) and a model for "social" aspects (UFO-C). UFO-A is where we start normally as it relates to the objects in our world.

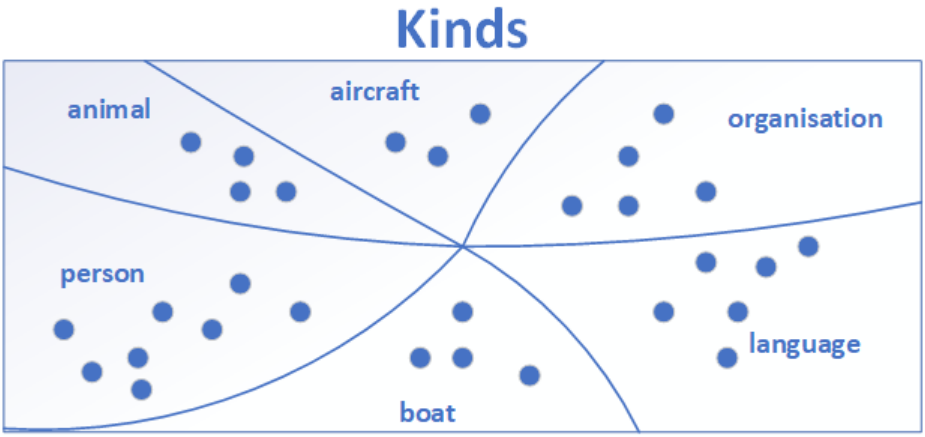

Kinds are a core UFO concept and reflect the objects we find around us like people, aircraft and animals. **Kind**s are objects that endure for their entire lifetime and define essential characteristics that instances must possess (called an identity principle). We can also distinguish between different instances of those **Kind**s: I am a person, you are a person and we are different people as we have different identities based on the properties we usually use to distinguish people (height, age, eye-colour, place of birth etc). Each 'dot' ● in the diagram indicates an individual instance of a particular **Kind**.

We might also talk about the **Roles** and **Phases** that **Kinds** might have during their lifetime. For example, a person may be a student for some part of their life (a **Role**), but will cease to be a student when they leave education. Similarly, a puppy is a **Phase** a dog will go through, which is transitory and only for the early part of their life. The following diagram shows how particular instances may exist in a particular **Phase** or **Role** as some point in their lifetime.

For example:

- An organisation may be dissolved at some point
- Some languages are no longer spoken as a native language by a community of people so are considered "Dead"
- A boat may at some point in its lifetime be used for hospitality (eg. parties, ceremonies etc.)

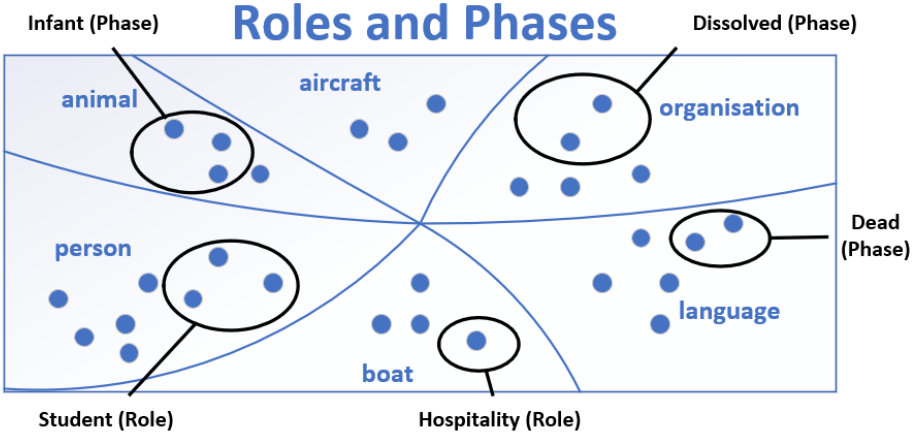

There are several other terms that become useful when modelling the world. One useful concept is that of a **Mixin**.

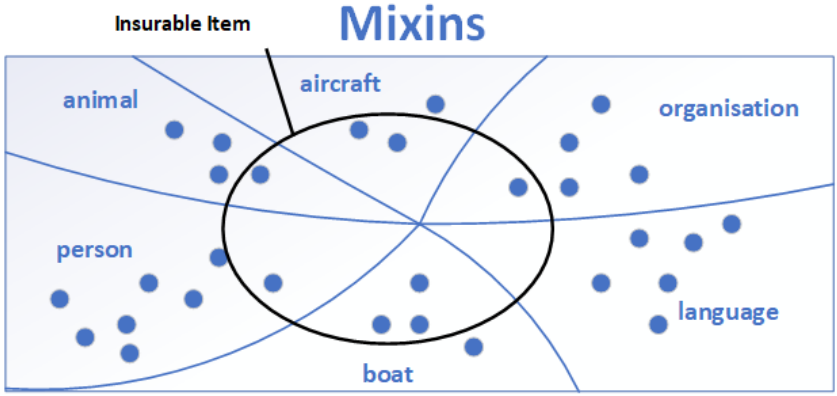

A **Mixin** is something we use to describe shared properties of objects that do not all belong to the same **Kind** and so can span multiple **Kind**s that do not share the same identity principles (a red boat does not have red hair colour, but both are red). The example shows a **Mixin** that some Kinds such as boats, aircraft, animals can be insured.

These terms along with the associated rules give us a powerful and consistent toolkit for describing important concepts and relationships in our world.

The Ontology Working Group (OWG)

The Ontology Working Group (OWG) has been set up to look specifically at developing an ontology of **Risk** and then subsequently **Value**. The OWG 2025 Objective is more formally:

Mission: Develop a clear, systematic terminology, and unambiguous model of risk and value concepts and their relationships.

The group will do this by using:

- The Unified Foundation Ontology (UFO)
- Visual modelling tools (Visual Paradigm)
- Ontologically consistent modelling notations (OntoUML)
- Monitoring the work of other working groups (eg. ISO & IEC)

The main reason for pursuing this is to provide a common framework for expressing both **safety and security risk**. Safety and Security Risk management processes are similar but different (or in other words, they are different perspectives of the world), so there is much confusion that can arise when the two groups need to collaborate on projects.

Working Group Developments

The following diagram shows some of the work in progress by the OWG:

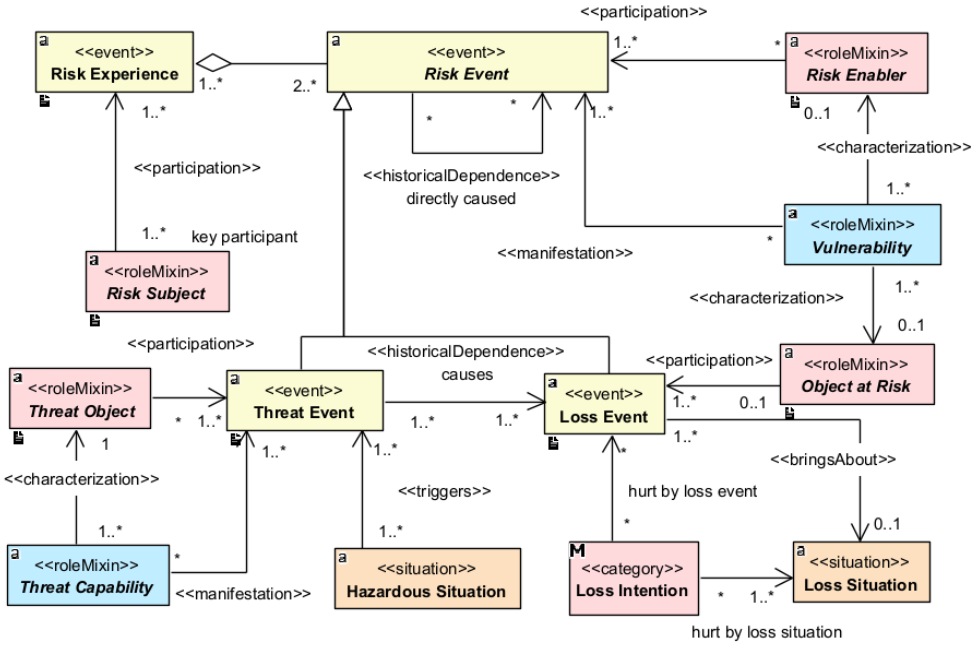

There have been many different attempts by other working groups and organisations to model risk, largely because, as we've seen, there are many different perspectives and ways of viewing the problem.

The OWG has assessed many different approaches but believe the most fruitful approach is through models that consider **risk** as an experience for participants and also to consider **value** as a key consideration when thinking about risk.

For example, any discussion around risk needs to really consider how participants *value* those items at risk, and indeed, how those who aim to exploit a vulnerability value *what they will gain* versus how they value *the liberties that may be lost* if they are caught.

It's hard to overstate the importance of this work, not only for the safety/security community where arguably ontologies *have* to be an essential aspect of the work we do. For example, it is difficult to see how a robust understanding of AI decision-making can be developed if not supported by formalised models of the concepts and relationships it operates on.

References

[1] T. R. Gruber, A Translation Approach to Portable Ontology Specifications, in Knowledge Acquisition, 5(2):199-220, Academic Press, 1993, http://www-ksl.stanford.edu/kst/what-is-an-ontology.html

[2] Air traffic control disruption hearing (British Airways), https://committees.parliament.uk/publications/42114/documents/209387/default/

[3] Unified Modelling Language, https://en.wikipedia.org/wiki/Unified_Modeling_Language

[4] Unified Foundational Ontology, https://en.wikipedia.org/wiki/Unified_Foundational_Ontology

[5] OntoUML, https://en.wikipedia.org/wiki/OntoUML

[6] OntoUML Anti-Pattern Catalogue (Relator Mediating Overlapping Types), https://ontouml.readthedocs.io/en/latest/anti-patterns/RelOver/index.html

The Ontology Working Group (OWG)

The OWG was set up in 2020 to help address issues with the lack of formalisation in terms and conceptual relationships used in standards and guidance documents. In particular, the similar but at times different, interpretations of terms and concepts used between safety and security domains is an ongoing cause of confusion and misunderstandings and presents barriers to more formal integration of these disciplines.

