

delivering benefits through science

source

pathway

receptor



Reliability in Flood Incident Management Planning

Final Report – Part B: Technical Report

Science project SC060063/SR2

The Environment Agency is the leading public body protecting and improving the environment in England and Wales.

It's our job to make sure that air, land and water are looked after by everyone in today's society, so that tomorrow's generations inherit a cleaner, healthier world.

Our work includes tackling flooding and pollution incidents, reducing industry's impacts on the environment, cleaning up rivers, coastal waters, contaminated land and improving wildlife habitats.

This report is the result of research commissioned by the Environment Agency's Science Department and funded by the joint Environment Agency/Defra Flood and Coastal Erosion Risk Management Research and Development Programme.

Published by:

Environment Agency, Rio House, Waterside Drive,
Aztec West, Almondsbury, Bristol, BS32 4UD
Tel: 01454 624400 Fax: 01454 624409
www.environment-agency.gov.uk

ISBN: 978-1-84911-130-0

© Environment Agency – October 2009

All rights reserved. This document may be reproduced with prior permission of the Environment Agency.

The views and statements expressed in this report are those of the author alone. The views or statements expressed in this publication do not necessarily represent the views of the Environment Agency and the Environment Agency cannot accept any responsibility for such views or statements.

This report is printed on Cyclus Print, a 100% recycled stock, which is 100% post consumer waste and is totally chlorine free. Water used is treated and in most cases returned to source in better condition than removed.

Email: fcerm.science@environment-agency.gov.uk

Further copies of this summary are available from our publications catalogue: <http://publications.environment-agency.gov.uk> or our National Customer Contact Centre: T: 08708 506506
E: enquiries@environment-agency.gov.uk.

Key Author(s):

Peter von Lany, Halcrow
Andy Barnes, Halcrow
Richard Dawson, University of Newcastle
Dennis Parker, Middlesex University

Dissemination Status:

Released to all regions
Publicly available

Keywords:

Flood Incident Management Planning

Research Contractor:

Halcrow Group Ltd, Burderop Park, Swindon, SN4 0QD, 01793 812479

Environment Agency's Project Manager:

Adam Baylis, Science Department

Theme manager:

Dr Suresh Surendran, Modelling and Risk Theme

Collaborator(s):

University of Newcastle, Middlesex University FHRC, JBA Consulting, University of Bristol

Science Project Number:

SC060063

Product Code:

SCHO1009BRDK-E-P

Science at the Environment Agency

Science underpins the work of the Environment Agency. It provides an up-to-date understanding of the world about us and helps us to develop monitoring tools and techniques to manage our environment as efficiently and effectively as possible.

The work of the Environment Agency's Science Department is a key ingredient in the partnership between research, policy and operations that enables the Environment Agency to protect and restore our environment.

The science programme focuses on five main areas of activity:

- **Setting the agenda**, by identifying where strategic science can inform our evidence-based policies, advisory and regulatory roles;
- **Funding science**, by supporting programmes, projects and people in response to long-term strategic needs, medium-term policy priorities and shorter-term operational requirements;
- **Managing science**, by ensuring that our programmes and projects are fit for purpose and executed according to international scientific standards;
- **Carrying out science**, by undertaking research – either by contracting it out to research organisations and consultancies or by doing it ourselves;
- **Delivering information, advice, tools and techniques**, by making appropriate products available to our policy and operations staff.



Steve Killeen

Head of Science

Executive summary

Flood incident management (FIM) plays an important part in reducing the adverse consequences of flooding, but it only does so when it is well planned and effective. Following the summer floods of 2007, the Pitt Review called for:

- improved planning and preparation for floods;
- a step change in the quality of flood warnings;
- a more resilient approach to managing floods.

This research focuses upon ways of improving the planning of FIM. New tools have been created recently that improve the way FIM planners can evaluate the likely reliability and performance of a FIM process in any given situation. These evaluations can be used to identify the components of FIM which should be addressed in the planning phase to improve the reliability and effectiveness of FIM actions.

Part A of the final report provides guidance on how to apply the various tools developed during this study. Part B describes the technical process undertaken to develop and test these tools. The test results provide the evidence for the guidance presented in this report.

Floods can be managed through structural and non-structural approaches. Structural approaches involve the use of physical structures to prevent, divert or mitigate the impacts of flooding. FIM aims to reduce the impacts of flooding upon society and the economy through non-structural interventions. The reliability of any approach, including FIM, has a direct influence on its effectiveness; this is why this project focuses on reliability.

Reliability issues arise from both the technical and human components of FIM. These technical and human components both introduce uncertainties which influence the reliability and performance of FIM systems. The overall performance of FIM systems depends on the reliability of a large set of individually linked, interactive components, which can act to either propagate or reduce uncertainty.

Part B of the Final Report describes the selection, development and testing of tools that can be applied at different levels.

- i. Overview level
At this level the tools provide an overview of FIM performance and help identify factors that could contribute to the reliability and performance of different aspects of FIM. The tools developed are a) performance matrices using a balanced scorecard approach, and b) root cause analysis employing fish-bone diagrams.
- ii. High level
A higher level of analysis is necessary where there is a need to model how specific processes contribute to overall FIM reliability, and to identify factors which make FIM vulnerable to uncertainty and risk of failure or underperformance. The tool developed is a Windows-based hierarchical process modelling tool called *Perimeta*.
- iii. Detailed or dynamic level
Dynamic modelling is needed where it is important to assess the dynamics of the FIM process and systems to show how the evolution of a flood event, technical systems and human behavioural processes may interact and

combine to influence FIM performance. The tool developed is agent-based modelling.

We used case study data provided by the Environment Agency for two separate flood events (July 2007 and January 2008) in the West Area of the Thames Region to test performance matrices and root cause analysis. Using the same flood events, we tested the application of the *Perimeta* modelling system to build hierarchical process models for two specific components of the flood damage avoided (FDA) equation. The form of the FDA equation currently used by the Environment Agency provides a means of evaluating the benefits of FIM, in terms of its contribution to avoiding flood damage. The evaluation of benefits, thus derived, can be used to help establish a business case for measures to improve FIM.

Each *Perimeta* model represented the sub-processes that contribute to specific aspects of FIM, showing how uncertainty and risks of failure could arise and propagate through these sub-processes, thus influencing the performance of FIM. Environment Agency staff from the West Area of the Thames Region were involved in developing, testing and validating the above models, and the insight they gained through the case studies is described in this report.

Agent-based modelling, a powerful tool that helps us to understand the behaviour of complex systems, particularly those tightly coupled with human behaviour, was tested through another case study: a flood event that occurred in Towyn, North Wales in 1990. This flood is a good example of a flood event requiring evacuation. An agent-based model was developed and tested to simulate the effect of flooding via breaches of the sea defences and the subsequent self-evacuation responses of the public.

The modelling tools developed are generic. They can potentially be applied to other forms of flooding and thus help evaluate FIM within the context of broadening Environment Agency responsibilities in the area of planning and implementing non-structural forms of flood risk management.

Acknowledgements

The authors would like to thank the Project Board (Adam Baylis, Kate Marks, Mike Steel and Suresh Surendran) for their valuable input and guidance throughout this project. We are grateful for the significant time and effort spent by Helen James and Kate Vincent (Environment Agency, Thames Region) in helping develop and carry out case studies. Their enthusiasm and knowledge were invaluable for this particular aspect of the project. We would also like to thank the Environment Agency staff who put aside time to review earlier drafts of this report and whose constructive comments helped shape this final report.

Contents

1	Introduction	1
1.1	Background	1
1.2	Study objectives	1
1.3	Purpose of Part B	1
1.4	Structure of Part B	2
2	Tools for assessing FIM reliability and performance	3
2.1	Introduction	3
2.2	Techniques considered	3
2.3	Selection of tools for further development and testing	9
3	Developing the decision-support tools	12
3.1	Introduction	12
3.2	Developing performance matrices and root cause analysis	12
3.3	Developing <i>Perimeta</i> models	13
3.4	Developing an agent-based model	15
3.5	Evaluating the benefits of FIM	17
3.6	Applying the tools to provide decision support	20
4	Testing the tools	25
4.1	Introduction	25
4.2	Applying the performance matrices and root cause analysis	27
4.3	Testing the <i>Perimeta</i> models	31
4.4	Testing the agent-based model	45
5	Outputs from the case studies	55
5.1	Introduction	55
5.2	Performance matrices and root cause analysis	55
5.3	<i>Perimeta</i> models	56
5.4	Comparing results with Environment Agency performance targets	60
5.5	The agent-based model	61
6	Discussion	71
6.1	Introduction	71
6.2	Evaluation of methods	71
7	Recommendations	75
7.1	Introduction	75
7.2	General recommendations	75
7.3	Specific recommendations	76

References	81
Appendix 1 <i>Perimeta</i> model data	84
Appendix 2 Rules of thumb for setting sufficiency and necessity	106
A2.1 Introduction	106
A2.2 Sufficiency	107
A2.3 Necessity	108
A2.4 Dependency	108
A2.5 Too much evidence – conflict, incoherence and confusion	108

1 Introduction

1.1 Background

The Environment Agency commissioned Halcrow, in collaboration with the University of Newcastle, the Flood Hazard Research Centre (FHRC) at Middlesex University, the University of Bristol and JBA Consulting to carry out research into improving Flood Incident Management (FIM) planning, taking into account its vulnerability to risk and uncertainty.

The project commenced in June 2007. It forms Science Project SC060063 under the joint Defra / Environment Agency Flood and Coastal Erosion Risk Management R&D Programme. It builds on the outputs from an earlier science study, Risk Assessment in Flood Incident Management – Phase 1 (SR11206) published by Defra / Environment Agency in 2006. The project was steered by a Project Board set up by the Environment Agency.

This research project has potentially far-reaching implications for the management and planning of non-structural flood measures. This is because it seeks to develop methods which evaluate the reliability of FIM, which integrates a number of non-structural approaches, so that reliability may be enhanced.

Although the project has developed reliability evaluation and enhancement tools to a level which can be used to good effect now, they will need to be refined in the light of application experience and learning for which there is no substitute, and subsequently further developed over time.

1.2 Study objectives

The overall objective of this study was to develop a set of tools to improve pre-event planning of FIM. The output from this research is intended to:

- enhance understanding of and to provide evidence on what influences the reliability of the different human and technical components of the FIM system, and how these components interact to influence the performance of FIM as a whole;
- identify techniques for evaluating FIM system performance (at a high level and at a more detailed level) in order to improve planning where improvements may be required;
- demonstrate how these techniques can be applied in the form of an analytic framework able to provide decision support to those planning and justifying improvements to FIM reliability and overall performance.

1.3 Purpose of Part B

The final report is presented in two parts. Part A is a guidance report and provides:

- background information on assessing the reliability of FIM planning;

- guidance on the tools which have been developed to evaluate and improve FIM planning;
- suggestions for ways in which these tools could be applied.

Part A also considers the current state of development of these tools, and recommends what additional science and R&D is required to enhance their potential capability and application.

This report, Part B, describes the science and technical work undertaken to develop and test these tools.

The approach used in this study was developed in response to the project brief and involved a number of steps:

- i. Defining the scope and nature of FIM following a workshop in October 2007 and a literature review.
- ii. Conceptualising the nature of performance measurement and management in the context of FIM (via a literature review).
- iii. Identifying techniques that could be used for assessing different aspects of FIM performance.
- iv. Selecting and developing tools to help assess and evaluate FIM performance, and testing these tools via case studies.
- v. Assessing the strengths and weaknesses of these tools in the context of modelling and managing the performance of processes within FIM.
- vi. Making recommendations for use and further development of these tools.

