Responsibility Modelling for Resilience

Gordon Baxter and Ian Sommerville
School of Computer Science, University of St Andrews, UK
{Gordon.Baxter, Ian.Sommerville}@st-andrews.ac.uk

Abstract. We describe a method for modelling responsibilities, which we define as duties that are discharged by agents, using some resources in a particular context. Responsibilities can be used to analyse socio-technical issues and represent the softer aspects of work. The relationships between responsibilities, agents and resources are captured using graphical responsibility models. These models help the stakeholders to identify vulnerabilities, so that they can then be appropriately managed, thereby making complex systems, such as critical national infrastructure, more resilient. We have also combined the method with a keyword approach to analyse tactical issues, such as alternative system configurations. Responsibility modelling is a lightweight method that we have used in areas ranging from contingency planning to system assurance. Based on our experiences we have developed checklist questions to guide the prospective analysis of complex systems. We are now using responsibility modelling to analyse aspects of critical national infrastructure in the UK.

1 INTRODUCTION

The discipline of resilience engineering is starting to mature, with the development of methods and tools to support the practice. The Functional Resonance Analysis Method (Hollnagel, 2004), for example, offers a way of analysing systems in terms of component functions and their inter-relationships. We have been analysing systems using responsibilities, rather than functions, based on our observations that systems often fail due to socio-technical, rather than technical reasons. We take a pragmatic view, and define a *responsibility* as:

> *a duty, held by some agent, to achieve, maintain or avoid some given state, subject to conformance with organisational, social and cultural norms.*

The key features of responsibilities are: that they are allocated to *agents* (usually people or organisations, but they could also be some technology or a software application); that *resources* (time, people, information or physical resources) are needed to discharge them; and that they are discharged in a context which includes factors that define (and

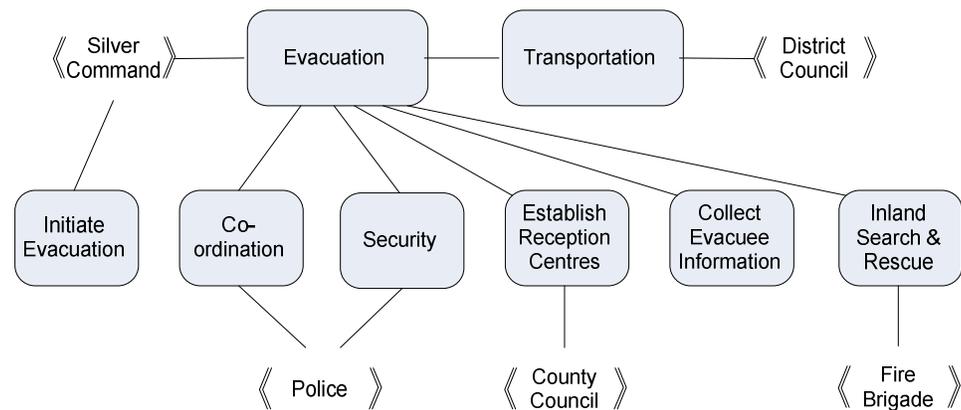possibly constrain) how the responsibility is discharged.

We will focus our discussions here on the resilience of critical national infrastructure (CNI). Using crisis management to deal with CNI failures is no longer sufficient; a resilient CNI is needed, which can quickly be restored to normal use in the event of a failure (Boin & McConnell, 2007). This is now widely accepted at national levels (e.g., The Scottish Government, 2011; The UK Government, 2011).

Relating responsibilities to CNI, we can see that the *given system state* that has to be attained or maintained is a safe and secure CNI; unsafe and insecure ones should be avoided. The applicable *norms* are often defined by local operating procedures, working practices, and constraints imposed by regulatory authorities and acts of Parliament.

The tools and methods that have traditionally been used to analyse technical systems and their failures cannot typically deal with failures that are linked to responsibility vulnerabilities. Technical systems are often described in terms such as tasks and function, which many system stakeholders find difficult to understand. On the other hand, stakeholders are more comfortable with the notion of responsibilities, which provide a more natural way to express the softer aspects of work.