The OWG has its origins in the Data Safety Initiative Working Group (DSIWG), which, between 2016 and 2019, explored the development of an ontology model for the concepts described in the Data Safety Guidance (DSG). However, it was realised that the work couldn't be completed without a wider ontological model for risk, and so the OWG was established and is currently working on developing a clear, systematic terminology, and unambiguous model of risk and value concepts and their relationships.

Seminar: How Safety Culture has to Change With AI

THE SAFETY-CRITICAL SYSTEMS CLUB, Seminar:

How Safety Culture has to Change With AI

Thursday 19 June, 2025 - London, TBA

This 1-day seminar looks at how an organisation's safety culture has to change when systems including AI are introduced.

The emergent application of artificial intelligence in safety critical systems raises many questions for safety culture. Who is accountable for safety - who has agency for safety - the AI developers or the operators? How does AI impact ownership of safety if operational decisions are made by AI? How does AI impact transparency in operational safety decision making? Does AI erode safety citizenship - a sense of self-determination - and disempower people? How do you design AI so as to allow people to intervene in a timely manner and retain effective oversight? Does AI erode human skills such that they can no longer truly retain operational responsibility? What if the person wrongly overrides the AI, how is this treated in a fair and just way? How do developers make decisions about the readiness of novel AI safety critical systems? And what are the opportunities offered by AI in safety critical systems? This seminar will hear from experts working at the cutting edge of AI and Human Factors about these and other issues, and discuss emerging principles and good practices.

Speakers and talks include:

Paul Leach, Head of Human Factors, Rail Safety and Standards Board - *Human factors principles for the design and operation of AI systems in rail*
Irene Ruiz-Gabernet, Airbus, Head of Operational Excellence and Business Management - *Human Factors and AI in Aerospace*
Sqn Ldr Kathy Syfret, Deputy MilCAM for A400M RAF platform - *AI and Aircraft Maintenance*
Carlan Carmen, System Safety Engineer - *AI and autonomous systems*

It will include speakers and inputs from the SCSC Safety Culture Working Group (SCWG).

It will be held in-person in central London.

Celebrating the IET Young Woman Engineer of the Year Awards 2024

The IET's winners and highly commended finalists of the Young Woman Engineer of the Year Awards with award winning computer scientist Dr Anne-Marie Imafidon MBE

The Institution of Engineering and Technology (IET) held an awards ceremony on 9th December 2024 celebrating the brilliant women showcasing their engineering excellence and truly engineering a better world. Louise Harney, SCSC Diversity Equity and Inclusion Lead, celebrates the success of these impressive women, who all work in industries dealing with safety-critical systems.

These prestigious engineering industry awards celebrate women working in modern engineering – and aim to help change the perception that engineering is predominantly a career for men by banishing outdated engineering stereotypes of hard hats and dirty overalls.

I caught up with the winners to discuss their careers – how they got into engineering, their challenges and achievements and what advice they might give to others in the early part of their careers.

Marisa Kurimbokus – IET Young Woman Engineer of the Year

I am a Chartered Engineer with a career spanning over a decade in product design and systems engineering within the automotive and power electronics industries, including Jaguar Land Rover, Triumph Motorcycles and Lyra Electronics.

I graduated from the University of Cambridge with a degree in Aeronautical and Aerothermal Engineering. I am currently the Head of Engineering for Raeon, designing and delivering custom batteries for a variety of applications, from automotive and marine, to robotics and drones.

I work within the Senior Leadership team to develop the company business strategy, plan resources and develop company processes. I am proud to have been listed in the Top 50 Women in Engineering: Inventors and Innovators in 2022 for my work in net-zero and green technology.

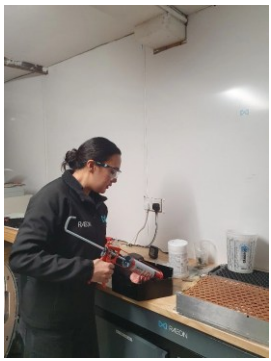

Safety should always be inherent in everything we do as engineers; from ensuring the way we work is safe, to designing products and systems that are safe for our users. Championing best practice in DFMEA and PFMEA in the design process is my way of capturing safety critical design features, and designing against failures.

I keenly support engineers at all levels, and STEM is a huge passion for me, fuelled by lack of guidance and inspirational role models when I was a student. I have set up STEM programmes at three of the companies I've worked for and I actively support developing engineers through their careers, providing mentoring for university students through to engineers applying for professional registration.

"Safety should always be inherent in everything we do as engineers; from ensuring the way we work is safe, to designing products and systems that are safe for our users"

Whilst the number of women in engineering still remains disappointingly low, I would urge all women to consider a career in engineering. The stereotypical "feminine" qualities of being caring and compassionate are too often trivialised in industry. However, these are the attributes that are so fundamental in developing safe, user-centric products and systems that are empathetic to all members of society. For diverse teams to be successful, individuals need to bring their own viewpoints, ideas and personalities.

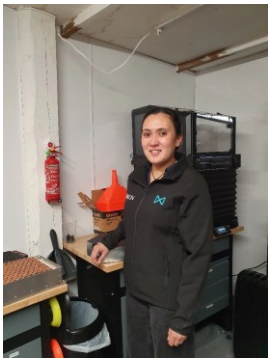

Natalie Parker –Women's Engineering Society Prize Winner

I am a Technical Specialist within Sellafield's Operational Technology Department and provide technical, process and assurance support to front line engineering teams and wider supply chain. I am also responsible for Technical training and line management of Control Systems Degree apprentices.

Control Systems at Sellafield provide automation for operations across Sellafield which is Europe's largest Nuclear facility. These systems range from basic to safety critical and therefore it is so important that I guide engineers on process and procedures to safely managing and maintaining these safety systems.

I never considered the field of engineering during my time at school as it was never part of the curriculum. I am not a person who is naturally clever and have always had to work hard to achieve my grades. During my struggle in my first year at sixth form I decided to explore apprenticeships and was lucky enough to secure a control systems apprenticeship and I have never looked back. I progressed from being a front line control systems apprentice on Europe's largest Distributed Control System to senior engineer and now technical specialist.

"Surround yourself by people who advocate for you and don't let those who don't push you down. You can achieve anything you set your mind to!"

My advice to others is to surround yourself by people who advocate for you and don't let those who don't push you down. You can achieve anything you set your mind to!

Alexia Williams, Mary George Memorial Prize for Apprentices Winner

From a young age, I was drawn to engineering and problem-solving, preferring remote-controlled cars and LEGO over traditional girls' toys. My interest in aerospace developed during my GCSE years, sparked by a visit to the Bath and West Show, where I learned about apprenticeships from a GKN graduate.

With my parents' support, I pursued Design and Technology Systems and Control at GCSE, becoming one of only three girls in my class. After A-levels, I completed a four-year Level 6 Aerospace Engineering apprenticeship at Rolls-Royce in 2022, transitioning into a full-time role as a Deployed Lifecycle Engineer. There, I solved shop floor issues and led improvement projects, including designing precise tooling solutions.

Motivated to further my education, I pursued a master's degree in through-life system sustainment as an apprenticeship through Cranfield University. With my manager's support, I became the first person at Rolls-Royce to complete an engineering master's apprenticeship and the first woman sponsored for the course in over 20 years.

In August 2024, I became a Through-Life Technical Lead, focusing on using data and digital tools to enhance in-service products. While often the only woman in my teams during my early career, I have seen more women join engineering roles over time, which I find encouraging.

I actively promote apprenticeships as Chair of the IfATE Apprentice Panel and as Apprentice Representative on the UCAS Apprenticeship Stakeholder group and Institution for Engineering and Technology Young Professionals Committee, advocating for diversity and inclusion in engineering. Over the last year I have spent over 200 hours promoting STEM careers via Air Shows, conferences, events and in schools. I have been recognised for my achievements via multiple Apprentice Awards.

"I would urge others to challenge stereotypes, recognise their worth, and inspire the next generation of engineers."

I advise aspiring engineers to explore their interests, embrace opportunities, and network early. Apprenticeships offer valuable rotations to discover what suits an individual best. Above all, I would urge others to challenge stereotypes, recognise their worth, and inspire the next generation of engineers.

Erin Lowe – Highly Commended for Mary George Memorial Prize

I am currently in my third year as an electrical apprentice at Yamazaki Mazak, specialising in the CV5-500 production line, where I focus on the electrical assembly and adjustment of our machines.

As I spend the majority of my day in a busy factory, PPE is something that I interact with constantly. Some PPE is required to be always worn in the factory, including Mazak uniform, safety glasses, and steel-toe cap shoes.

During my second year, a group of us apprentices came together to establish the Mazak Apprentice Charity Committee. Through this initiative, we have organised over a dozen events and successfully raised more than £1,700 in the past year to support various causes.

Through Mazak, I have had the opportunity to attend numerous career fairs, give many presentations, and many in-house events too (mostly factory tours). This year, I have normally had at least one event every fortnight.

"STEM outreach events are very beneficial to the community, but also to my own personal growth "

Part of being a Mazak apprentice is being an apprentice ambassador. This is not compulsory and so some choose not to participate however, I think that STEM outreach events are very beneficial to the community, but also to my own personal growth.