This report is aimed at those with an interest in what influences the performance of FIM. It will be of specific interest to those who wish to gain insight into:

- evaluating the reliability and performance of FIM in particular;
- non-structural flood risk management (FRM) measures in general.

Part A: Guidance Report should be read in conjunction with this report by those looking for broad guidance on the application of these tools to plan improvements in FIM.

1.4 Structure of Part B

Chapter 2 outlines the tools that have been selected and developed to help assess the reliability and evaluate the performance of FIM.

Chapter 3 describes how these tools have been developed within the context of FIM.

Chapter 4 details the design and construction the tools and their testing using case studies. The outputs of the case study work are described in Chapter 5.

Chapter 6 discusses the case study work and outlines the insights that the case studies provide regarding the evaluation of FIM reliability and performance.

The report concludes with Section 7, which contains recommendations for the use and further development of these tools.

2 Tools for assessing FIM reliability and performance

2.1 Introduction

This study has identified decision support tools that can help to assess the reliability and evaluate the performance of FIM at three levels:

- i. **Overview level**
The tools can be used across a range of possible flood incidents to provide an *overview* of FIM performance. This will help to identify the root causes of, and contributory factors to, good, adequate and inadequate performance. Two tools have been developed: performance matrices using a balanced scorecard approach, and root cause analysis using fish-bone diagrams.
- ii. **High level**
The tools can analyse quantitative and qualitative information on reliability and uncertainty in order to determine how these factors contribute to overall FIM performance and identify how FIM is vulnerable to uncertainty, risk of failure or underperformance. The tool developed is a Windows-based hierarchical process modelling tool called *Perimeta*.
- iii. **Detail**
The tools can focus on modelling the dynamics of FIM processes or systems and show how the evolution of a flood event, technical systems and human behavioural processes may interact and combine to influence FIM performance. The tool developed uses agent-based modelling.

These tools can be applied using evidence relating to FIM performance that is:

- *retrospective* (i.e. from post-flood reviews and/or simulated flood exercises);
- *prospective* (i.e. on proposed options for improving FIM processes, resources or delivery mechanisms).

The purpose of these tools is to assess performance and to help identify and evaluate measures that can be adopted to maintain or improve FIM performance.

This chapter outlines the techniques considered for application within FIM, and summarises those selected for further development and testing by this study.

2.2 Techniques considered

Several techniques were considered for use, in the context of FIM, at each of the three levels of application outlined previously.

2.2.1 Overview level techniques

Techniques capable of providing an overview of factors that influence the reliability and performance of different aspects of FIM include tools such as:

- i. **Performance summaries**
Information on past performance is presented in a number of ways, including performance tables, shaded matrix cells, spider diagrams or on radial charts.
- ii. **Performance matrices**
The performance of systems can be assessed from different perspectives against an agreed set of criteria.
- iii. **Balanced scorecards**
Performance is evaluated against multi-criteria indicators assigned to measuring different aspects of performance.
- iv. **Root cause analysis**
Used to help identify fundamental or root causes of problems, or what factors need to be in place to ensure good performance.

Performance matrices

The development and application of performance matrices in the context of improving FIM reliability was considered in an earlier study by the Environment Agency (2006). Two performance matrices have already been developed: one to assess the performance of 'FIM processes', the other to assess performance of 'FIM outcomes'.

The structure of performance matrices is based on the concept of balanced scorecards, used extensively in business, industry, government, and non-profit organisations worldwide as a strategic planning and management tool to:

- to align business activities with the vision and strategy of the organisation;
- improve internal and external communications;
- monitor organisation performance against strategic goals.

As described in Part A: Guidance Report, the balanced scorecard approach originated in the 1990s with the work of Robert Kaplan (Harvard Business School) and David Norton. As a performance measurement framework, it adds strategic non-financial performance measures to traditional financial metrics, thus giving managers and executives a more 'balanced' view of organisational performance (Kaplan and Norton, 1992;1993;1996).

The scorecard is typically constructed to assess performance from four main perspectives, each with its own set of agreed performance indicators. If required, scores for performance in each aspect can then be combined in a weighted manner to give a score for overall performance.

The four perspectives typically used in the balanced scorecard are:

- i. **The customer perspective**
To be successful, how should an organisation appear to its customers and key stakeholders?
- ii. **Internal processes**
To be successful, which processes should an organisation be good at?

- iii. **Continuous improvement**
To be successful, how should an organisation sustain its ability to learn and to improve?
- iv. **Finance**
To be successful, how should an organisation appear to those who provide its financial resources?

Each of these perspectives is also relevant in an assessment of FIM performance.

Root cause analysis

Root cause analysis is driven by the belief that failures and associated risks are best managed by dealing with their fundamental causes of failure rather than by responding, as a matter of expediency, to their symptoms. But in systems that are complex, dynamic and/or inherently uncertain it may be difficult to identify, with confidence, single root causes; failures and their associated risks may arise from a combination of causes (some deeper than others) that interact with one another.

Root cause analysis is used, for example, by the National Patient Safety Agency of the National Health Service (NHS) to determine the root causes of incidents that affect patient safety. A toolkit to guide the application of root cause analysis following incidents that have affected patient safety can be found on the NHS website of the National Patient Safety Agency (2009).

The process of carrying out a root cause analysis can be helped by using 'cause-and-effect' diagrams such as the 'fish-bone diagram'. This form of diagram is also known as the Ishikawa diagram (see Figure 2.1), after Kaoru Ishikawa, who pioneered quality management processes in the Kawasaki shipyards in the 1960s and became one of the founding fathers of modern production management (Ishikawa, 1990). In this form of 'cause-and-effect' diagram, causes are typically grouped into six main categories of factors that can influence process reliability and performance: 'equipment'; 'process'; 'people'; 'materials'; 'environment'; and 'management'. Within each of these categories, primary and secondary causes of failure can be identified.

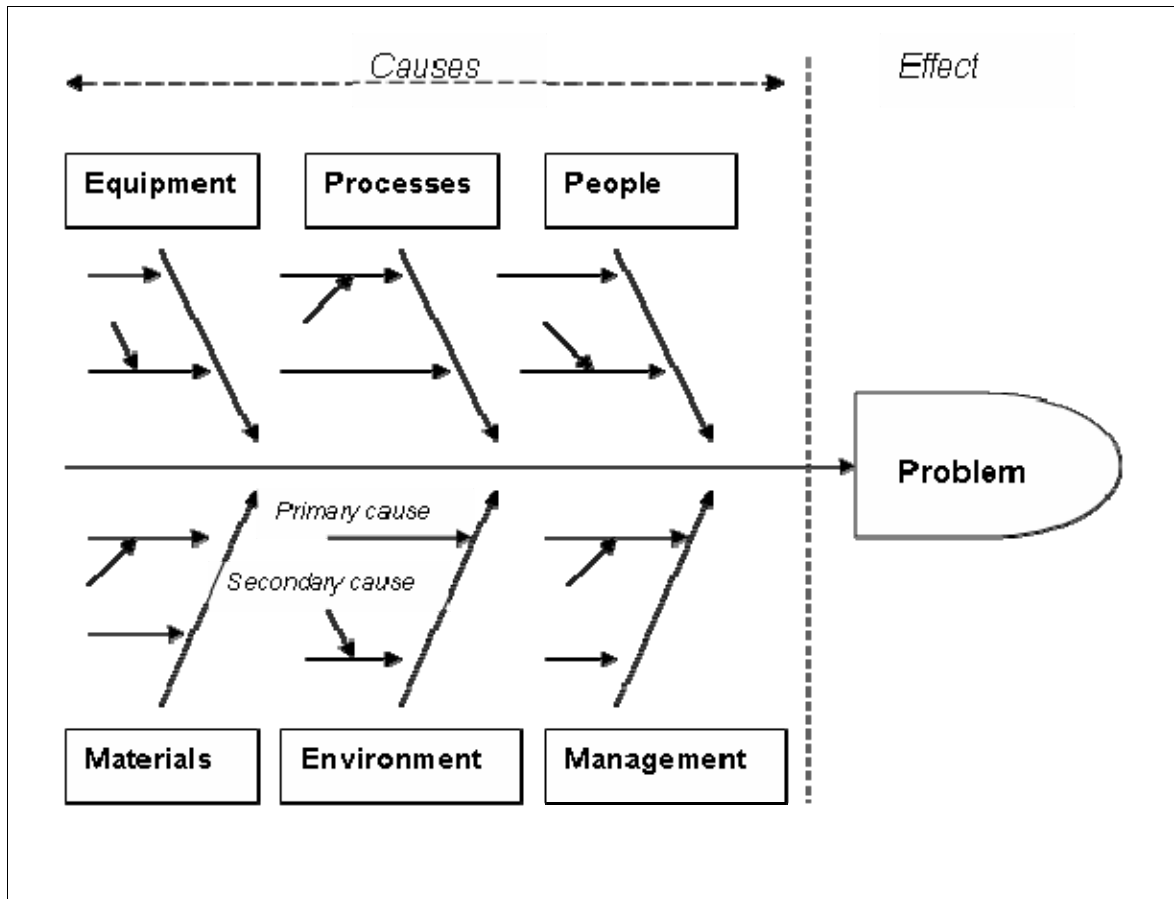


Figure 2.1 Ishikawa (fish-bone) diagram.

2.2.2 Network and hierarchical models of systems

Hierarchical models of systems

A *hierarchical* model is a special case of a *network system* model. In a hierarchical model, the system under consideration is described at a range of different levels of detail. At the top of the hierarchy is a general description of the overall system. This is then broken down into increasingly detailed descriptions of subsystems. For example, policy makers will be interested in the performance of the system as a whole, whereas process teams will be interested in the performance of individual sub-systems.

A hierarchical approach provides a linkage between different levels of decision-making so that the influence that detailed decisions have on high-level performance can be demonstrated and evaluated.

A limitation of many hierarchical modelling techniques is that the hierarchical structure limits the capability of the model to represent key connections and dynamic feedback within a system. There is often more than one reasonable way to decompose a system, but model outputs can be sensitive to the hierarchical structure used.

Hierarchical approaches include:

- i. **Fault trees**
A logical diagram showing all the failure or partial failure mechanisms that

contribute to the failure of a system, in which each event can be assigned a probability in order to calculate the probability of total system failure;

ii. **Event trees**

These diagrams provide a logical representation of the events that may follow an initiating event.

Fault and event trees are inherently hierarchical in structure. Human events/faults can be incorporated (e.g. 'call out team unavailable') into these trees, as has been demonstrated for operational barriers in the Thames Estuary.

Other hierarchical approaches include:

i. **Failure mode element and criticality analysis (FMECA)**

FMECA combines event trees with a risk register to produce a location-cause-indicator diagram that can be used to rank each failure mode according to the combined likelihood of occurrence, consequence and confidence.

ii. **Perimeta**

This tool is a Windows-based hierarchical process modelling tool (developed by the University of Bristol) that allows reliability and uncertainty to be measured and assigned to weighted indicators of performance for each part of a system. These measures can be combined and propagated up the hierarchy of the system, in order to develop an overall measure of system reliability and associated uncertainty.

iii. **Analytic hierarchy process (AHP)**

This structured multi-criteria technique deals with complex decisions by providing a framework for representing and quantifying the elements of a decision, relating those elements to overall goals. The process can evaluate alternative solutions.

In applying the AHP, decision-makers compare each element of a system with every other system element (i.e. a systematic comparison of element pairs). To make their comparisons, decision-makers can use their judgment or concrete data about the elements' relative meaning and importance. The judgments can be converted to numerical values that can then be processed, evaluated and compared over the entire range of the problem.

