

System safety and the law

Dr Chris Elliott FREng

System engineer and barrister, Director Pitchill Consulting Ltd, UK
chris.elliott@pitchill.com t+44 (0)1483 273793 m +44 (0)7836 217901 f +44 (0)1483 273465

Keywords: Law, safety, complexity.

Abstract

Engineers face prosecution, and their companies compensation claims, when safety-critical systems fail and cause harm. The law reduces the complexities of systems to simple linear and hierarchical structures, risking injustice and obstructing investigation.

A three-fold solution is proposed, involving a more constructive approach to investigation, empowering system architects and corresponding contractual frameworks, and adapting the criminal and civil law to hold the system architects to account.

1 Two cultures – lawyers and engineers

The legal view of engineering

Criminal and civil law both impact on engineering. Both concern the failure to discharge duties that are imposed by society, and most legal duties apply both to natural persons (human beings) and legal persons (companies), but the two arms of the law have different goals and effects:

- criminal law regulates the relationship between persons and society. It imposes specific obligations and seeks to punish those who fail to discharge them
- civil law regulates the relationship between persons. It imposes general obligations and seeks to compensate the victims of those who fail to discharge them.

Criminal law

The crimes that are most likely to be alleged after a serious failure of a safety-critical engineering system are:

- breach of the Health and Safety at Work etc Act (HSWA)
- manslaughter (individual or corporate)
- environmental or hygiene offences under numerous statutes.

The criminal law seeks a binary decision – the defendant is either innocent or is guilty beyond reasonable doubt. Although more than one defendant may be found guilty for

the same event, each of them individually must be guilty; guilt is not shared.

The circumstances under which a defendant can be found guilty vary greatly between crimes. Three broad classes are, in order of increasing ease of conviction:

- the defendant intended his actions to cause their effect or at least have been reckless as to their consequences (eg manslaughter)
- the defendant failed to take reasonable care to prevent the consequences of his actions (eg breach of HSWA)
- the defendant is guilty if the consequences occurred, independent of his actions or intentions (strict liability, eg many environmental crimes).

It is a fundamental of English criminal law that the defendant is innocent until proven guilty. That principle applies equally to criminal charges that might arise for a failure of a safety-critical system but it is not always obvious. For example, if an accident has occurred at work and the prosecution alleges that the employer could have prevented it, section 40 HSWA says

.... it shall be for the accused to prove .. that it was not practicable or not reasonably practicable to do more than was in fact done ...

In other words, the defendant has to prove the circumstances that establish his innocence.

Manslaughter has been particularly difficult to prove where the defendant is a company. The traditional legal doctrine is that there must be an identifiable “controlling mind”, an individual who was personally responsible for the corporate failing that led to death. This is almost impossible in a large organisation, so the only successful prosecutions have been of small companies (see for example [1]).

The current government is committed to introducing a new offence of corporate manslaughter. A Bill is anticipated for autumn 2006 that takes account of the extensive criticism of the draft legislation. The indications are that this will go some way towards recognising that the corporate culture is as important as the nominal procedures but will still look narrowly for a single defendant and not for a defective system. Also it is unlikely to expose the company or its managers and directors to any greater penalty than the unlimited fines that may be imposed for the breach of HSWA

that would almost inevitably also be provable if corporate manslaughter had occurred.

Civil law

Duties under civil law can arise from statute (such as the Occupiers' Liability Acts) or through the historical evolution of the common law (such as trespass) and increasingly as a result of a combination of these (such as statutory nuisance). The common law duty not to be negligent is typical of these civil law duties and of great importance for the legal view of engineering.

A person can be found liable in negligence if:

- he owed a duty of care to another person, and
- he breached that duty by conduct that fell below the standard of a reasonable practitioner, and
- the breach has caused compensable harm to that other person.

In practice engineers owe a duty of care to everyone who might be affected by their actions and the test of breach, as for any other profession, allows honest mistake and for innovation. It is the third test that is important for this paper, summed up as

One is never simply liable; one is always liable for something ...[2]

The civil law has great difficulty deciding whether an action caused harm (ie is there something for which one is liable) in complex situations where there are multiple possible causes of the harm. Where:

- it is not possible to determine with a certainty greater than 50% which of several events caused the harm; and
- a different person was responsible for each event

then it is not possible to find any of them liable to compensate the victim [3]. If the harm itself might have taken several different forms and can only be estimated statistically, the law is even more reluctant to find any liability [4]. The only exception is where the different causes are of the same kind and it is more likely than not one of them caused actual, not statistical, harm. This is a recent development of the law, applied to the claim of a victim of asbestos inhalation who had worked for two employers, both of whom were equally likely to have caused his condition [5].