Final Words

It was great to catch up with these truly inspirational women, not only for their achievements, but for developing careers in an industry still subject to perceptions that it is career solely for men wearing hard hats. Their experiences show that there are many different routes into engineering and I have noticed since my school days that apprenticeships are both more common and more respected throughout engineering.

On behalf of the SCSC, we commend the winners for their focus on STEM, engineering education and providing inspiring role models for the future generation of engineers.

Further Information

If you would like further information about the IET Young Woman Engineer of the Year Awards, check out the website here: https://youngwomenengineer.theiet.org and do feel free to contact us to support you on DEINetwork@scsc.uk.

Louise Harney, SCSC Diversity Equity and Inclusion Lead

Top Image credit: The Institution of Engineering and Technology (IET)
Marisa Kurimbokus images © Marisa Kurimbokus.
Natalie Parker images © Natalie Parker.
Alexia Williams images © Alexia Williams.
Erin Lowe images © Erin Lowe.

Seminar: Deployment, Operations and Maintenance of Safe AI Systems

The "Deployment, Operations and Maintenance of Safe AI Systems" seminar was held at the Cumberland Hotel in London, on 26th September 2024 and hosted by Mike Parsons. The speakers presented a number of perspectives on bringing AI into service, including safe operation, adequate maintenance, human integration and organisational readiness.

The six speakers presented insightful and compelling angles on development and using frameworks for safe AI systems, covering a range of topics such as:

- Bob Oates and Carolina Sanchez (Cambridge Consultants): Organisational readiness for AI: Governance and Assurance – how we ensure that our organisations are ready and willing to integrate AI.
- Kate Preston (University of York): A human-centred assurance framework for deploying autonomous systems based on previous research and demonstrator projects – case studies and some early steps towards defining a human-centred assurance framework.
- Karin Rudolph (Collective Intelligence and the Ethical Technology Network): Responsible Deployment of AI Systems: What businesses need to know – the AI regulatory landscape, socio-technical systems, AI standards and the important societal impacts of AI.
- Jane Fenn (University of York): A New Approach to Creating Clear Operational Safety Arguments – how do we argue about operational safety and why is this not yet good enough, the complicating influence of AI.
- John McNicol (Nova Modus): Systems for Safety When Automated Vehicles Go Wrong! – AI in vehicles, new and emerging standards, the concept of minimal risk events and how considering remote operation safety can enhance AI vehicle systems.
- Richard Hillman (Horiba Mira): Safety assurance methods for automated vehicle deployments: learnings from the Harlander project – challenges of assuring automated vehicles and an introduction to the Harlander project.

Bob Oates and Carolina Sanchez

A double act from Bob Oates and Carolina Sanchez was a timely reminder of the fact that AI integration is not just about the technology: the organisations, stakeholders, individuals and human processes are also of paramount importance. Bob and Carolina asked us to consider the seven AI assurance pillars, and how AI amplifies existing safety challenges such as ensuring data integrity, adequate testing, and reliable software supply chains. The motivating example – a lawyer in the US using ChatGPT to perform legal research – resulted in a very lively response from the audience!

Bob and Carolina then moved on to discuss the topic of governance of AI systems. The emphasis here was very much on providing a grounding for this in best practice, and also on an examination of how organisational responsibilities for AI implementation don't always align with the way the technology itself is developed.

Bob and Carolina also challenged us to consider the processes and policies which organisations need to put into place to ensure responsible adoption of AI: a statement of ethics, AI usage policy and a system for AI risk management. The final takeaway was a very relevant observation here that we, as safety engineers, have to use assurance foresight so that we don't in future end up having to use the word hindsight: definitely a thought to chew on for the rest of the day!

Kate Preston

Kate Preston followed, with a whistle-stop tour of some of the existing research that considers human factors and AI ethics, including work from the Assuring Autonomy International Programme (AAIP), which the SCSC was lucky enough to hear about at a previous seminar. However, as Kate observed, many significant challenges still remain! Real-world AI integration has not yet made adequate use of the frameworks and guidance to ensure safe and ethical development, while it still remains unclear who takes responsibility for the AI.

Kate took us through some of the reasons for this, including the fact that the socio-technical aspects of AI are still being underconsidered: the message here is that the people are as important as the technology! A human-centred approach would apply human factors to ensure that the appropriate consideration is being given to how humans might interact with the AI-enabled systems, as well as how external systems and organisations might affect AI use.

Next, Kate introduced us to how such a human-centred approach might be developed, by a two-pronged attack consisting of a scoping review and a systematic analysis of existing research performed under the auspices of the AAIP. The initial findings for these were impressive; I think the extent of existing research surprised many of us, and provided an inspirational note on which to end Kate's presentation.

Karin Rudolph

Ethics and human factors remained firmly in the spotlight for Karin's presentation, looking specifically at responsible deployment of AI. Karin took us through the tangled web of legal instruments mentioned in the EU AI Act: a daunting reminder of just how much legal infrastructure is needed to begin to assure that AI is being adequately regulated. As Karin noted, we're currently in the "era of risks", and adequate AI governance is essential.

The topic of the environmental impact of AI came up here again, with many of us surprised to learn just how much worse AI is for the environment than flying, car manufacturing and fuel consumption! A timely reminder indeed that risks go beyond safety, and Karin emphasised this with a concise and helpful introduction to the National Institute of Standards and Technology (NIST) Framework for AI risk management. However, legislation is only as good as the thinking behind it, and Karin soon challenged us to don our thinking caps by asking for an exact definition of "ethics": we all know what we mean by it, but the importance of establishing a shared definition quickly became clear!

Karin also invited us to consider some of the fundamental socio-technical aspects of systems, with a thought-provoking example of how people can form emotional attachments to AI systems, and come to depend on these beyond their simple functional performance. The landscape of risk is wide ranging, and it was rapidly becoming clear that adequate management of all of these risks is necessary if we are to claim that we really are developing "responsible" AI.

Jane Fenn

After a break for lunch, Jane Fenn picked right back up on one of the major themes for the day: how do we ensure operational safety of AI systems? Safety has certainly changed in the past thirty years, and Jane took us through a number of significant safety events, from Piper Alpha to the Haddon-Cave report. It was rapidly becoming clear to all that operational safety has not always been considered with as much attention as it requires!

Jane highlighted some of the reasons for this, ranging from a lack of understanding of risk on the part of the operators, a general preference for "box-ticking" and simple rules rather than engineering judgement, and a lack of understanding of how the safety mitigations truly contribute to mitigations of hazards. A daunting task, but Jane emphasised that clear and explicit identification of operational safety requirements would go a long way towards addressing these. She proposed a separate – although linked – argument for operational safety, which has the advantage of hiding confusing design details while at the same time ensuring that the link between operational safety and design remains clear.

Of course, this applies to all systems, but Jane encouraged us to consider just how AI – as

a new and evolving technology – can enhance this challenge. It isn't always clear with an AI system exactly what needs to be monitored to provide assurance of operational safety, and to date there have been too many examples where operational safety is provided by simply handing control back to the human operator, regardless of how alert or situationally aware they may be!

John McNicol

John refocused our thoughts firmly on the automotive domain, with a presentation focused on automated vehicles and system safety. With mechanical failure, power failure and software failure all potentially in the offing, John introduced us to the new BSI Flex standard 1888 for considering just these eventualities. BSI Flex 1888 discusses both minimal risk manoeuvres and minimal risk conditions, and much of the material in this new standard has come out of the SafeMRX project exploring these.

John also reviewed the state of existing standards in this sector, noting in particular that these fail to adequately consider minimal risk manoeuvres or minimal risk conditions, in particular where busy roads or automated buses and trucks are involved. It was rapidly becoming clear that something new was definitely required, and John's presentation also offered a way for audience members to get involved via the industry consultation stage in November 2024.

This wasn't the only standard John introduced us to, though, with the presentation finishing up via a look at BSI Flex 1886, considering systems for remote operation safety of Autonomous Vehicles (AVs). Version 2 of this had just been published, and audience members very much appreciated the chance to walk through and see what was new in terms of remote operation systems, whether these be perception systems, communication systems or control systems. Coming to the end of a talk focused primarily on what could go wrong with AVs, the knowledge that these standards are out there felt definitely like a step in the right direction!

Richard Hillman

Richard continued the automotive focus with the final presentation of the day, an in-depth look at the Harlander project, and what we can learn from this about safety assurance methods for AV deployment. The Harlander project was new for some of the audience, who were intrigued to hear about the use of Belfast Harbour Estate for two autonomous vehicles, timed to link specifically with the rail schedule. Richard cautioned us about some of the ways in which the real-world setting can introduce challenges for AVs, however, with edge cases such as confusing marking and signage, and even the presence of animals on the road: surely the first time an SCSC presentation has considered runaway seals!

Moving on, Richard emphasised the need to consider best practice before executing any testing. This is particularly the case where AI is involved, as machine learning creates a large black box that means the traditional methods of cascading system level requirements to subsystems are made much harder. Looking further into specific forms of testing, Richard's presentation introduced the audience to three categories of scenario-based testing: functional, logical and concrete scenarios. While functional scenarios require simple description in natural language, concrete scenarios quickly require statistical sampling methods to help control test volume: with a reasonable real-world number of 30 parameters and 20 levels each, this generates $1x10^{39}$ test cases for a single logical scenario!

Having thoroughly convinced us of the need for intelligently-applied best practice, Richard then discussed how multi-pillar testing combines different forms of test evidence, but still is not sufficient to assure Safety of The Intended Functionality (SOTIF): rather, design analysis evidence is also needed. With SOTIF itself only being a tiny fraction of the overall safety case, this presentation had been a convincing way to motivate us to consider how safety assurance requires so much more than testing!

