SSS'23 Event Report



In February 2023, the Safety-Critical Systems Club's annual Symposium returned to a fully in-person event hosted at the Principal Hotel in York. With over 100 attendees, the event enjoyed the type of attendance not seen since pre-Covid years and hosted over three days, the event once again, provided an exceptional and diverse range of talks. Paul Hampton, SCSC Newsletter Editor, provides highlights of the event.

While Covid has not entirely disappeared, there was a definite feel this year that we had returned to the type of event we'd enjoyed in the past, with the high attendance giving the event a real buzz and with lots of new features such as "Room 101", the crossword and roving video interviews, the event was one of the best we've seen in my memory. The talks were also of the highest quality covering many of the pressing challenges we have in the industry today. The keynote speeches are summarised below.

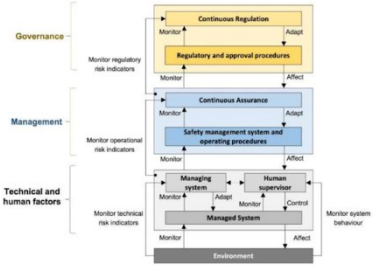
Safe, Ethical and Sustainable: Framing the Argument

Professor John McDermid, Director of Assuring Autonomy International Programme at the University of York, opened the conference with a talk developing some of the ideas he'd presented at the previous year's symposium where he introduced the mantra of "Safe, Ethical and Sustainable", and aimed to link three strands of thought:

- Ideas on how to govern safety of complex systems and their emergent properties
- How to develop an ethics assurance argument drawing on tenets of biomedical ethics
- The mantra – that broadens safety considerations to include society and environmental impacts

Governing Safety

John first covered the governance of safety and introduced a model showing a hierarchy of three layers. Technical and human factors at the bottom covering day to day safety management, a management layer looking at the owning organisation & supply chain and then the governance at the regulator level. There are feedback loop aspects with continual iterations but John stressed the process needs to operate a lot faster to manage the complexity and emergent technology.



Safe, Ethical and Sustainable

While the ambition of the mantra is clear, John said that there are challenges in how one actually frames an argument balancing these three quite diverse areas, as there will be trade-offs required between them.

John first covered ethics with an illustration of the type of assurance case drawn from biomedical disciplines. In this argument, along with standard non-maleficence concerns, beneficence and human autonomy also need to be balanced to give an overall "Justice Assurance" argument.

John then covered the sustainability aspect by describing associated models such as the UN Sustainable Development Goals (SDG), which describes 17 "goals" although he noted they are really a mix of different things. He also introduced the concept of "Doughnut Economics" that show a middle area that is sustainable, with poverty in the middle and overconsumption on the outside. There are other versions that enumerate specific concerns. These models help identify and inform the range of concerns that have to be considered – there will be trade-offs, but the aim is to remain in the green band.





John also mentioned a worrying aside where he noted that although electric vehicles are good for the environment at point of use, they have a bigger carbon impact than conventional vehicles in their construction and require the use of metal and rare earth metals that have limited supply.

An illustration

John illustrated framing such an argument through an example where a local authority was considering replacing a fleet of 45 city buses with autonomous vehicles. John showed that the choice of vehicle had different impacts on the environment and utility to different types

of travellers. For example, large buses can carry more people but have poorer safety performance against smaller vehicles because of their mass, shape and risk to vulnerable road users. From a different perspective, smaller vehicles are good and efficient for lone travellers but not good for larger families possibly with additional equipment such as buggies for small children.

John concluded that:

- The work is illustrative and not put in practice yet, but other work such as Rawls reflective equilibrium could be used to help make the justice argument
- There are many ethical principles. The biomedical principles seem to be appropriate as it spans most of the concerns but needs interpretation for a particular context
- For sustainability models, the SDGs are quite different and have surprising omissions such as "human autonomy"
- John said that no one can be expert in all concerns and multi-disciplinary teams are essential. The role of the safety engineer is changing to be more of a facilitator as well as an expert in hazard and risk analysis

Rising to the Challenge of Certifying Automated Vehicles

Jamie McFadden presented the second keynote talk of the day. Jamie is from the Vehicle Certification Agency (VCA), which is responsible for the oversight of vehicle type certification for use on UK roads.