A numerical weight or priority is derived for each element of the hierarchy, allowing diverse and often incommensurable elements to be compared to one another in a rational and consistent way. However, AHP uses arbitrary scales, based on expert elicitation, without any underlying statistical theory. The process has some fundamental theoretical flaws, although these can generally be taken into account if the AHP is designed and implemented with care.

Network models of systems

Network models are often required to model more complex systems that are poorly represented by hierarchical models, as is often the case where there is feedback within the system or multiple connections between different branches of a hierarchical structure.

Dynamic networks can capture time-changing conditions by simulating stochastic processes. Bayesian networks and influence diagrams (or belief networks) are probabilistic networks that represent a set of variables and their (probabilistic) interdependencies. An influence diagram is a generalisation of a Bayesian network where

both probabilistic inference and decision-making problems can be addressed together (the influence diagram also includes decision variables, probabilistic variables, objective variables and functions).

2.2.3 Social simulation

Social simulation is a research field that applies computational methods to examine issues in the social sciences. Social simulation has been used to model a number of 'social systems' in areas that include political science, economics and geography. Social simulation aims to bridge the gap between the descriptive approach frequently used in the social sciences and the formal, often quantitative approaches used in the physical and natural sciences.

There are a number of different approaches to social simulation, briefly outlined below. It should be noted that many real applications of social simulation have developed hybrids of the described approaches.¹

Micro-analytic simulation (micro-simulation)

Micro-simulation is a modelling technique that operates at the level of individual units such as persons, households, vehicles or firms. Each unit is assigned a set of associated attributes (e.g. people may be assigned age, sex, marital and employment status; vehicles may be assigned origins, destinations and driving speeds). A set of rules are identified and applied to these units which lead to changes in their state and behaviour. These rules may be deterministic or stochastic (probability ≤ 1), such as the chance of a person dying, marrying, giving birth or moving within a given time period. The change to the unit over the time scale of interest can then be assessed.

Multi-level simulation

Multi-level simulation is an extension of the 'micro' approach; it enables multiple scales of 'unit' (e.g. individual, household, total population) to interact. For example, population attributes depend on aggregated individual attributes and individual attributes depend on the population attributes. For example, gender distribution depends upon the number of individual male and females in the total population, whilst the rate at which new individuals are born depends upon population size and gender balance. Micro-simulation and multi-level models are not spatially explicit.

System dynamics

System dynamics models are used to understand the behaviour of complex systems over time by modelling the circular, interdependent and sometimes time-delayed relationships among its components. A system's properties and dynamics are described in terms of equations which derive the future state of the system based on its current state. System dynamic models deal with internal feedback loops, flows of stocks (which can include materials and information) and time delays that affect the behaviour of the entire system.

¹ An excellent and thorough description of these models is given by Gilbert and Troitzsch (2005).

Discrete event models

Discrete event models (often referred to as queuing models) describe a system in terms of a sequence of events which mark a change in the system state. Unlike system dynamic and micro-simulation models (which change the modelled system state in discrete time-steps), in discrete event models the state only changes when an event has occurred. In these models, events typically change only a part of the system's state – between two events nothing changes, not even implicitly. For example, the event 'surge in North Sea' would schedule the next event 'alert Thames Barrier staff', which in turn might schedule a 'close barrier' event if necessary.

Cellular automata models

In cellular automata models, the system domain is divided into a grid; each cell of the grid has a set of attributes (e.g. population, topography) and a finite number of states. The cells do not have to be square, and can be defined in more than two dimensions; however, each cell is subjected to the same set of rules. The behaviour of each cell is governed by local interactions between neighbouring cells. At each time-step every cell in the model is updated. The 'Game of Life' is a famous example of a simple cellular automata model. Cellular automata focus on local interactions so are best applied in scenarios where communication is spatially constrained.

Agent-based models

In an agent-based model, individuals and organisations can be represented as 'agents' who are governed by rules that determine their general behaviour, interactions and spatial location. They can be multi-level and spatial. Agents can be both reactive (i.e. they are influenced by other agents) or proactive (i.e. they actively seek to perform a task). Their behaviour can evolve throughout a simulation.

With just a few rules for a number of agents and locations, complex system behaviour can emerge from the simulation, providing useful insights into vulnerabilities to processes, components and systems.

Agent-based models can simulate many of the aspects of the other modelling methods described above. For example, they can capture the dynamics of system dynamics models, the local interactions of cellular automata, multiple levels of system modelling and the individual attributes of micro-simulation.

2.3 Selection of tools for further development and testing

Each technique described in Section 2.2 was assessed to see if it met a number of criteria agreed with the Environment Agency. The criteria stipulated that a technique should:

- be able to represent specific characteristics of the FIM processes and systems;
- remain effective in situations where information about the state of the FIM system is incomplete or when a precise definition of all elements of the system is not possible;

- be able to make good use of information on reliability and performance from different sources and in a variety of formats;
- provide spatial information on reliability and performance, so that FIM measures could be prioritised at a broad scale (e.g. flood warning areas / Environment Agency regions / catchments).

Performance matrices and **root cause analysis** fulfilled all these criteria and were selected for further development and testing.

Science Report 11206/SR, published by the Environment Agency and Defra in June 2006, recommended the development of a hierarchical process modelling tool to help with the assessment of FIM reliability and performance. The *Perimeta* software is better than other network and hierarchical systems models because:

- it is encoded within well developed software that has in the past been applied to a variety of problems (see below);
- uncertainty can be explicitly represented;
- the influence of reliability *and* uncertainty on performance can be modelled;
- it accepts numerical as well as qualitative input data;
- it provides numerical output, using an interval probability to represent uncertainty.

Perimeta has already been applied to a variety of different problems, for example:

- the management of flood defence system assets by Hall *et al.* (2004) and Dawson *et al.* (2004);
- dealing with uncertainty and risk in engineering systems by Davis and Hall (2003);
- the management of performance in the Highways Agency by Harding *et al.* (2003).

The social simulation approaches (briefly reviewed above) are limited in their ability to capture dynamic responses, but agent-based modelling is one of the few practical methods able to provide this simulation capacity (see Table 2.1).

Research shows that human and organisational responses to flood risk and flood warnings is strongly related to previous experience of flooding and flood warnings, and also to a learning process. Agent-based modelling is therefore well suited to modelling the influence of these kinds of system dynamics on emergency responses such as, for example, the susceptibility of evacuation routes to overcrowding (as introduced in Part A: Guidance Report, Section 3.4).

Table 2.1 Summary and comparison of social simulation methods.

Technique	Simulation levels	Agent interaction	Complexity of agents	Number of agents	Spatially explicit
Microsimulation	2	No	High	Many	No
Multi-level simulation	>2	Potentially	High	Many	No
System dynamics	1	No	Low	1	No
Discrete event	1	No	Low	Many	Yes
Cellular automata	2	Local	Low	Many	Yes
Agent-based modelling	>2	Local and non-local	High	Many	Yes

3 Developing the decision-support tools

3.1 Introduction

This chapter describes the step involved in the development of the techniques selected to assess the reliability of FIM processes. The techniques chosen for incorporation into assessment tools were:

- performance matrices and root cause analysis (to provide an overview of FIM performance);
- *Perimeta* based hierarchical modelling of FIM processes;
- agent-based modelling for dynamic simulation of emergency response aspects of FIM.

3.2 Developing performance matrices and root cause analysis

3.2.1 Performance matrices

The development and application of performance matrices in the context of improving FIM reliability was discussed in an earlier Science Report (11206/SR) published by the Environment Agency and Defra in June 2006. Two performance matrices were developed for this prior study: one to assess the performance of 'FIM processes', another to assess performance of 'FIM outcomes'.

By drawing on the idea of assessing performance from different perspectives, as in the case of a balanced scorecard, a third performance matrix, covering 'FIM planning and readiness', was developed during the course of this current study thus broadening the ability of the performance matrices to provide an overview of FIM performance.

The study also introduced a further refinement to the tool, making it possible to categorise each element within each matrix as 'inadequate', 'adequate' or 'good' (i.e. exemplary and worth sharing as an example of best practice) when compared with the expected performance of the FIM system given the characteristics of a particular flood event. By using relative rather than absolute descriptors of performance, the tool takes into account the circumstances of the flood event (as mitigating factors) when assessing the level of performance achieved. This approach accepts that expectations of FIM performance are not absolute and may vary with the potential severity of flood risk; the design standard and reliability of flood defences; the possible flood warning lead time; and the spatial scale of flooding.

The performance matrices tested during the case study are shown in Figure 4.2 to Figure 4.4 (Chapter 4).

3.2.2 Root cause analysis

A 'fish-bone diagram' was used for root cause analysis during the case study work of this project. This form of the 'fish-bone diagram' is based on a generic form of diagram developed by Ishikawa (1990).

The generic form of diagram (Figure 2.1) shows factors that influence the reliability and performance of processes. They are grouped into six main (typical) categories: 'equipment', 'process', 'people', 'materials', 'environment' and 'management'. These category headings were adapted to suit the FIM context. The adapted categories are:

- effectiveness of available staff and resources (i.e. people);
- effectiveness of embedded procedures (i.e. 'process');
- reliability of technical systems (i.e. 'equipment');
- effectiveness of communications and information management (i.e. 'materials');
- good leadership promoting effective team-work (i.e. 'management');
- quality of the working environment (i.e. environment).

These categories could apply to any FIM process.

The fish-bone diagram was adapted for these new categories (see Figure 4.5). This new fish-bone diagram was then assessed to see if it could provide a structured framework within which specific factors influencing the performance of individual FIM processes could be identified and examined.

3.3 Developing *Perimeta* models

For *Perimeta* the performance of the FIM system is represented in the form of a hierarchy of interconnected processes. This model structure make it possible to assess how 'lower' processes within the system exert influence on the performance of 'higher' processes. The evidence on risk and uncertainties is combined at each level of the hierarchy and propagated upwards through the FIM process.

In this manner *Perimeta* shows how risks and uncertainty could first arise and then propagate thorough the various FIM activities, eventually affecting the performance of the top process within the FIM system.

Uncertainty in performance is derived using interval probability theory. The application of this theory can be used to derive a 'figure of merit' for a system (Figure 3.1). In a figure of merit evidence to support a situation (in this case, 'no property flooding') is shown as green, evidence against the situation (in this case, 'property flooding') is shown as red; uncertainty in the evidence is shown as white. The figure of merit is colloquially described as the 'Italian flag'. A related 'fragility curve' is used to assess evidence on the reliability of individual elements of flood defence systems.

Figure 3.1 is a graphical representation of an interval probability in the case of forecast water levels. An s-shaped value function has been used to represent the forecast water level and the uncertainty associated with the forecast. In this case, the green zone in the 'Italian flag' represents the evidence that properties will not flood ($S_n = 0.25$), the red zone represents the evidence that properties will be flooded ($1 - S_p = 0.28$) and the remaining white section represents the uncertainty ($S_p - S_n = 0.47$). For the sake of

simplicity in this example, it is assumed that there is no uncertainty in describing the flood warning threshold.




Figure 3.1 Example of a ‘figure of merit’, mapping uncertainty onto an ‘Italian flag’.

The key steps involved in setting up and applying a *Perimeta* model are outlined below:

- i. Identify processes within FIM for a *Perimeta* model that could be built to help assess performance;
- ii. Build the structure of the *Perimeta* model to represent those processes, identified above, that are thought to influence performance;
- iii. Identify the evidence requirements for each sub-process within the model where direct evidence on performance can be used as input data for the model;
- iv. Set weights within the *Perimeta* model to define the level of influence that each (child) process has on the process above it (parent process);
- v. Test the model using data on reliability and uncertainty for different flood events or for different flood risk management areas;
- vi. Validate the models against data on performance that is (ideally) from sources independent of the data used as inputs to the model.