Mike Parsons – Wrap Up

Mike Parsons wrapped up the session by presenting the answer to a very timely question: why can't we simply ask ChatGPT to solve this for us? Having asked ChatGPT the query How can I safely deploy an AI system into service? Mike obtained what looked like a reasonable answer that focused on such aspects as design for safety, robustness, failures; testing and validation; scenario planning for failures etc. However, the audience were quick to realise that these statements hide some very difficult and complex challenges, and ChatGPT doesn't get us any closer to solving these!

At the end of a day focused on the challenges of AI deployment and operation, the audience were unanimous in agreeing on the necessity for human judgement, clear thinking and safety assurance that demonstrably considers the people, organisation and real-world challenges of AI systems. A thought-provoking day, and a great discussion to finish!

Report by Catherine Menon.

Catherine Menon is a principal lecturer at the University of Hertfordshire. Her research focuses on the intersection of trust, safety and ethics in public-facing autonomous systems. She has a PhD in category theory and has been involved in the development of regulatory standards and guidance for safety, security and ethics assurance within autonomous systems.

Seminar: Safe Autonomous Transport – the Good, the Bad and the Ugly

The "Safe Autonomous Transport - the Good, the Bad and the Ugly" seminar was held at the Eurorstars Book Hotel in Munich, Germany on 28th November 2024 and hosted by Carmen Carlan and Mike Parsons. The event looked at progress in autonomous and highly automated transport and the problems encountered in implementing systems that are sufficiently safe.

This was the very first SCSC seminar held outside of the UK and was the inspiration of Carmen Carlan who felt that an event in Germany would better facilitate attendance from those, for various reasons, who were not normally able to travel to the UK. Carman was instrumental in organising the event, arranging the speakers and helping select a suitable venue. She also co-hosted with Mike Parsons, who, along with Alex King, Brian Jepson and myself, made the journey via train from the UK.

Destination: Munich!

Rather than fly separately, we decided to meet up at St Pancras in London and take the train as there was quite a bit of high-value AV equipment and we would be able to collaborate more closely travelling together.

The only problem was that the train journey entailed covering the best part of 1,000 miles for us all. Even if the trip had gone smoothly, we each had about 16 hours of travelling ahead of us! Unfortunately after

meeting up at St Pancras and a relatively smooth journey to Brussels, we arrived to find our Frankfurt train cancelled (we were scheduled to change there for the final leg to Munich). With not much information from staff on the ground and conflicting digital-based information on various European websites and apps, we decided to ignore the attendant's advice to travel to Aachan (just inside the German border) and wait for the next Frankfurt train. After lunch and feeling pleased with ourselves, we boarded the fast intercity ICE train at Liege Guillemins only to be told that the trip was going to be delayed 90 minutes and we would be better getting the regional train from Aachan! A further change at Cologne left us with a 5-hour journey to Munich and we eventually arrived at the hotel close to midnight, three hours late!

Markets, Soup and Museums!

The team did have some time the next day to recover and see a few of the sights before setting up the room for the seminar the following day. Munich itself was very pleasant with a vibrant cosmopolitan feel enhanced by the Christmas market stalls that were marvellously lit in the evening. The stalls of course abounded with stollen, bockwurst, chocolate, cheese and Christmas stocking fillers, but there were also market stalls selling fruit and veg, giving the city a more local village feel, which was delightful. We eventually found a marvellous Suppenküche selling soups of course, but also chilli-con-carne and other dishes. Standing room only but very good and we rounded the meal off with some excellent coffees.

The rest of the day turned into a mini "Tech Trip" for Mike and myself. The Deutsches Museum sits next to the river Isar and is reputedly the world's largest museum of science and technology. We only had a few hours there, but it did not disappoint with many wonderful exhibits many being hands-on.

Particular interactive favourites for me were the quantum experiments such as the 'double-slit' and the 'information-eraser'. The museum also had a room dedicated to bridge building and architecture, which was fascinating and informative and the models used to illustrate building processes were mini works of art in their own right!

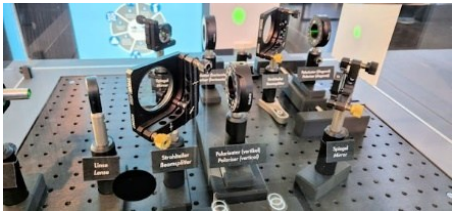

A model railway with a snowbound theme filled half of one room and there was a whole floor dedicated to aviation with more German-oriented craft such as the Messerschmitt and the Junkers with their distinct 'corrugated' bodywork.

We rounded the day off in the Ratskeller a very large and popular restaurant in the wonderful undercroft of the New Town Hall ('new' is perhaps a misnomer as it is over 150 years old now!) No meal is complete in Munich of course without a Bavarian Weissbier or two!

Trains, Planes and Automobiles

The seminar took place on the Thursday and was very well-attended with 32 in-person and 2 online. The Eurostars Book hotel was a very comfortable location with plenty of space for delegates and the refreshments and lunch were ample and of a high standard.

As the name suggests, the hotel has a literary theme with a marvellous Trompe L'oeil on the wall in the reception and many classic titles decorating the wall.

The six speakers presented a wide range of transport-related AI topics, touching on the domains of autonomous trains, aircraft and road vehicles:

- Mario Trapp, Safe Intelligence, Fraunhofer IKS, TUM - Assuring Safety in the Face of the Unpredictable
- Alex Haag, Futurail - How Safe is Safe Enough for Autonomous Trains? Analyzing Human Performance as a Reference System
- Levi Lúcio and Christoph Neuböck, Airbus - Lessons Learned from Enabling Model Based Systems Engineering for a Large Autonomous Aircraft Programme
- Torben Stolte, Volkswagen - ADMT's Approach Toward Arguing Safety for Automated Driving - An Introduction
- Simon Burton, University of York, UK - Safety Under Uncertainty: Automotive Standards for AI Safety and Research Perspectives
- Henrik Putzer, Cognitron - Developing & Assessing for Trustworthiness in AI

Mike Parsons introduced the event with an introduction to the SCSC and highlighted some of the perceived challenges of AI; not only in how we can assure systems, but how AI might need to be used to give us additional tools to help address the emerging challenges.

Mike then handed over to Carmen Carlan who introduced each of the speakers in turn.

Mario Trapp

Prof. Mario Trapp was our first speaker and started by describing the "Economic Safety Gap" where the excitement around the utility that AI can bring is exponentially driving technological innovation, but leaving an increasingly large gap between use and safe use. Safety engineers are also in a very small minority and so there is opposition around those that raise safety concerns.

Mario asked, what can we do to give safety a voice? He said we are a linear discipline now operating in an exponential world. Mario proposed that we need to lift safety to the next level as part of the next generation of safety. This he called 'Safety NxT' – a 5-point framework as summarised below:

Nos 1: To Change the Narrative

Traditionally, safety is in the nay-sayers domain and we are not good at storytelling. We need to tell our story in a more motivational and constructive way. Safety is a very small minority compared to all other stakeholders and so it's important to give one consistent voice for safety.

Nos 2: Design

Safety needs to be part of the design process from the beginning. A key point is the *value* that the system might bring and the value we want to create.

Nos 3: Speed

There is a desire to increase the speed of release of updates as technology is evolving exponentially. For example, Tesla has approximately 40 releases a year, so we can no longer take months to assess a release and so we need to be quicker than we are today.

Nos 4: Adaptability

Things that worked in the past may no longer work in the future, especially with AI. Typically in automotive, there is a 14-day cycle to recertify any change. This does however offer an opportunity to try things. For example, we can explore missing evidence gaps through gathering small pieces in the field and provide micro-evidence.

"We need to think about the value: the more capability we need, the harder it is to guarantee safety – if it is not safe there will be no value"

Nos 5: Adaptivity

What can systems do themselves to adapt? Mario's area of research is looking at how systems can be taught to adapt to an unpredictable context. Traditional safety assurance usually considers the worst case outcome for a hazard, but considering the current operational situation as part of a dynamic risk assessment at runtime will give freedom to optimise utility without violating safety.

Alex Haag

Alex was our second speaker. Alex works for the startup company Futurail, which aims to advance autonomous train technologies and AI products are seen as paving the way to full autonomy.

Alex asked why have autonomous trains at all and said the team were not targeting high volume trains like ICEs (less of a business case for relatively low number of vehicles), but more the thousands of regional trains where there is an acute driver shortage (5 drivers are required for each train due to working hour constraints). Replacing the drivers reduces costs not only in salaries, but also in the high cost of driver training and it can also increase line capability by optimising travelling speed.

Alex showed the high-level architecture of their Objective Detection System, which uses a combination of optical and LIDAR sensor data to detect objects. A 'structure gauge' is overlaid to show the safe space to be occupied by the train. He showed a video of the detection challenges and the need for filtering, for example, the train shouldn't brake if a pigeon flies in front of the train, and fog will introduce noise into the data. Alex said their systems are installed on a train in Belgium and have already capture 30,000km of very diverse data (eg. tunnels, narrow track, different seasons and climate conditions).

Vegetation detection is important as well as the rail itself and platforms can be challenging as, by design, they are very close to the structure gauge and so there is a need to distinguish things that are expected to be close and those that are not like a pedestrian.

Alex said their aim is to have a SIL 2 certified system by the end of 2026 in time for live train operations.

One question is how safe is safe enough? And what does SIL2 mean in this context and there are still discussions in these areas. The traditional methods to demonstrate safety (from EN50126) cannot be used:

- **Explicit risk estimation**: requires precise data that is not available (obstacle type, frequency in all environmental conditions)
- **Standard practices**: no Rail Object Detection Systems and AVs are all in US and China
- **Reference systems**: object detection is all human-based. There is accident data available but not avoided accidents

It is important therefore to estimate obstacle occurrence from unavoidable accidents, typically at night or in curves. Based on some assumptions, the team have calculated from accident data from trains running in South West France, that humans avoid 30% of accidents in these scenarios and so puts the target for an AI system into perspective, for example, maybe 90% detection rate for AI may be acceptable. Alex concluded by saying that this is still work in progress and more diverse data is being captured to help validate the system.