The principle behind this confused areas of law that has only been most briefly summarised here was summed up by Lord Hoffman:

... the law regards the world as in principle bound by the laws of causality. Everything has a determinate cause, even if we do not know what it is.[6]

That principle should not worry any engineer but it is applied without appreciating the concepts of *a priori* uncertainty or of expectation value.

The engineering view of the law

The engineer's view of the law is, for this author, summed up by the experience of researchers in the Health and Safety Executive. They had developed great understanding of the risk of explosion arising from industrial processes and had codified their knowledge in a software tool. They were keen to make this available for industry, rightly believing that its use would help to achieve HSE's goal of ensuring that risks to people arising from work activities are properly controlled. They were however concerned that HSE might be liable if the software were published and subsequently found to have contributed to an accident.

The research managers sought specialist (and expensive) legal advice to answer the question "is there a risk that HSE might be liable for faults in the tool?". The answer they received was equivocal – there were circumstances under which liability might arise (subsequent case law has established that HSE can be liable⁷ in negligence). Although the researchers were dissatisfied with the answer, it reflects an ill-founded question. Had they asked "is the risk that HSE might be liable small enough to be outweighed by the benefits to society of this tool?" they might have received a useful answer. The lawyers could no more give an unequivocal answer to the question posed than the researchers could to the question "is this process safe?".

Engineers often encounter lawyers, if not the law itself, when negotiating contracts. They see the lawyers' job as helping them to achieve a contract which simultaneously minimises their potential liabilities and maximises their potential returns. Although that sounds attractive in the short term, in the longer term a more cooperative and balanced relationship with customers and suppliers may be better. Which project is seen as a greater success – the Wembley stadium being built by a contractor that prides itself on its tough approach to subcontractors or the Eden project and T5 Heathrow founded on the cooperative contractual principles of the Egan report [8]? Engineers, like many of their legal advisers, have difficulty distinguishing their rights from their interests.

Engineers also encounter the law when acting as expert witnesses. It is this author's experience that they rightly see this more as a gladiatorial contest than a scientific exercise in truth-finding. This and their experience of the law in their private lives can lead then to hold the law in awe, fear and contempt because it is so powerful and yet so apparently unable to understand the practical constraints that are inherent to engineering.

Law and engineering – a summary

In 1959 C P Snow delivered his infamous lecture on "*The two cultures and the scientific revolution*" in which he identified a "gulf of mutual incomprehension" between the sciences and the arts. The title of this section of this paper is a nod to Snow, but is there really a gulf between engineers and lawyers?

Engineers and lawyers both work by following rules and processes, complemented by creative thinking to address novel situations. Both have a large body of theory that has to

be applied pragmatically to every problem. Both are seeking to deliver to their clients the least-bad trade-off between conflicting and often irreconcilable constraints. Both look for certainty from the other, whilst being unable to deliver it themselves.

And both have a deep blind spot where they overlook the realities of the behaviour of systems:

- Criminal and civil law are founded on the principle that every person is either guilty/liable or not and that, where actual harm has been caused, someone is to blame. Law has difficulty with the concept of failure arising from the interaction of elements of a system, rather than from the failure of one of those elements.
- Engineering is often taught and usually perceived as an essentially linear and logical process, where complex problems may be rendered simple by decomposition into constituent parts, each of which may be solved in isolation and integrated to create the whole. This is difficult in practice even in the most mature and pedestrian fields of engineering; it is far from true in the engineering of safety systems, especially those that must cope with the complexity and uncertainty introduced by reliance on software.

2 Systems – engineering in the real world

What is a system?

A *system* is a set of elements or components which, when brought together, exhibit properties which were not present in the elements alone. A trivial example is a system consisting of a battery, a bulb and two wires. When brought together, they create light. The light is an *emergent property* of the system. A system truly is “more than the sum of its parts”.

Systems can also have unintended and undesirable emergent properties. For example, the specification of what a component will do often has some implicit and unstated assumptions of how it will be used. If the supplier’s assumptions are different from those of the user, something unanticipated might occur when it is built into the system. That unanticipated consequence need not be a property of the component itself; it could arise solely from its interactions with the system. It might not emerge immediately. A component may work satisfactorily at first but problems emerge later, either because the component changes (for example because of incorrect maintenance or inappropriate modification) or because of changes to other components with which it must work.

