

School of Computing Science,
University of Newcastle upon Tyne



Human compensations for undependable systems

Denis Besnard and Gordon Baxter

Technical Report Series

CS-TR-819

November 2003

Copyright©2003 University of Newcastle upon Tyne
Published by the University of Newcastle upon Tyne,
School of Computing Science, Claremont Tower, Claremont Road,
Newcastle upon Tyne, NE1 7RU, UK.

HUMAN COMPENSATIONS FOR UNDEPENDABLE SYSTEMS.

Denis Besnard

CSR, School of Computing Science
University of Newcastle upon Tyne
Newcastle upon Tyne NE1 7RU
United Kingdom
denis.besnard@ncl.ac.uk

Gordon Baxter

Department of Psychology
University of York
York YO10 5DD
United Kingdom
G.Baxter@psych.york.ac.uk

Abstract: Randell's (2000) dependability fault-error-failure model was originally designed with the objective of describing the propagation of faults in technical systems. Conversely, Reason's (1990) swiss cheese model was intended to describe the organisational facet of systems' failures. However useful these two views have been, there has not been a lot of effort devoted to highlighting their common features. Moreover, these two models say little about the positive human contribution to the delivery of an acceptable service with undependable systems. The investigation of these two aspects forms the main focus of this paper.

Our first objective will therefore be to integrate the two models. In doing so, we will also provide an answer to the problem of scale in the description of events in complex settings: organisational factors and pure technical causes could be integrated in the same descriptive picture. Our second objective will be to show that the dependability of the service of socio-technical systems is often a matter of human compensations for poorly designed systems. This claim will be supported by three concrete examples where human compensations have permitted a partly-automated system to deliver an acceptable service.

Keywords: Dependability of service, socio-technical system, adaptations, workarounds, violations.

1 INTRODUCTION

It has been a classical view that catastrophes in large socio-technical systems (STSs) in general, are caused by some sort of human error. As a consequence, there are still societal pressures for the identification of one single person responsible for a given event. This has been shown to be the rule in medical domains, for instance (Reason, 2000; Svenson *et al.*, 1999). A somewhat more modern conception is the idea that failures in STS are the result of a combination of factors meshed into a complex causal network spread over several hierarchical levels within an organisation (Reason, 1990; 1997). In this view, the responsibility lies in the upper hierarchical levels where flawed decisions can have dramatic knock-on effects at lower levels, as happened in the explosion of the Challenger space shuttle (Rogers, 1986). In parallel to this evolution, the idea that became predominant was that perfection could not be achieved, neither by designers nor users. Somewhat surprisingly this idea did not have much impact on the classical dependability model. This is one area where this paper is hoping to make some progress.

Dependability can certainly be assessed at a purely technical level (e.g. probability of failure per demand for a given software module). However, computer systems and technical artefacts in general are being more and more tightly integrated with human activities. It follows that the service delivered goes far beyond the mere technical correctness. So much so, in fact, that it is now worth taking a fresh look at the concept of *service* dependability. In other words, one needs to consider the use of the technology and how this use moderates the level of dependability of the final service. Without such a view, technical arguments will have a limited validity. For

instance, the certification of a piece of code at a probability of 10^{-3} failures per demand only captures a tiny portion of the dependability of the final socio-technical system.

Within this context, we believe that human operators play an essential role in the final dependability of a service. The latter can obviously be degraded when, for instance, front-end operators err, perform dangerous violations (Reason, 1990) or cannot cope with exceptional circumstances because of a design flaw (e.g. see the Plugger accident, Loeb, 2002). But this angle only reflects one side of the coin. As Amalberti (1996) states, these errors are the side-effects of a cognitive system that achieves a correct performance most of time. A part of humans' positive role has to do with compensating for unanticipated or adverse conditions. There is nothing really new, here. This notion has been around in the cognitive ergonomics community since the work of Bainbridge (1983). What is a more recent development is that dependability must not be about just the technical components. Instead, it must be about the service, as delivered by the joint human-machine system. Using this perspective allows the human compensating role to be captured and factored into the dependability equation.