Once the model is validated it can be used to test and visualise how improvements (to increase reliability and/or to reduce uncertainty) in the performance of individual sub-processes affect the performance of the entire system.

A description of each of the above steps, which we carried out with guidance from Dr John Davis (University of Bristol) who has experience of applying *Perimeta* to a range of performance management problems, is included in Chapter 4 of this report.

The development of the *Perimeta* model was an iterative process. It involved consultation between the project team and the Environment Agency staff who assisted with the case study work. This joint participative approach, in which a number of people contribute to developing model structures and assigning model parameters, is fundamental to building a good *Perimeta* model.

In developing the *Perimeta* models we accessed several sources held by the Environment Agency to gather information on the structure, performance and reliability of FIM processes structure. These included:

- Flood Incident Management Benefits Roadmaps (v1.7);
- Work instruction: 'Flood Warning Performance Measures' (Version 3, Issued 14/03/05);
- Work instruction: 'Flood Warning Levels of Service' (Version 2, Issued 05/05/06);
- Flood Risk Management process / activity diagrams;
- 2007 post flood review reports:
 - Internal Debrief Report: Thames Region Floods July 2007 (Final report, September 2007)
 - Summer 2007 Floods – Thames Region: Multi Agency Forward Look Event (Draft report and action plan, November 2007)
 - Review of Thames Region Summer Floods 2007: Technical report (Final draft, Version 2, 23rd November 2007)

3.4 Developing an agent-based model

An agent-based model is a computational method for simulating the actions and interactions of autonomous decision-making entities in a network or system, with the aim of assessing their effects on the system as a whole. Each agent individually assesses its situation and makes decisions according to a set of rules. Agents may execute various behaviours appropriate for the system component they represent (for example, members of the public and fire and rescue services will have different objectives during a flood event). At the simplest level, an agent-based model consists of a system of agents and the relationships between them. Even a simple agent-based model can exhibit complex behaviour patterns because a series of simple interactions between individuals may result in more complex outcomes than could not have been predicted just by aggregating individual agent behaviours.

3.4.1 Background to agent-based modelling

Agent-based modelling is well suited for certain studies in social science. The social sciences seek to understand not only how individuals behave but also how the interaction of many individuals leads to large-scale outcomes. Understanding a political, economic or social system requires more than an understanding of the individuals that it comprises. Moreover, it is necessary to understand how the individuals interact with each other, and how the outcomes can be more than the sum of the parts. Agent-based modelling is extremely useful when the actions and interactions of agents is dependent upon their past experience, and especially when the agents continually adapt to that experience. Mathematical analysis tends to be limited in its ability to capture dynamic consequences; agent-based modelling is one of the few practical methods of analysis that enables this to be captured.

Research shows that human and organisational responses to flood risk and flood warnings is strongly related to their prior experience of flooding and flood warnings, and also to a learning process. Agent-based modelling is therefore well suited to

simulating these kinds of system dynamics. Other approaches (that are briefly reviewed in Section 2.2.3) may provide some useful insights into FIM and are worth exploring in more detail. They may inform certain types of policy question, but they are limited in their ability to capture dynamic responses.

Not only can agent-based model capture dynamic responses, but they also have a good pedigree for testing the effectiveness of dissemination mechanisms for warnings and alerts and the susceptibility of evacuation routes to overcrowding in simulations of fire and terrorist incidents (Still, 1993; Galea *et al.*, 1996; Wong and Luo, 2005) and situations of 'panic' (Helbing *et al.*, 2000; ZARBOUTIS and MARMARAS, 2005). The use of agent-based models for these scenarios suggests that they would be appropriate for simulating flood emergencies.

Axelrod and Tesfatsion (2006) have described specific purposes for which agent-based modelling may be used:

- i. **Enhancing empirical understanding**
Why have particular large-scale irregularities evolved and persisted even when there is little attempt to manage them? Examples might include the evolution and persistence of informal, unofficial flood warning systems.
- ii. **Enhancing normative understanding**
How can the design of a flood warning or FIM system, or a FIM system component, be improved? Examples might include whether designs for flood warning policies, institutions or processes will result in socially or economically desirable system performance over time.
- iii. **Enhancing heuristic understanding**
How can greater insight be attained about the fundamental causal mechanisms in social systems? Examples might include whether or not communities adopt temporary (i.e. flood warning dependent) measures to bolster the resilience of properties, and whether the adoption of such measures will emerge from individual choices and adaptation. Is there a tipping point when individual adaptations become sufficiently common for the community to want to adopt further collective resilience measures?

Agent-based modelling is also applicable when the granularity of the model is beyond the reach of mathematical modelling. Granularity should not be confused with geographical scale, however, because agent-based modelling may be used for intensive small-scale modelling (e.g. a flood problem in a coastal settlement, as used in this study) as well as for large-scale modelling of flood policy interventions at a national level (see later). Large-scale applications need not necessarily be more data hungry than small-scale ones because in many situations groups of people, or institutions, may be represented by a single agent.

Bossomaier *et al.* (2005) describes many successful agent-based modelling simulations of complex systems in the physical and biological sciences, from climate modelling to fluid dynamics, and from genetic regulatory networks to ecosystems. However, quantitative modelling of human social systems is still in its infancy. A key challenge lies in the difficulty of modelling human behaviour. Nevertheless, as larger data warehouses of individual preferences and actions are built, our ability to model collections of people, at least in aggregate terms, will continue to improve.

Fortunately, responses to flood risk and flood warning is one of the areas of human behaviour in which the Environment Agency and others (e.g. the European Commission) have been investing to accumulate such data using social surveys.

3.4.2 Developing an agent-based model for FIM

As far as we are aware there have been very few applications of agent-based modelling to aspects of flood management or FIM; this work is high experimental and work in progress.

The research of Brouwers and Verhagen (2003) and Brouwers *et al.* (2009) has developed a geographical explicit dynamic model. Its main purpose is to investigate the possibilities for a national Hungarian flood insurance programme. Agents within the model included the government, insurers and individuals; different levels and types of insurance cover were modelled. Using one set of assumptions the insurer agent performed well and avoided insolvency, government expenditures rose because of increased compensation payouts, and the more vulnerable individuals could not afford flood insurance. Basic assumptions regarding insurance premiums and coverage were then altered, leading to diminished income for insurers, increased government expenditure and extremely adverse effects on some individuals. The modelling demonstrated how system-scale changes emerge from seemingly small differences.

Of more direct relevance to this project is the MassVac evacuation model (Hobeika and Jamei, 1985) and the Life Safety Model (Johnstone *et al.*, 2005) which were developed to estimate casualties under dam-break scenarios. As part of the EC-funded FLOODsite project, the Life Safety Model was applied in Canvey Island (London) and calibrated to reproduce the observed casualties in 1953 (Lambruso *et al.*, 2008).

The application of agent-based modelling to FIM has been implemented through the following processes:

- i. Identifying the key processes, infrastructure, components and actors associated with FIM.
- ii. Identifying the relationships and feedbacks between the different FIM components to develop a generic conceptual (spider) diagram describing FIM systems.
- iii. Collating evidence for quantifying the relationship between each link in the spider diagram.
- iv. Coding the model and implementing a case study demonstration using the outputs of the above steps to input the correct parameters into the model.
- v. Analysing the model and demonstrating the risk analysis methodology.

Steps i–iii above are shared with the construction process of the *Perimeta* model. Given the resources assigned to this component of the project, it was not feasible to build a complete agent-based representation of the FIM system and all its processes. Therefore a subset of FIM components (predominantly flood warning lead time and evacuation) was selected to demonstrate the capabilities of the agent-based model and a risk-based approach to FIM.

3.5 Evaluating the benefits of FIM

FIM is recognised as one component in a range of activities carried out by the Environment Agency to manage flood risk. Flood risk planners and managers recognise that they are dealing with dynamic '*risk producing-risk response*' systems; natural and man-made components interact to generate floods and respond to flood risks. Flood risk management is best achieved through well integrated combinations of structural and non-structural measures, which offset each other's disadvantages and are matched to the specific requirements and characteristics of particular locations.

FIM helps to reduce the physical, economic and environmental damage caused by flooding through a variety of activities, including issuing flood warnings. Flood warnings are the most significant element of FIM because they enable timely actions to be taken to increase the amount of flood damage avoided through:

- *structural measures* such as the operation of flood defence systems by the Environment Agency (and others);
- *non-structural measures* such as encouraging the recipients of flood warnings to take appropriate action to reduce flood damage.

The benefits generated by flood warnings can be measured using a form of 'flood damage avoided' (FDA) equation (see below). This form of the FDA equation is taken from *Flood Warning Performance Measures* (Version 3, Issued 14/03/05); and *Flood Warning Levels of Service* (Version 2, Issued 05/05/06). The origins of the equation in this form, and its application using data collected from surveys of flooded households, is described by Parker *et al.* (2007).

The equation provides a means of converting the level of performance (in each of a set of FIM processes) to FIM benefits in terms of flood damage avoided:

$$FDA \text{ (flood damage avoided)} = AAD \text{ (annual average damage)} \times DR \text{ (damage reduction)} \times C \text{ (coverage)} \times r \text{ (service effectiveness)} \times RA \text{ (availability)} \times PR \text{ (ability)} \times PE \text{ (effective action)}$$

Of particular interest to this study is the use of *Perimeta* to assess the non-structural flood risk reduction benefits of FIM (i.e. the influence of FIM on flood damage reduction through flood warnings).

Chapter 4 describes in detail the case study tests in which *Perimeta* was used to derive evidence-based values for the performance of selected terms in the FDA equation (Table 3.1). Whilst *Perimeta* could in principle be applied to calculate values for many of the terms of the FDA equation, we focused our attention on those terms in the equation that are directly influenced by Environment Agency FIM activities: *coverage* and *service effectiveness*. The remaining terms (*availability*, *ability*, and *effective action*) are all indirectly affected by the Environment Agency FIM activities. Assessing these terms would require data collected by the Environment Agency to be combined with data from other sources.

Table 3.1 Potential for applying *Perimeta* to the terms of the FDA equation.

Terms in FDA equation	Definition	Evidence from	Influenced by Environment Agency FIM processes?	Potential for applying <i>Perimeta</i> to model each term in the FDA equation
Annual average damage (AAD)	Potential flood damage	National Flood Risk Assessment (NaFRA) and other studies	Directly, via operation of flood defence systems	Data derived from other sources
Damage reduction factor (DR)	Variation of % potential flood damage avoided with lead time of flood warning provided	Studies to relate flood warning lead times with potential to reduce flood damage	Directly, by providing timely and effective flood warnings	Data derived from other sources
Coverage (C)	Proportion of 'properties at risk' 'offered a flood warning service'	Area data and estimates	Directly, by offering a flood warning service to flood prone areas	Yes
Service effectiveness (r)	Proportion of 'flooded serviced properties' sent a flood warning	Post-event data; forward looking assessments	Directly, by providing effective flood warnings	Yes
Availability (RA)	Proportion of properties sent a flood warning that received a warning	Confirmed receipt of flood warning	Indirectly, by raising public awareness	Yes – provided sufficient input data is available
Ability (PR)	Proportion of serviced properties able to respond to flood warnings	Public opinion surveys	Indirectly, by raising public awareness	Yes – provided sufficient input data is available
Effective action (PE)	Proportion of serviced properties that took effective action	Post-event response survey	Indirectly, by raising public awareness and emergency responses initiated	Yes – provided sufficient input data is available

3.6 Applying the tools to provide decision support

The tools were tested on different types of flood incident – the *overview* tools and *Perimeta* used data for fluvial flood events, and the agent-based modelling was tested in the context of a coastal flood incident. The case studies provided insight into how these tools can support and complement one another, as described below.