These arguments apply equally to physical components, such as a bolt or power station, or a virtual component, such as an operating system or method of working.

Safety is an emergent property of a system, not a quality of the components or something that can be bolted on. A poorly designed or managed system is intrinsically unsafe; it remains unsafe however well the components (things or people) work.

Conventional management practices are oriented towards dealing with issues that are *complicated*, that is, *rich in detail*. Hierarchical, confrontational management operates well when the purpose and behaviour of the structure can be used to determine what is needed from even its smallest components and the interactions between the components are simple and predictable.

Systems are *complex* because they are *rich in structure* - the components are, for any significant system, interconnected and interdependent. As a result, the emergent properties are often hard to predict or manage. If safety is to be an emergent property, the system needs a robust self-correcting management mechanism that encourages cooperation between those responsible for each of the components.

3 System failures

Systems fail in many different ways, illustrated by these examples of systems that have behaved in a way that was not intended. The consequence may have been an accident, financial loss or poor performance.

Failure of a simple system

A small factory automation system, designed in the early 1980s, used an HP85 computer to control a stepper motor drive and a parallel I/O port. These external devices were made by two different companies, neither of which was HP. Communications was via HP’s proprietary standard HP-IB. This had proved to be so valuable that the IEEE had codified it as IEEE-488, a recognised international standard. The HP computer and both of the external devices had certificates of compliance with IEEE-488.

The HP computer could control either of the devices perfectly but, when both were connected, the system did not work. After extensive and expensive investigations, it was found that the IEEE-488 standard did not define the timing of one of the control signals. The designers of the two external devices had made different, and incompatible, assumptions about this signal.

The client sought someone to blame (and sue) for the considerable cost and delay. There was no one. Both of the devices and the HP computer complied with every aspect of IEEE-488, an international standard issued by the leading US standards body.

The problem would probably not have occurred had the two external devices been made by the same company, especially if that company had been HP who made the computer. The timing of the signal would still not have been defined but the “normal way of doing things” in the company would have meant that all its products were based on the same assumptions.

The customer incurred substantial financial losses for something that was not his fault but civil law did not allow him to recover those losses from any party. The only party which might have been negligent was the IEEE standards

committee, but the law would almost certainly find it too distant to be held responsible.

Failure of a complex system

The space shuttle Columbia was lost when it broke up on re-entry. The direct cause was that the heat shield on one of its wings had been damaged by a piece of insulating foam that had broken away from the External Tank during the launch. The Columbia Accident Investigation Board conducted an exhaustive formal inquiry. Chapter 7 of its report is entitled "The accidents organizational causes" and starts with the following text:

The organizational causes of this accident are rooted in the Space Shuttle Programs history and culture, including the original compromises that were required to gain approval for the Shuttle Program, subsequent years of resource constraints, fluctuating priorities, schedule pressures, mischaracterizations of the Shuttle as operational rather than developmental, and lack of an agreed national vision. Cultural traits and organizational practices detrimental to safety and reliability were allowed to develop, including: reliance on past success as a substitute for sound engineering practices (such as testing to understand why systems were not performing in accordance with requirements/specifications); organizational barriers which prevented effective communication of critical safety information and stifled professional differences of opinion; lack of integrated management across program elements; and the evolution of an informal chain of command and decision-making processes that operated outside the organizations rules.

Appearing before the Senate Committee on Commerce, Science and Transportation, the Chairman blamed NASA's system, not any individuals, and said there was "not one person responsible."

No-one could be found guilty in criminal law for the accident. None of the engineers or managers were individually to blame and the criminal law could not be applied to those who had, at the highest level, established NASA's culture and permitted it to persist.

Failure of a non-engineering system

During the 1990s concerns were raised about the care of children receiving complex cardiac surgery at the Bristol Royal Infirmary (BRI). A Public Inquiry was conducted from October 1998 to July 2001, chaired by Professor Ian Kennedy. Its report [9] drew important conclusions on the behaviour of BRI, and more generally of the National Health Service, as a system and on the behaviour and culpability of individuals within that system. Some relevant paragraphs from the Summary of that report:

3 The story of the paediatric cardiac surgical service in Bristol is not an account of bad people. Nor is it an account of people who did not care, nor of people who wilfully harmed patients.

