



David Slater takes a look at how risk assessment has evolved and argues that today's more complex, interdependent systems need a new 'top down' approach

Risk assessment: from the TOP

Successful risk assessment methodologies have been developed over the years to 'model' the behaviour of increasingly more complex systems, from space shuttles to nuclear power plants. They look at the effect of failures of individual components and human interventions on the overall behaviour of the system. This reflects the natural way we look at accident scenarios (ie. cause and effect).

One of the first formal methodologies (failure mode and effects analysis – FMEA) simply listed the effects for every failure mode that could be thought of, at a component level. This approach is even used by classical Boolean fault tree analysis techniques. First a set (not necessarily complete) of top events ('failures') is identified and it then requires the complete underlying fault 'tree' structures for their analysis. These 'bottom up' methods also often used a probabilistic (mainly from the historical incident and failure record) basis for estimating the reliability and risk of failure of these top events.

As systems have become more complex, this approach becomes more difficult to use and unreliable in its predictions and insights (if it hasn't happened, we have no data and no rigorous basis for identifying all the 'top events'). Nowadays, with the emphasis increasingly on the interactions, behaviours and interdependencies of 'systems of systems', the limitations of the classic 'bottom up' approach means that the simplifying short cuts and assumptions needed to make the construction of the extensive logic networks tractable seriously challenge the integrity of any results obtained.

It's therefore timely and necessary then to look again and see if there is a better way. Instead of trying to build suitable logical and quantitatively accurate descriptions of complex systems of systems from the bottom up (component by component, system by system) possibly being unaware of missing, parallel branches and whole areas of interdependencies (common failure modes), Cardiff University start up company Intradependency has developed an alternative 'top down' approach which identifies only the criteria for successful operation first and then drills down (logically and quantitatively) into the underlying structural complexity to critical and common dependencies only as far and wide as required (a just-in-time analogy).

Our analysis method then 'stress tests' these dependencies on a *what if?*, *how much?* and *most likely?* basis for discovering critical dependencies, failure modes, their probabilities and knock-on effects.

α brief history of risk

Predictions of risk have been around as early as 3200 BC, where they relied on reading 'signs from the gods' but the first reference to more rational risk assessment goes back to Ancient Greece, where, during the Peloponnesian War in 431 BC, the Athenian leader, Pericles, gave a funeral oration which, among a list of Athenian virtues, included, "we are capable at the same time of taking risks and estimating them beforehand."

In more recent times, the more formal discipline of estimating risk (risk analysis) evolved originally from military logistics and its demands for predicting reliability. These techniques were also adopted for similar purposes by NASA and the Nuclear Industry (in the 1950s), where the first report on nuclear power plant accidents (*WASH-740*) was issued. The consequences predicted were often unacceptable, but it was believed that the probability of such an accident was very small. Many of the classic techniques such as failure modes and effects analysis (FMEA), fault tree analysis (FTA) and event trees, as well as significant sources of actual failure rate data, evolved from the demands on the nuclear industry to demonstrate safe operation.

These techniques concentrated on failure likelihoods and neglected the consequences, presumably on the basis that any failure was undesirable and to be designed out. The Flixborough incident of 1974 (where an unconfined release of cyclohexane vapour exploded and levelled the NYPRO nylon works near Scunthorpe, UK, killing 28 people) underlined the importance of the extent of consequences and highlighted an industry-wide under-developed approach to formal risk assessment and explosion modelling. Some 36 years later, the explosion mechanism at Buncefield is still, apparently, a mystery!

The Netherlands government was the first to try and remedy this lack of expertise, and embarked on a programme to standardise risk assessment techniques and also to provide a standard toolkit of consequence models. This was incorporated into an approved methodology for planning and community safety legislation and proved very successful in regulating hazardous installations and transportation in a spatially-challenged environment.

After the 1980 Kjelland accident (where a pentangle offshore drilling rig broke up and sank in rough seas with no survivors), the Norwegian government regulator (Norwegian Petroleum Directorate, NPD), also enshrined a quantitative 'residual risk' methodology in law as the required basis of offshore safety cases for obtaining operating licences in the North Sea.

In the 80s, the growing recognition that there was more to reliability than equipment failure generated a substantial 'human factors' contribution.

Finally, in the UK in the early 80s, the Sizewell B inquiry was exercised on the safety of the pressurised water reactors (PWRs) derived from submarine experience. And the Health and Safety Executive (HSE) published its view on the *Tolerability of Risk*.

In the 80s, the growing recognition that there was more to reliability than equipment failure generated a substantial 'human factors' contribution. The problem was that analysts were still looking for numbers to quantify the 'and/or' gates and branch probabilities of their logic trees, rather than being open to the more complex subtleties being demonstrated.

It was therefore largely on this incompatibility of hard and soft techniques and the increasingly significant contributions from social scientists, psychologists and ergonomists, that the development of purely mathematical logic trees peaked – and stuck.

The resources needed for rigorous consequence modelling and frequency estimation are significant, and the value of the extra precision over more qualitative approaches began to be questioned. This led to the evolution of a logic diagram approach, which, although it had the FTA and event tree wings, did not need their detailed internal structures. These 'bow tie' diagrams became increasingly adopted for offshore safety cases, producing a clear picture of the issues at significantly less resource and expertise level requirements.

The problem with non-numerical risk assessments was underlined during the UK's BSE crisis of 1992–97, where the Medical Spongiform Encephalopathy Committee had only qualitative expert judgements to offer the regulators, who needed quantitative criteria in order to control the abattoirs.

