

The Risks Of LSCITS: The Odds Are Stacked Against Us

John McDermid¹ OBE FREng

¹ Department of Computer Science, University of York,
Deramore Lane, York, YO10 5GH

Abstract. Complex IT Systems are often used in applications which can pose a risk to their owners or to the public. Many of these are subject to extensive risk assessment before they are deployed and operated yet, despite this, undesired events do arise, leading to financial loss or loss of life. This paper investigates the role of existing risk assessment methods and draws the conclusion that they do not effectively predict the causes of actual loss events. The paper then suggests an alternative approach, which has the potential to offer a unified approach to risk assessment across a number of domains, and across different system properties, e.g. safety and financial risk. It concludes with observations on similar methods and research results, especially from accident analysis, and makes suggestions for future research directions.

Keywords: Large Scale Complex IT Systems, Risks, Safety, Security.

1 Introduction

Many Large-Scale Complex IT Systems (LSCITS) are used in roles where organisations depend on them for a key aspect of their business. As a consequence, these systems may be safety, security or mission (business) critical. It is common to assess such systems in terms of the risks they pose – whether to the organisation that owns them or to third parties – but different approaches to risk analysis are used in different domains. This paper analyses some “loss events” associated with a range of LSCITS (and one comparatively simple system) then uses the “signatures” of these events both to question current approaches to risk analysis and as a source of ideas and inspiration for an alternative model.

Serious failures of the more “critical” LSCITS are relatively rare, and that might suggest that “all is well” in terms of our ability to design and assess such systems. However a cursory assessment of a range of “loss events”, e.g. accidents or financial losses, suggests that the current approaches to risk assessment do not throw much light on the actual causes of the loss events. The paper considers a range of “loss events” which illustrate safety, financial and availability issues. It shows that the risk assessment methods used, explicitly or implicitly, in these different domains do not provide a good basis for gaining an understanding of these events.

This analysis also shows that the events studied have remarkably similar “signatures” in the sense that the “confluence of events” which leads to the loss are very similar in nature, although they are in different technologies, systems and domains. As LSCITS are increasingly depended on for multiple critical properties,

e.g. safety and security, the availability of a “unified” approach to risk assessment has the potential to be valuable in the design and assessment of future generations of LSCITS.

It should be noted that these initial findings are tentative, and need further and more rigorous study. Observations on methodological issues are presented in the next section, followed by a brief, textual, analysis of five loss events. The “signature” of each loss event is discussed, and some observations are made which are intended to help in developing a methodological approach to assessing LSCITS risk. This is followed by a discussion of risk analysis. This discussion first outlines the risk analysis processes typically used in several different domains, and then discusses how well these processes reflect the signatures of the loss events described previously. Next, the paper outlines an alternative perspective on risk analysis, believed to be suitable for assessing LSCITS. The paper then considers related research, before drawing conclusions and proposing directions for future work.

2 Methodological Remarks

It is not possible to do an exhaustive, scientific, analysis of LSCITS and their risks. In some domains the allowable failure/loss rates are so low that the expectation would be that there would no failures in the operational life of the system – and that is many decades. Further, there are far too many systems, and the numbers deployed are increasing at a rate that defies analysis.

Thus our approach reflects an approach developed in social and management sciences, e.g. by Van der Ven [1], which provides a framework for observing, modelling and (ultimately) intervening in real-world applications. This approach is outlined in Fig. 1 overleaf.

The framework can be “entered” anywhere but, for our purposes, it is easiest to think of it starting with observations of reality (the bottom of Fig. 1) to produce a problem formulation. From the problem formulation and the observations, it is then possible to develop theories and models which provide explanations of the observed phenomena (which are better than current theories in this regard). The framework then proceeds to build a “research design” enabling an intervention – changing reality – which can then be re-assessed to seek to confirm or refine the theory and model. At this stage in our work, we are firmly in the (early) stages of theory and model building.

Van der Ven also uses the nature of the research, and the degree of engagement, to refine his research framework; this is outlined in Fig. 2 overleaf. In terms of this model, we are working in the “describe/explain” part of the framework, and within that mainly in the “detached/outside” research perspective (although one of our examples below is in the “attached/inside” quadrant, as it is a personal experience).

Thus, at this stage, the criteria for assessing the theory and model are relevance and validity; we would also say that they should give greater explanatory power than current models of risk.

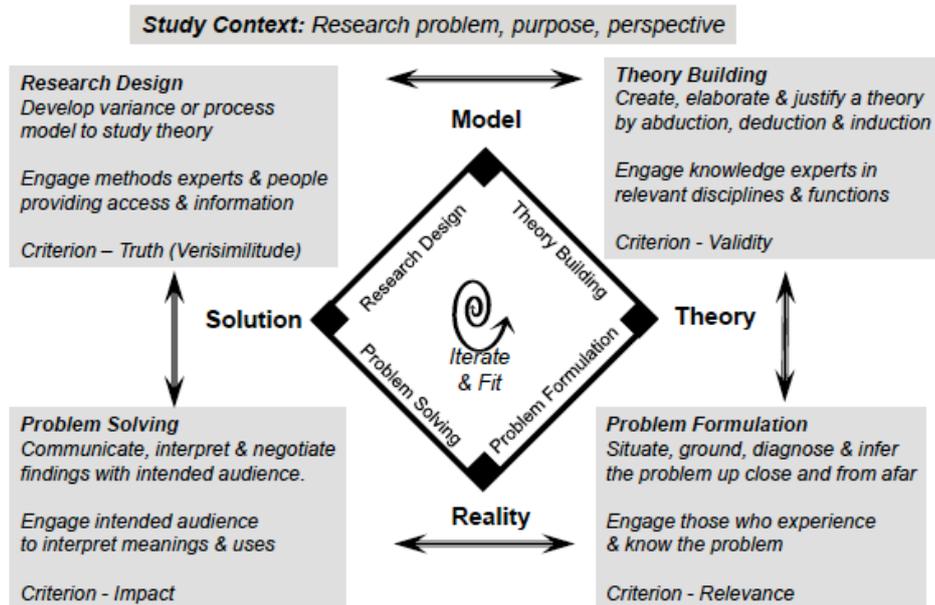


Fig. 1. Model Of Engaged Scholarship (from Van der Ven [1])

Research Question/Purpose

		To Describe/Explain	To Design/Intervene
Research Perspective	Detached Outside	Basic Science With Stakeholder Advice 1	Policy/Design Science Evaluation Research For Professional Practice 3
	Attached Inside	2 Co-Produce Knowledge With Collaborators	4 Action/Intervention Research For a Client

Fig. 2. Types Of Research (from Van der Ven [1])

This work is undertaken as part of the LSCITS programme [2]. The original LSCITS proposal divided the space of concern for LSCITS into four layers in a “stack”, viz:

- PSS – Predictable Software Systems – the development and application of the most advanced scientific principles to large-scale computing problems;
- HISE – High Integrity Systems Engineering – rigorous approaches to dealing with the design and assessment of systems beyond the reach of the PSS methods, including Systems of Systems (SoS);
- STSE – Socio-Technical Systems Engineering – the analysis of systems and their failures where the causes of the difficulties arise in the interaction between technology and users, both individuals and organisations;
- CiO – Complexity in Organisations – focusing on the problems of large-scale organisations, and how they influence system success.

Also, the LSCITS programme includes orthogonal work on non-standard computational approaches to complex problems, and work on the cloud.