4 It is an account of people who cared greatly about human suffering, and were dedicated and well-motivated.

Sadly, some lacked insight and their behaviour was flawed. Many failed to communicate with each other, and to work together effectively for the interests of their patients. There was a lack of leadership, and of teamwork.

5 It is an account of healthcare professionals working in Bristol who were victims of a combination of circumstances which owed as much to general failings in the NHS at the time than any individual failing.

10 And it is an account of a system of hospital care which was poorly organised. It was beset with uncertainty as to how to get things done, such that when concerns were raised, it took years for them to be taken seriously.

21 We adopt a 'systems' approach to analysis, by which poor performance and errors are seen as the product of systems which are not working well, as much as the result of any particular individual's conduct.

26 Bristol was not unusual in having problems. It was, after all, managing the transition from the known (the old NHS) to the unknown (Trust status)....

30 At a national level there was confusion as to who was responsible for monitoring quality of care. The confusion was not, however, just some administrative game of 'pass the parcel'. What was at stake was the health, welfare, and indeed the lives of children. What was lacking was any real system whereby any organisation took responsibility for what a lay person would describe as 'keeping an eye on things'...

The tragedy behind these observations is that, although the problems of BRI were becoming known by the end of the 1990s, the Inquiry's experts' statistical analysis concluded that between 30 and 35 children died between 1991 and 1995, over and above the number which would be expected if the Unit had been 'typical'. More generally, the report says that:

80 Around 5% of the 8.5 million patients admitted to hospitals in England and Wales each year experience an adverse event which may be preventable with the exercise of ordinary standards of care. How many of these events lead to death is not known but it may be as high as 25,000 people a year.

The Inquiry report, although not exonerating the individual doctors and administrators, makes it clear that their failings have to be seen in the light of the weaknesses of the system in which they were working and that they also did much that was good. Tony Giddings, who is a consultant surgeon, a non-executive Director of the National Clinical Assessment Authority and Surgical Advisor to the NHS Modernisation Agency, said at a recent seminar on "Humans in complex engineering systems":

There was a watershed in the publication of the Kennedy Report. This is an enormous document which will take you a weekend to read, but which is well worth it if you are interested in the area. What we learned from Kennedy was the whole systems approach to medical safety and surgical performance. It identified shortcomings in all the

systems involved, in the infrastructure and indeed in the individuals concerned. In his words, 'good people doing bad things'.

The legal point is essentially the same as that made by the Columbia example – that it is not possible to pursue the people who permitted the system to remain defective and not right to pursue the people who worked within it.

Systems and fragmentation

All of the examples so far have concerned systems that were constructed with interfaces, rather than as a single element. There are additional complications when a previously integrated system is fragmented.

An example of the human consequences of fragmenting an integrated system appears in one of the leading textbooks on system engineering [10]. Although most of the examples given in that book are technical systems, it includes the following:

Historians have remarked that one of the more serious consequences of European colonisation was the drawing of maps showing national boundaries. Up to the time of colonisation, the political structure had been that of city-states in which authority of the state decreased with distance from the city. Midway between the cities, their authorities were the least. ... Boundaries had little meaning. With the drawing of maps and line boundaries, authority was abruptly made absolute and total right up to the line on the map. Distance no longer separated the states and new kinds of wars resulted.

The boundaries between two elements of a system may be ill-defined but this is of little importance if they are not subject to political or contractual disputes, or if there is little interaction between the neighbours. Once hard boundaries are defined, there must be absolute clarity as to the relationships between the neighbours.

This is particularly relevant where human systems had been integrated and are in transition to a fragmented set of interacting components. This was the case in BRI (NHS to Trust) and Columbia (NASA management to contractor) and clearly underpinned much of the challenge for UK railways during the 1990s (BR to private companies). It is remarkable how well people and organisations bear up under the pressure of change. The accident rate for UK railways was lower after privatisation than before (despite what one might read in the Press) and life expectancy and general health appear to have continued to rise in the fragmented health service.

The difficulties of fragmentation are compounded if accompanied by privatisation or competitive targets. Former public servants who are thrust into the competitive marketplace tend to over-react. They assume that being commercial means being aggressive and confrontational at all times. After a few years they realise that private industry is much more about cooperation than confrontation – today's competitor may be tomorrow's partner in a large project, and the customer and supplier have a shared interest in success the moment that the ink is dry on the contract. However, until

they appreciate this, the very time at which cooperation is needed to smooth the transition into fragmentation is when it is least available.