In section 2, we highlight the compatibility between two classical models of system failure: Reason's (1990) swiss cheese model and Randell's (2000) fault-error-failure model. In doing so, we will try to convey the idea that technical and organisational issues need to be simultaneously considered to capture the causal mesh leading to mishaps and catastrophes. We then describe some examples of human compensations (section 3) and suggest that they should be incorporated into dependability models (section 4).

2 TWO MODELS OF SYSTEMS FAILURES

In this section, we develop an integrative representation of Reason's and Randell's models. The intention is to focus on the dual consideration of technical and organisational aspects when analysing socio-technical system failures. We begin with a brief recap of Reason's and Randell's models.

2.1 Reason's swiss cheese model

In Reason's (1990) conception, accidents occur when the imperfections (represented as holes) that are inevitably present in the functional layers of a system, are put in alignment, leading to a dangerous system's state to breach its defences (see Figure 1). One of Reason's major contribution to the understanding of catastrophes is that events propagate from the highest organisational layers (managers) down to the lowest layers (front-end operators). In Reason's view, managerial decisions and policies force the workers to find illegal ways to do the job (see the JCO example in Furuta *et al.*, 2000). In doing so, they act in an unprotected manner, leaving room for an accident. When several actors implement this dangerous -yet unavoidable- strategy, they implement safety breaches that will stay latent in the system until the required accidental conditions happen, i.e. when a series of holes in different layers become aligned.

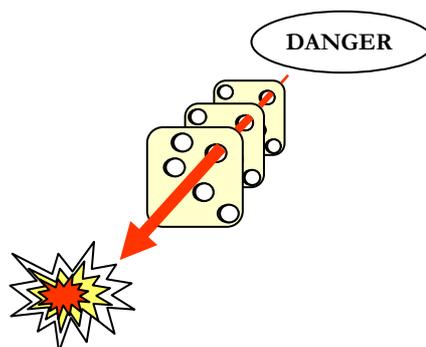


Figure 1: Reason's (1990) swiss cheese model.

One of Reason's concerns has been to focus on the organisational causality of mishaps. He opts for a holistic view in which the legal perspective of blaming the front-end operator (see Svenson *et al.*, 1999) is seen as an obstacle to safety improvements (Reason, 2000). One of the key attributes of this model is that each of the contributing factors is seen as necessary but not sufficient on its own to cause the occurrence of a mishap.

2.2 Randell's fault-error-failure model

Randell's (2000) dependability model uses a similar architecture to Reason's. A fault is inserted in a system during its creation. This fault then creates an undesirable system state (error) which remains dormant in the system. A failure occurs when the error is involved in the performing of a given function. This simple model can be stretched back in the causal chain as far as is needed, and whether a given piece of data has to be considered as a fault or a failure is only a matter of scope. For instance, a computer program crashes (failure) because a bug (error) was not fixed (fault). If one needs to, the presence of the bug itself can be considered as a failure to provide a bug-free program, the cause of which can be traced back to bad programming practice (fault), and so on.

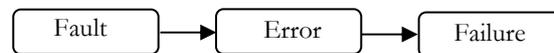


Figure 2: Randell's fault-error-failure model.

Roughly summarising Randell, for failures not to happen, systems must exhibit a number of dimensions (availability, integrity, security, reliability, confidentiality and maintainability). Ideally, a system exhibiting an appropriate implementation of all these dimensions is said to be dependable. The main criticism of this model is that the human and organisational aspects are hard to describe and analyse within this framework. Instead of attempting to directly refute this criticism, it may be of some value to try to find a way to back the model with an existing framework and attempt to extend the scope of the fault-error-failure model beyond its initial area of application. This is attempted in the next section.

2.3 Integrating Reason's and Randell's models

There are strong common ideas between these two models.