More recently, health and safety regulations have made the demand to perform risk assessments a legal obligation. This ubiquitous requirement has meant that the methodologies recommended for use have to be over simplified to the point where they are simply a checklist for awareness raising – a reasonable goal in itself, but at the expense of the status of the more rigorous and searching risk disciplines?

For example FN curves have morphed into coloured risk matrices, 'semi'-quantified,

on dubious mathematical grounds. As responsible engineers, we should not continue to oversimplify because we can't get a handle on the complexities and counter intuitive behaviour of complex systems. We need to address formally, rigorously and if possible, quantitatively, all those difficult and 'human' factors that complicate the sums.

it all depends!

Many existing methods of risk analysis also need the analyst to think through what could go wrong. This approach is often unproductive and hard to carry out, because thinking of what can go wrong is inherently counter-intuitive, not to say depressing, especially for projects' proposers and managers.

There is considerable debate about the definition of risk, anyway. In the classical approach it needs to combine and convey the important aspects (dimensions) of likelihood and consequences. Since then, almost every new field opened up to a risk-based approach has reinvented its own version. When you add the myriad colloquial (non technical) and traditional uses, it does seem to be pointless in insisting academically on one, and only one, definition; and we make no apology for introducing yet another vernacular terminology.

Most (non academic) people would agree with a meaning of risk as 'the degree to which things outside our control threaten the things we hold dear', and the way we propose to use it is merely a more precise rendering of that sentiment.

The term 'goal' is thus used to indicate something we hold dear or want to achieve. It may include something we already have and the goal might be to keep things that way.

There are major goals and minor goals. We can speak of top-level or strategic goals such as keeping the economy afloat, maintaining the national infrastructure or establishing a trade agreement with China. More tactical goals might be keeping a given power station online, ensuring terrorists don't gain access to a particular building, or maintaining supplies of vaccine in a certain geographical area. Lower-level goals might include ensuring a particular vehicle is maintained, ensuring that a key worker has an understudy, or ensuring that an insurance

policy is renewed. There's no real lower limit to the level of goals we can define and we could consider the integrity of a certain electrical fuse or the serviceability of a particular mobile phone as goals too.

The insight on which our 'top down' approach pivots is that each of these goals depends upon other things which (if we take care of our definitions) are nothing less than other goals, often at a 'lower level'. From a risk standpoint, the whole structure of any organisation, from a sovereign state down to a village fête consists of nothing less than a fractal-like network of goals each depending on lesser goals which depend on even lesser goals, and so forth all the way down as far as we care to delve.

the problem with complex interdependent systems

These individual systems (such as the National Infrastructure) are "so complex that no detailed dynamic simulation exists" (Fisk, 2006). Fisk reviews an increasing literature on the societal implications of increasing the scale and complexity of systems from ecology to civilisation.

Our 'top down' approach, however, models only the results of critical dependencies and how these are pragmatically affected by the status of other systems. It does not need to understand the mathematical intricacies of individual system stabilities, merely the outcomes of interactions.

As Karl Weick pointed out, complex systems always retain the capacity to produce novel, or surprising events. Reliability or redundancy among safety-critical elements is not sufficient to achieve safety. You also need "some flexibility and responsiveness to pick up things that haven't happened before, ie. the inevitable and constantly-emerging anomalies".

Accidents, using a 'system of systems' approach, can therefore be thought of as the result of flawed processes, involving interactions among system components, or between interdependent systems themselves. These can include people, societal and organisational structures, engineering activities and physical systems and components. These interactions can be spatial, temporal, or systemic; designed, or inadvertent, associations. To prevent these occurrences, it consequently becomes vital to understand how systems of

systems can experience a 'drift' (Rasmussen Drift) in performance by changes in these interactions with time. This can be due to external pressures (vulnerabilities), or influences (dependencies) acting at various levels throughout the entire socio technical, regulatory and natural environments. The recent global financial crisis has been analysed as a classic example of this drift.

Unexpected perturbations (accidents) in complex (stiff) systems have been described as 'normal' accidents (Perrow, 2000), or 'non-linear accidents' (Hofnagel, 2008). They can occur from interactions (stochastic resonance) among components, each of which is functioning as intended, but where the independent decisions and organisational behaviours interact in dysfunctional ways. (Hofnagel's Functional Resonance Accident Model).

α solution

It's clear that in contrast to the bottom-up individual component focussed cause and effect emphasis of early risk approaches, it's more important to understand this wider picture of systemic interactions and escalations.

Our method borrows heavily from dependency analysis, bayesian networks, failure mode analysis, database management and geographic mapping techniques. It determines disaster scenarios and numerical measures of risk. It is scalable and as it's based on (positive) goals and objectives rather than (negative) weaknesses and disasters directly, it takes account of the critical dependencies of those objectives.

The approach demands a way of modelling dynamically, on a just-in-time basis, the implications of the (real time), dependent behaviour of interacting complex systems. Not by looking at local system failures, but by asking for each system the question *what does our successful operation depend upon?* It should be noted that the emphasis here is on the goal or objective of the successful operation of the system and not on the potential for component failures, ie. *what if this fails?* It is aimed at facilitating and encouraging a more rigorous approach to predicting and assessing risk; and to give valuable insights into hitherto counterintuitive and perplexing behaviours of modern engineering systems of systems.

We're currently in the process of developing software to implement this approach – so watch this space. **tce**