Fragmentation inevitably brings new players to the system. All industries, and companies within those industries, rely on tacit understanding – the unstated knowledge that everyone who works in the industry should have about how things are done. It is the knowledge that is given to apprentices or trainees by people with more experience. It becomes particularly important for safety if there is an unwritten understanding that something will not be done. The overview of decision-taking published by Rail Safety and Standards Board on behalf of the rail industry [11] gives a simple example:

... there might be a tacit understanding that a high tensile steel bolt is used where a mild steel bolt is specified. It could become normal practice to hang greater loads on that bolt than intended because "everyone knows that it's strong enough". A new company that joins the industry, knowing nothing of that tacit understanding, might decide to use a mild steel bolt since that is all that the specification demands.

System failures – a summary

Some very clear messages emerge from all of the examples:

- the true cause of a failure is often not that an individual or component has failed; the failure may be a symptom of a poorly designed or managed system and that system may be human or technical
- in other words, the failures occurred primarily because of the incorrect interactions of the components of the systems, not because the components themselves were defective
- it should be possible to avoid such failures by defining properly the interfaces between the components but even the best and simplest interface specifications may not be correct and exhaustive
- an exhaustive interface specification is less necessary when the same organisation is responsible for the components on both sides of it
- fragmentation of an existing system is particularly fraught, especially when subject to financial pressures and to a culture that discourages open communications.

4 Ownership – who owns the risk?

Nils Diaz, Chairman of the US Nuclear Regulatory Commission, said [12]:

I would like to start my remarks today with a story that you may have heard before. It has several versions but the point of the story always remains the same. It is a story about four people named Someone, Anyone, Everyone, and No one. There was an important job that had to be done - Someone should have done it, Anyone could have

done it, and Everyone thought that Someone would do it. In the end, No one did it. At the beginning, the responsibility was not assigned.

In the nuclear arena, we cannot afford the luxury of allowing a job to go undone or be poorly done.

In a seminar [13] on the safety of rail systems at the Health and Safety Executive's headquarters in London, the author presented the following slide:

Liability

- a system may fail even though all of its components work to specification
 - who is responsible for the failure?
 - who is liable for the failure?
 - who had the power to prevent the failure?

There was an awkward, shuffling silence. No member of the audience had a satisfactory answer. Let us look at each line separately.

Responsibility for the failure

Health and safety law in the UK is largely based on the seminal report of the committee chaired by Lord Robens in 1972 [14]. That report argued that complicated prescriptive standards should be replaced by a duty on each employer to strive to eliminate risks to workers and others, so far as is reasonably practicable. However, the report states in paragraph 182:

We accept that transport safety is a vast study in its own right, involving many technical problems of a highly-specialised nature. Provisions for the safety and health of those engaged in flying aircraft, driving trains, lorries and so on clearly cannot be considered in isolation from a whole complex of special considerations such as the constraints imposed by the design of transport vehicles; the circumstances in which they operate which include many eventualities beyond the control of an employer; and the predominant need—in terms of numbers at risk—to safeguard the travelling public and the public generally. We accept that these matters must be dealt with within transport legislation.

Paragraph 475 of the report summarises the conclusion:

The legislation [the Health and Safety at Work etc Act 1974] should not apply to the normal use of the highway, to domestic service, or to transport workers whilst actually engaged in transport operations.

Lord Robens and his committee understood that it was not appropriate to hold one person responsible for failures of a system over which he does not have control. It is often

impossible to isolate a single cause from the set of interacting elements of, for example, a transport system.

Liability for the failure

The law seeks to find people to blame, either so that they may be punished or so that they may compensate the victims. This is not a new approach – Robert Hubert was executed for “causing” the Great Fire of London in 1666, even though it was subsequently shown that he was not even in England at the time and the true cause lay in the overhanging wood-framed buildings and narrow streets. Blaming individuals under the circumstances of a system failure is unlikely to be just (although the prevailing jurisprudence of the 17th century would see no injustice in executing a Frenchman).

In a review of error management in aviation and medicine, Helmreich [15] concludes: “... *there is seldom a single cause, but instead a concatenation of contributing factors*”. He cites a fatal medical error which, on its face, arose because of negligence by the anaesthetist but in fact was the result of nine sequential failings of the surgical system.