Jamie said that the fullscale deployment of autonomous vehicles is a challenge, but the first steps have already been taken with, for example, the introduction of the Mercedes "Drive Pilot" in 2021 of automatic lane keeping on certain sections of autobahn in 13,000km of pre-approved German motorway.

Jamie said that the type approval process comprises:

- Around 70 tests covering emissions, crash, lighting and tyres
- Conformity of production: can the vehicle manufacturer maintain the type certification approval in full production

Jamie talked through some of the key safety and security aspects and noted that functional safety is part of the type certification approval process.

Two case studies were described that galvanised the views of regulators on the effectiveness of type approval:

- In Sep 2015, Volkswagen were found to have installed software on vehicles that could detect that they were being tested for emissions and allow non-compliant vehicles to pass the test. This impacted 11m vehicles worldwide, and was known as "Dieselgate".

This called into question the regulatory approval practices and resulted in revisions to the regulatory framework brought into force in 2018.

- Researchers exploited a security vulnerability in an internet connected infotainment system to remotely brake and steer a moving Jeep vehicle. 1.4m vehicles were affected and this is known as the "Jeephack" and was a wakeup call for the industry to the risks of connection to vehicles.

This has resulted in improvements in industry working practices and standards and increased security of software and the need to keep software up-to-date.

In 2015, the need for a new set of technical regulations was identified by the UN economic community for Europe and culminated in three new regulations in 2021: UNEC 155 (cyber security), UNEC 156 (software updating) and UNEC 157 (automated lane keeping systems).

Jamie then summarised the regulatory requirements:

- The system shall be free from unreasonable risk for the occupants and other road users

Jamie broke this down into:

- Monitor the design domain environment and driver
- Execute longitudinal and lateral control
- Monitor for failures and degraded operational conditions

This shows the more open and objective-based approach to achieve approval rather than the historical test-based approach.

- The system shall not cause any collisions that are reasonably foreseeable and preventable

Guidance is provided in the regulation to cover the types of issue to be addressed, as a starting point, such as road users and objects blocking the lane, vehicles cutting in and the reveal of objects when vehicle peel away.

- Activated systems shall comply with traffic rules including responding to emergency and enforcement vehicles

There are no explicit list of scenarios to comply against. Regulators want to see a robust assessment process that manufacturers have performed to understand and characterise the traffic rules the system needs to be developed against, the implementation of that system in the design and how these are maintained as rules change.

These might need to consider aspects such as Smart Motorways, a wide range of enforcement vehicles (not just police) and filtering motorcycles etc.

There is a need for collaboration between organisations as no one organisation can know all traffic rules for all markets of Europe.

- The manufacturer shall demonstrate their Cyber Security Management System across the vehicle's lifecycle

Approval also needs cyber security approval and the need for a Cyber Security Management System for Development, Production and Post-production phases. Focus not only on the security of the vehicle but how the manufacturer has gone about securing the vehicle over the lifecycle of the vehicle.

- The manufacturer shall identify critical elements and shall identify risk appropriately

This involves the assessments of the manufacturer's risk assessment process and evidence.

The Automated and Electric Vehicle Act of 2018 that states that the Secretary of State must prepare a list of vehicles that can safely drive themselves. Critical to this will be the monitoring and control test:

Can the system safely control the vehicle.