3.6.1 Spatial scale

All three methods can in principle be applied at a range of spatial scales, from an operational area used in flood risk management down to an individual flood risk zone within such an operational area. In the case studies the *overview* and *Perimeta* tools were tested for an entire operational flood risk management area, whereas the agent-based modelling was applied at a more local level.

The most appropriate scale of application depends on:

- the decisions that are to be informed by the modelling work;
- the processes that must be represented to adequately model the relevant phenomena.

The overview methods can be applied over the widest range of spatial scales; *Perimeta* is probably best applied at the scale of an operational area. Agent-based models are probably most appropriate for modelling flood risk areas. However, as described previously, agent-based models could in principle be constructed to represent different processes, geographical scales and decisions.

The most appropriate scale of application of each of the tools is influenced by:

- the scale at which the particular FIM planning decisions are taken;
- factors relating to the effect that the failure of FIM processes has on FIM performance (e.g. the failure to deliver flood warnings in time to a flood warning area will affect performance at a local rather than at a catchment level);
- the level of detail at which we can represent specific risks and uncertainties that can influence FIM performance – the selection of too coarse a model resolution may mask the influence of specific factors;
- the level of detail at which indicators and data on FIM performance are typically measured;
- the amount of input data required – applications that require data at a high resolution will be limited in terms of the spatial area over which they can be applied effectively.

3.6.2 The complementary use of the tools

The *Perimeta* case study demonstrated the benefit of using performance matrices in advance of building the *Perimeta* model. The performance matrices helped to confirm the FIM processes that had caused concern during recent flood events. The builders of the *Perimeta* models were then able to ensure that these processes were represented at an appropriate level of detail in their models. The structure of the fish-bone diagram

proved to be useful in helping develop an appropriate structure for the *Perimeta* models.

Figure 3.2 shows that:

- performance matrices can be used to gain an overview and identify which of the many processes within FIM are performing adequately well or not, thereby indicating the FIM processes (or sub-systems) which should be assessed in more detail;
- specific causes of inadequate performance can be identified using root cause analysis, which identifies more clearly those processes (or sub-systems) that are contributing to poor (or good) performance;
- the insight gained from the applying performance matrices and root cause analysis can also be used to structure and provide information for the construction and application of *Perimeta* and the agent-based models;
- outputs from *Perimeta* and agent-based models can be used in a complementary way, via a suitable value function, to test improvements to FIM and to evaluate the benefits of improved performance.

The outputs from *Perimeta* models provide information on the reliability and performance of flood warning. This output can be used as input data to an agent-based model that simulates how the evacuation of members of the public is affected by variations in flood warning performance (i.e. lead-time and accuracy). The response of agents (in the agent-based model), and the outcomes of the flood incident, depend on the quality and coverage of the flood warnings as shown by the Towyn case study.

The subsequent behaviour of the agents can also be influenced by factors such as their *availability* to receive flood warnings; their *ability* to respond; and the *effectiveness* of their response actions. These factors are all represented in the FDA equation and influence the benefits derived from good FIM (see Section 3.5).

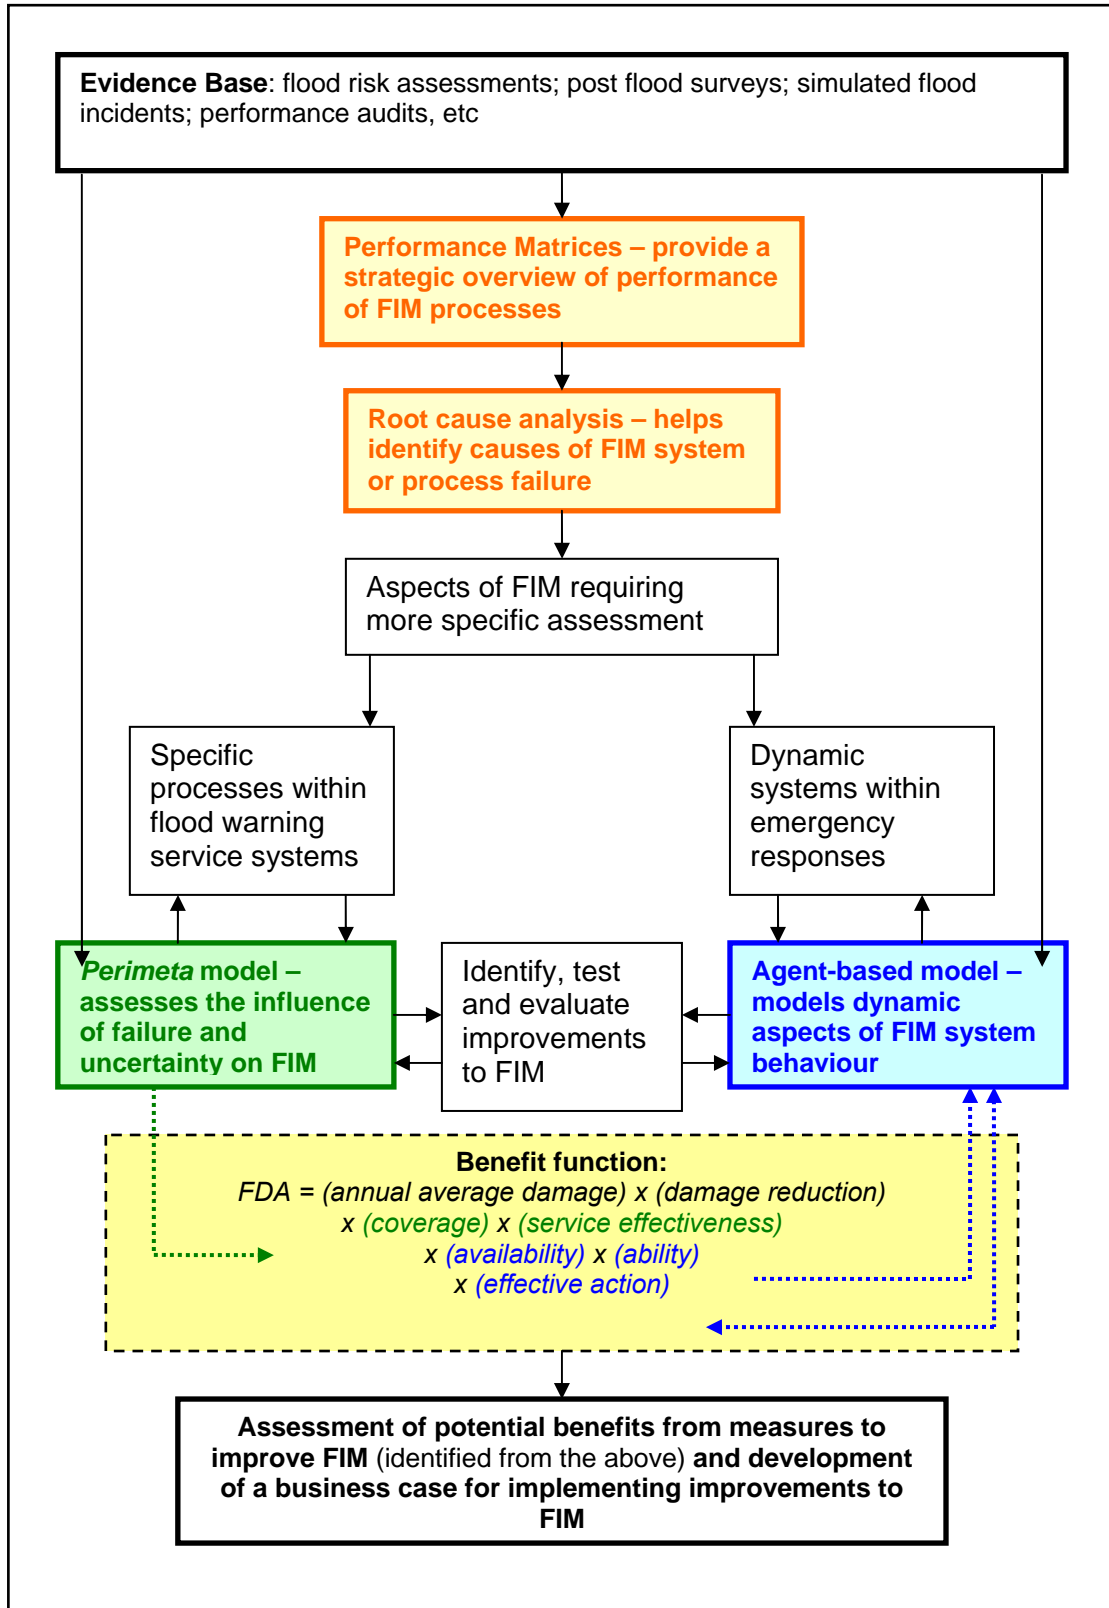


Figure 3.2 Conceptual integration of the tools within a decision-support framework.

3.6.3 Integrating the results from *Perimeta* and agent-based modelling

Outputs from the agent-based model can be used to inform performance indicators in the *Perimeta* model; this enables the detailed representation of certain processes to be better enumerated in *Perimeta*. Conversely *Perimeta* performance indicators can be used to help parameterise an agent-based model, as illustrated below.

Figure 3.2 indicates how results from *Perimeta* and agent-based modelling can provide parameter values for specific terms in the FDA equation (see Part A: Guidance, Section 4.7). *Perimeta* provides information on the terms relating to flood warning 'coverage' and 'service effectiveness' terms.

On the other hand, the agent-based modelling makes specific assumptions regarding the availability and ability of agents to respond to flood warnings, and simulates their subsequent action(s) –the model can use information on agents' 'availability' and 'ability', and provide insight into the 'effectiveness of the subsequent actions' of agents.

The outputs from these tools will be probabilistic and indicate uncertainty. In a subsequent benefit (or business-case) analysis, they can be converted to 'expected' values. This informs a deterministic estimate of the flood damage avoided. Alternatively they can be used in a probabilistic form for a probabilistic evaluation of flood damage avoided. In either case, the incremental benefit of improvements to FIM can be estimated by subtracting the flood damage avoided 'after improvements to FIM' from the 'base case' flood damage avoided.

In principle *Perimeta* could be used at a higher level to integrate results on performance for each term in the FDA equation. *Perimeta* could combine evidence from different sources including agent-based models and post-flood surveys. This use of *Perimeta* was proposed in the GENESIS (Generic Process for Assessing Climate Change Impacts on the Electricity Supply Industry and Utilities) study described by Walsh *et al.* (2007).

The above approaches can also be applied when using an improved benefit equation such as the Flood Warning and Response Benefits Pathways (FWRB) model, as described in Section 4.7.1 of Part A: Guidance. In this case, the outputs from *Perimeta* would provide information on the 'contents moved or evacuated' component of the FWRB model. An agent-based model set up to explore the consequences of, for example, self-evacuation by road following receipt of a flood warning can be used to inform the 'evacuation' term of the FWRB model.

The combined use of *Perimeta* and agent-based modelling could generate outputs that could be represented graphically in the form of 'Italian flags' (see Figure 3.3). The figure shows how the performance of an emergency response, measured in terms of the number of people exposed to flood water depths of greater than 20cm, varies in relation to the percentage of the population at risk that receives an effective flood warning.