The longer version of Nils Diaz's story has a further line:

Everyone blamed Someone when No one did what Anyone could have done

It is this blame culture that the Columbia Accident Investigation Board and the BRI Inquiry sought to avoid, because they recognised that it would mask the true causes of the accident.

Liability can only be assigned if, in the management of critical systems including those that are safety-critical, every risk has an “owner” and every task is defined. Where they are not, the law fails (and, arguably more importantly, so does the management of the system). Criminal law has difficulty finding someone guilty for a system failure unless the guilt is determined by strict liability, either in theory or in practice (the practical consequence of s40 HSWA). Civil law similarly finds it hard to find liability, as illustrated by the cases in negligence with multiple causes and the failure to find anyone to sue for the defective factory automation system.

Power to prevent the failure

It would appear that, in many systems failures, no-one had the power to prevent it because no-one owned the system, or more specifically the hazards that lie in the interactions between its elements. This paper now shifts its direction, from cataloguing the woes of engineering and the law to proposing a course of action for engineers and lawyers that should lead to fewer system failures by ensuring that someone does own those hazards.

5 Safe system management

There are three critical steps that engineers and lawyers can take to reduce the risk of failure of safety-critical systems:

- learning from experience by conducting proper investigations after accidents and incidents that might have developed into accidents

- appointing and empowering system architects and holding them to account
- aligning the legal framework with the system and the responsibilities of the players.

Investigation - inquiry or inquisition?

There is a Spanish proverb – only a donkey makes the same mistake twice and even he learns by the third time. Why do we hold an investigation after a failure of a safety-critical system – to find out who was to blame or to prevent it happening again? Learning from experience requires a proper inquiry. The report of the investigation into the loss of Columbia had no doubt has this should be done:

Many accident investigations make the same mistake in defining causes. They identify the widget that broke or malfunctioned, then locate the person most closely connected with the technical failure: the engineer who miscalculated an analysis, the operator who missed signals or pulled the wrong switches, the supervisor who failed to listen, or the manager who made bad decisions. When causal chains are limited to technical flaws and individual failures, the ensuing responses aimed at preventing a similar event in the future are equally limited: they aim to fix the technical problem and replace or retrain the individual responsible. Such corrections lead to a misguided and potentially disastrous belief that the underlying problem has been solved. The Board did not want to make these errors. A central piece of our expanded cause model involves NASA as an organizational whole.

Similarly, Professor John Overveit [16] analysed medical incidents and showed that 85% are due to organisational failures and 15% due to individual failures. The individual failures are because of inattention to detail, lapses of memory, negligence, forgetfulness and carelessness. This was corroborated by Wilson [17] who also agreed with Overveit's overall findings that remedial action often suited the hierarchical management style. Wilson found that 98% of remedial action was concentrated on the individual failures and only 2% focused on the organisational failures.

There are systematic tools to help to analyse the underlying causes of system failures. One such, SMART, has been applied to a wide range of systems, from TCAS airborne anti-collision systems to the public water system in the province of Ontario, which contains computers, hardware, human operators, management decision-making, and government regulatory components. Leveson [18] described the need for SMART as follows:

Traditional approaches to hazard analysis and safety-related risk management are based on an accident model that focuses on failure events in static engineering designs and linear notions of causality. They are therefore limited in their ability to include complex human decision-making, software errors, system accidents (versus component failure accidents), and organizational risk factors in the analysis. These traditional accident models

do not adequately capture the dynamic complexity and non-linear interactions that characterize accidents in complex systems, i.e., what Perrow called system accidents. System accidents often result from adaptation and degradation of safety over time: The move to a high-risk state occurs without any particular decision to do so but simply as a series of decisions or adaptations (asynchronous evolution) that move the system into a high-risk state where almost any slight error or deviation can lead to a major loss.

To handle this more comprehensive view of accidents, risk management tools and models need to treat systems as dynamic processes that are continually adapting to achieve their ends and to react to changes in themselves and their environment. Leveson's new accident model, STAMP (Systems-Theoretic Accident Modeling and Processes), provides the foundation for such a risk management approach by describing the process leading up to an accident as an adaptive feedback function that fails to maintain safety constraints as performance changes over time to meet a complex set of goals and values.

Constructive investigation is the first step on the road to proper management of safety critical systems, a step shared by the engineers and the lawyers. The engineers have the expertise and insight to conduct the analysis, either informally using their general understanding or formally with a tool like SMART. The lawyers have the forensic skills and experience of disentangling complex arguments which, together with the framework of legal process, allows robust conclusions to emerge and be authenticated.