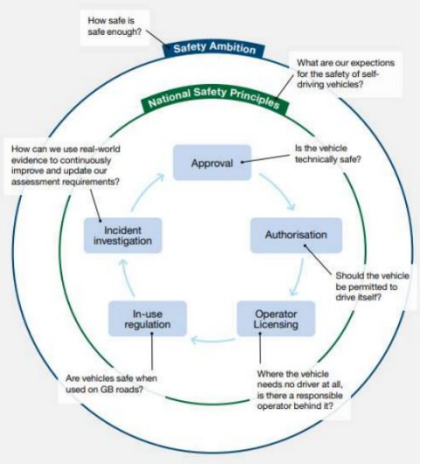
Can it be shown the vehicle can act without a human monitoring it:

- comply with traffic rules
- avoid collisions that a human can, treat others with reasonable consideration
- avoid putting itself in a position where it can cause a collision
- recognise when it is leaving its operational design domain

Jamie concluded by saying there are two significant challenges as they are not truly addressed by the regulation.

- Safety assurance of machine learnt systems bias, interpretability of results and robustness of decision making
- Appropriate use of modelling and simulation

The UK government published its "Connected and Automated Mobility 2025" document and expects that self-driving vehicles should be held to equivalent standard of behaviour as a human driver: competent and careful. There is currently a consultation that will result in a set of national safety principles and a safety framework.



Jamie said that regulators are aiming to address these challenges through:

- Objective based requirements for system safety, machine learning and modelling and simulation
- A foundation of management systems to establish confidence in the system development, production and deployment
- Application of risk assessment to appropriately treat risk
- A new legal framework as an enabler

What the First 100 Years of Railways Can Teach Us …

Tom Anderson presented the last keynote speech of the day with possibly the longest titled talk ever! "*What the First 100 Years of Railways Can Teach Us about the first 10 years of self-driving road vehicles".* Tom discussed what lessons can be learned from previous railway accidents and how these can be read across into the emerging autonomous automotive vehicle industry. Tom referred to various books such as "Red for Danger" by LTC Rolt, which provides a history of British railways disasters and from these, 15 accidents were selected to:

- Understand why the accidents occurred, what the consequences were and what lessons could be learned for today's systems
- Emphasise particular key issues for autonomous automotive vehicles.

Tom discussed 3 accidents:

Parkside (1830) This covered one of the first ever fatalities on a railway when a number of VIPs were being transported during the opening of a line for the Liverpool & Manchester Railway. When the train stopped to take on water, many passengers disembarked and walked onto the track of an oncoming train. One individual, William Huskisson, an MP and former member of the Cabinet, failed to clear the track in time and was hit and mortally wounded by the train. Tom said that one of the key lessons learnt from this accident was that people don't expect new hazards from new technologies, some may ignore safety practices and reduced mobility of a system user can cause additional issues.

Winchburgh (1862) In this accident, maintenance work required single track operations and a train was being used as a pilot engine acting as a token. This pilot engine would normally be a smaller distinct vehicle but had been swapped with a larger engine as the smaller pilot engine had other duties at the time. A ballast train was then mistaken for the pilot engine and was let through to collide with another train. Lessons from this accident where:

- Safety procedures must be understood and followed correctly
- Maintenance situations offer different environments (eg. single track working) that need equal consideration to the norm. Autonomous vehicles will also need to adapt to maintenance situations.

Quintinshill (1915) The Quintinshill signal box controlled two passing loops, one on each side of the double-track. At the time of the accident, both passing loops were occupied with goods trains, and a northbound local passenger train was standing on the southbound main line.

Neglect of the rules by two signalmen, such as failure to apply locking collars, led to the presence of the local passenger train that had been reversed onto the southbound line was overlooked by the signalmen. A southbound troop train was cleared for passage and collided with the stationary train. A minute later the wreckage was struck by a northbound sleeping car express train travelling north. Over 200 people were killed in the accident.

Tom said, lessons from this accident were:

- Accidents arise from multiple causes and failure to adhere to safety practices
- If an AV sees an accident, it could report this to reduce the risk of further accidents

Tom concluded by saying that there were 75 individual items that could be read across into the autonomous automotive industry and recommends that an independent body is required to investigate AV accidents.

Room 101

Roger Rivett hosted a new event called 'Room 101' for System Safety: What would you get rid of? This mirrored the popular UK TV show of the same name where individuals are invited to discuss their pet hates and provides an argument for why they should be 'put in the bin'. The audience then votes on whether the particular gripe should indeed be disposed of. This was an entertaining session with a number of items agree by consensus to be consigned to the dust bin!