To help understand the risks of LSCITS, it would be possible to classify the example loss events in terms of the LSCITS stack. However, it is sometimes hard to make distinctions between the four different concerns, and as several of the loss events studied have complex causal factors, it has been decided instead classify the events on a “scale”, viz:

- Pure technical – there is a clear technical cause of the loss event, and the interaction with, and behaviour of, the organisation is much as intended;
- Socio-technical – the causes of the event include erroneous interaction between the system and users, and may also include individual human errors or technical failures;
- Pure organisational – there is a clear organisational cause of the event, e.g. failure to implement separation of duty, and the system behaved according to intent (and requests from users).

This is intended to be a “sliding scale” not a hard categorisation, and it is used to “locate” the primary causal factors in each loss event on the technical-organisational axis.

In analysing loss events there is always a risk of hindsight bias – looking for evidence to prove the author’s hypothesis. In part we have sought to address this risk through consideration of events which span the range from highly technical causes to those whose origins are largely organisational. Further we are seeking to build a theory and model, not to prove one. However there is always a risk of such biases, and we return to this concern in the discussion.

3 An Analysis of Some Loss Events

In order to shed light on the issues of risk assessment we consider five “loss events”. Four are documented in the literature, to varying degrees; one is a personal experience. Many more examples could be chosen, but the rationale here has been to choose events which span the range from technical to organisational, and cover

safety, financial risk and availability (integrity) of private data. Of course there would always be benefit in considering more loss events; see the conclusions for a discussion of future work.

It would be possible to analyse the events using methods such as Why-Because Analysis (WBA) or Why-Because Graphs (WBG) [3]. We have chosen not to do so here, partly for brevity, partly so that we can emphasise what we perceive to be key points, and partly as there is not a body of work to draw on showing how to apply the techniques outside the safety domain. However if this work is to be taken further, then it will be necessary for the analysis of these individual events to be put on a more rigorous footing (although this might require extensions to techniques such as WBA).

In each case, the description contains a brief overview of the event, primarily to set context. This is followed by a descriptive analysis of the event(s) leading to the loss, and ends with an assessment of the “signature” of the event. Following the discussion of these five loss events, some observations are presented, followed by the problem formulation, as suggested by Van der Ven’s framework [1].

3.1 A Syringe Pump

This example relates to a syringe pump, that is a medical device which delivers liquid drugs on hospital wards, or anaesthetics in operating theatres. A desired delivery rate is set by a nurse, or an anaesthetist, then an electric motor drives the syringe plunger to deliver the fluid at the set rate. The flow rates set may be very low.

For safety, the initial design had diverse mechanisms for measuring movement of the plunger: a linear interference grating directly measuring plunger movement and a quadrature system measuring the rotation of the motor shaft. In initial use, the device suffered a lot of “spurious trips” where the pump stopped because the linear grating system detected inappropriate movement. Investigation showed that this was spurious, and was due to backlash in a gearbox, which was significant given the low rates at which the pump was meant to operate.

It was decided that the device should be modified so that the linear grating was not enabled until the quadrature system had confirmed that the syringe plunger was moving. Two patients were killed when the plunger emptied the syringe at high speed (it emptied a 250ml syringe in a matter of seconds). A more detailed analysis of the control loop is presented in [4], but the reference does not describe the context of use which was deemed to be sensitive at the time.

Note that this is a simple example, and can not be considered to be an LSCITS, but it is included as it is possible to give quite a detailed technical exposition of what happened, and it also shows the problems of change – which are a causal factor in many accidents and incidents.

Description of the Event. The quadrature system had a reference square wave signal (Fig. 3 a). A sensor on the motor shaft generated a square wave signal; if the rising edge from the sensor was before the rising edge on the reference signal that indicated movement in one direction (Fig. 3 b); movement in the opposite direction was indicated by the edges being in the other order (Fig. 3 c).

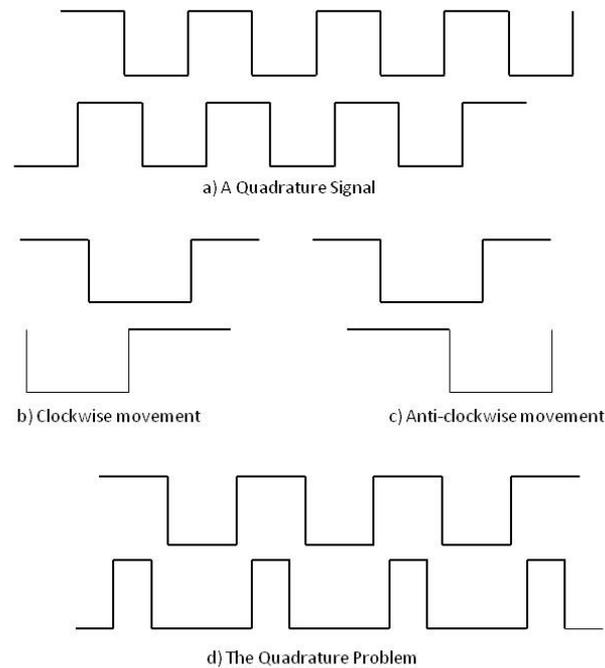


Fig. 3. Critical Waveforms (adapted from [4])

The signal coming from the sensor was not “clean” and was “squared up” by means of a Schmitt trigger. Schmitt triggers vary in performance; most produce a signal which is close to square (with a mark-space ratio of 1), but some produce signals with a much higher mark-space ratio. For some devices at the limit of the manufacturing variability, the resultant signal was very distorted and the complete positive pulse was within the positive part of the reference signal (Fig. 3 d).

The logic in the software was designed on the basis that the complete overlap of signals shown at Fig. 3 d) was impossible; unfortunately rather than flagging this as an error and stopping the motor, the software looped, kept power on the motor, read the next set of inputs, and kept going waiting for a “valid” input showing that the motor was moving. Following the change, the motor control software did not enable the linear grating until it had a “valid” input, so the protection system did not stop the motor either.

Signature: The key causal factors in the loss event are:

- Intrinsic flaw (software “bug”) exposed by the circumstances;
- The protection system was disabled due to a single point failure (the Schmitt trigger at the limit of tolerances); thus it was a common cause failure;
- Opportunities for further protection, i.e. detecting invalid inputs and/or stopping the device after a time had elapsed without detecting a “valid” input, were missed.

This is viewed as largely a technical loss event as there was a “bug” in the software (incomplete coverage of possible inputs) and an electronic device which was at the limit of manufacturing variability. However there was a socio-technical element – the change to improve availability – however this change only uncovered a basic design flaw, rather than contributing directly to the accident.

3.2 The Cloud

This example is a personal experience which is largely an integrity problem, but also shows problems of loss of service availability. Cloud systems can be viewed as LSCITS in themselves – although the particular application which caused the problem is quite simple (it certainly wouldn’t be viewed as complex by today’s standards).

Description of the Event. I was one of the early adopters of cheaply available cloud services, and decided to migrate my diary and contacts database to the cloud so that they could be updated by certain colleagues, and viewed by many more. For a few months this worked well, and the cloud diary became the master copy (my “back up” was no longer updated). Then I started to have problems in updating the diary (I could make changes, but they were only transient, and they did not update the stored diary).