Proper inquiries that are designed to find causes, not scapegoats, can shed light and learn lessons and are within the scope of the two professions.

System architects

Systems will only be safe if they are designed, not if they emerge by chance. Someone must be responsible for establishing the system architecture – deciding what is the problem that must be solved, how to decompose it into its parts and how to put them back together again. Those are three distinct tasks.

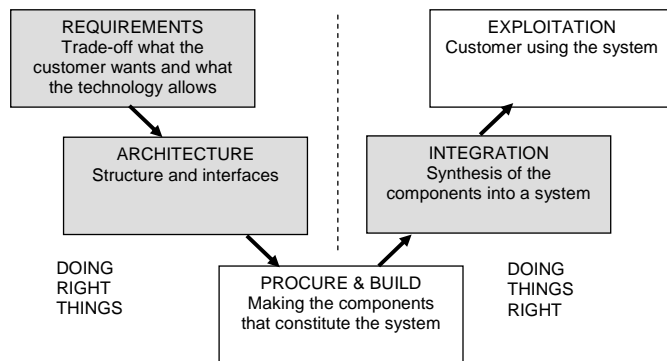
The first is necessary to help the customer get away from specifying his requirements in terms of candidate solutions. NASA realized that Apollo astronauts would need to write notes so specified a pressurized ball point pen that would work without gravity. It was built at great expense, only for NASA to realize that the Russian cosmonauts used pencils. Working out the requirement is an iterative task, helping the customer to get the best out of credible technology without aiming so high as to be unaffordable, in time or money.

The second task requires clear insight into how the task should be decomposed. Lilleniit [19] described the system architect as:

... the person who can look at a jewel and hit it just the right way so it falls into the right number of pieces. It is that ability to decompose in just the right way

The architecture is defined by its interfaces – a robust design has clear and logical interfaces that minimize the interactions between the components on the two sides. The architecture is as far as possible an overview of the system, free from details of the implementation.

The third task requires deep understanding of how the technologies that have been used in the components interact. It therefore needs good knowledge of the technologies, wide experience of how they can be used and, crucially, how they can fail, and the personal skills to manage the sensitivities of the engineers responsible for each part.



The three tasks are illustrated in the simplified version of the classic V-diagram above, where the grey boxes are the responsibility of the system architect. The drawing is not to scale; the lion's share of the work and expenditure still lies in the white boxes but the architect's contribution has a disproportionate effect on system performance and hence safety.

The RAEng report from which the Lilleniit quotation was taken continues:

This role has been forced on projects due to the sheer breadth of hardware, software and communications options now available. Only a truly experienced and knowledgeable individual can harmonise the selections into an effective whole. A skilled architect will also produce a design which is robust, scalable and evolvable. Experienced architects should be able to incorporate sufficient flexibility to accommodate the changes in specification that generally arise during the course of projects without introducing unnecessary complexity which could compromise the integrity of the design. Ideally the scope for evolution should extend beyond the project in hand to encompass future projects or products.

Systems architects possess the exceptional conceptual skills required to translate a business vision into a technical blueprint, which should ideally be expressed in a notation that supports formal analysis and reasoning. Moreover, systems architects must have the breadth of human and organisational understanding to address the underlying organisational and motivational issues that can critically impact on project success.

A competent system architect given the resources and authority to act can transform a system from one that is intrinsically unsafe (in the sense that no amount of effort by individual players can ensure safety) to one in which human, organisational and machine weaknesses are recognised, managed and accommodated.

Legal framework

Lawyers too have three crucial roles to play in ensuring system safety, in defining the contractual relationships, avoiding inappropriate blame and holding system architects to account.

The first task lies in adopting contracts that encourage cooperation and in which contractual interfaces align with the logical engineering interfaces. The need for cooperation emerged from the report of the Egan Task Force [8], which found that the fragmentation of the construction industry was reinforced by routine use of competitive tendering for procurement and frequent recourse to the courts to settle disputes. The relationships between the companies (along the supply chain and between customer and suppliers) that need to work together were essentially confrontational but, as a result of Egan, new forms of contract are being adopted which are built on partnership. They recognise that the customer and supplier have a shared interest in the success of the project and that a long term relationship provides a greater incentive to succeed than the threat of litigation. Although Egan's recommendations were driven primarily by commercial considerations, safety is equally improved by a cooperative relationship.