Working Group Meetups

Day two started with an early morning get together of working group members. This included a demonstration in the main hall of the "RADISH" tool for helping to manage data safety risks. More details of the RADISH tool are given in the Feb 2023 edition of the newsletter scsc.uk/scsc-182

Special Events

The morning then continued with a number of special events:

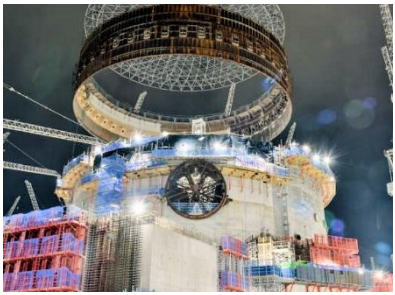
- Visit to the Institute of Safe Autonomy
- Guided City Walk
- York Minster Tour
- York Railway Museum

These were hugely popular with an additional trip to the Railway Museum being organised to accommodate everyone.



How Nuclear New Builds Incorporate Lessons Learnt

The second day's keynote was given by Tom Hughes from EDF discussing the Hinkley C new nuclear power station that is being built in Somerset. The focus of the talk was how the build was incorporating lessons learnt from other accidents and incidents in the nuclear industry.



Tom first described the Hinkley Point Power Station that will produce 3260MW of electrical output from the world's largest turbine generators and account for 7% of the UK's electricity.

Tom described the main characteristics of nuclear plants – the "3 C's"

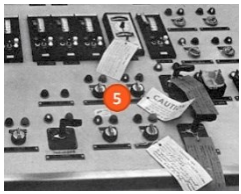
- Control: the rate at which stored energy is released needs to be controlled
- Cooling: even when the reaction is shutdown the fuel continues to produce decay heat so needs to be cooled
- Containment: by-products are radioactive so need to be contained to protect personnel, the public and the environment

Tom gave three specific examples to illustrate the application of lessons learnt:

Three Mile Island (TMI)

On 28th March 1979, a cooling malfunction caused part of the TMI-2 reactor core to melt – the reactor was destroyed.

Tom explained how a mixture of operator and latent errors in the system conspired to cause the accident. The lessons learnt included:



- **Being intolerant to defects & poor quality**. Issues that seem inconsequential now could later compromise layers of the 'defence-in-depth' model
- **Performing work to the highest standards** and in accordance with the approved designs / instructions / procedures – whilst maintaining a questioning attitude
- **Embracing learning** (internal and external) – review relevant Operational experience and close actions meaningfully, ensure training programmes equip personnel for all situations (combination of skills and knowledge)

Chernobyl (1986)

The accident occurred as a result of a failed safety test. Operators mistakenly (and somewhat unknowingly) placed the reactor into a highly unstable state. However, the reactor design was fundamentally flawed. Several factors led to a huge power surge that destroyed the reactor core and released a radioactive plume. Causes were both human and technical in nature.



Tom noted that there are significant design differences in cooling and containment between the Chernobyl RBMK system and Hinkley Point C (HPC), making HPC much safer. The key lessons learnt were:

- **Culture**: Whilst the HPC environment is very different, cultural failings that played a huge part at Chernobyl (Time Pressures, Cost Pressures, Management Blindness and (lack of) Questioning Attitude) could affect any plant, anywhere

Fukushima Dai-Ichi (2011)

In 2011, the Fukushima Dai-Ichi (a 6-Unit Boiling Water Reactor (BWR) power plant) was struck by a magnitude 9.0 earthquake (larger than the plant was designed for). Units shut down automatically (as expected) and shutdown cooling systems started. However, the subsequent tsunami overwhelmed sea defences and flooded numerous sensitive plant areas. Electrical supplies were completely lost, resulting in loss of cooling to nuclear fuel and core damage.