In this graph the red portion indicates poor performance, the green portion shows good performance and the white band indicates the uncertainty in the assessment. A sensitivity analysis of this type, using the agent-based model, provides information of relevance to the *Perimeta* model. It relates a decision of interest to flood incident managers (e.g. what coverage of the floodplain do we need?) to a measure of a successful (or not) outcome (e.g. the number of people exposed to a dangerous depth of floodwater).

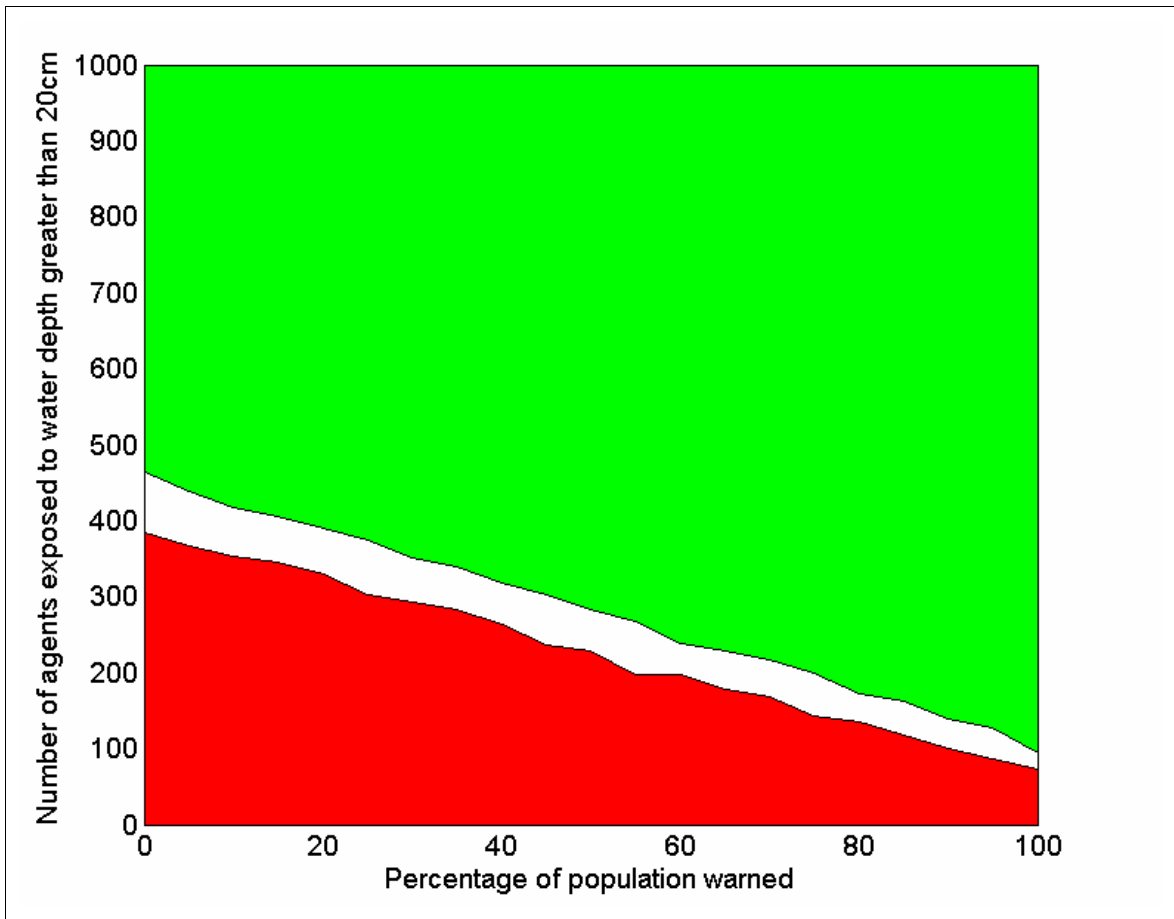


Figure 3.3 An 'Italian flag' showing the extent to which an outcome (i.e. the number of agents exposed to flooding) is conditional on the coverage of the flood warning system (i.e. percentage of population warned).

4 Testing the tools

4.1 Introduction

This chapter describes how the tools were tested using case studies from two locations. The Thames Region case study relates to fluvial flooding; the Towyn case study is concerned with coastal flooding in North Wales.

Several factors were considered to select the most appropriate case study locations and associated flood events, including:

- the availability of sufficient data and information on the flood event and FIM performance;
- the availability of Environment Agency staff from the local area to assist with the case study development (if necessary);
- the relative severity of flood events (testing events of different magnitudes will help demonstrate the flexibility of the approaches);
- any prior knowledge within the project team of the area and/or events.

The Thames Region, and the West Area in particular, was severely affected by flooding in summer 2007 (Figure 4.1). The severity of the flood event and the fact that it is a relatively recent event meant that there was a large amount of evidence available. JBA Consulting was involved in the post-flooding reviews, and many of the Environment Agency staff who were involved in the management of the event were still working in the same roles at the time of this study. There was therefore a large amount of information and knowledge available for this particular flood event, making it an excellent event on which to base one of the case studies.

A smaller flood event affected the same area in January 2008; this was selected as the second case study event to provide contrast to the larger summer event in 2007. By using the same area for each case study it was possible to test the models on floods of differing severities and assess how FIM performance had changed over time. Performance matrices and root cause analysis were examined and the *Perimeta* models were tested using data and information on FIM performance.

The Towyn case study covered a site in North Wales. It was selected as a good example of an area where evacuation is an important feature of emergency response planning. The project team already had access to the key data and flood models necessary to set up an agent-based model.

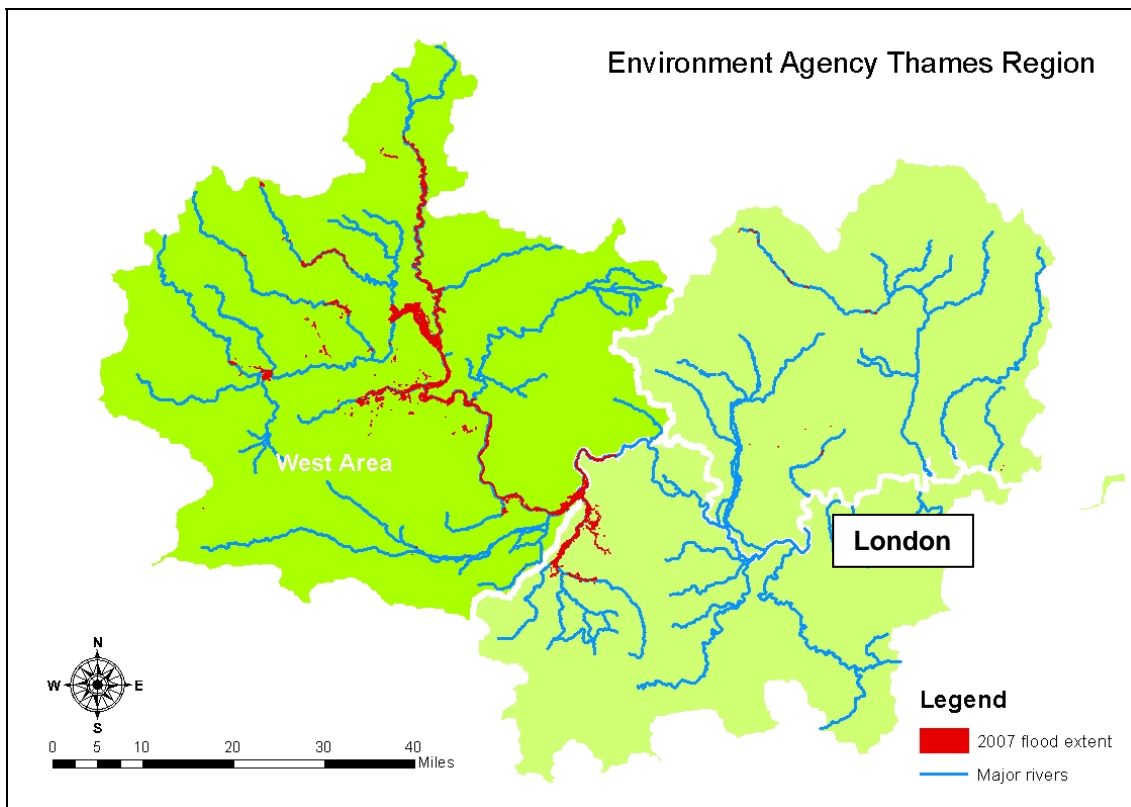


Figure 4.1 Location of West Area (Thames Region).

Towyn FIM policy issues

Since the 1990 flood in Towyn, the North Wales Flood Defence Group has recognised that flood risk management must go beyond primary prevention and are only satisfactorily achieved when internal structural and non-structural measures are deployed in a balanced effort to increase community resilience.

Evacuation was a very real issue in 1990; many people were picked up from water around or above waist height. In one case, an upstairs flat acted as a temporary safe haven for 20 people for five hours. Outside of Towyn, overtopping caused smaller local scale flooding of roads and property.

The local authority is therefore seeking to build wider community resilience through spatial planning, building regulation, warning systems and evacuation contingency plans. The evacuation planning is taking into consideration the vulnerability of the population (e.g. age, mobility) as well as access and exit roads. Options under consideration are to expand the capacity of key roads out of the floodplain and to construct additional multi-storey buildings in strategic locations within the floodplain to act as temporary flood shelters.

The local council has identified the potential of the agent-based model to inform its community resilience planning. The council has been in discussion with the development team for the agent-based model regarding a more detailed study and testing of FIM options.

4.2 Applying the performance matrices and root cause analysis

The two overview methods (performance matrices and the root cause analysis fish-bone diagram) were tested during the initial stages of the case study work using data for the West Area of the Thames Region.

4.2.1 Performance matrices

We tested the application of performance matrices for three aspects of FIM, namely:

- 'planning and readiness';
- 'process – operational performance';
- 'outcome – operational performance'.

In particular we:

- reviewed the structure and descriptions of individual components of each matrix;
- adopted the terms 'inadequate', 'adequate' or 'good' (i.e. exemplary and worth sharing as an example of best practice) when performance was compared with reasonable expectations of FIM performance given the characteristics of the flood event;
- tested the matrices with some recent post-flood data.

We created the performance matrices shown in Figure 4.2 to Figure 4.4 to assess post-event flood data for the widespread July 2007 flooding in the West Area of the Thames Region and the smaller flood event in the same area in January 2008. The ellipses on the performance matrices show the assessed level of performance of the different elements of FIM processes during each flood event.

The results indicated how well different aspects of FIM performed in each event. They also indicate elements of FIM where improvements could be considered. Insight gained through using the performance matrices is discussed in Chapter 5.

Element	"Inadequate" FIM	"Adequate" FIM	"Especially Good" FIM	Elaboration and categorisation
Availability & readiness of trained staff	Poor state of availability and preparedness	Good availability & readiness levels	Excellent availability & readiness levels	Categories: staff FIM training; staff rotas; holiday rotas; experienced staff availability
Embeddedness of procedures	Some key procedures yet to be fully embedded	Minor lack of embeddedness	No lack of embeddedness at all	Categories: key procedures (e.g. communication with HQ) are codified and embedded in practice
Availability & readiness of appropriate equip	Significant availability, serviceability & readiness problems	Minor issues	All equipment is appropriate, serviced and ready to go	Categories: demountables; watercourse maintenance equipment, vehicles etc.
In-house simulated performance	Significant performance issues	Minor performance issues	Negligible performance issues	Categories: local, regional, national rehearsals and performance in each, desk-top/other etc.
Relations with professional partners	One or more key relationships is problematic	Minor issues	Excellent - no issues at all	Categories: relations with Category 1 and 2 responders
Multi-agency exercise performance	Significant performance issues	Minor performance issues	Excellent - no performance issues at all	Categories: local, regional, national rehearsals and performance in each, desk-top/other etc.
Readiness of flood risk & socio-economic database	Major data gaps; no GIS based system operational	Minor data gaps; GIS based system operational but improvements needed	No data gaps; excellent GIS system operational	Categories: locational data; Agency assets; community assets; vulnerable people etc.
Health and safety	Significant H&S issues	Minor H&S issues	No H&S issues	Categories: Office-based staff, Operations staff, staff Involved in emergency actions

Key: 2007 2008

Figure 4.2 'Planning and readiness' performance matrix.