Aligning the contractual and technical interfaces seems obvious but is too often overlooked. Parties should be contractually liable for that which they can control, without relying on complex interactions across a contractual and technical boundary.

The second challenge for lawyers, both prosecuting authorities and litigators, is to refrain from pursuing individuals for system failures. Criminal prosecutions are unlikely to succeed but cause great distress and encourage "defensive engineering" where the engineer chooses the course of action that minimises his chance of criminal charges rather than that which is best for the client. Civil claims become mired in extensive and expensive actions; in the majority of UK civil actions the legal costs exceed the damages eventually awarded and much greater benefit can result to both parties by using mediation rather than litigation.

The third challenge is to develop the law so that both criminal and civil actions can be directed against those who have the power to prevent systems from being unsafe – the system architects and their political masters or clients. This may be uncomfortable for engineers but it is the logical consequence of calling for system architects to be empowered. Society wants safe systems and it wants to be able to hold those who cause a lack of safety to account. System architects, it was argued earlier, are the people who can deliver system safety, so must accept the legal consequences.

This is no more onerous a burden than that placed on any professional whose actions directly and culpably affect the safety of the public. Like dentists, structural engineers and lawyers, they would need to carry insurance for their civil liabilities; they would remain criminally liable for their actions.

The analysis earlier in this paper touched on the difficulty that civil law has with finding liability for harm arising in complex situations, where the cause and the type of harm may only be attributed statistically. Although two of the five Law Lords in *Greg v Scott* [4] believed that it was possible to reinterpret the existing law effectively to look at expectation values, the majority found that this was not possible. Lord Hoffman believed that this would be “a radical departure from precedent” and concluded:

But a wholesale adoption of possible rather than probable causation as the criterion of liability would be so radical a change in our law as to amount to a legislative act. ... I think that any such change should be left to Parliament.

A proper legal regime for the regulation of safety in systems therefore requires fundamental rethought and might be best addressed by the Law Commission.

6 Summary of conclusions

Both engineers and lawyers have difficulty ensuring system safety. They both have a role to play in constructive investigation of system failures. Specialist engineers need to act as system architects and lawyers need to reinforce their role by aligning contracts with system interfaces and holding system architects to account.

These are grand and far reaching actions that cannot be easily achieved but they provide a framework for identifying incremental improvements that will reduce injustice and, most importantly, improve system safety.

References

- [1] *R v OLL Ltd*, 1993
- [2] *Kuwait Airways v Iraqi Airways* [2002] WLR 1353 at 1388 per Hoffman L
- [3] *Wilsher v Essex Area Health Authority*, HL
- [4] *Gregg v Scott* [2005] UKHL 2
- [5] *Fairchild v Glenhaven* [2002] UKHL 22
- [6] Hoffman L in *Gregg v Scott* at para 79
- [7] *Thames Trains v HSE* [2003] EWCA Civ 720
- [8] “Rethinking construction”, DETR, July 1998
- [9] “Learning from Bristol: the report of the public inquiry into children's heart surgery at the Bristol Royal Infirmary 1984 -1995” Command Paper: CM 5207, July 2001
- [10] “Systems architecting: creating and building complex systems”, Reichtin E, Prentice-Hall 1991, ISBN 0-13-880345-5
- [11] “How safe is safe enough?”, RSSB, February 2005
- [12] Nils K. Diaz, Chairman, US Nuclear Regulatory Commission, CNRA Regulatory Industry Forum, Paris, 17/6/04
- [13] Elliott, C J, “Safety-critical railway systems - a technical, legal and management perspective” HSE Chief Scientist’s Seminar, 17 December 1999
- [14] “Safety and Health at Work”, Report of the Committee 1970-72 Cmnd. 5024, July 1972
- [15] Helmreich R L “On error management: lessons from aviation” *BMJ* 2000; 320:781-785
- [16] Overveit J, “Health Service Quality”, Brunel University, 1998
- [17] Wilson J “Learning from healthcare mistakes”, HCRR Sept 2000 p14
- [18] Leveson N G “Model-Based Analysis of Socio-Technical Risk”, Technical Report, Engineering Systems Division, Massachusetts Institute of Technology, June 2002 <http://sunnyday.mit.edu/papers/stpa-tech-report.doc>
- [19] Lilleniit H, quoted in “The Challenges of Complex IT Projects”, RAEng ISBN 1-903496-15-2