Although the UK EPR system already had design improvements from the older-generation reactors at Fukushima, following the accident, a number of additional modifications were made to the UK EPR design including:

- **Seismically protected water storage building** with additional diesel driven pump capability
- **Increased runtime** of all diesel generators and severe accident batteries
- **New plant configurations** that ensure provision of cooling water to the reactor and spent fuel pool at all times.

Tom concluded with some key points:

- It is nearly always better to 'engineer out' the hazard or error potential than rely on individuals to follow complex safe systems of work perfectly
- People must be trained to work to the highest standards and in accordance with the approved designs / instructions / procedures – whilst maintaining a questioning attitude
- Accepting poor conditions or poor behaviours makes these the 'norm'
- Operational Experience is of huge value – review relevant experience, conduct investigations, set good actions and close them properly
- Make safety & quality personal – never work on the assumption that another 'layer' of the defence-in-depth model will catch an error. What seemed trivial today could cause a problem tomorrow

Poster Session



The days talks were followed by a poster session in the main exhibition area where members of the SCSC Working Groups and others presented posters summarising their work on large presentation boards.

See page 33 for more details of the poster sessions and for copies of the posters as presented.

Banquet

The day concluded with the symposium's banquet and Les Hatton from Oakwood Computing gave an amusing after dinner speech.



Making the Water Visible: A methodology for exploring Systemic Change

A fish is swimming along one day when another fish comes along and says: "Hey, how's the water?" The first fish stares blankly at the second fish and then says, "What's water?"

Kania J, Kramer M, Senge P, 2018, The Waters of Systems Change, FSG

Gill Kernick, Transformation Director at ARUP University, gave a moving account of how she witnessed first-hand the Grenfell Tower fire of 2017. At the time of the fire, she was living in an adjacent block and having previously lived on the 21st in Grenfell Tower, she knew several people, some children, who died that night. In response to the fire, Gill was committed to ensuring an event of this nature never happened again. She thought the desire to improve would be shared, but found this to be far from the case, and this led her to ask why this failure to learn makes sense.

The Grenfell Tower Inquiry concluded that:

"... each and every one of the risks which eventuated at Grenfell Tower, on that night, were well known by many and ought to have been known by all who had any part to play. As a result, you will be able to conclude with confidence that each and every one of the deaths that occurred in Grenfell Tower on 14 June 2017 was avoidable."

Gill said the heart of the matter was a failure to understand the nature of complexity and a piecemeal versus systemic approach to change.

Complexity: many of the traditional top-down bureaucratic and mechanistic ways of thinking are grounded in mythical cause-and-effective narratives and are becoming redundant – and yet, responses to disasters largely come from this paradigm.

Piecemeal v Systemic Change: Gill provided a comparison between these two different approaches to change:

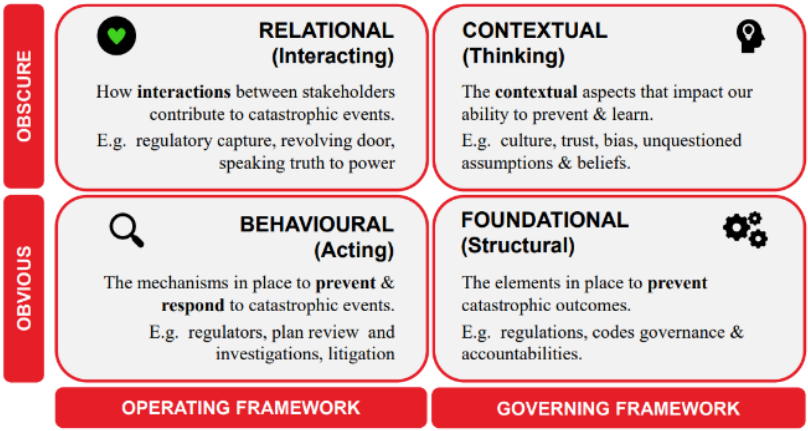
	Piecemeal Change	Systemic Change
Intent	Solving a piecemeal issue	Shifting the conditions holding the status quo in place
Question	What's wrong with a discrete part of the system?	What is the system perfectly designed for?
Assumption	Controllable, predictable world	Complex, emergent world
Access to change	Fix what is wrong	Make the Water Visible: grapple with the messy kaleidoscope
Approach to change	Technical Solutions *(if I do x, y will happen)*	Disrupting the status quo, experimenting *(if I do y, what will happen?)*
Leadership Style	Bureaucratic, command and control, rules based	Organic, emergent, values and principles based
Requires	Traditional Expertise	All Stakeholders - tacit expertise *(validating different ways of knowing)*