After a period of time on email and interactive chat with the cloud service provider it was concluded that the data had been corrupted, and that it was not practicable to “fix” the problem (they had no tools which could repair the data). The support team agreed that they would “package up” my diary and email it back to me (so I could import it into another diary/calendar tool). Again, after an extended exchange with the cloud service provider it became clear that this “resolution” wouldn’t work either – the corruption which prevented update also prevented an export being produced!

Fortunately, the diary could still be displayed and printed, thus it was possible to print out my diary (about a year and a half ahead) and to type it all back in. In total the process took 2-3 weeks during which my diary was not up to date, and it cost a significant amount of time in discussion with the service provider and in retyping the diary. (Contacts had changed very little, so there was a minimal amount of effort needed to restore them.)

Signature: The key causal factors in the loss event are:

- Intrinsic flaws (software “bug”) exposed by the circumstances;
- A single point failure (the data corruption) disabled both the primary function (diary update) and the protection system (the ability to export data);
- Other protection systems, e.g. the replication of the data in the cloud, were rendered worthless as the data was corrupted not “lost”.

Overall this is largely a technical failure, and clearly a detailed software design (data dependency) issue. There was a socio-technical element – the point at which the service provider decided they would stop trying to solve the problem – but as the annual charge for the system was under \$100, this was understandable (they will have made a loss on my account given the amount of time they spent in helping me).

3.3 The “Flash Crash”

On May 6th 2010, there were a number of significant anomalies in the financial markets both in New York and in Chicago. Perhaps the most significant event was the drop in the Dow Jones Industrial Average by almost 10%. Unlike some accidents leading to loss of life, e.g. in the aerospace sector, the problem is not well-described in the literature, and the description here draws heavily on [5] and contemporaneous press reports, e.g. [6].

Current financial markets are operated through a mixture of highly autonomous algorithmic trading systems, referred to as algo traders, and human traders. As well as trading directly, the humans set parameters for the algo traders. The changes in the way markets work have been rapid and significant. In 2003 the human traders on the New York Stock Exchange (NYSE) handled about 80% of trading volume of stocks listed on the exchange. By the end of 2009, the proportion being traded “manually” had fallen to 25% with much of the trading moving to electronic-trading platforms, such as Direct Edge and BATS, which execute trades in milliseconds.

Further, markets are linked and some of the trading is done on “spot differences” between markets, e.g. between New York and Chicago. This sort of trading tends to be automatic (algo) as the computer systems have the speed (the millisecond trades) to capitalise on small differences in prices, by trading huge volumes of shares or other commodities.

Description of the Event. On May 6th 2010 the Dow Jones Industrial Average plunged by nearly 1,000 points, with most of the losses occurring between 2.40pm and 3.00pm, see Fig. 4. It was the largest single day decline in the market’s history. Some well-known stocks, such as Accenture, briefly traded for as little as a cent. The market later rebounded, to close down by 348 points, although it was “off” by 9.2%, and over \$800 billion, at worst.

There has been considerable speculation about the cause or trigger of the “Flash Crash”, and it seems that what happened was a combination of general nervousness (about the state of the Greek economy and the UK general election results) and some specific trading actions. Automated systems certainly played a very big role in the rapidity at which events occurred, but human traders also influenced the markets.

First, there is evidence to suggest that human traders were active and significant participants in the market during the big drop. Also, it seems that some human traders were experiencing serious delays in their data feeds caused by the huge volume of trades being executed, so they issued orders in good faith but on the basis of bad (stale) data, and that just made things worse.

Second, humans had “rigged” their algo trading systems to get around some regulations without actually breaking the law. The regulations require that traders always offer two prices: one to buy and one to sell shares. If the traders don’t want to take business, then they would offer to buy at \$0.01 (1 cent) and to sell at \$99,999 (the allowed limits). Whilst the human traders may not have used these prices directly, they were encoded in the algo traders, and these prices were used during the event.



Fig. 4. The Change in the Dow Jones Industrial Average during the “Flash Crash”

In setting these values, the trading houses didn't consider that in a big panic like the “Flash Crash”, that many of the traders would get out of the market, cancelling their existing (sensibly-priced) bids and offers, and so the extreme prices would then be left exposed as the best offer and bid prices in the market. At that point, other algo traders transacted at these prices because they had been programmed to automatically deal with the best bid or offer price, regardless of its absolute value (and whether or not it was sensible). Thus the “Flash Crash” is what can be viewed as an emergent property of a complex set of interacting “systems” (an SoS) – both human and automated.

Several companies such as Tradeworx, a hedge fund with a high-frequency trading business, shut off their systems. Manoj Narang, the CEO of Tradeworx said he did this when he “noticed the prices were erroneous”, because he knew exchanges would cancel those trades. Many of the trades were cancelled, and the share values returned to near normal, however many companies suffered sustained losses as the trades went through before “limits” were reached where the trades were later cancelled.

Signature: The key causal factors in the loss event are:

- Intrinsic flaw (algorithmic trading at the “best price” regardless of the actual price) exposed by the circumstances;
- Protection (requirement to set buy and sell prices) rendered ineffective by setting of extreme values;
- Other protection systems (cancelling of trades) did at least partially rectify the problems, but some traders did suffer lasting damage (losses on trades which were upheld, as the prices were not deemed “erroneous”).

Although much of the “damage” was done by automated trading systems this is a socio-technical issue, as human traders were still operating during the drop, and they set the algo parameters which so significantly contributed to the event.

3.4 Überlingen

In July 2002, two aircraft collided near Überlingen over the Bodensee (Lake Constance) [7]; the description here draws heavily on [8]. One aircraft was owned by DHL and was carrying freight; the other was a commercial aircraft carrying passengers, and operated by Bashkirian Airlines.

One of the roles of air traffic control (ATC) is to monitor flights and to offer guidance or instructions to aircraft so they maintain safe separation. In this case the primary control centre was Zurich. Many aircraft are also fitted with a Terminal Collision Avoidance System (TCAS) which is a “last resort” system which gives pilots “advisories” if it detects that there is another aircraft on a collision course. The TCAS systems coordinate their advisories so the two aircraft take diverging paths. Both aircraft were fitted with TCAS.

Description of the Event. In July 2002, a DHL-owned Boeing 757 aircraft collided with a Tupolev 154 operated by Bashkirian Airlines. All passengers and crew were killed. The trajectories are shown in Fig. 5, where the Bashkirian Airlines aircraft is moving South West, and the DHL aircraft is moving almost due South.

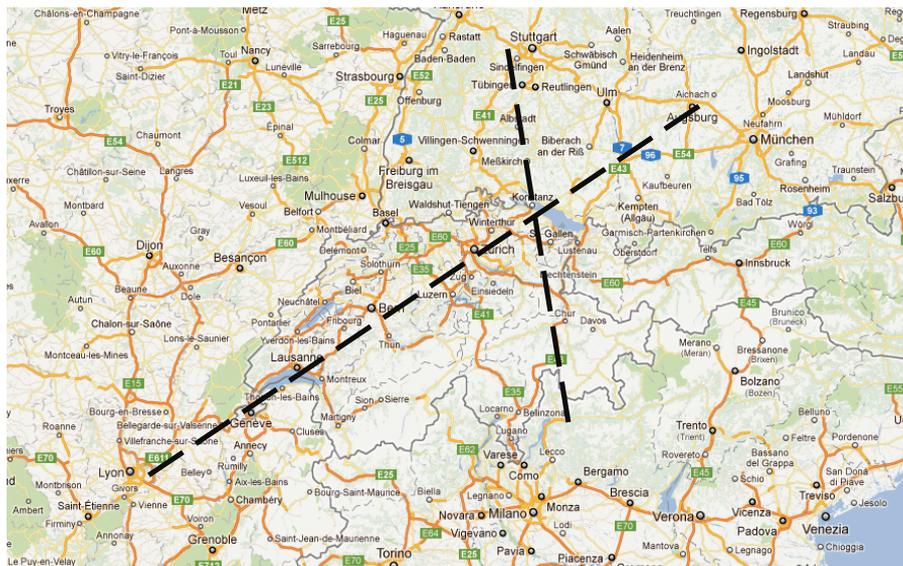


Fig. 5. Überlingen Accident

The two aircraft were initially on a collision course at 36,000 feet, and were first made aware of each other when their TCAS systems issued a traffic warning. Soon afterwards both aircraft received collision-avoidance instructions from TCAS — the B757 was to descend and the Tu-154 was to climb. Shortly after, the Tu-154 received an instruction from Zurich ATC to descend to avoid traffic. According to the cockpit voice recorder on the Tu-154, the pilot originally chose to follow the instruction from