2 RESPONSIBILITY MODELLING

We have developed a graphical representation technique called Responsibility Modelling (RM) for analysing systems and organisations, which we have used to analyse major event contingency plans (Lock et al., 2009a). These plans, which are drawn up in advance, are often long unwieldy documents, possibly written by several people at different times, and hence tend to be inconsistent and incomplete. Delivering a resilient response to a major incident, however, is highly dependent on the way responsibilities are allocated and discharged. Fig. 1 shows an example model for a flood contingency plan for Carlisle, England.
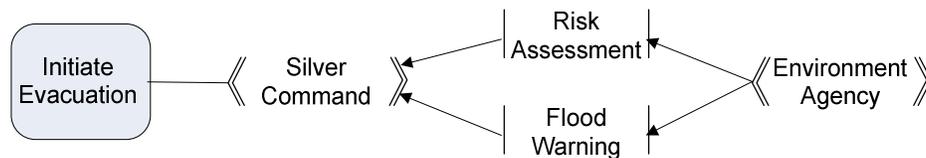

Fig 1. High level responsibility model for flood evacuation plans in Carlisle, England.

Responsibilities are shown as shaded rounded rectangles (e.g., *Co-ordination*), and the agents are shown in angle brackets (e.g., *Police*)

The links between responsibilities and agents show where responsibilities have been allocated. Each level of responsibility can be represented by its own responsibility model diagram(s). For large socio-technical systems, such as those providing and supporting CNI, it is easier to present the system as a set of related responsibility models, reflecting the different levels of responsibilities, rather than try to model the whole system on one diagram. These models are usually organised hierarchically.

In Fig. 1, the responsibility *Collect Evacuee Information* has not been allocated to any agent. Drawing the responsibility model highlighted that the contingency plan had not defined which agent (*Police*, *District Council*, etc.) should collect this information.

Normally, responsibility models (and particularly lower level models) will also show the resources needed to discharge each responsibility (these are shown as a pair of vertical lines with a label in between). Fig. 2, for example, shows which resources are needed by *Silver Command* to discharge the *Initiate Evacuation* responsibility from Fig. 1.

Fig. 2. Responsibility model showing resource associations. (the responsibility for the *Environment Agency* to produce the *Risk Assessment* and *Flood Warning* information resources has been omitted for the simplicity)

The arrows on the links between agents and resources show how those resources are used. In Fig. 2, the arrows point from the *Environment Agency* into the resources *Risk Assessment* and *Flood Warning*, indicating that these resources are produced by the Environment Agency. Similarly, because the arrows point from these resources into *Silver Command*, this indicates that these resources are read by Silver Command, and used to initiate an evacuation.

3   RESPONSIBILITY VULNERABILITIES

When failures occur in socio-technical systems they can manifest themselves in different ways. Some will arise during operation, such as when people interact with the technology; when processes do not work as documented; and when people try to provide data (or information) to the users. In CNI it is obviously important to minimise any adverse consequences that may follow on from any failures.

System failures can arise through a class of vulnerabilities that is associated with responsibilities. Sommerville (2007) identified six main types of responsibility and agent-related vulnerabilities that can give rise to failures:

1. Unassigned responsibilities: a responsibility has not been assigned to any agent.
2. Duplicated responsibilities: different agents believe that they have been assigned a particular responsibility and all try to discharge it.
3. Uncommunicated responsibilities: a responsibility is assigned to a role, but the agent that is assigned that role is not told about the responsibility.
4. Misassigned responsibilities: the agent does not have the capabilities or resources to discharge the assigned responsibility.
5. Responsibility overload: the agent does not have sufficient resources to discharge all its responsibilities in a timely manner.
6. Responsibility fragility: a responsibility is assigned to an agent, but if that agent is not available, there is no back-up agent who can discharge that responsibility.

Whilst duplicated responsibilities are a vulnerability, a system may be made more resilient if back-up agents are assigned to particularly important responsibilities. The secondary agents must understand, however, how and when they should act.

In addition, a misassigned responsibility may not necessarily be discharged ineffectively (if at all). The inherent flexibility and adaptability of people often means that some way will be found to discharge that responsibility, using workarounds as appropriate.

One more vulnerability can be added to the list: conflicting responsibilities, which can arise where an agent fulfils several roles. If, for example, a person has the responsibility for managing a project (making sure it is delivered on time and within budget), but also has the responsibility for technical assurance of that project (making sure that the application or system has been rigorously developed, tested and so on), these two responsibilities are likely to come into conflict.