Levi Lucio and Christoph Neubock

Christoph presented remotely from Toulouse discussing MBSE in Airbus for a large autonomous aircraft programme running with a team of over 300 people. Christoph showed the V-model system engineering approach and showed the scope of the MBSE covering requirements definition through to design definition. The model covers all levels from the entire aircraft right down to equipment level.

Christoph said that because of the size of the team, there was a need for MBSE Model management to manage the combined efforts and contributions of 300 people.

This, for example, involves the assessment and measurement of the quality of the models. Even with everyone using MBSE, there is still a need to produce documents, including data and configuration management and there are security and IP rights to consider.

Levi then took over as lead presenter and said things have evolved since starting with MBSE and he shared some of the problems encountered, which occurred at many different levels:

Cultural Challenges: people were reluctant to discard the legacy processes that they were familiar with, there was missing knowledge/skills/ability for MBSE developments, missing understanding and still a heavy dependence on document-based ways of working.

Lessons learnt: value legacy (don't throw away as people are attached to it).

Scalability: there is a need to manage models at scale. Vendors' product features are not always as expected and it is hard to turnaround changes quickly. Levi said that they ended up writing software and tools themselves.

Lessons learnt: start with big bang but be flexible and make and introduce rules over time. Handle data exchanges and variants as first-class citizens.

Integration Challenges: There were issues with the integration of organisations, processes, methods and tools.

Lessons learnt: build the community on small and constant successes.

Software Tools: don't expect tooling to give you everything you need, but if you write it yourself, any software development has to be professional.

Lessons learnt: Mainly a matter of education on quality and insufficiency of external tools.

There are now, however, achievements of note gained through much pain!

Levi concluded by describing the road ahead. The first objective is to make sure the product can be architected (get the plane to work!) He thought they still need much more automation as there are still a lot of manual activities. More also needs to be integrated and many more lessons are expected to be learned!

Torben Stolte

ADMT's is VW's **A**utonomous **D**riving **M**aaS (mobility as a Service) and **T**aaS (Transport as a service). ADMT is their approach to arguing the safety of Level 4 Autonomous Vehicles, that is, AVs with no human drivers but operating in a limited area such as shuttle services and goods transport.

Torben said that we want AVs to be safe but we need a proper definition of "safe" as there is no universally agreed specification for the term. VW use the ISO26262 definition: "safety is the absence of unreasonable risk" but Torben asked how is unreasonable risk defined?

An engineering perspective might simply use a severity versus occurrences model but as Phil Koopman has said, 'safety is more than counting injured or dead people' and is much more nuanced. For example, there is biasing that may occur from particular road users such as cyclists. How also would targets be defined and how would these be modulated by the class of driver (eg. young, drunk, elderly etc.)

There is a need for safety argumentation but there are challenges:

- Proving, with statistical significance, events that are very rare would require many (billions) of miles to be driven, which is impracticable. Statistical proofs are not reasonably achievable prior to deployment.
- There is uncertainty about the safety level expected by the public, the perceived level of safety compared to the engineered level of safety, and perceptions are likely to evolve over time
- Regulatory demands are very coarse-grained at the moment

The approach is therefore to use safety indicators that collectively combine to form a claim that the system is free from unreasonable risk. Prior to operation, these will be prospective and include aspects such as good design and post-deployment they will be retrospective indicators using more statistical argumentation.

The Safety Assurance Case being developed is based on Goal Structuring Notation (GSN) and Torben provided a high-level GSN diagram of the argument but said structuring the argument is a multidimensional challenge: there are assurance domains over the system lifecycle but also cross-cutting topics that also need to be considered such as Safety Culture.

A key research question is how to express uncertainty in the Safety Assurance Case. For example how to address aspects such as qualitative versus quantitative approaches, argumentation and evidence uncertainty and the sheer diversity of topics.

Torben concluded with some observations:

- Safety argumentation must fit to the organisation
- Generalised argument patterns are hard to define due to diversity of topics
- Safety argumentation is an evolutionary process
- Top level safety claim formulation does not matter too much
- Terminology and consistent use of language is important

Simon Burton

Simon started his talk by stating that AI/ML models will eventually fail sometimes under certain conditions and asked a number of questions: how safe is safe enough? What measures should be taken to reduce risk? How do we argue an AI-based system is safe? Are there limits to what we can reasonably achieve and how do we interpret existing standard and regulations?

Simon said he was asked by ISO to lead the standardisation effort for AI so he first looked at existing standards. He found that ISO26262 focusses on malfunctioning behaviour but doesn't particularly address risks with non-malfunctioning but primarily erroneous behaviour. He also looked at ISO21448 (SOTIF), which requires the absence of unreasonable risk from functional insufficiencies and so might be a better standard for AI safety.

He then introduced ISO PAS 8800: "Road Vehicles – Safety and Artificial Intelligence". The intention was not to repeat requirements and guidance already held elsewhere so, for example, PAS 8800 does not cover functional safety arising from malfunction.

Simon provided an overview of the standard and said the draft was approved on 16th October 2024 and full release is scheduled for 13th December 2024.

He shared his thoughts on what was next for AI and safety and research directions and said one of the biggest open questions is:

Are we able to make definitive statements about the safety of AI?

He concluded that "it depends!" Simon defined uncertainty as "deviation from the unachievable ideal of completely deterministic knowledge of the relevant system" and showed a diagram that illustrated how there are a number of layers that compound uncertainty.

He said there was a relationship between the complexity of a system and our ability to make definitive statements about the safety of the system; the assurance argument therefore gets harder as the system gets more complex. He then showed how the standard could be applied by using the example of a construction sign detection system.

"Are we able to make definitive statements about the safety of AI? ... it depends!"

Simon then presented his view on the limitation of safety assurance providing examples of applications as they sit on a graph of complexity of environment/task/data/model versus assurance uncertainty with some "Frontier AI" for safety-critical applications now being suggested such as using LLMs (self-trained on the Internet) for AV driving tasks.

He then listed the fundamental challenges that need to be overcome:

Can AI still help in managing the complexity of systems? Should we really use LLM which are simply "guessing the next word" in safety systems or can they be used in a more hybrid fashion? Simon concluded with a summary of the foundations of an AI safety argument (found in the standard but still evolving).

Henrik Putzer

The final presentation was from Henrik covering Developing and Assuring for Trustworthiness in AI. He structured his talk around the elements in the seminar's title: The Good, The Bad and The Ugly.

The Good

AI is quick to market and there is commercial imperative to include it in future products from company's like Bosch. He said there is a lot of AI going on and it will undoubtedly be part of our systems in some form in the future.

Henrik asked what is so special about AI and looked at the history of technology usage from mechanics, electronic and software and now arriving at AI/ML/DL. AI can make decisions about larger unseen datasets based on only partial samples of training data. However, it also introduces uncertainty related failures and Henrik gave some of example of 'AI Fails'.

"We shouldn't be afraid of AI but afraid of those that don't use it properly"

The Bad

There is Mysticism about AI, which has been seen in research – changing the order of training data sets or quantifying a neural net model can suddenly make everything work. We need to change Mysticism using structured processes.

Henrik then looked at the timeline of functional safety and asked whether the old risk-based approaches can be used in AI? Henrik thought so, but with some tweaks, in particular there is a need to consider fault models in the AI.

In summary AI still sits within the system architecture and covered by standards like ISO 26262. A risk-based approach can still be adopted so it's important to understand the AI failures through a fault model and there needs to be a structured approach, so a V-model for AI including specific methods to be applied.

Henrik then discussed a standard that introduces a structured approach to developing Deep Neural Network (DNN) elements (VDE-AR-E 2842-61 part 5). The approach is based on a V-model. The left-hand side has the Design phases where the training dataset is built. After hyperparameter optimisation there is then a series of verification activities on the right-hand side with training, design, functional and semantic verification.

The Ugly

Henrik said "The Ugly" side of AI from an engineer's perspective is when the assessors come along! However, by following the standard the situation should not be so bad as embedded in the V-model are the general process areas that an assessor will be looking for.

In summary Henrik said that he believes there will be (and is already) beneficial (and safe) AI but we have work to do and there needs to be system (of systems) engineering to support it. We shouldn't be afraid of AI but afraid of those that don't use it properly.

Mike Parsons – Wrap Up

Mike wrapped up the seminar with a general discussion session, which allowed the attendees to discuss their thoughts on specific AI-related questions such as:

> *Question: Do we have the whole ecosystem we need for AI covering all aspects such as development, evolution, regulation, maintenance and documentation?*

Mike then opened the floor for any further discussion topics.

Conclusions

Despite some of the epic train journeys required to bring the event to mainland Europe, the seminar was hugely successful with fantastic speakers in a great venue. The event being so well-attended led to some great discussions during the breaks and it was a fantastic opportunity to share ideas and current best practice in the challenging area of AI in transport systems.

Feedback from delegates on the event have been very positive and so we hope this will be the first of many more seminars to be held on mainland Europe.