TCAS. However, his co-pilot, a senior company executive who was on board in order to assess the pilot's performance, overruled him, and the aircraft began to descend; this was in accordance with company procedures and the Tu-154 operations manual.

At no point did Zurich ATC give any instructions to the Boeing 757 pilot, although the pilot did tell the ATC that he was descending, shortly before the collision. The two aircraft descended to 35,400 feet where they collided.

The entire accident, from the first TCAS traffic warning to the collision, took slightly less than a minute. Neither pilot was aware of the precise location of the other aircraft until a few seconds before the collision. Zurich ATC was not operating at full effectiveness on the night of the incident. Only a single controller was working, rather than the usual two, and he had to cover two frequencies and two radarscopes. In addition, upgrade work on the Zurich radar processing system meant that the system's performance was severely impaired. In particular, the STCA (Short Term Conflict Alert) function was not available. Further work on the ATC telephone network meant that it was unavailable. There was a backup line, but it was effectively useless due to technical problems.

The impending collision was noticed by a number of ATC centres in neighbouring regions, but they were unable to contact Zurich because of the telephone problems.

Signature: The key causal factors in the loss event are:

- Intrinsic flaw in that one airline took ATC as primary and the other took TCAS as primary, exposed by the circumstances;
- Protection systems (TCAS and ATC) rendered ineffective by the intrinsic flaw, and by reduced staffing and equipment problems in the Zurich ATC;
- Other protection systems, e.g. STCA and communications, rendered ineffective by the technical status of equipment at the Zurich ATC centre.

There are technical, socio-technical and organisational elements to this accident. It can be viewed as further along the spectrum towards an organisational accident, by comparison with the "Flash Crash" for at least two reasons. If DHL and Bashkirian Airlines had treated TCAS as primary, then the accident would have been averted. Further, Zurich ATC operating under such constrained conditions – low staffing, inoperative equipment – can be seen as an organisational failing.

3.5 Société Générale

In January 2008 Société Générale (SocGen) discovered that one of their agents, Jérôme Kerviel (JK), had been building up fraudulent trading positions over a number of years. The positions built up by JK amounted to about €50 Billion. These were "unwound" by SocGen resulting in a net loss of €4.9 Billion for the bank [9]. The actions taken by JK led to a court case and his being given a custodial sentence.

Description of the Event. The root of the problem came from "massive directional positions" [9], i.e. transactions assuming a massive movement of an asset's price in one direction (without any hedging); JK's activities went on over a number of years, and it was only towards the end that these positions became "massive". JK used a

number of methods for hiding these positions, including a significant number (nearly 1000) of fictitious trades which both hid his fraudulent positions, and altered various parameters which were monitored by the bank to detect excessive risk-taking. JK also used intra-monthly provisions (adjusting information at month end) which hid his position. There appears to have been some collusion with a trading assistant (who would normally make such intra-monthly provisions) although this does not seem to have been proven in the Court.

The positions went undetected partly due to JK's activities to conceal them, but also apparently due to poor supervision (although the decision by the Court could be seen as exonerating SocGen in this regard). For example, in 2007 JK was without an immediate superior for about two-and-a-half months, and no effective provisions for monitoring his activities were put in place during this period. Also, the new manager coming into post in April 2007 was weak [9], and the new manager was not given much support in taking on his new role.

A further factor, related to weak supervision, is the failure to act on the numerous alerts generated by systems which monitor positions and trades, for undesirable/suspicious activity. For example, in January 2007, an unusually high number of trades were marked as pending or with no counterparty; these were in fact fictitious, so the alert was a clear sign of the issue. In many cases the alerts were direct evidence of the fraudulent activity; it seems that they were followed up, but explanations from JK were accepted, and issues not escalated to superiors. The internal investigation [9] showed some 64 alerts which were directly linked to the fraudulent behaviour (and several more which were indirect).

A number of other factors, e.g. monitoring the growth in JK's share of the trades and profits in his division, and running a number of computer-based monitoring tools, could have helped to detect the problem. Also, it would have been possible to design the system so that JK could not make some of the trades, and his assistant would have had to, but this only increases the personal risk which would have been taken through collusion, rather than preventing the loss.

Signature: The key causal factors in the loss event are:

- Fraudulent behaviour, together with fictitious transactions which (to a degree) hid the inappropriate transactions;
- Failure of supervision, meaning that many of the systems put in place to detect such anomalous activity were either inoperative or not acted upon;
- Failure to investigate adequately alerts which indicated that fraudulent activity was taking place.

This is the closest to a "pure organisational" problem of the five examples reviewed here. Although there were technical systems which could have been used to help detect the fraud at an earlier stage than actually occurred these all appeared to work, if not perfectly, at least well enough to provide alerts and hence warnings of problems. The underlying "weakness" is that these systems were not used, for several reasons, including leaving JK without an immediate supervisor for a period of time, or because the warnings were not adequately handled.

3.6 Observations

The description of the signature of the above loss events is focused on protection, or barriers to accidents or other loss events (in the SocGen case, the supervision and alerts act as protection or barriers). One of the reasons for starting with two very simple examples is that the role (and inadequacy) of the barriers is reasonably obvious and unequivocal. Inevitably, for the more complex events, the choice of key factors is rather more selective (subjective) as there are many causal factors. Thus there is a risk of hindsight bias – but protection/barriers are introduced for a purpose, and it is thus worthwhile at least as part of our investigation considering why they were not effective, in these cases. There are some other factors that support the focus on barriers.

First, in some domains, e.g. nuclear, there are very clear design principles, e.g. [10], which are based around the idea of layers of protection. Here, the notion of protective layers and “defence in depth” seems to be fundamental to system design and risk control.

Second, even where the standards are not so explicit about protection, e.g. aviation, analysis of real system designs [11] indicates that the degree (number of levels) of protection varies with criticality. Thus it seems that design engineers “naturally” seek to introduce layered protection systems, even where this is not formally required.

Third, financial regulation also supports the idea of layers of protection, with measures both intended to reduce the likelihood of a loss event occurring, and to ameliorate a problem if it does arise [12].

Fourth, as is hopefully obvious from these examples, the notion of barriers is quite general and can apply to technical systems, to the interaction between technical systems and people (i.e. in the socio-technical space) and in organisations. Thus it seems to be a useful unifying concept.

We would thus argue that the focus on barriers is both relevant and valid as the basis for a “theory” and model of risk and loss in LSCITS (see below).