Gill used the analogy of the conversation between the two fish presented earlier to demonstrate the desire to "make the water visible" and set out the following principles:

- Traditional bureaucratic, linear, cause and effect ways of thinking and leading are ineffective
- Piecemeal solutions don't cause systemic change and can have unintended consequences
- The system is functioning perfectly to produce what it is producing
- Systemic change requires disrupting the status quo

Gill then set out an approach and framework for exploring systemic change. The approach is grounded in "Enquiry" as distinct from a more formal Inquiry, inhabiting the space of the not knowing, and "Sense-making", distinct from analysis, interpretation or exploration and not demanding neutrality. The key enquiries being:

- Why does our failure to learn make sense?
- What will it take to enable systemic change?



Gill presented the supporting framework against which the following questions can be posed:

- What are the known issues?
- What myths are contributing to our inability to learn?
- What are the conditions holding our inability to learn in place?
- How might we enable systemic change?

Gill concluded by discussing the transformative power of grief and the hope offered through the democratising of change where ordinary individuals are more empowered to effect change through purposefulness and connectedness with one another.

The Language of Risk and the Risks of Language

A pre-recorded talk from Catherine Menon from the University of Hertfordshire was then shown. Catherine said she works with assistive robotics and ethically-informed safety cases, but is also a published author. Catherine showed the different perspectives of risk language drawing out the differences between the language used in formal safety argumentation and the more common layperson's perception of the language. For example, a lay person has a very different understanding of what an "argument" is from a safety engineer.

She said we need to understand how our use of language will be perceived to our readership and the dangers of using metaphors.

She gave an example of the Byzantine General problem where a number of generals need to coordinate a decision to attack a city or retreat with the knowledge that one general may be a traitor and deliberately give false instructions.

Catherine discussed the results of a study where members of the public (a group of actors) where asked what was important about the Byzantine General problem. For safety engineers, "consensus" was important but for the actors it was that the "generals want to attack" showing quite different perceptions of the problem.

Catherine said while metaphor is the main mechanism through which we comprehend abstract concepts and perform abstract reasoning, our language choice can govern our thinking. For example, using the expression "**facing** significant risk", presents risk as an adversary. Safety arguments can also be presented in different ways depending on the language – it can be adversarial or more a form of partnership with the reader.

Catherine concluded with the idea of skilful storytelling with examples of how even a six-word story can have connotations:

"For sale: baby shoes. Never worn."

She warned that it is possible to be too convincing and that word choice, argument structure and explanation of context are important and asked how we assess the risk associated with ourselves.

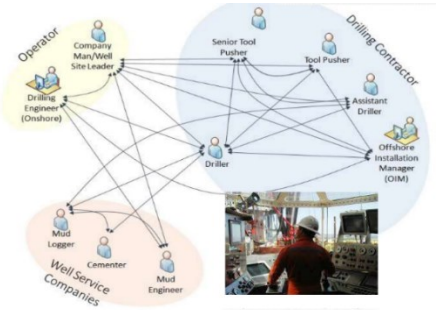
Resilience in Safety Critical Systems – Offshore FRAM

David Slater from Cardiff University starting off his talk by discussing what we mean by "safety" and argued that nothing is absolutely "safe" as freedom from accidents cannot be guaranteed. He suggested that (as proposed by Erik Hollnagel) we should stop using the word "safety" and start using the word "safely".