It is therefore important, when using RM to analyse a system that vulnerabilities are considered. The results of the analyses can be used to identify strategies for appropriately managing vulnerabilities (in some respects, the process for dealing with vulnerabilities is similar to risk management).

4 RESOURCE VULNERABILITIES

There are also vulnerabilities associated with resources, so how they are provided, used, and accessed has to be taken into account. If an agent continuously consumes resources in discharging a responsibility, for example, then appropriate consideration needs to be given to the supply of those resources. This could potentially lead to the identification of a new responsibility for another agent to produce those resources.

If information resources have to be used in emergency situations, then consideration has to be given to how these resources can be accessed by the appropriate agent. If, for example, a network connection is lost, and the information resource is only stored on-line, this could mean that the agent cannot discharge a particular responsibility. Similarly, consideration needs to be given to how physical resources can be accessed if there is a flood, for example, or if the road network becomes grid-locked.

5   TACTICAL CONSIDERATIONS

RM can be combined with keyword-based approaches, like Hazard and Operability Studies (HAZOPS; Kletz, 1999), to analyse the vulnerabilities and risks of particular system configurations. This involves applying standard keyword phrases such as "Too much", "Not enough", "Too early", "Too late" and so on, to the different elements of the responsibility model. The sorts of questions that might be asked are "What if there are not enough resources available to discharge the responsibility?" and "What if this responsibility is discharged too late?". The results of this analysis can be used to define appropriate prevention and mitigation strategies to deal with those situations.

Where agents are assigned several responsibilities, they usually deal with them by prioritising them in some way (or there may be standard operating procedures that define the priorities of some responsibilities). The timing of the discharge of responsibilities that have a specific deadline has to be carefully considered: if a responsibility is discharged too early, it rarely has adverse consequences; if it is not discharged by its deadline, however, consideration needs to be given to whether that responsibility should be abandoned. Sometimes there may be merit in carrying on after the deadline, particularly if it means that the system keeps operating, albeit at a degraded level, when abandoning the responsibility would cause the system to cease operating.

In the ideal world, responsibilities should be discharged by agents with the appropriate combination of knowledge, skills and attributes. When incidents happen, however, they have to be handled using the available agents and resources, so agents may be assigned responsibilities that they are not able to fully discharge. As noted above, there may still be merit in partially discharging some responsibilities, but there will be others that should just be abandoned, rather than assigned to the scarce supply of agents and resources who might be better deployed elsewhere.

The synchronisation of responsibilities can also lead to potential problems. If an agent is responsible for ensuring that an information resource is kept up to date, for example, and another agent uses that resource to discharge its responsibility, the second agent may need to be kept informed of updates to that resource. If they are not, they could potentially be using resources that are out of date, which may lead to time and effort being expended unnecessarily, and a responsibility being discharged ineffectively.

6   USING RESPONSIBILITY MODELS

There are two basic ways of using responsibility models. The first is to describe an existing system. This involves looking at descriptions of that system, either based on existing studies, or by using the documentation—design documentation, operating manuals and so on—that describes how the system should work. The collected data is then analysed, and one or more responsibility models are created. As noted above, we have already used RM to examine contingency plans in this way (Lock et al., 2009a).

The second way is to use the models in a prospective fashion. In other words, the models are used to draw up system descriptions or plans before a system exists, to help ensure that the system will be less vulnerable to responsibility failures when it is built. The

resulting responsibility models can be incorporated into the system documentation.

Whether responsibility models are used retrospectively or prospectively, the final step is the same. Once the models have been created, they are taken back to the appropriate stakeholders so that they can verify that the models reflect their understanding of how the system should work. Any feedback is used to refine the models appropriately.

By using RM, the responsibility vulnerabilities that could lead to system failures can be identified and highlighted. This means that responsibility models can be used to develop scenarios for use in system evaluation, through table-top exercises and simulations, for example, to see how the system would cope. Also, because responsibility models help to identify and describe the softer aspects of work, they can supplement the process of risk management to increase assurance about the resilience of a particular system or service.