However, there is an apparently contradictory or countervailing issue which arises from standards and regulations – that is the requirement to evaluate risk, usually quantitatively. For example there are numerical targets for aircraft of 10^{-9} per flight hour for catastrophic events, and of 1.55×10^{-8} per flight hour for ATC induced accidents (e.g. mid air collisions). In other domains, e.g. financial markets, the notion of risk targets is less explicit, but there is still an expectation that risk is evaluated quantitatively (see below).

To simplify the issue, we can state that designers are often required to quantify risks “before the event”. However it is less clear how useful this quantification is “after the event” (here we are thinking about it as an explanatory mechanism; it is clearly not meaningful to talk about the probability of an event arising after it has occurred).

Thus this initial assessment of these loss events leads us to a problem formulation (in the sense meant by van der Ven [1]): *what is an appropriate risk assessment method for LSCITS?*

4 Risk Analysis

The term risk is used in many different ways, but with broadly similar meanings – the chance of harm or loss. We briefly set out some of the key principles of risk analysis below then use these principles in considering risk in the five loss events described in section 3. This approach is adopted in order to throw light on the problem formulation set out above. This then leads on to the suggestion of a theory for risk in LSCITS – the next step in Van der Ven’s model (see Fig. 1) to help us to reach the point where we might define models which can be evaluated via experiments and interventions.

4.1 Risk Analysis Principles

In its simplest form, risk is normally represented as follows:

$$\text{Risk} = \text{likelihood} \times \text{severity}$$

Where the likelihood is the probability of the loss event, or the frequency of the event, and the severity is the extent of the loss. This allows the risk of different events to be compared. For example, consider two risks, A and B, where:

$$\text{Risk A} = 10^{-7} \text{ per hour} \times 10 \text{ deaths}$$

$$\text{Risk B} = 10^{-6} \text{ per hour} \times 1 \text{ death}$$

Both have the same risk – an expectation of one fatality in a million hours, on average. Similar calculations can be done in terms of financial risk, e.g. expected loss of \$10M pa.

Some models of risk don’t quantify severity, but rank it qualitatively, e.g.: catastrophic, major, minor, and then risk is evaluated in terms of the probability in each risk class. In some cases, the probabilities are grouped into classes as well; when this is done, risk is evaluated via a matrix, see for example MilStd 882D [13].

In some circumstances, other factors are introduced, e.g. exposure to the risk, or the controllability of the risk by the operators. It is not uncommon for the exposure to be used to modify the probabilities, and factors such as controllability to be used in determining risk categories. Although there are many variations on a theme, the notion that risk is fundamentally a combination of probability and severity of loss is fairly universal, and that will be the focus in our analysis.

Finally, it should be noted that we are always interested in predicting or estimating risk to answer questions such as “is this system safe enough to deploy?” Even when making *post-hoc* decisions, e.g. “is this system now too insecure to continue using?”, we are making predictions of future behaviour based on knowledge of the past.

4.2 Risk Analysis of Loss Events

As may be apparent from the loss event descriptions above, it is not always easy to evaluate risk. The approach taken here is to seek to identify, in broad terms, what would need to be done (or known) to evaluate risk quantitatively in each case. An assessment is made of what risk might have been estimated before the events, and what might have been estimated with hindsight. This analysis is then used to inform a discussion of an approach to risk assessment for LSCITS.

Syringe Pump: The safety risk of a device such as the syringe pump would normally be evaluated using a tool such as fault trees [14] which enable accident probabilities to be evaluated based on data about the failure probability and failure modes of basic components, e.g. motors. To the author's knowledge this wasn't done, but a rough estimate of risk can still be made. The intent was that there was triple redundancy: the motor control, primary protection (quadrature system) and secondary protection (linear grating) would all need to fail for the device to fail in a hazardous manner. A failure rate of 10^{-3} per hour for each element is not unreasonable (a "rough" figure for commercial electronics); thus the accident probability might have been estimated at circa 10^{-9} per hour. With 10,000 devices in operation, this suggests 100,000 hours, or about 11 years between accidents.

However this estimate was not appropriate, in the circumstances. Two critical factors in the syringe pump accidents were the software which ignored "impossible" inputs rather than detecting them and taking safe actions, and the Schmitt triggers which could produce "impossible" inputs, at the extreme of their manufacturing tolerances.

To estimate the likelihood of any syringe pump containing a Schmitt trigger with the undesirable behaviour requires a model of the manufacturing distribution, and hence what proportion of the production would have the "dangerous" behaviour. Based on informal data on the system and the accident, this is about 100-1000 ppm (parts per million), or one in 10,000 to one in 1,000.

The likelihood that this erroneous behaviour would give rise to the accident was:

- ~0 prior to the modification to the code which disabled the start of the linear grating checking for movement, until it was detected by the control subsystem
- 1 after this modification

Note that the post-modification probability could also have been made 0, with defensive design of the software. However, without that design change between one in 1,000 and one in 10,000 of the devices would have been flawed, giving rise to an accident rate of one-ten per annum. Assuming that the Schmitt triggers "reliably" produced poor signals, then the accidents would occur early in operational life, and the actual accident rate per operating hour would be many orders worse than the estimate of 10^{-9} per hour.

The optimism in the estimated risk arises because the model used for risk estimation did not adequately reflect the way in which the devices (syringe pump software and the Schmitt triggers) worked (and failed).

The Cloud: In the case of using the cloud to store calendars, a very informal risk evaluation was undertaken. In essence a view was taken that cloud services were highly resilient (gave good availability) and if the service proved poor, the diary could be "repatriated" to a PC without too much difficulty. Also, an informal view of security was taken – that the calendar data wasn't too sensitive (although it would allow someone to determine travel arrangements) thus password protection was sufficient. However the terms of service say "you assume all risks and costs ..." , so it should have been apparent that there were risks! Further, the terms of service do say "does not guarantee or warrant that any content you may store or access through the

service will not be subject to inadvertent damage, corruption or loss”. However this was viewed (perhaps naively) as an “escape clause”, not a “real warning” so, informally, the view was that the risk of unauthorised access to data was low, and the risk of “losing” the data was effectively nil.

As was the case in the syringe pump example, the model used for risk estimation was inappropriate. As it turned out, the real cost of the failure was in the time to retype the calendar into a different tool (and this wasn’t even identified as an issue) and the failure mechanism, i.e. inability to re-export the calendar, was not considered either although, arguably, the wording of the terms of service should have sensitised me to this possibility.

Flash Crash. Financial markets have long understood the concept of “market risk”, see for example [15], and related concepts such as credit risk. These ideas are also at the basis of bank regulation; under the “Basel 2” arrangements banks have to hold reserves based on the notion of the “Value-at-Risk” (VaR). At its simplest, the requirements are for banks to maintain a level of capital which covers VaR at the 99.9th percentile confidence interval [16]. Whilst the details are complex, as many of the traders involved in the “Flash Crash” will have used hedging techniques (buying options to enable adverse movements in the price of assets to be offset), the majority if not all of the organisations involved will have undertaken some form of market risk analysis.