- *Systems can be decomposed into layers.* Each layer represents a sub-system, a state or an actor that has an impact on the functioning of the entire system.
- *Failures wait for calling conditions.* Some unstable conditions can be present in a given system without having any immediate effect. A failure, from this point of view, is an unlikely combination of a number of contributing factors. The analysis of major catastrophes (Wagenaar, 1987; Mancini, 1987) support this view.
- *Events propagate.* Accidents are not caused by the occurrence of sudden unfavourable circumstances. Instead, they are generated by early design faults which, under certain conditions, trigger an undesired event.
- *Events escalate.* A combination of local failures accounts for the breakdown of full systems.

As depicted in Figure 3, we believe that each organisational layer invariably contains one or more holes which can be attributed to the occurrence of fault-error-failure chains during its creation or functioning. One then gets an elementary failure generation chain for a hole in a given system's layer.

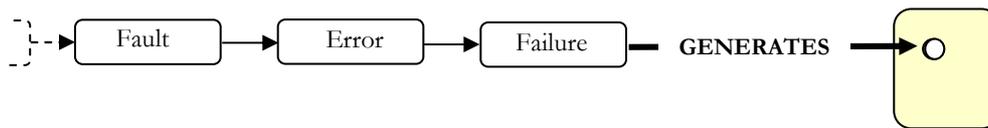


Figure 3: An elementary fault-error-failure chain generating a hole in a given system's layer.

There may be many ways to represent the mechanisms leading to holes in system layers. One of them is having further swiss cheese models orthogonal to the master one. Instead, we believe that a fault-error-failure chain provides an interesting angle where each single hole of a faulty system's layer can have an identifiable causal path. In other words, we advocate a mapping between failures and holes in the system's layers. In this view we have reached a step where Reason's and Randell's models can be turned into compatible representations of systems failures. This opens up a new area of application for Randell's model, that of socio-technical system failures. Equally, it allows Reason's model to connect to technical causal paths of failures in systems. We will attempt to demonstrate the validity of this integration in the next subsection where the Therac-25 socio-technical failure is described. This case has been studied extensively and widely documented (see e.g. Leveson, 1993) and therefore provides an unambiguous basis for discussion.

2.4 Extension to large systems: THERAC-25

THERAC-25 was an X-ray treatment machine designed to kill tumours in deep body tissues. Radiation overdoses happened between 1985 and 1987 and several patients died from subsequent injuries. The machine was recalled in 1987 for extensive design changes, including hardware safeguards against software errors (see Leveson, 1993).

When reasoning about an event in STSs, there are several levels of abstraction that need to be considered. These levels range from individual activities up to entire sections of a company. In the case of THERAC-25, a number of contributing factors were involved, including the regulation authorities, the company who developed the system (AECL) and the programmer who wrote the code. Each of these stakeholders failed, in some sense, in contributing to the system's dependability as a whole. Following Reason's conception, we would say that each contributing layer to the full THERAC-25 system contained flaws (holes).

We now analyse the failures in the STS of THERAC-25 using our integrated model. If one accepts the idea that systems are made of layers (e.g. in the THERAC-25 case: the programmer, the company and the regulators), then a series of fault-error-failure chains can be used to account for the superimposition of the weaknesses contained in each of these layers (see Figure 4). For the sake of exhaustiveness, there should be as many fault-error-failure chains enumerated as there are holes in each of the THERAC-25 layers. However, for simplicity, we will only describe one of the many chains for each of the layers considered in our example.

- *The programmer* did not take all of the system's real-time requirements into account (fault). This led to the possibility of flaws in some software modules (error) which degraded the reliability of the software (failure).
- *The company* did not perform all the required tests on the software (fault). This resulted in bugs in some modules remaining undetected and hence unfixed (error), thereby triggering exceptions when the given modules were called (failure).
- *The regulation authorities* did not thoroughly inspect the system (fault). This led to some flaws remaining undetected (error). In turn, these flaws caused injuries and deaths when the system was used (failure).