Element	"Inadequate" FIM	"Adequate" FIM	"Especially Good" FIM	Elaboration and categorisation
Immediate preparedness	Poor state of preparedness	Good levels of preparedness	Excellent state of preparedness	Categories: awareness raising; training; pilot exercises; appropriate documentation; contingency plans
Forecasting	Flood forecast very poor in terms of accuracy	Some forecasting accuracy problems	Floods forecast accurately in terms of time and extent	Categories: Detection; forecasting; dissemination; response; review
Warning and promoting response	<50% of warnings received in time for sensible response	75% of warnings received in time for sensible response	95% of warnings timely for sensible response	Categories: Internal warning messages; dissemination; response; review
Other communication	Major communication problems	Some communication problems	No major communication problems	Categories: Inside Agency; outside Agency; Met Office; between responders; to/from public; from gauges; to sirens
Co-ordination	Distinct lack of co-ordination	Good co-ordination	Excellent co-ordination; no problems at all	Categories: Bronze, Silver, Gold Command; DEFRA /COBRA; professional partners; voluntary bodies
Media management	Media poorly managed	Good media management	No major media management problems	Categories: National TV; local TV; national radio, local radio; printed press locally /nationally; internet; others
Equipment provision	Clear equipment shortcomings	Some equipment lacking	All relevant equipment available and in place	Categories: Detection; data processing; communication power; flood fighting (demountables); other equip.
Equipment reliability	Significant reliability problems	Some reliability problems	Equipment reasonably reliable	Categories: Detection; data processing; communication power; flood fighting (demountables); other equip.

Key: 2007 2008

Figure 4.3 'Process – operational performance' performance matrix.

Element	“Inadequate” FIM	“Adequate” FIM	“Especially Good” FIM	Elaboration and categorisation
Environmental damage	Significant environmental damage	Minor environmental damage	Negligible environmental damage	Categories: flora; fauna; landscape; archaeological remains; others
Economic damage (external)	Greater than Multi-Coloured Manual damage	Less than Multi-Coloured Manual damage	Much less than Multi-Coloured Manual damage	Categories: Residential; non-residential; public buildings; infrastructure damage (utilities); beyond floodplain affects; others
Damage to Agency’s assets	Significant damage to Agency assets	No more than to be expected	Much less than expected	Categories: Agency installations (gauges; depots; stores; offices); communication equipment; flood defences; other
Injuries	Greater than 5 in 1000	Equal to or greater than 1 in 1000	None	Categories: Immediate and directly flood related; medium term effects (direct and indirect)
Loss of life	Equal to or greater than 1 in 1000	None	None	Categories: Age; gender; Social class; local/visitor; prior health status
Victim trauma	Greater than 25 in 1000 people affected	Equal to or greater than 5 in 1000 people	No people affected	Categories: Age; gender; Social class; local/visitor; prior health status
Reputation of Agency and others	Reputation damaged	No effect on reputation	Reputation enhanced	Categories: Public nationally; public locally; politically (MPs; DEFRA etc) with Professional partners

Key: 2007 2008

Figure 4.4 ‘Outcome – operational performance’ performance matrix (refers to those outcomes influenced by FIM).

4.2.2 Root cause analysis and the use of the fish-bone diagram

Root cause analysis is typically used to investigate the root causes of critical failures in systems. The West Area case study did not provide any examples of such failures in FIM during the floods of 2007 and 2008 so the fish-bone diagram (see Figure 4.5) could not be tested in this way.

Fish-bone diagrams can also be used to help the structuring of *Perimeta* models of individual FIM processes (the fish-bone diagram is a form of hierarchical network diagram). Its hierarchical structure was found to be particularly helpful in constructing the template for the *Perimeta* ‘Service Effectiveness’ model, highlighting processes which may otherwise have been missed out or duplicated.

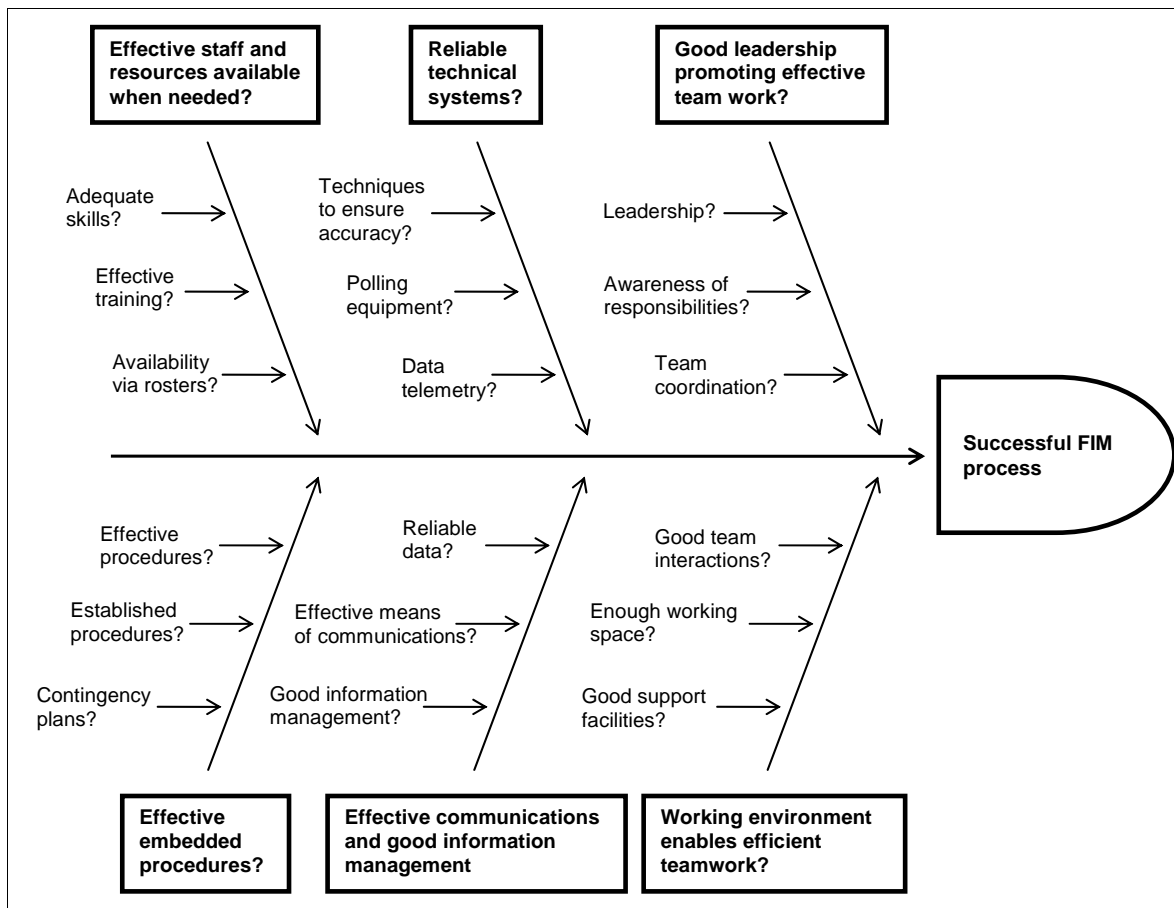


Figure 4.5 Illustration of a generic ‘fish-bone diagram’ that can be applied to analysing failures in a process within FIM.

4.3 Testing the *Perimeta* models

Table 4.1 summarises the meetings held with Environment Agency staff during the development of the case studies. As already explained, a number of internal project team meetings were held during the course of the model development and case study testing. These meetings supplemented the formal meetings detailed in the tables as part of the iterative model building process.

Table 4.1 Case study meetings with the Environment Agency.

Session	Date	Purpose
1	07/10/08	Provide the Environment Agency with background information to the project Explain what assistance and information we will require from the Environment Agency during case study work Arrange dates for further case study meetings
2	11/11/08	Introduce the Environment Agency to the FIM project, including background and key aims Introduction to <i>Perimeta</i> as a modelling tool Assessment of the overview level methods (performance matrices and fish-bone diagram) Initial views on <i>Perimeta</i> model structure for ‘coverage’ and ‘service effectiveness’ Discussion of key performance measures used to assess FIM performance and how these could potentially be used within the model structure Data / information availability for providing evidence within <i>Perimeta</i> Decision on the spatial scale of model application and selection of case study areas
3	19/11/08	Iterative development of <i>Perimeta</i> model structure
4	27/11/08	Agreement of spatial scale at which <i>Perimeta</i> models will be applied, as well as specific events for case study focus
5	18/12/08	Revision of necessity and sufficiency values Incorporation of evidence into <i>Perimeta</i> models
6	30/01/09	Summary of model refinements since previous case study sessions Discussion of <i>Perimeta</i> model application
7	10/02/09	Discussion of key findings of the case study work and lessons learnt Model validation and sensitivity

An overview of the process followed to develop the *Perimeta* models is given in Figure 4.6.