However what happened in the “Flash Crash” was not a market risk, but a systemic risk (or, perhaps better, the systemic issues meant that the market risk analysis was not accurate). The concept of systemic risk in financial markets is not new. In 2008 Long-Term Capital Management (LTCM), a US hedge fund, lost about 90% of its capital in about 9 months, for example losing \$1.8 Billion in August 2008 alone. It was “rescued” as there was a concern that it could collapse and cause significant consequential business failures [17]. The root cause of the “Flash Crash” was not the same as with LTCM – instead it was a socio-technical problem caused by a combination of the use of algo trading and the way certain trading parameters were set. However the critical point here is that the classical market risk analyses were not good predictors of events – again the underlying model of risk was inappropriate.

Überlingen. The safety of air traffic management in Europe is subject to Eurocontrol regulations, specifically ESARR 4 [18]. ESARR 4 sets a quantitative target for catastrophic accidents, which includes mid-air collisions, in European controlled airspace of 1.55×10^{-8} per flight hour (the figure is derived from historical achievement). It also requires “use of a quantitative risk based-approach in Air Traffic Management when introducing and/or planning changes to the ATM System” (section 1.1). In other words, providers of ATM services are required to provide a quantified risk assessment which shows that the risk of accidents, such as that at Überlingen, are less than 1.55×10^{-8} per flight hour.

Due to the way regulation is carried out, the services at Zurich will either have been subject to this regulation, or evaluated based on similar regulations which require a quantitative risk assessment. Thus there was a belief, prior to the accident, that the risk per aircraft was of the order of 10^{-8} per flight hour. As there are many

accumulated flight hours in Europe, the occurrence of this one accident does not mean that this average accident rate has been exceeded, however it is very unlikely that the models on which the risk assessment was carried out will have reflected the circumstances which arose at Überlingen.

In particular, the risk assessment models would have assumed proper staffing, working telephones, working STCA, etc. – or perhaps more accurately, the models would have assumed that where there were such deficiencies appropriate means would have been taken to mitigate risks, e.g. calling on neighbouring centres. TCAS is viewed as an aircraft system, not part of ATM, so it is unlikely that the ATM risk analysis would have considered TCAS. Further, it seems very improbable that the risk analysis would have considered the fact that ATM might have, in effect, rendered TCAS ineffective by giving instructions which over-rode this “last line of defence”. So, once more, the model (which almost certainly would have been) used for the risk calculations was not representative of the situation that arose.

Société Générale. SocGen will have carried out market risk analysis but, as with the “Flash Crash”, what happened was “outside” the models used to assess risk. However, what occurred at SocGen would generally be classified as operational risk, rather than systemic risk. There are, nonetheless, similarities with the “Flash Crash”. The type of problem seen was not unprecedented; for example work by the Federal Reserve Bank of Boston [19] states that the “capital charge for operational risk will often exceed the charge for market risk”. Put another way, the VaR for operational issues may well be greater than that due to the market.

As the causes of the SocGen issues were largely organisational, an effective risk model would have to address these issues. Some work has been done in this area, e.g. using Bayesian approaches to modelling operational risk including fraud in insurance [20], but this remains a little explored area, to the author’s knowledge.

4.3 Risk Analysis for LSCITS

The examples given above show that “classical” analyses of risk do not shed much light on the causes of the loss events. Implicitly, system safety engineering methods (which apply to the syringe pump and Überlingen) assume that physical failure mechanisms reflect aleatoric or aleatory uncertainty, i.e. “randomness”, which can be characterized by a stochastic model. Further, we implicitly assume ergodicity – i.e. that past failure behaviours are good predictors of the future. Based on these assumptions we can use probability density functions (PDFs) and often we approximate those functions by point probabilities, e.g. the mean of the PDF, in evaluating risk. Such approaches are good ways of modelling processes such as the tossing of coins, and the failure of simple components, e.g. resistors. They underlie the most common quantitative models of system safety, e.g. the calculations supporting fault tree analysis. Similar assumptions underlie the processes of modelling market risk (and in some approaches to software safety [21]).

However, in many cases we face epistemic uncertainty, i.e. limited knowledge of the system model or of the stochastic model. In other words we do not know the shape of the PDF, nor can we estimate its mean. In the cases above, whether sophisticated

risk analysis was carried out, or it was very informal, as in the cloud example, the loss events are much better explained in terms of epistemic uncertainty – or to put it simple, the wrong model was used.

There is a further factor in some cases – that the models need to change (or be changed). In other words even if the right model was used in the initial assessment of risk, the system structure and thus the model which is used for assessing risk changes as the system operates and evolves. As markets and trading systems evolve rapidly, it is almost inevitable that, in situations typified by the “Flash Crash”, any analysis done before introducing a new trading system would rapidly become inaccurate. Further, in the SocGen case, the risk controls assumed a model of the organisation with people filling key roles – the risks were very different when JK’s superior left and was not replaced for over two months.

Returning to Van der Ven’s framework, we can propose an *explanatory* theory: *risks and loss events in LSCITS are better explained via analysis of epistemic uncertainty than aleatory uncertainty.*

Note that this is not to say that the techniques based on aleatory uncertainty are worthless – indeed it can be argued it is because they are so effective that the epistemic factors dominate in actual loss events. However, even if it is accepted that this is a plausible explanatory theory, it does not really help us towards a model which can be used to analyse LSCITS, so we now consider some aspects of the LSCITS “stack” and consider how we might use this to build a generative theory, as a step towards a model (in Van der Ven’s terms).

In an analysis of Australian defence avionics systems [11] it became clear that there are “layers of protection” against systematic (design) faults in systems, which vary with criticality – the worse the outcome the more the layers of protection. Further the “innermost” layer of protection was concerned with either avoiding or containing any systematic causes of hazards, at source, and the outer layers were concerned with detection and mitigation (of hazard causes). Barriers can be seen in all the five examples discussed above; in some cases these are technological, and in several of the cases they are organisational. This leads us back to the LSCITS “stack” – or something like it.

We can think of “barriers” in the following ways:

- Prevention of problems, at source, or managing them to a low and controlled probability of occurrence – the province of PSS, and other techniques, e.g. Six Sigma, where the processes are human and organisational, not technical;
- Detection and mitigation of technical problems – (in part) the province of HISE, especially considering the interaction of peer components (in a system or SoS);
- Prevention of socio-technical problems, and detection and mitigation of problems through socio-technical means – the province of STSE which is both concerned with good socio-technical system design, and with handling errors arising at this level;
- Prevention of organisational problems, and detection and mitigation of problems through organisational means – the province of CiO which is both concerned with good organisational design, and with handling errors arising at this level.

In general the “barriers” can be characterised in the following ways:

- Their detection and handing of failures (undesirable behaviours) which arise from lower levels;
- Their detection and handing of failures (undesirable behaviours) which arise from peer systems;
- Internally generated failures (undesirable behaviours);
- Failures (undesirable behaviours) “exported” to higher levels.

If this is an appropriate way of looking at LSCITS, in this context, then a means of evaluating risks is needed. We briefly discuss this below, but first set out a further “theory” in the sense of Van der Ven’s framework. We propose a *generative* theory: *risks and loss events in LSCITS are best controlled (and risks estimated) via the design and analysis of barriers.*

In order to proceed from the above theory towards a model, in Van der Ven’s terms, we need to produce means of identifying the need for barriers, for “designing” them, and for evaluating risk. For brevity, we assume here that barrier identification is possible, e.g. by using adaptations of current methods, which do identify barriers in both technical systems and organisations, and focus on risk evaluation. There are at least three possible approaches:

- Qualitative approaches, e.g. the use of tabular ways of expressing the “depth of defence” against particular potential causes of loss events – these can then be evaluated based on loss event severity, to assess the adequacy of risk controls (this is essentially a generalisation of the approach used for safety in MilStd 882D [13]);
- Quantitative approaches, perhaps by extending the Fault Propagation and Transformation Analysis (FPTA) method [22] developed in part through the LSCITS programme, to consider fault propagation between barriers;
- Quantitative approaches, building on Bayesian approaches such as those proposed for operational risk [20].