Report by Paul Hampton SCSC Newsletter Editor

Images: © Paul Hampton and Mike Parsons

Recent Safety Publications

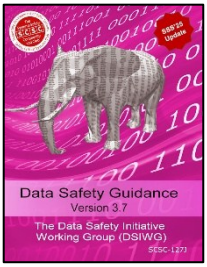

Guidance on the management of data safety risks (V3.7)
scsc.uk/scsc-127J

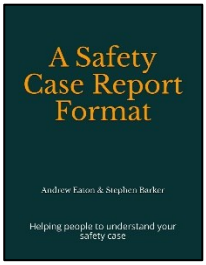

A Safety Case Report Format Andrew Eaton & Stephen Baxter
www.amazon.co.uk/dp/B0DPL3C84P

Safety-Critical Systems eJournal vol.3 no.2
Summer Issue 2024
scsc.uk/scsc-196

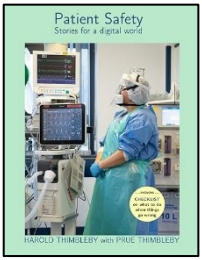

Patient Safety – Stories for a digital world
www.amazon.co.uk/dp/1399975420

Safety Assurance Objectives for Autonomous Systems.
scsc.uk/scsc-153C

Proceedings of the 32nd Safety-Critical Systems Symposium.
scsc.uk/scsc-188

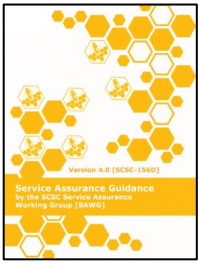

Service Assurance Guidance Version 4.0. Jan 2024
scsc.uk/scsc-156D

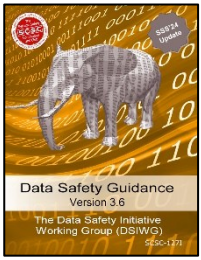

Guidance on the management of data safety risks (V3.6)
scsc.uk/scsc-127I

Safety-Critical Systems eJournal vol.3 no.1
Winter Issue 2023
scsc.uk/scsc-191

Book Review

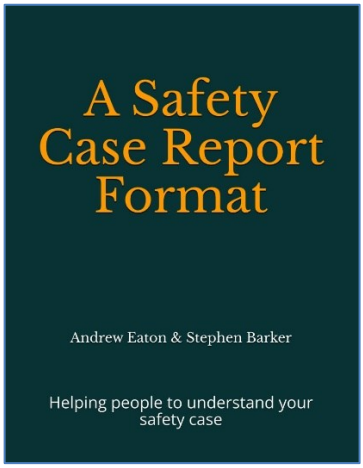

The size and complexity of safety cases is usually such that they are not easy to deliver and difficult to understand. A Safety Case Report can therefore be used to present the key parts of the safety case, which can then be more easily delivered and understood – but how should such as report be structured and presented?

In their book 'A Safety Case Report Format', Andrew Eaton and Stephen Barker provide a generic format for such a Safety Case Report to help anyone responsible for justifying and communicating a safety case to other parties.

Steve Kinnersly provides a review of the book and shares his views on how well it meets its objectives.

Having agreed to write this review, I asked a friend who has decades of experience in safety matters whether he had any thoughts on Safety Case Reports (SCRs). *"I've read hundreds of them*," he replied, "*that's why I'm now a cynical old #*#*#*!*"

In case you are wondering what these seemingly hazardous documents are, an SCR is a document that summarises the arguments and evidence of a safety case. It is not the safety case itself but rather a high-level guide and a gateway into the full safety case. Interim SCRs may be used to document progress against a project's safety programme.

As an Independent Safety Assessor (ISA) for projects that have often lasted years, I have seen many grow from initial outline to full and final fruition. From this has emerged a personal wish list for an SCR. It should be:

1. Clear – in both writing and structure
2. Concise – not verbose, avoiding what is not necessary
3. Comprehensible – for the intended readers
4. Comprehensive – addressing all relevant aspects of the system (or service etc.) and its intended uses
5. Complete – addressing all the evidence and argument needed to reach its conclusions
6. Coherent – logical and balanced
7. Correct – a faithful summary of the safety case

If it fails on one or more of those points then it might not be as convincing as the author hoped – and who writes a SCR and does not want it to be convincing? Anything that helps SCR authors to satisfy these seven points must surely be a good thing.

'A Safety Case Report Format' aims to "*help people to understand your safety case*". It mainly addresses the argument and evidence needed for an SCR to be comprehensive, complete and coherent. The focus throughout is technical; it is not a style guide so has little to say about points 1-3 in my wish list. A generic, high level contents list for an SCR is given and used as a framework for a thorough, detailed, reasoned and logical presentation of what an SCR should address and how it should be structured. As befits a book that is aimed at the full range of possible SCRs, the emphasis is on recommendations and guidance: the writer of an SCR must decide whether and how they apply to their situation.

"… the emphasis is on recommendations and guidance: the writer of an SCR must decide whether and how they apply to their situation"

The authors, Andrew Eaton and Stephen Barker, each have more than two decades experience working for the UK Civil Aviation Authority (CAA), chiefly on the safety of air traffic control systems. Their background is reflected in how they have approached the book. It is written from an assessor's top-down viewpoint of what needs to be in an SCR in order for them to be convinced, rather than from a project safety engineer's bottom-up viewpoint of how to summarise all the available safety argument and evidence in a single document. Put simply, an SCR produced in accordance with 'A Safety Case Report Format' should provide an assessor with all the information needed for their assessment. Indeed, the authors have used as a reference point the CAA document CAP 1801 'Assessment of Change Safety Cases' (for which they were joint authors), which provides guidance for assessing a generic (not aviation specific) safety case.

'A Safety Case Report Format' claims to be "*generically applicable, not specific to any regulatory regime or legislation*". The range of possible safety cases and thus SCRs is, of course, huge. This challenge is addressed by including everything which might be needed for what the authors consider to be the most complex situation for a safety case to address, specifically a change rather than a new build since "*a change ... has the most topics to address*". The user must then tailor the contents to their situation by selecting and adapting. While this is a reasonable approach in principle, it remains to be seen how well it works in practice. The authors note that a user needs to have *"adequate competence in safety assurance and the subject of the safety case*". The absence of examples does not help here – perhaps something for the authors to address in a Part 2?

Given that the book aims (and appears) to be comprehensive, a user intending to write an SCR is faced with an important question: since an SCR must be a faithful summary of the safety case, what if something the book identifies as needed is not in the safety case itself? The answer given in the book is that it must be decided whether this "*can be justified or whether the safety case needs to be enhanced before the SCR can be produced*". It is hard to disagree with that. It is also hard to avoid concluding that the book is as much about what needs to be in the safety case itself as it is about the corresponding SCR. Why not use the book at an early stage in a project to help design a convincing safety case and thus ensure that all necessary evidence is produced by the safety programme? There should then be no gaps when summarising the arguments and evidence in the SCR.

A key question for a 'how to' book such as this is how well it works in practice. There are no examples in the book and, being newly published, there is not yet any evidence from 'real world' application. From personal experience, the recommendations and guidance given in the book appear to be reasonable and consistent with good practice in safety. As an ISA, I would have been happy to see a project following them (suitably interpreted and only if they apply to that specific project, of course). However, I would want to see that the project treated them as they are intended: prompts for careful consideration rather than instructions to be followed blindly.

The authors' background in air traffic control safety is no doubt responsible for some idiosyncrasies which seem at first sight to limit the book's scope of application. Air traffic control safety is primarily concerned with the safety of a service. This is presumably why the book is written in terms of an SCR for the safety of a service. Someone responsible for an SCR for something other than what is normally considered a service might reasonably wonder whether this book is for them. This is addressed in the book as follows:

1. 'Service' is defined very broadly as *"an output from a functional system that is intended to be of use",* the output "*may be either intangible… or tangible (e.g. a product, commodity or function)*" and protection against harm "*is itself a service*".
2. Harm is not restricted to what might be caused by a service itself, it also includes harm that might be caused by the system which delivers the service (called 'local harm' in the book).

The authors note that this "*may require some adjustment of perspective for some people".* Whether it is sufficient to cover all situations for which an SCR is needed remains to be seen – probably yes, but this reviewer has struggled here and is still hiding in the long grass!

A substantial part of the book is concerned with what needs to be in the SCR for each transition stage during a change. This might reflect the fact that an essential part of air traffic control safety is maintaining safety during the transition stages involved in implementing a change. Someone producing an SCR for something that is not a change to an existing system or service might reasonably wonder why that part of the book is relevant to them. The answer is that in the book a transition stage is regarded as any stage in which the system is operating in state that is not a final operating state. This implies that, for example, commissioning tests might be a transition stage. My advice to a potential user is to think carefully about whether you can dismiss this part of the book: you might actually know transition stages by a different name.

A background in air traffic control safety might or might not be responsible for a third idiosyncrasy. As an ISA, I always expected an SCR to include a clear and prominent statement of and justification for top level safety requirements which then flow down into lower level requirements (derived requirements etc.). It was therefore a surprise to find that this did not seem to be required here. A trip to the Index and Glossary was needed to find the reason. Top-level safety requirements are 'Risk acceptance principles', 'safety criteria' (etc.), with 'safety requirement' being used only for lower level requirements. This appears to be a matter of terminology rather than substance. However terminology is important. The reader is advised to read the introductory material and glossary carefully to ensure that they understand the meaning of terms and concepts as used in the book rather than assume that it follows conventions used in their own domain.

In conclusion, 'A Safety Case Report Format' is a thorough, logical and reasoned treatment of an important topic. It provides comprehensive recommendations and guidance on what should be considered for inclusion in an SCR. As such, it can help not just with writing an SCR but also with designing and developing the safety case itself.

"'A Safety Case Report Format' is a thorough, logical and reasoned treatment of an important topic"

The aim of being generic and thus applying to any and all domains and possible safety cases appears to have been met. However, a consequence is that a user must exercise considerable judgement in interpreting, selecting and tailoring the contents to their own situation. This is therefore a book for experienced safety professionals to use. Indeed, it is arguable that it might best be used as a basis for developing domain-specific versions that can be used with less tailoring.