Where responsibility models are particularly useful, however, is in the description of work situations where several groups or agencies have to collaborate and co-operate, such as in the provision and maintenance of CNI. In multi-agency situations there are often discrepancies between the way that the different agencies view responsibilities, and there may also be emergent responsibilities that arise out of the collaborations between the different agencies. RM can help here in at least two important ways. The first is that responsibility models make it easier to identify any responsibilities that have not been assigned to any agent. This may arise in situations where all the agencies know about a responsibility but each assumes that it has been allocated to one of the other agencies. The second is that there may be different interpretations of what a particular responsibility means. The RM process helps to identify these too, and provides a basis for discussions aimed at achieving a common level of common agreement about what the responsibility entails, who it should be assigned to, what resources are needed, and who provides those resources.

7 SUMMARY AND FUTURE WORK

We have applied RM retrospectively to several application areas ranging from contingency planning to system procurement. We are now starting to extend the method so that it can be used prospectively in the analysis of CNI, to help facilitate resilient responses to events (Boin & McConnell, 2007). We have been developing checklists that can be used to ensure that responsibilities are assigned, appropriately resourced, and made visible to all the agents involved, so that they can be discharged in a timely and effective manner. We have started to apply these ideas to analyse the resilience of elements of CNI within the UK

RM is intended to be a relatively lightweight method that requires little in the way of training to learn. The checklists associated with each responsibility occupy about one side of A4 paper. A standalone web-based tool for creating responsibility models is planned for development in 2011.

The main strengths of RM lie in its capability for describing some of the softer aspects of work, and in the fact that responsibilities are a concept that is readily understood and can be discussed by stakeholders. The method has been successfully applied to the analysis of the emergency flood plans for Cumbria (Lock et al., 2009a), to help in

identifying information requirements for a system that co-ordinates agencies to deal with emergencies (Sommerville et al, 2009), and to investigate some of the problems with the e-counting systems used in the 2007 Scottish Elections (Lock et al., 2009b). The inherent ability to deal with systems involving the use of multiple agencies makes RM particularly suitable for analysing systems that involve the support and maintenance of CNI to identify potential vulnerabilities.

RM has been designed to be used alongside the process of risk management, and in parallel with other techniques that can be used to identify vulnerabilities in software and hardware. Using RM, however, can help to identify a class of vulnerabilities that are not addressed by these other methods. Implementing appropriate strategies to deal with these vulnerabilities will help to increase the resilience of the overall socio-technical system.

As we build larger and more complex systems, and systems of systems (including CNI), we need to find new ways to analyse them. We have found responsibilities to be a useful level of abstraction from this point of view. RM has proved useful both to us and to the agents and agencies involved in identifying and addressing potential responsibility vulnerabilities, such as unallocated responsibilities, and resource problems.

REFERENCES

Boin, A. and McConnell, A. (2007). Preparing for critical infrastructure breakdowns: The limits of crisis management and the need for resilience. *Journal of Contingencies and Crisis Management*, *15,* 50-59.

Hollnagel, E. (2004). *Barriers and accident prevention.* Aldershot, UK: Ashgate.

Kletz, T. (1999). *HAZOP and HAZAN: Identifying and assessing process industry standards* (Fourth edition). Rugby, UK: Institution of Chemical Engineers.

Lock, R., Sommerville, I., & Storer, T. (2009a). Responsibility modelling for civil emergency planning. *Risk Management, 11*, 179-207.

Lock, R., Storer, T., Sommerville, I., & Baxter, G. (2009b). Responsibility modelling for risk analysis. In *Proceedings of ESREL 2009*, (pp. 1103-1109).

The Scottish Government (2011). *Secure and Resilient: A Strategic Framework for Critical National Infrastructure in Scotland*. Edinburgh, UK: The Scottish Government.

Sommerville, I. (2007). Models for responsibility assignment. In G. Dewsbury and J. Dobson (Eds.), *Responsibility and dependable systems* (pp. 165-186). London, UK: Springer.

Sommerville, I., Lock, R., Storer, T., & Dobson, J. (2009). Deriving information requirements from responsibility models. In *Proceedings CAiSE 2009: 21st international conference on advanced information systems engineering* (pp. 515-529). London, UK, Springer: 515-529.

The UK Government Cabinet Office (2011). *Keeping the Country Running: Natural Hazards and Infrastructure*. London, UK: The UK Government Cabinet Office: © Crown copyright, 2011.