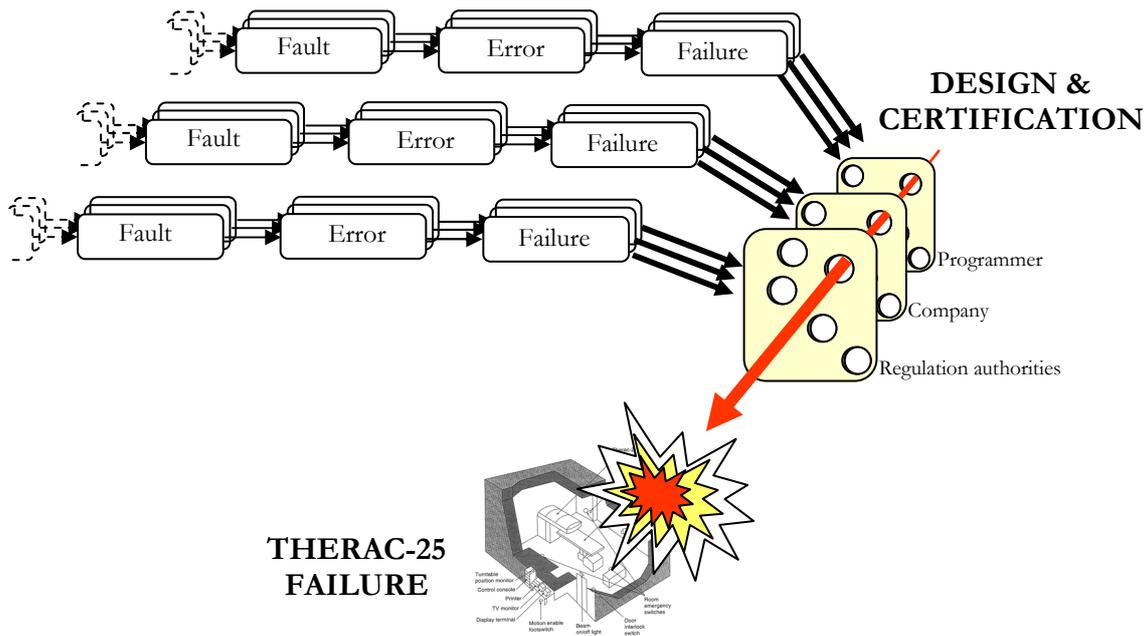


Figure 4: Faults, errors and failures generating impairments in some of THERAC-25 systems' layers.

We have now evaluated the usefulness of an integrated dependability model. The latter can be applied to the description of socio-technical systems and constitutes an attempt to capture both technical and organisational aspects of large failures. To this respect, the integrated model we are proposing here supports the recent idea that dependability must encompass the human actors in systems. This is the focus that we believe should be given to the second part of the paper, since nothing has been said about the users yet. As the latter are core stakeholders, we need to explain where they fit in our conception and how they can affect the dependability of the service delivered by the STS. This is an important question since users often are the last barrier before accidents occur. Indeed, in many cases humans actually improve the dependability of a system by developing ad-hoc compensation strategies. In this respect, there seems to be an issue that is not fully addressed either by Randell's or Reason's models: humans can enable undependable systems to deliver an acceptable level of service. It usually involves such actions as ad-hoc adaptations, workarounds and safe violations. Ensuring the technical dependability of a complex system is not enough to guarantee the quality of service it will deliver. Therefore, some effort is needed to understand the human contribution to the dependability of the delivered service. To this end, we now consider some cases where an undependable system was made to deliver an acceptable service.

3 HUMAN COMPENSATIONS IN SYSTEMS

A system delivers a service when there is an overlap between what the users expect and the system's specification. This overlap can be more or less complete and its size (see Figure 5) represents the quality of the service. The larger the overlap the more acceptable the service. In contrast to the traditional dependability view, we will consider service as a fundamentally non-binary notion.

