Chapter 6

Resilience Engineering

Gordon Baxter, University of St Andrews

Summary

Resilience engineering is concerned with building systems that are resilient to change. In other words systems that continue to work, often through the results of human endeavours, when faced with adverse situations (both anticipated and unanticipated). The work grew out of the safety engineering community around the end of the 20th century. There is a strong relationship between resilience and dependability, with resilience being described as the persistence of the dependability of a system (or organisation) when facing changes. The discipline of resilience engineering focuses on three main areas: developing tools and techniques to assess how organisations achieve resilience in their particular domain; on improving organisational resilience; and on modelling and predicting how organisational change and decision making affects risk and resilience.

Background

Historically, safety engineering has focused on the negative aspects of systems, and tried to achieve failure rates that are as low as reasonably practical (such as 1 failure in every 10,000 events or 10-4). In this view, a system is perceived as being made safer if the number of adverse events is reduced. This approach to safety engineering analyses what goes wrong, looking for failures and malfunctions, and then tries to prevent recurrences by eliminating causes and putting appropriate barriers in place. Resilience engineering was developed to take a

more positive view of safety. In the resilience engineering view, safety is regarded as the ability to succeed under varying conditions. Systems succeed, far more often than they fail: a failure rate of 1 in 10,000 events, means a success rate of 9,999 in 10,000 events. The resilience engineering approach therefore analyses why things go right, and uses that as a basis to understand what counts as normal performance, so that work can be made better and safer. Work situations are invariably underspecified and therefore not completely predictable, so resilience engineering looks at issues to do with performance variability which is not only necessary (to deal with the changing situations) but also inevitable because of the inherent variability of people, organisations, contexts and technology.

The Trade-off Between Efficiency and Thoroughness

The Efficiency-Thoroughness Trade-Off (ETTO) Principle was formulated to help explain why things that often go right can sometimes go wrong. It is not really a new principle, it is more a way of integrating lots of similar work together under a single unifying umbrella. Examples of the ETTO appear to be ubiquitous. Efficiency is achieved when a particular goal (or objective) is attained at minimum cost (time, effort, resources and so on); thoroughness involves carrying out a detailed analysis that allows one to be confident that the current conditions will lead to some desired activity being successful and having no unwanted side-effects. When the balance between efficiency and thoroughness is achieved, successful performance results. If the balance tips too much towards efficiency this can lead to wrong actions being performed (through lack of analysis of the situation); if the balance tips too much towards thoroughness this can lead to actions not being performed too late to be effective, because so much time has been spent on analysing the situation. In order to manage safety, it is important to understand how the balance between efficiency and thoroughness is realised. ETTOs can happen for several reasons:

- Scarcity of resources, particularly time, or uncertainty about the amount of time.
- The inherent human trait of following the line of least effort.
- A need to keep something in hand (reserves of resources, or time) to handle unanticipated situations.
- Peer pressures to do things in a particular way or to meet a specific deadline.

- Organisational pressures, such as the conflict between priorities (safety first) and practices (be ready on time).
- Individual characteristics, such as personal priorities, working habits, and personal ambitions.

The trade-off is a heuristic one that applies to people and to organisations. It can only be made by machines when it has been included in their implementation (embedded in the software), and in such cases is algorithmic, rather than heuristic.

The Functional Resonance Analysis Method (FRAM)

The FRAM was developed as a way of describing the performance of socio-technical systems. It regards variability as being inherent in normal performance, and uses this to explain why accidents happen: performance variations can lead to positive as well as negative outcomes. Shortcomings in performance, however, are linked to variability in complex relationships, so they cannot be adequately described using simple linear models. Some adverse events can be attributed to the breakdown in normal functions, but generally they are best understood if considered as the result of the combination of several sources of variability in human performance. The FRAM is built on four principles:

1. **The principle of equivalence of successes and failures**. People and organisations continually have to adapt to the current conditions. When these adjustments are made correctly and when failures and potential harms are correctly anticipated, this leads to success; when this ability to correctly make adjustments is absent, failures can result.
2. **The principle of approximate adjustments**. Work situations are invariably underspecified and hence partly unpredictable. Individuals, groups and organisations have to adjust their performance to suit the prevailing conditions. These adjustments are approximate because resources (time, information etc.) are scarce.
3. **The principle of emergence**. The variability in several functions can combine in unanticipated ways, giving rise to consequences that are disproportionately large, and produce non-linear effects. Performance (successful or otherwise) is emergent, rather than resultant.

4. **The principle of functional resonance**. When the variability of several functions resonates, this can cause the variability of one function to exceed its normal limits. These consequences can dissipate through tight couplings rather than well-defined cause-effect links.

Cause-effect models have traditionally been at the core of safety engineering. These structural approaches (such as MTO: (hu)man, technology, and organisation) can be used to provide analyses of complex situations, but the analyses are necessarily linear because they are based on simple direct relationships (cause-effect). The FRAM focuses on the system dynamics (and variability) rather than modelling individual failures, and hence can be categorised as a systemic analysis approach (like Nancy Leveson's Systems-Theoretic Accident Models and Processes, STAMP). These more holistic approaches describe events as coupled functions, with links between functions showing dependencies, rather than cause-effect relations.

Retrospective

The field of resilience engineering is still a relatively new one, and it is continuing to develop. The move towards a more systemic, functional approach to understanding system performance, rather than a structural approach reflects the need to find new ways to deal with the ongoing rise in complexity in systems. Models like the FRAM (and STAMP) appear to offer much promise in this area, and have been successfully used in domains such as healthcare, aviation, and finance. The models continue to be developed as they attract more and more users. As organisations and nations increasingly focus on critical nature of networked (and national) infrastructures, the need for resilience engineering methods and tools continues to grow. There is an active and growing resilience engineering community, centred around MINES, ParisTech (in the south of France), which significantly overlaps with the cognitive systems engineering community. They run regular conferences on resilience engineering, and a training school for use of the FRAM.