Finally, would 'A Safety Case Report Format' have stopped my friend becoming a "cynical old #*#*#*!" ? Ah, that is not for me to say…

Steve Kinnersly B.Sc. D.Phil. MRSC

Steve Kinnersly has spent most of his career in various aspects of safety. This has ranged from modelling accidents at nuclear power plants to independent safety assessments of complex, high hazard systems involving hardware, software, people and procedures. His work has included software development and validation, technical management and consultancy. Steve was a member of the Independent Safety Assurance Working Group from its inception and led the development of its Code of Practice for Independent Safety Assessors (ISAs). Now retired, he maintains an interest in safety and related matters.

60 Seconds with ... Aimee Avrill

Aimee is a qualified Project, Programme and Change Manager with experience working across various industries, including Defence, Public Sector, Logistics, Banking and Technology. She's a member of the Association for Project Management and Chartered Management Institute.

With a passion for diversity, equity & inclusion, Aimee joined BAE Systems in 2022 as a Diversity, Equity & Inclusion Manager, leading the company's Employee Resource Group Programme across the UK, Kingdom of Saudi Arabia and Australia. Over the past two years, Aimee has led the expansion and growth of the programme with a membership base of 12,500+ globally.

For her recent work, she was recognised as a finalist for the Inspiration of the Year Award at the LGBTQ+ in Defence Awards 2024.

Aimee is originally from Jersey, Channel Islands, where she served as a Police Officer in the States of Jersey Police Force. She moved to Devon in 2016 where she lives with her partner, Sophie. In her spare time, she enjoys going to festivals, concerts and dining out.

You are a champion of Diversity, Equity & Inclusion (DE&I) – in a few sentences, how would you describe DE&I?

DE&I is fundamental to creating an environment in which everyone – regardless of their race, gender, age, sexual orientation, ability, background, or any other characteristic – can thrive. We can all help to create a place where people feel they belong and are respected and accepted for their authentic selves.

What first attracted you to working in DE&I?

It stems from my personal experiences in previous workplaces, where I witnessed people being overlooked or undervalued because of certain aspects of their identity. This isn't only morally wrong, it's bad for business! I saw the difference it makes when people feel included, respected, heard and accepted. That inspired me to join an Employee Resource Group and be part of the change, before securing my current role within the DE&I Team.

What aspect of your career are you most proud of?

A few years ago, I made the decision to leave an organisation that didn't align with my values. The role was everything I'd ever wanted, but the workplace culture wasn't. This particular company valued presenteeism, extended working hours and failed to address bad behaviours, all of which impacted my happiness, mental health and productivity. I'm proud of the decision I made because now I work for a company where I feel respected and valued. It's a company where my voice is heard and I'm encouraged and supported to maintain a positive work-life balance.

What advice would you give to yourself age 12?

I'd encourage myself to take risks, ask questions, and not be afraid to challenge the status quo. Up until the past 10 years, I always felt like everything had to be perfect. As time went on, I realised that growth often happens through those moments when things don't go as planned. We shouldn't be afraid of mistakes, but embrace them.

How important is DE&I to a discipline such as system safety?

Diversity of thought is crucial to spot business opportunities and address business challenges. But to harness diversity of thought, organisations must create an inclusive culture that encourages participation and values different perspectives.

What's your most favourite quote or motto?

"*One day you'll leave this world behind, so live a life you will remember*" – Avicii, The Nights. Great song, great quote.

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

I'd love to learn British Sign Language (BSL). Communication is one of the most powerful ways to connect with others and being able to sign with people who are deaf or hard of hearing would allow me to have meaningful conversations with friends and colleagues who use it as their primary language.

How do you see the rise of Artificial Intelligence affecting the objectives of DE&I?

Broadly speaking, I'm excited by the rise of AI, but do believe it will present both opportunities and challenges. On one hand, I foresee AI being a powerful tool in advancing DE&I – enabling data-driven decision making, enhancing accessibility options and helping to create personalised pathways for employees. But if not carefully managed, it could also reinforce biases and inequalities. It's important that AI systems are designed, implemented and continuously monitored to identify and mitigate these risks.

Connect

The Newsletter and eJournal

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. There are now two publishing vehicles for content – shorter, more informal content, can be published in the Newsletter with longer, more technical peer-reviewed material more suitable for the eJournal. If you are interested in submitting content, then get in touch with Paul Hampton for Newsletter articles: paul.hampton@scsc.uk or John Spriggs for eJournal papers: john.spriggs@scsc.uk

The SCSC Website

Visit the Club's website thescsc.org 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.

Facebook

Follow the Safety-Critical Systems Club on its very own Facebook page.

www.facebook.com/SafetyClubUK

X/Twitter

Follow the Safety-Critical Systems Club's X/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 or event you would like to advertise in the Newsletter? The SCSC Newsletter can reach out to 1,000's of individuals 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

Safe System Architecting

The SCSC Safe System Architecting Working Group (SSAWG) is a joint initiative between the INCOSE UK Architecture Working Group and the Safety Critical Systems Club, with a vision of addressing systems safety through the promotion of choosing, using, and developing appropriate architecture. Two focal areas are defined as:

- Safety-driven architecting – what safety drivers are likely to exist when architecting is being considered.
- Architecting system safety – what architectural factors may realise, enable, support, or preclude the achieving of safety considerations.

The groups Mission Statement is:

- To produce practical guidance on addressing system safety during system architecting.
- To encourage and facilitate collaboration between INCOSE and the club.

Statement of the Problem

The design of a system architecture is a critical first step toward addressing safety requirements. With the community moving toward digital engineering to include model-based systems engineering, digital twins and digital threads, there is an opportunity to advance current best practices in architecture to:

- Embed safety concepts into architecture concepts, as opposed to capture safety solely as an architectural view.
- Develop meaningful metrics for the evaluation of architecture against safety goals, which will feed into the overall architecture trade-offs.
- Generate (digital) architectural artefacts that contribute to the construction of the safety case.

It is envisaged that the outcomes developed in this joint working group will be beneficial to both the systems engineering community and the system safety community.

WG Operation

The WG will meet every 6 weeks online via Teams. Note that the INCOSE UK AWG meets every quarter.

If you would like to be involved in the group, please contact the co-chairs.

Cochairs: **Siyuan Ji** s.ji@lboro.ac.uk and Jane Fenn jane.fenn@baesystems.com

SCSC Working Groups

The Safety-Critical Systems Club is committed to supporting the activities of 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

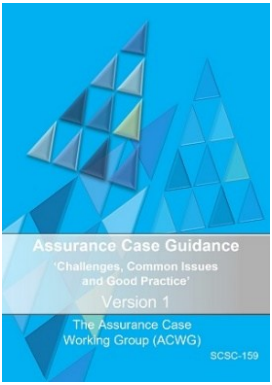

In Aug 2021, the group published v1.0 of the Assurance Case Guidance: scsc.uk/scsc-159

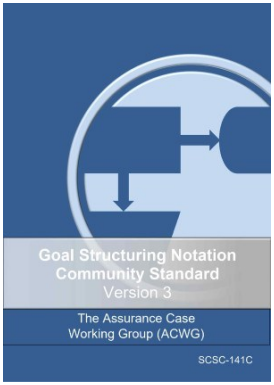

One of the working group's activities is the maintenance of the Goal Structuring Notation (GSN) Community standard.

See scsc.uk/gsn for further details.

In May 2021, the group published v3.0 of the standard: scsc.uk/scsc-141C

Lead Phil Williams phil.williams@scsc.uk

SCSC Working Groups

Security Informed Safety

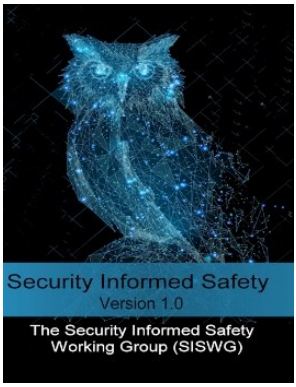

The Security Informed Safety Working Group (SISWG) aims to capture cross-domain best practice to help engineers find the 'wood through the trees' with all the different security standards, their implication and integration with safety design principles to aid the design and protection of secure safety-critical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

Data Safety Initiative

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

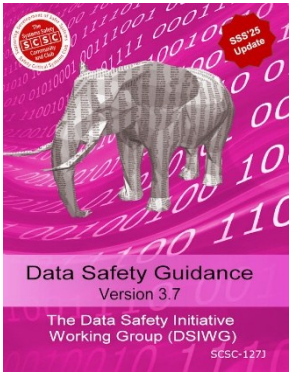

It is acknowledged that data has been a contributing factor in several incidents and accidents to date and there is foreseeable harm that can arise from Machine Learning and Large Language Models' use of data by Artificial Intelligence (AI) systems that are subject to issues of biasing, interpretation and, arguably, falsification. 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.

An update to the guidance (v3.7) was published in Jan 2025: scsc.uk/scsc-127J

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

SCSC Working Groups

Safe AI

The SCSC Safe AI Working Group (SAIWG) aims to capture cross-domain best practice and guidance on key topics within the design, evaluation, assurance, and approval of safety systems that use or are developed using AI, bringing together emerging standards and key results from the incredible amount of research being conducted into AI safety.

The working group was kicked-off at SSS'24 and is led by Alan Simpson. The SAIWG will conduct regular meetings, workshops, and publications to share knowledge and experience on various topics related to AI and safety systems, such as coordination of safety with other disciplines, evaluation of risk, and mapping of terminology and language.

Lead **Alan Simpson** alan.simpson@ebeni.com

Safety of Autonomous Systems

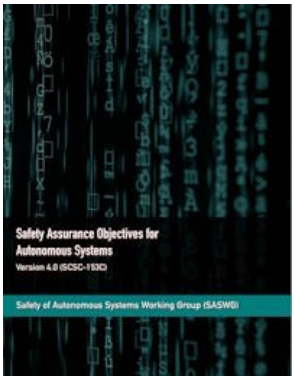

The specific safety challenges of autonomous systems and the technologies that enable autonomy are not adequately addressed by current safety management practices and standards.