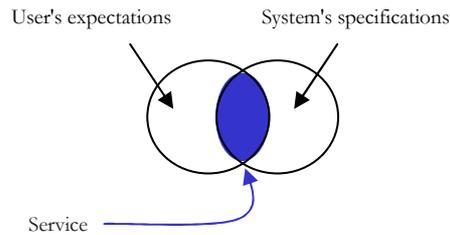


Figure 5 Description of a service as an overlap between a system's specifications and users' expectations.

A service can fluctuate in quality. However, human users often accommodate these fluctuations and find value in an imperfect service. Let us take the simple example of a leaking bucket. A user can tolerate the leak as long as the objective (keeping some water in over a given amount of time) is met. If one now attempts to keep the contents of the bucket over a length of time that is long enough for all the water to leak, then the empty bucket can be considered as a failure. So what matters here is not the specifications of the system as a watertight container, but what the user wants to do, combined with the awareness of the system's limitations. From this standpoint, the system specification is a secondary issue. The correctness of the specification helps define the artefact and helps to identify a minimum level of service. However, the reality is that incorrectness is paramount in our everyday dialogue with technology. This incorrectness is nonetheless widely tolerated and compensated for by the users.

The examples in the below respectively deal with adaptations, workarounds and violations. These mechanisms capture a continuum in the departure from prescribed work whilst still yielding a positive outcome. The examples illustrate how ad-hoc forms of control can increase the quality of service from a nominally undependable system.

3.1 Adaptations

Clarke *et al.* (2003) show how humans adapt their work to local contingencies. An ethnomethodologically informed ethnographic study was conducted of a steelworks factory producing steel slabs of varying sizes. The study focused on the use of a computer-driven roughing mill. Slabs are produced to a required size by rolling large metallic rolls in a series of passes. Because there are several qualities of steel and a variety of ways to reach the final slab's dimensions, operators have developed various work strategies. Some of them override the computer's control. For instance, operators sometimes shift to manual control mode for the final passes on slabs of a particular thickness. In doing so, they reduce the number of passes in order to avoid slabs taking a U-shape, an occurrence which is not always avoidable under computer mode. As quoted from Clarke *et al.* (2003):

"...because the computer, at less than 45, pisses about...does 4-5 passes...that's what's causing turn-up."

The work reported here highlights how operators develop strategies that compensate for flaws in the automation. The latter may be fit-to-purpose under nominal work settings but adaptations are required for any other case. In this example, the adaptations performed by the operators prevent the occurrence of undesired outcomes. If such adaptations did not take place, the production would be (at best) much longer due to the necessary corrections to malformed slabs. Another interesting point is the anticipation skills exhibited by the roughing mill operators. They often pro-actively adjust the quality of the slabs they produce to the requirements of the next processing stage (finishing mill). This is a sign of expertise in piloting the system and constitutes another example of human skills compensating for the inadequacy of the computer function. This steelworks factory study shows how an acceptable level of dependability is achieved in the

service through ad-hoc collaboration between human operators and an imperfect piece of technology.

3.2 Workarounds

Voß *et al.* (2002) reported on a field case study carried out at EngineCo, a plant producing mass-customised engines. In unpublished observations, Voß *et al.* analysed workarounds performed by the production line operators. EngineCo relies on a software tool order to track orders, deadlines, special customisations for particular customers, etc. The software drives the entire supervision of the process from the management of stock, all the way down to delivery dates. This software is designed in such a way that all the parts needed for an engine have to be in stock before assembly can begin. This appears to be an unworkable constraint for the operators who can still work on areas of an engine for which parts are available. However, because the software system does not allow this sort of ad-hoc adaptations to contingencies, operators create items in the stock for the parts that are missing and begin the assembly. Later, when the missing parts get delivered, they are mounted on the engine and the stock is set back to zero. This type of workaround shows how operators overcome technical limitations and rigid designs in order to maintain the production targets, even when all the conditions needed are not present. It also highlights how users can adapt their work to tools that do not always support their practice (Randell, 2003). This issue has been widely discussed by Gasser (1986).