Accepting that a system cannot be entirely free from risk, David said that the natural response is defensive – to put up "barriers" in depth so that a number of natural or designed defences need to fail before there is an accident. He noted however, that these can be subject to common cause failures with barriers failing as a group or one failure triggering the pulling down of others like dominoes.



David then focussed on the Macondo Blowout (aka Deepwater Horizon explosion) of April 2010, where an offshore drilling rig experienced a blowout from a drilling well that led to a fire, many fatalities and loss of the platform. He first described the sequence of events that led to the accident and showed that all the barriers somehow failed on demand. He said that traditional forms of analysis such as Bow Tie that assume sequential and entirely independent failures can give quantitative failure probabilities on demand that are extremely low suggesting an accident is all but impossible, but an accident occurred nevertheless.



David explored the reasons for this by making the distinction between "Work as Imagined" and "Work as Done". David noted that the operators had all the information they needed to tell them something was wrong but they rationalised it and assumed they could "normally" work around abnormal indications.

David modelled the "Work as Done" that illustrated the complexity of the interactions between the various organisations such as the Operator, Drilling Contractor and Well Service Companies.

He observed that:

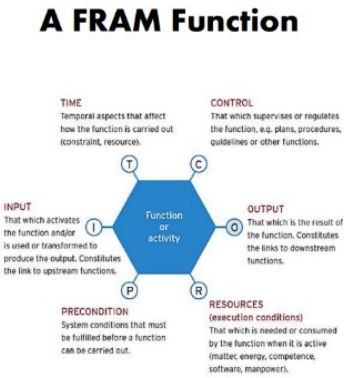
- The interaction complexities would be difficult to model using Fault Trees
- The ownership of particular barriers is unclear
- Is the Driller really in charge/control of all the Key Variables?

And he asked how all the critical interdependencies can be shown/traced.

David said we need to move forward to establish a new paradigm in thinking from "Determinism" to coping with uncertainty and to:

- Operate in the real world, not an idealised, standardised environment
- Move away from expecting linear, predetermined, prejudiced, sequential, "as designed" outputs and behaviours
- Live with Nonlinear, interdependent, unexpected, as necessary, emergent outputs and behaviours
- Establish better models to support the safety engineer aspects

David then looked at what modelling tools might be more suitable to support this new way of thinking. He referred to Leveson's System Theoretic Accident Model (STAMP) and said it was a conventional, (hard wired) "control" diagram of the system. He then detailed Hollnagel's proposal that envisages a system as a collection of interacting, interdependent "Functions" (The Functional Resonance Analysis Method - FRAM) and gave an example of the model from a medical Trauma team.



He said the FRAM method allows you to build a model of an actual complex sociotechnical system and then showed a FRAM model of the Macondo case showing the instantiation for the procedure being operated.

He noted that:

- System behaviours "emerge", they are not predetermined
- The supply of services is a crucial aspect of almost every "safety" function

David's assessment of the case study listed several proposed contributing factors:

- There was no single, simple "Root Cause", component failure
- The system was too complex (pushing the boundaries: record water depth, etc.)
- BP was not directly in control due to a culture of contracting out risks
- The "services", equipment suppliers and maintenance activities that the drillers relied on, were "variable" in quality
- Perhaps complacency assured the drillers that normally, all these minor variabilities could be worked around

He said that when all these normal variabilities "resonated", the system became unmanageable.

David finally introduced the idea of "resilient" safety-critical systems where behaviours are more predictive rather than reactive, that is, not so much monitoring and responding, but learning and anticipating. He said that to be more responsibly resilient we need to develop proactive design approaches and tools and offered FRAM as a possible solution.

Crossword Competition

The SCSC Safety System crossword was first published in the Oct 2022 edition of the newsletter and was also given to delegates to complete during the symposium. There were 5 correct entries, and the winner was drawn from the 'hat' by arguably the SCSC's youngest member – Charisse Lin (aged 7), who also produced the remarkable piece of artwork on the back of this edition of the newsletter!



The answers to the crossword are given below.