It may be practical to combine these approaches in particular ways, or to learn from them, e.g. using the scenario testing approach proposed in [20] to validate FPTA models. In practice, it might be that the quantitative approaches are best thought of as means of ranking designs (sets of barriers), than evaluating risk in the aleatory sense, or in the sense of loss per unit time, which is the underlying measure in safety and in financial risk. In practice, the idea of scenario testing may prove to be vital, as the only practicable way of handling system complexity.

Several of the examples discussed above, e.g. the financial ones and Überlingen can be viewed as SoS. A characteristic of an SoS is that the constituent systems – its configuration – changes over time, and typically faster than individual systems can be redesigned. If any change violates assumptions made about the SoS then there can be undesired behaviour – such changes can be thought of as inflection points [23]. No SoS or system design can be robust against all potential changes, but perhaps it might prove possible to use scenario testing on barrier models to demonstrate robustness against epistemic uncertainty – or at least to identify what classes of change bring about undesirable inflection points. Of course, this will only be as good as the underlying models.

5 Discussion

There has been work, particularly in the safety community, focused on the modelling and analysis of accidents. We review this work here, and draw some distinctions with the approach which we have outlined above. We then make a few further observations about the difficulties of quantification of risk for high criticality systems.

Peter Ladkin in Bielefeld has developed Why-Because Analysis (WBA) [3] as a “rigorous technique for causally analysing the behaviour of complex technical and socio-technical systems”. Whilst it is also intended to assist in analysing safety requirements, to the author’s knowledge it has found greatest utility in accident analysis, where its flexibility enables it to be used to address relevant causal factors. Our experience with WBA, for example [8] which analyses the Überlingen accident, and our as yet unpublished work on the Wenzhou train crash, shows its utility. Indeed, one possible step for making the ideas set out above more rigorous would be to analyse all the loss events using WBA. However our work on Wenzhou suggests that WBA is not good at dealing with influences, rather than causes, thus there may be merit in seeking to extend WBA before analysing all the above loss events.

Further, we are not aware of cases where WBA has been used proactively to drive designs and we do not see how it would help in identifying barriers, although we note that [3] refers to the use of WBA to identify requirements. As we understand it, WBA does not help to evaluate risk (at least quantitatively) although again one can envisage ways of extending the method to do this.

Nancy Leveson at MIT has developed STAMP [24] as a means of analysing both socio-technical and organisational causes of accidents – thus it gives a framework for analysing the type of loss events discussed earlier. One of the great attractions about STAMP is that it gives a generic model of factors in accident causation from low-level technical issues through organisations, up to political institutions. A number of examples using STAMP have been published. However our experience, to date, has been that it is hard to apply, and that the guidewords in the method for assessing deviation from intent do not seem to be sufficiently comprehensive. For example, one of the issues in the Wenzhou accident is that the Ministry of Railways (MoR) was both the operator of the trains and the regulator; although the STAMP model identifies operators and regulators there is no obvious way to reflect the conflict of interest (potential single point of failure) due to MoR’s dual role, in that framework. As one of the key factors in some of the five loss events discussed above was single point failures which undermined multiple protective barriers, this at present seems to be a limitation of STAMP (this must be viewed as a tentative assessment as our work on Wenzhou is ongoing).

Like WBA, we have yet to see STAMP used proactively in system design although there is nothing intrinsic in the method which should prevent this. Again, like WBA, STAMP does not appear to provide a basis for evaluating risk in the sense investigated here although, again, extensions might be possible.

Recent work on resilience engineering [25] has a stronger influence on the ideas set out herein. Both in his publications on resilience engineering, and in prior work, Hollnagel emphasises the importance of designing barriers, and the need to assess human behaviour and cognitive processes, in designing systems and barriers. As we

develop the ideas set out above we need to draw on the insights from resilience engineering, but note that the scope of our endeavour is broader – seeking to take a unified view of critical systems, rather than the focus on safety in Hollnagel’s work.

Other work in LSCITS is addressing issues relevant to the approach outlined here, for example the use of responsibility modelling as an aid to risk analysis in socio-technical systems [26]. As currently defined, this work would most naturally form part of the qualitative risk analysis approach identified above (indeed we have used it this way in our Wenzhou analysis).

There is some literature, for example [27], which is casting doubt on the validity of quantitative risk assessment. This can be read two ways: as supporting our analysis here, by confirming that real-world risk assessments are often flawed, or contradicting it by implying that trying to quantify risk is impractical. We hope, in time, to be able to support a third view; that taking an approach, informed by quantitative analysis, can lead to more robust designs (e.g. better and better-placed barriers) and more resilience to changes in models, than achieve by current approaches. Separately, we are working on approaches to assessing whether or not risk predictions are valid, or trustworthy.

As indicated earlier, there is a risk of hindsight bias, including finding examples which confirm the author’s hypothesis. Also identifying “root causes” of a loss event is always judgemental – in other words, when do you stop looking for prior causes of events? In the cases considered, several are the subject of existing public domain analyses, so this helps avoid hindsight bias. Further, barriers are intended to stop the propagation of faults and errors – so it is not biased to observe that they weren’t effective, once a loss event has occurred.

Further, the author’s “foresight bias” was that the loss events would be explained by change – in the technical system, in usage, etc. In some of the examples, e.g. the syringe pump (technical) and Überlingen (organisational) there are clear changes (if only temporary in the case of Überlingen) which have a causal influence on the loss event, but the other cases are less clear-cut. Arguably, they all involve change – with the “cloud” example it was moving a calendar, with the “Flash Crash” and SocGen there were changes in behaviour. However these can be viewed as changes in usage within design parameters, not a change in the intended usage. Partly for this reason, and also because change can be thought of as one of the possible reasons why the models used for analysis do not reflect the system (in the broadest sense) as used, it was decided to treat epistemic uncertainty as the primary factor in the explanatory theory. Whilst we cannot prove that there is no hindsight bias, the fact that this is an explanatory theory, and it is not used directly to produce the generative theory and any solutions, makes the problem of hindsight bias less of a concern than it might otherwise be.

Finally, we believe that the observation we have made about the limitations of risk analysis because the causation of loss events is based more on epistemic than aleatory uncertainty to be a distinctive, if not unique viewpoint. There are, for example, criticisms of ESARR 4, e.g. [27], which challenge the underlying safety models in ATM (especially for setting targets), but this, and all the other examples we know, focus on a particular system or scenario. However, if nothing else, this analysis of ESARR 4 serves to show how important it is to analyse the models behind standards, as well as systems designs, to ensure that they are effective in their role.

6 Conclusions

There are growing numbers of LSCITS in operation, many of which are critical, e.g. those supporting ATM and the financial markets. Also more “classical” safety-critical applications are becoming more extensively networked. The failure or misbehaviour of such LSCITS could lead to harm, be it in terms of loss of life or financial impact. This paper has sought to demonstrate, by means of examples, that classical approaches to assessing risks of critical systems have severe limitations in practice, and do not seem to be effective for LSCITS. In general this is because the basis on which the risk assessment is done is not representative of the causal mechanisms in actual loss events.