3.3 Safe violations

Although given some attention by Reason (1990), violations that improve system's safety need more attention for they demonstrate the importance of the accuracy of operators' mental models in acting reliably in adverse conditions. This is the angle taken by Besnard and Greathead (2003) on the DC-10 that crash-landed at Sioux City Airport, Iowa (reported in NTSB, 1990). The aircraft was forced to land after the tail-mounted engine disintegrated and damaged the lines of the two wing-mounted engines, resulting in a complete loss of hydraulic control. The damage to the hydraulic lines meant that the crew had no control over the ailerons, rudder, elevators, flaps, slats, spoilers, or steering and braking of the wheels. The only control they had was over the throttle controls of the two, wing-mounted engines. By varying the throttle controls, they were able, to a limited extent, to control the trajectory of the aircraft. As both the pilot and the co-pilot were struggling with the yoke, they could not control the throttles. Fortunately, another DC-10 pilot was onboard as a passenger and was brought to the cockpit. This second pilot could then control the throttles allowing the pilot and co-pilot to control the yoke and the co-pilot to maintain communication with the ground. This is, understandably not common flying practice and several flying procedures were obviously violated on this flight.

By performing these violations to basic flight practices, the crew were able to reach the airport and about 60% of the passengers survived. This event shows that some violations can be beneficial to system safety when they are coupled with a valid mental model. They allow operators to implement *ad hoc* control modes and to some extent, cope with unknown piloting conditions. The pilots' accurate mental models helped them to define viable boundaries for possible actions and allowed them to restore some form of control on the trajectory under strong time pressure and high risks. Controlling the aircraft on the basis of such a model allowed the implementation of positive desirable violations. The latter, although contravening existing procedures, nevertheless exhibited a high degree of relevance. From a human-machine interaction point of view, the pilots avoided a complete disaster and brought the safety level of the damaged aircraft back within acceptable boundaries. Again, we are in the face of undependable -however exceptional- settings being largely compensated for by human intervention.

4 ACCOMODATED UNDEPENDABILITY

The important point here is that the human compensations for technical imperfections can be achieved using actions that are, strictly speaking, illegal. What we have described as adaptations, workarounds and violations are acts that do not belong to the best practice, as described by operating procedures, for example. However, the reality of the work as it happens on the ground supports the well-known discrepancy between these procedures and practice. It follows that procedures must be seen as a resource for action (Wright *et al.* 2000; Fujita, 2000) rather than a comprehensive prescription of the work. The reason is that there will always be situations for which rules do not work and that will require some kind of adaptive response from a human operator.

A solely technical approach to the dependability of sub-parts of machines with which humans do not interact directly (e.g. the hard disk of the computer that drives the roughing mill described in section 3.1) appears to be adequate. However, where human-interaction is needed, one should consider a bigger picture. Therefore, the dependability of the *service*, as the result of the interaction of humans with the automation, is believed to be an appropriate concept because it takes into consideration the human operators, the automation and the context in which the human-machine system is embedded. In the following refinement to our initial model (see Figure 6), we suggest that a) dependability should be regarded as a socio-technical issue and that b) the abovementioned compensating actions should be part of the big picture of the socio-technical dependability.