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.

The group will work closely with the new Safe AI Working Group (SAIWG) as AI is used extensively in Autonomous Systems.

The group published v4 of its guidance Safety Assurance Objectives for Autonomous Systems, in Feb 2024 scsc.uk/scsc-153C

Lead **Philippa Ryan** pmrc@adelard.com

SCSC Working Groups

Multi- and Manycore Safety

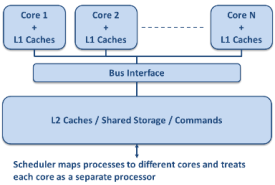

It is becoming harder and harder to source single-core devices and there is a growing need for increased processing capability with a smaller physical footprint in all applications. Devices with multiple cores can perform many processes at once, meaning it is difficult to establish (with sufficient evidence) 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 **Lee Jacques** Lee.Jacques@leonardocompany.com

Safer Complex Systems

The Safer Complex Systems Working Group (SCSWG) builds on the IET/RAE work already done in this area. It is recognised that the RAE work is ongoing and collaboration is encouraged. The group's mission is to produce practical guidance on developing, managing and assuring complex systems throughout their lifecycle (so as to achieve and justify their safety).

The following provides a statement of the problem:

- It is acknowledged that complex systems are becoming more prevalent with more opportunities to cause harm
- There are also new complex systems arising from using combinations of existing systems and services which are then used for safety purposes
- Complex systems are not sufficiently addressed in current safety management practices and standards
- In particular, complex interactions and emergent behaviours are not currently assessed and managed sufficiently
- There could be benefit in developing new analysis, tools and techniques to manage complex system risks
- There are clear business and societal benefits, in terms of reduced harm, reduced liabilities and improved business efficiencies, in improved management of complex systems risk

The group is aiming to produce draft guidance for SSS'25.

Contact the group lead if you would like to attend future meetings or find out more.

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

SCSC Working Groups

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.

The group published guidance v4.0 in Jan 2024: scsc.uk/scsc-156D

Lead Kevin King kevin.king@baesystems.com

SCSC Safety Culture

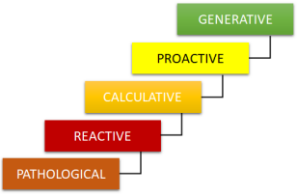

The Safety Culture Working Group (SCWG) has been established to provide guidance on creating and maintaining an effective safety culture. The group seeks to improve safety culture in safety-critical organisations focussed on product and functional safety, by sharing examples and latest approaches collated from real-life case studies.

Meetings provide an opportunity to discuss any particular aspects attendees are interested in taking forward, and to help set future directions for the group.

In Dec 2023, the group published a position paper for assessing and managing safety culture. https://scsc.uk/r189:1

Lead Michael Wright michael.wright@greenstreet.co.uk

Systems Approach to Safety of the Environment

The Systems Approach to Safety of the Environment Working Group (SASEWG) is a new group intending to apply Systems Safety practices to systems that are embedded within the natural environment, while focussing on that environment.

The group aims to produce clear guidance on how engineered systems should be developed and managed throughout their entire lifecycle so as to preserve, protect and enhance the environment.

Please get in touch with the working group lead if you would like to join or find out more about this group.

Lead James Inge james.inge@scsc.uk

SCSC Working Groups

Ontology

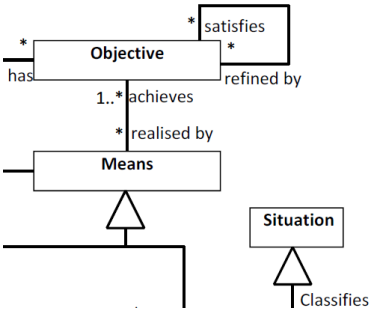

The Ontology Working Group (OWG) develops ontologies that will form the basis of SCSC guidance, as well as having wider industrial and academic applications.

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

The Safety Futures Initiative

The Safety Futures working group meets once a month to bring together fresh perspectives, innovative ideas and insights.

The Safety Futures is diligently working on developing a comprehensive roadmap to guide new safety professionals in exploring career paths across various industries. To supplement this, they have an upcoming initiative to create a University 'League' Table to rank universities based on their safety engineering courses and the career pathways they support.

They are also looking to start a reverse mentoring programme where we can match more experienced people with new professionals for information exchange.

The group encourages anyone to get involved to help shape the future of safety engineering careers.

Lead **Khadijah Khatun** khadijah.khatun@scsc.uk

SCSC Membership

The SCSC provides a range of services to the System Safety community including seminars, tutorials, leadership events, specialist topic working groups, the annual symposium and a comprehensive body of publications. Membership brings many valuable benefits such as free access to online events, the SCSC Newsletter and access to presentations and other resources from events.

Individual Membership

To become an individual member of the SCSC please register on the SCSC website using the icon at the top right of any page and select "Register". Complete and save your account registration and then verify your email address. Once registered and logged in click the link "why not join the SCSC..." inviting you to become a member at the top right of the page or select "Pay membership" from the icon.

Individual membership can be paid online using a credit/debit card through our secure payment partner Realex Global Payments or contact Alex King for other payment methods. For student or retired member rates please contact Alex King to get your account status changed.

Corporate Membership

Your company contact with the SCSC should arrange the membership and any renewals for your organisation. To join as a member covered by a corporate membership, register as per the instructions for an individual member and then contact Alex King to confirm your affiliation.

Renewing Membership

You should be notified by email when your membership is almost expired or shortly after it has expired. These notifications will contain a link to the online renewal page or you will be able to renew when logging onto the website through the 'click to renew' link.

Membership Fees

The following fees are applicable from January 2025 for new and renewing members:

- 1 year Individual Membership: £149
- 2 year Membership: 8% discount: £275
- 3 year Membership: 16% discount: £375
- 1 year SFI Membership: FREE for first year, £35 for years 2 & 3
- 1 year Membership, retired member rate: £35
- For Corporate Membership discounts contact Alex King

A one-month Publication Pass is also available for £15. This allows access to all SCSC website publications in a particular calendar month.

Contact Alex King using office@scsc.uk

The SCSC Steering Group

Current Members


Stephen Bull
stephen.bull@scsc.uk


Dewi Daniels
dewi.daniels@scsc.uk


Paul Hampton
paul.hampton@scsc.uk


James Inge
james.inge@scsc.uk


Khadijah Khatun
khadijah.khatun@scsc.uk


Mark Nicholson
mark.nicholson@scsc.uk


Wendy Owen
wendy.owen@scsc.uk


Davy Pissoort
davy.pissoort@scsc.uk


John Spriggs
john.spriggs@scsc.uk


Carmen Carlan
carmen.carlan@scsc.uk


Jane Fenn
jane.fenn@scsc.uk


Louise Harney
louise.harney@scsc.uk


Brian Jepson
brian.jepson@scsc.uk


Alex King
alex.king@scsc.uk


Yvonne Oakshott
yvonne.oakshott@scsc.uk


Mike Parsons
mike.parsons@scsc.uk


Roger Rivett
roger.rivett@scsc.uk


Sean White
sean.white@scsc.uk

Honorary Members


Tom Anderson


Robin Bloomfield


Dai Davis


Graham Jolliffe


Tim Kelly


Felix Redmill


Phil Williams

Club Positions

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

Managing Director

Mike Parsons 2019-

Tim Kelly 2016-2019
Tom Anderson 1991-2016

Steering Group Chair

Dewi Daniels 2024-

Roger Rivett 2019-2024
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-

Honorary Solicitor

Dai Davis 2022-

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-

eJournal Editor

John Spriggs 2021-

Administrator

Alex King 2016-

Joan Atkinson 1991-2016

Diversity, Equity and Inclusion (DE&I) Lead

Louise Harney 2024-

Wendy Owen 2023-2024

Safety Futures Initiative Leads

Khadijah Khatun 2024-

Zoe Garstang 2019-2024
Nikita Johnson 2019-2021, 2023-2024

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		

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	30

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

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	31

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	30

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

4-6 Feb 2025
SCSC Symposium
York, UK

Safety Critical Systems Symposium SSS'25

scsc.uk/e1099

27 March 2025
Conference
Cambridge, UK

CWIC 2025: CW International Conference 2025

cambridgewireless.co.uk/event-calendar/conferences/cwic.html

8-11 Apr 2025
Conference
Lisbon, Portugal

EDCC 2025: 20th European Dependable Computing Conference

edcc2025.campus.ciencias.ulisboa.pt/

1 May 2025
SCSC Seminar
London, UK

Safe Agile Developments

scsc.uk/e1154

10-13 June 2025
Conference
Paris, France

AEiC 2025: 29th Ada-Europe Int. Conf. on Reliable Software Technologies

scsc.uk/e1114

15-19 June 2025
Conference
Stavanger, Norway

ESREL SRA-E 2025: 35th European Safety and Reliability Conf.

esrel2025.com

19 June 2025
SCSC Seminar
London, UK

How Safety Culture has to Change With AI

scsc.uk/e1156

23-26 June 2025
Conference
Naples, Italy

DSN 2025: 55th IEEE/IFIP Int. Conf. on Dependable Systems and Networks

dsn2025.github.io/

4-6 Aug 2025
Conference
Chania, Crete, Greece

IEEE Int. Conf. on Cyber Security and Resilience

www.ieee-csr.org

9-12 Sept 2025
Conference
Stockholm, Sweden

SafeComp 2025: 44th Int. Conf. on Computer Safety, Reliability and Security

safecomp2025.se


thescsc.org/membership
Supporting Development of Safer Systems
The
Systems Safety
SCSC
Community
and Club
Safety Critical Systems Club