Our approach in this paper has been influenced by Van der Ven’s approach to research in social sciences, building research problems and theories from empirical observations. Although this is perhaps unusual, it seems justified in that the social sciences deals with very complex situations where experimentation (in the classical scientific sense) is not possible – and the same problems exist in assessing the effectiveness and risks of LSCITS. It is our intent to take this on further, to build models from which we can then plan and conduct experiments to help refine our ideas. To do this requires at least three areas of exploration:

- Assessment of the signatures of a larger set of loss events;
- More rigorous assessment of the causal structures and signatures of a number of loss events, e.g. using WBA;
- Construction of a prospective model for system risk analysis and design refinement, perhaps based on work on FPTA and Bayesian approaches to risk analysis (to rank risks, if not to evaluate them accurately).

An underlying assumption in the approach we have sketched here is that the concept of barriers is a useful abstraction in LSCITS. It has several merits:

- The concept is already used in technical systems, e.g. aviation and nuclear, and in organisations, e.g. the financial sector, and is one of the underlying principles in resilience engineering;
- The concept applies independent of implementation technology;
- It offers a significant abstraction away from the detail of an LSCITS;
- Analysis of “integrity of barriers” may give a way of assessing the continued robustness and resilience of a system (or SoS) following change.

In extending this work, one of the key challenges will be to demonstrate that the concepts are effective in the presence of change, especially in SoS, as this is central to the challenges of constructing and assuring LSCITS.

Finally, the LSCITS principals have recently set out their views on the engineering of LSCITS [29], and identified several challenges. It is hoped that the work set out here will contribute to providing solutions to two of these challenges: 5 (“how can systems be designed to recover from failure?”) and 6 (“how can we manage complex, dynamically changing system configurations?”). If we can do this, then we will have made a significant contribution to the understanding of how to design and assess LSCITS for critical applications.

Acknowledgements

In producing this paper I have been influenced by the work of colleagues in York and in the LSCITS partner Universities, and by discussions with other collaborators.

I am grateful for input made by my LSCITS colleagues, especially Dave Cliff and Ian Somerville. In York I have benefited from discussions and contributions from Rob Alexander, Georgios Despotou, Gaocheng Xe, Tim Kelly, Andrew Rae, Derek Reinhardt and Niu Ru.

I have also had some useful and stimulating discussions with Mark Connelly and Mark Rodbert of Neural Insights, and Maurice Perks of IBM. I am particularly grateful to Mark Connelly for identifying some relevant literature on practices in the financial sector. Further, I am grateful for the brief but helpful discussions with Robert Cowell of City University, the lead author of [20], during a visit to York.

Finally, I should like to acknowledge the support to this work by the EPSRC through the LSCITS programme, ref. EP/F001096/1.

References

1. Van der Ven, A.H., Engaged Scholarship: A Guide for Organizational and Social Research, Oxford University Press, 2007.
2. LSCITS research programme, see <http://lscits.cs.bris.ac.uk/research.html> (last accessed, 3rd February 2012).
3. Ladkin, P.B., Why-Because Analysis, see <http://www.rvs.uni-bielefeld.de/research/WBA/> (last accessed, 4th February 2012).
4. Clarke, S.J., Coombes, A., McDermid, J.A., The Analysis of Safety Arguments in the Specification of a Motor Speed Control Loop, YCS 136, Department of Computer Science, University of York, 1990.
5. Cliff D., Private Communication, January 2012.
6. The Economist (on-line edition), A Few Minutes of Mayhem, May 13th 2010.
7. Bundesstelle für Flugunfalluntersuchung (BFU: German Federal Bureau of Aircraft Accidents Investigation), Accident on 1 July 2002, Near Überlingen/Lake Constance, Germany Involving Boeing B757-200 and Tupolev TU154M, Investigation Report AX001-1-2/02, May 2004.
8. Alexander R., Hall-May M., Modelling and Analysis of System of Systems Accidents, DARP/TN/2003/19, University of York, February 2004.
9. Société Générale, General Inspection Department, Mission Green, Summary Report, May 20th 2008 (English version, translated from the French).
10. Health and Safety Executive, Safety Assessment Principles for Nuclear Facilities, Revision 1, 2006.
11. Reinhardt, D.W., McDermid, J.A., Assuring against Systematic Faults using Architecture and Fault Tolerance in Aviation Systems, in Proc. Improving Systems and Safety Engineering, Brisbane, Australia, August 2010.
12. The Basel Committee on Banking Supervision of the Bank for International Settlements, see: <http://www.bis.org/bcbs/about.htm> (last accessed 4th February 2012).

13. US DoD, MilStd 882D Standard Practice for System Safety, 2002.
14. Roberts, N.H., Vesely, W.E., Haasl, D.F., Goldberg, F.F., Fault Tree Handbook, Systems and Reliability Research Office of U.S. Nuclear Regulatory Commission, Washington, DC, 20555, 1981.
15. Alexander, C., *Market Risk Analysis*, Volumes I-IV, Wiley, New York, 2008.
16. Basel Committee on Banking Supervision, International Convergence of Capital Management and Capital Standards (Basel II), Bank for International Settlements, 2004.
17. US General Accounting Office, Long-term Capital Management: Regulators Need to Focus Greater Attention on Systemic Risk, GAO/GDD-00-3, October 1999.
18. Eurocontrol Safety Regulatory Requirement (ESARR) 4, Risk Assessment and Mitigation in ATM, Eurocontrol, 2001.
19. Fontnouvell, P. de, DeJesus-Reuff, V., Jordan, J., Rosengren, E., Using Loss Data to Quantify Operational Risk, Federal Reserve Bank of Boston, April 2003.
20. Cowell, R.G. Verrall, R.J. Yoon, Y.K., Modelling Operational Risk with Bayesian Networks, *Journal of Risk and Insurance*, Vol. 74, No. 4, pp795-827, 2007.
21. McDermid, J.A., Risk, Uncertainty and Software Safety, in Proc 28th International System Safety Conference, Vancouver, Canada, International System Safety Society, 2008.
22. Ge, X., Paige, R.F., McDermid, J.A., Probabilistic Failure Propagation and Transformation Analysis, in Proc Safecom 2009, LNCS 5775, B Buth, G Rabe, T Seyfarth (Eds), 2009.
23. Perks, M., Private Communication, February 2012.
24. Leveson, N.G., A New Accident Model for Engineering Safer Systems, *Safety Science*, vol. 42, no. 4, pp. 237-270, 2004.
25. Hollnagel, E., Woods, D.D., Leveson, N.G., Resilience Engineering: Concepts and Precepts, Ashgate Publishing, 2006.
26. Sommerville, I., Lock, R., Storer, T., Responsibility Modeling for Risk Analysis, Proc. ESREL 2009. Prague, September 2009.
27. Hansson, S.O., Seven Myths of Risk, *Risk Management*, vol. 7, no. 2, pp. 7-17, Jan. 2005
28. Brooker, P., Air Traffic Management Accident Risk, Part 2: Repairing the Deficiencies of ESARR 4, Cranfield Research report PB/5/05, May 2005.
29. Sommerville, I., Cliff, D., Calinescu, R., Keen, J., Kelly, T.P., Kwiatkowska, M., McDermid, J.A., Paige, R.F., Large-Scale Complex IT Systems, *Communications of the ACM*, Vol. 55 No. 7, Pages 71-77, June 2012.