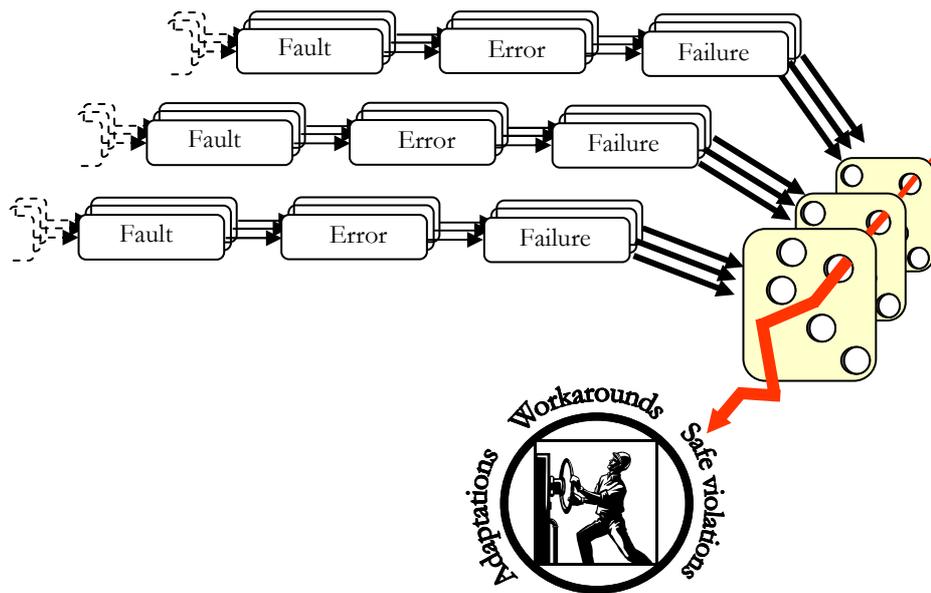


Figure 6: Adaptations, workarounds and safe violations contributing to the dependability of socio-technical systems.

From this standpoint, the nature of an action (i.e. whether it is legal or not) does not matter. As acknowledged by Besnard & Greathead (2003), what is of importance are dimensions such as:

- the knowledge that operators have about the limitations of the system;
- how compatible their mental model is with the functioning of the system;
- the extent to which future events can be anticipated;
- the understanding of the consequences of the actions performed.

Our aim is not to encourage operators to disregard the rules. We are not claiming either that design should not take into account the human physical and cognitive characteristics. Instead, we

wish to promote a view in which ad-hoc changes to prescribed practices can help to mitigate, or even alleviate a system's undependability. The examples described in section 3 show how necessary these corrections can be, in the face of system flaws.

We have addressed a number of points that we believe contribute to the progress of understanding the dependability of STSs. These points can be summarised as follows.

- Randell's technical standpoint must be reconciled with more organisational views of failures, since what matters is the service - as provided by the dialogue between humans and the automation - which is delivered.
- Beyond very low levels of complexity, perfect design is a dream (see for instance Kleinman *et al.*, 2001, on a faulty pace-maker design) more than a reality. There will always be unexpected conditions for which human compensations will be needed (Bainbridge, 1983). Therefore, dependability has to be thought of in terms of a non-binary concept emphasizing the human adaptive use of technology.
- The view according to which systems must be made dependable by designers must acknowledge that systems are often made less undependable by users. This obviously calls for a greater involvement of end-users in iterative and evolving designs. These are issues tackled by a new approach called *co-realisation* (Hartswood *et al.*, 2002).

5 CONCLUSION

We have assessed the compatibility of two different models, both dealing with the propagation of faults in systems. We have reached a position where these accepted models of failures are thought to be both compatible and incomplete: Reason (1990) and Randell (2000) describe different fault propagation aspects which belong to two orthogonal views. In this paper, we have attempted to combine these models as a way of describing failures across the whole socio-technical system. The resulting integrated model offers a richer description of socio-technical failures by suggesting a mapping between sequences of events (a fault-error-failure chain) and holes in the layers of a system.

We believe that our approach has some intrinsic interest since it constitutes a step forward in reconciling technical and organisational views on failures in socio-technical systems. In doing so, we position ourselves within a stream of research where dependability can no longer be accepted as a sole technical issue. When considering how computers are used in socio-technical systems, it is obvious that the scope of dependability models need to be expanded to encompass human compensations to technical flaws.

6 REFERENCES

- Amalberti, R. (1996). *La conduite de systèmes à risques*. Paris : P.U.F.
- Besnard, D. & Greathead, D. (2003). A cognitive approach to safe violations. To appear in *Cognition, Technology & Work*.
- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19, 775-779.
- Clarke, K., Hughes, J., Martin, D., Rouncefield, M., Sommerville, I, Gurr, C., Hartswood, M., Procter, R., Slack, R. & Voss, A. (2003). Dependable red hot action. In Kuutti, K., Karsten, E. H., Fitzpatrick, G., Dourish, P. & Schmidt, K. (Eds.) *Proceedings of the European Conference on Computer Supported Cooperative Work*, Helsinki. Dordrecht. Kluwer (pp. 61-80).
- Fujita, Y. (2000). Actualities need to be captured. *Cognition, Technology & Work*, 2, 212-214.
- Furuta, K., Sasou, K., Kubota, R., Ujita, H., Shuto, Y. & Yagi, E. (2000). Human factor analysis of JCO criticality accident. *Cognition, Technology & Work*, 2, 182-203.

- Gasser, L. (1986). The integration of computing and routine work. *ACM Transactions on Office Information Systems*, 4, 205-225.
- Hartswood, M., Procter, R., Slack, R., Voß, A., Buscher, M., Rouncefield, M. & Rouchy, P. (2002). Co-realisation: Towards a principled synthesis of ethnomethodology and participatory design. *Scandinavian Journal of Information Systems*, 13, 7-20.
- Kleinman, B., Baumann, M., Andrus, C. (2001). Faulty design resulting in temporary pacemaker failure. *Chest*, 120, 684-685.
- Leveson, N. (1993). An investigation of the Therac-25 accidents. *IEEE Computer*, 26, 18-41.
- Loeb, V. (2002). Friendly fire deaths traced back to dead battery. *Washington Post*, March 24.
- Mancini, G. (1987) Commentary: Models of the decision maker in unforeseen accidents. *International Journal of Man-Machine Studies*, 27, 631-639.
- NTSB (1990). *Aircraft accident report. United Airlines flight 232. Mc Donnell Douglas DC-10-10. Sioux Gateway airport. Sioux City, Iowa, July 19, 1989.* National Transportation Safety Board, Washington DC, USA.
- Randell, R. (2003). User customisation of medical devices: the reality and the possibilities. *Cognition, Technology & Work*, 5, 163-170.
- Randell, B. (2000). Facing up to faults. *The Computer Journal*, 43, 95-106.
- Reason, J. (1990). *Human error*. Cambridge University Press, Cambridge, UK.
- Reason, J. (1997). *Managing the risks of organisational accidents*. Aldershot, Ashgate.
- Reason, J. (2000). Human error: Models and management. *British Medical Journal*, 320, 768-770.
- Rogers, W. P. (1986) *The presidential commission on the space shuttle challenger accident report*. Available online at <http://science.ksc.nasa.gov/shuttle/missions/51-l/docs/rogers-commission/commission.txt>
- Svenson, O, Lekberg, A. & Johansson, A. E. L. (1999). On perspective, expertise and differences in accident analyses: Arguments for a multidisciplinary approach. *Ergonomics*, 42, 1567-1571.
- Voß, A., Slack, R., Procter, R., Williams, R., Hartswood, M. & Rouncefield, M. (2002). Dependability as ordinary action. Proceedings of *SafeComp2002*, Catania, Italy (pp. 32-43).
- Wagenaar, W. A. & Groeneweg, J. (1987). Accidents at sea: Multiple causes and impossible consequences. *International Journal of Man-machine Studies*, 27, 587-598.
- Wright, P. C., Fields, R. E., & Harrison, M. D. (2000). Analyzing human-computer interaction as distributed cognition: The resources model. *Human-Computer Interaction*, 15, 1-41.

7 ACKNOWLEDGEMENTS

The authors wish to thank Cliff Jones, Michael Jackson, Tony Lawrie and Carles Sala-Oliveras for useful inputs into earlier versions of this work. The authors also wish to thank anonymous reviewers for useful comments. This study was written as part of the DIRC project¹ which is funded by the EPSRC.

¹ Visit DIRC at www.dirc.org.uk