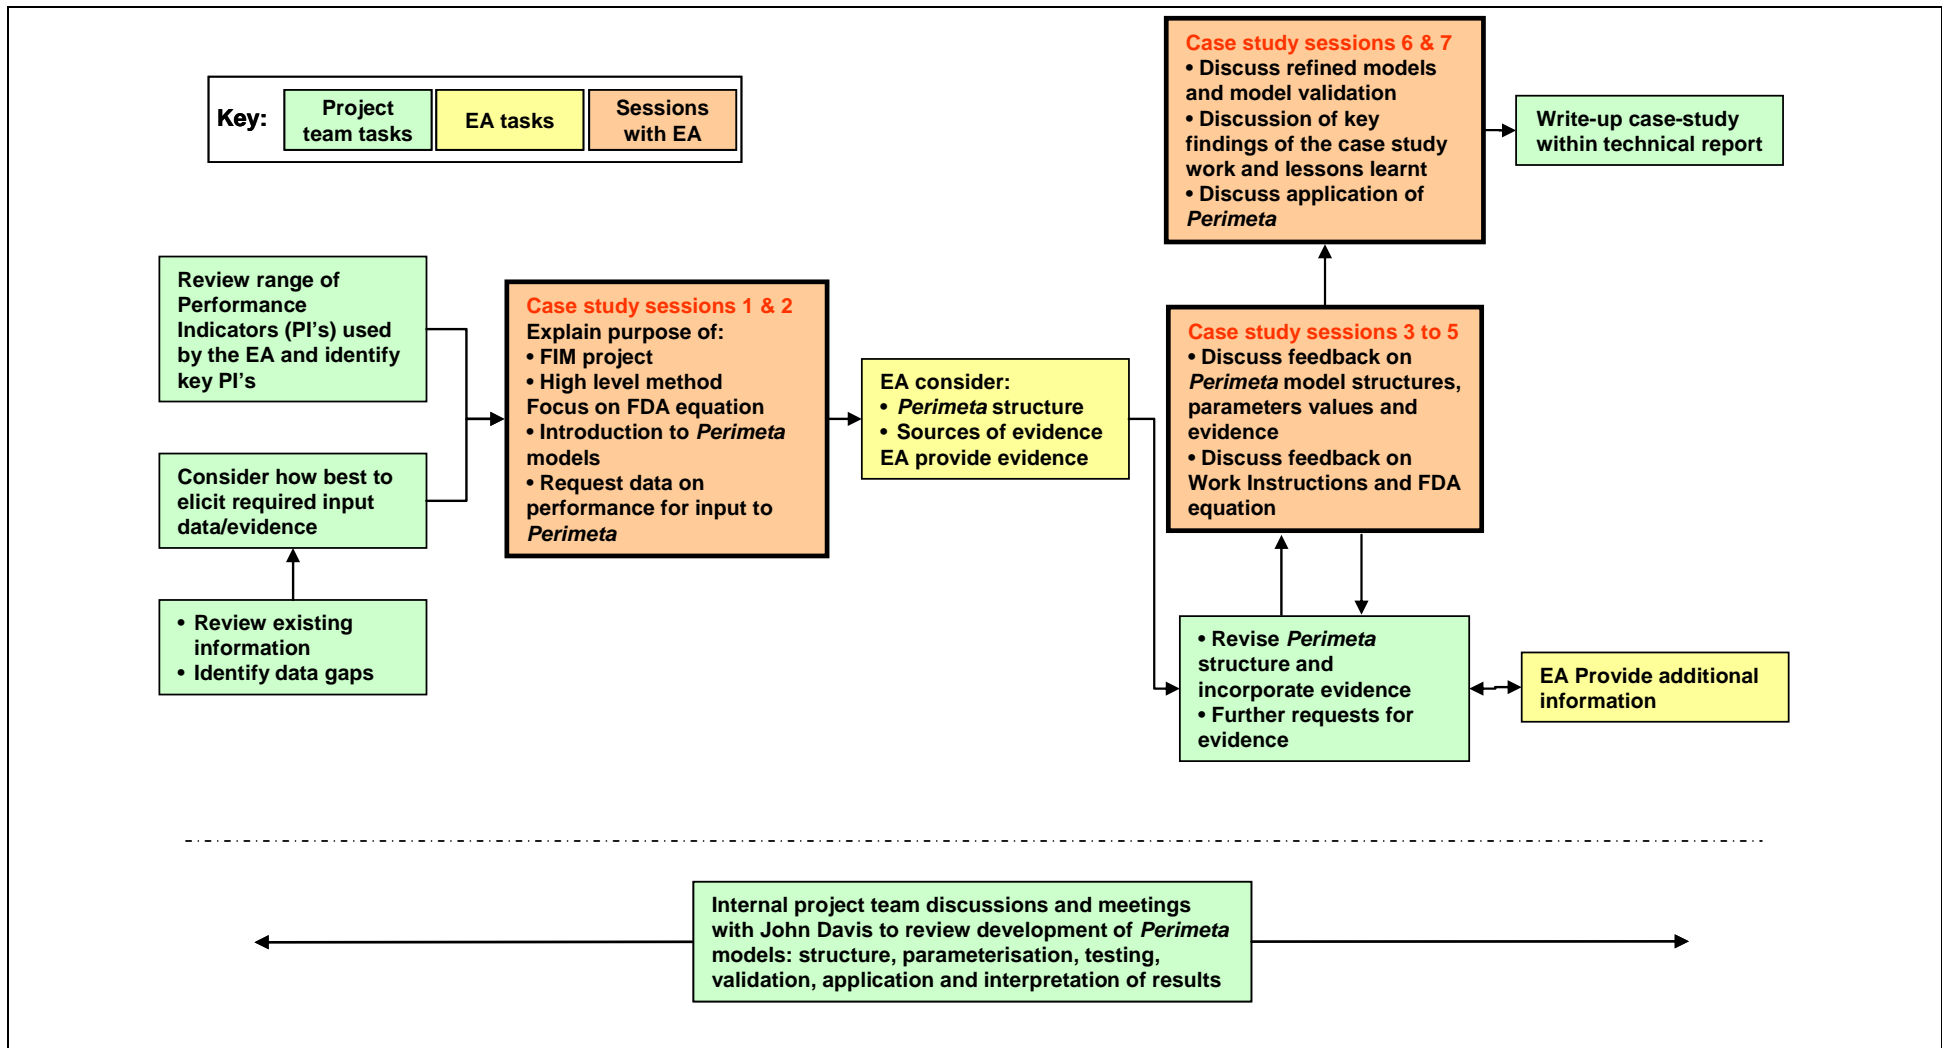


Figure 4.6 Approach used for the *Perimeta* case study.

4.3.1 Building the Perimeta model structure

Constructing the 'Coverage' and 'Service Effectiveness' *Perimeta* models was an iterative process which required considerable input from both the project team and Environment Agency staff. Discussions during several model-building sessions helped shape the models and produce a structure in which everyone involved had confidence.

Initially the models are best constructed in a 'top-down' manner, working from the overall aim of the model down through the processes that contribute to the achievement of this aim. Once a working template is built, evidence can then be entered into the models. This is best carried out with a 'bottom-up' approach as evidence is input via child processes at the bottom of the model and then propagated up through the model hierarchy.

Initial analysis of the processes involved suggested the 'Coverage' model would be the least complex, so this model was constructed first. It also made sense to assess the 'Coverage' element of the FDA equation prior to the 'Service Effectiveness' element because in practice the service cannot be effective if there is no coverage.

Structure of the 'coverage' model

The Flood Warning Performance Measures work instruction document defines the 'Coverage' term within the Flood Warning Service Limit, as

$$(Number\ of\ serviced\ properties) / (Number\ of\ properties\ at\ risk)$$

The top (or 'root') process of the coverage model indicates the strength of evidence in support of the fact we can provide 'coverage' in line with the target levels of service (LoS) for flood warning. In other words, it represents the degree of confidence that we have in being able to meet the performance target we are aiming to satisfy in terms of flood warning 'coverage'.

The structure of the coverage model (as shown in Figure 4.7) is based on the AMS Flood Warning Levels of Service guidance (Version 2, Issue date 05/05/06); this document contains performance measures for the FIM processes that make up the coverage element of the service. The coverage model indicates how successful the FIM processes are at achieving the Environment Agency's coverage targets. The model breaks this aim down into two aspects:

- i. **Having the ability to provide a flood warning service (FWS)**
This measure the ability to provide accurate, timely and reliable flood warnings.
- ii. **Having the ability to communicate these warnings**
Communication must also be accurate, timely and reliable – and in compliance with LoS targets.

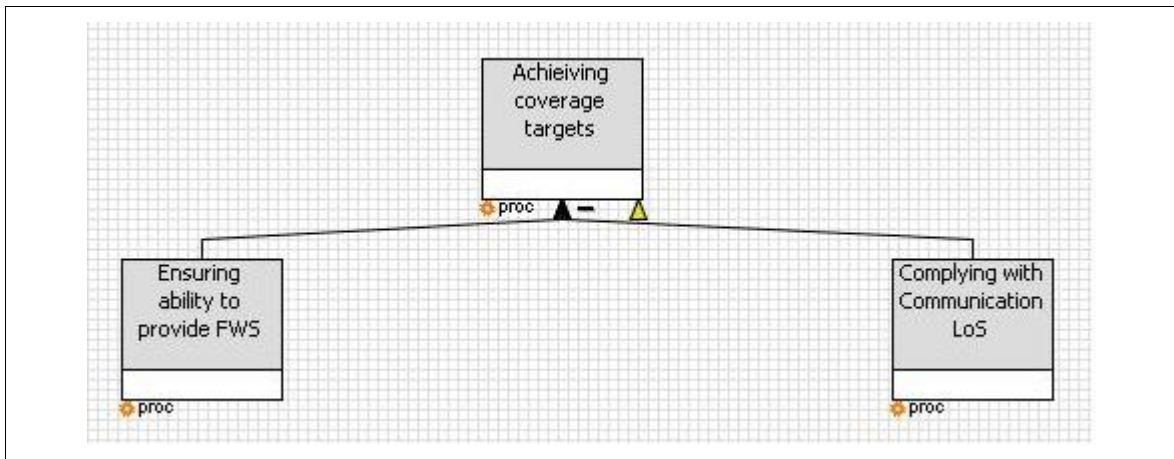


Figure 4.7 Aim of the coverage model.

The first of these aspects (i.e. provision of a FWS) is divided into three elements, namely:

- flood detection;
- flood forecasting;
- warning dissemination.

This division is made on the assumption that if FIM complies with the LoS requirements for each of these elements it will be possible to provide accurate, timely and reliable flood warnings. The model provides a further level of detail which incorporates the actual evidence to show the level of compliance. The structure of the three elements is illustrated in Figure 4.8.

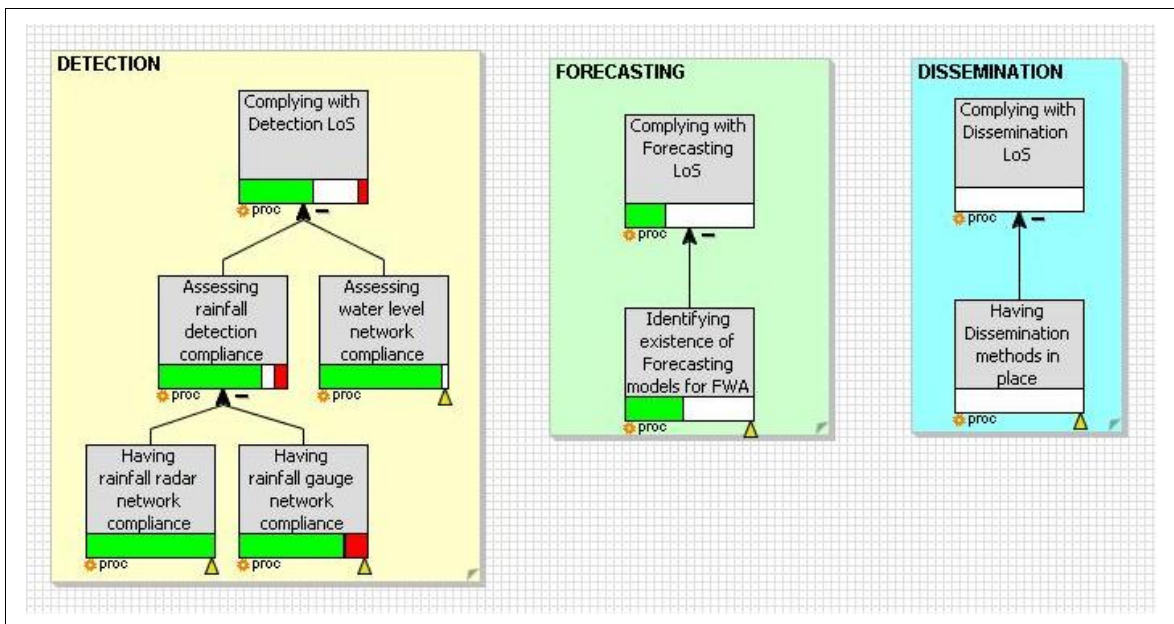


Figure 4.8 Coverage model: detection, forecasting and dissemination elements.

The modelling of the second aspect (i.e. communication of warnings) focuses on the communication of the flood warning to its intended recipients. This does not have the separate elements of the first aspect, but does contain the same level of detail at which the actual evidence is inserted (see Figure 4.9).

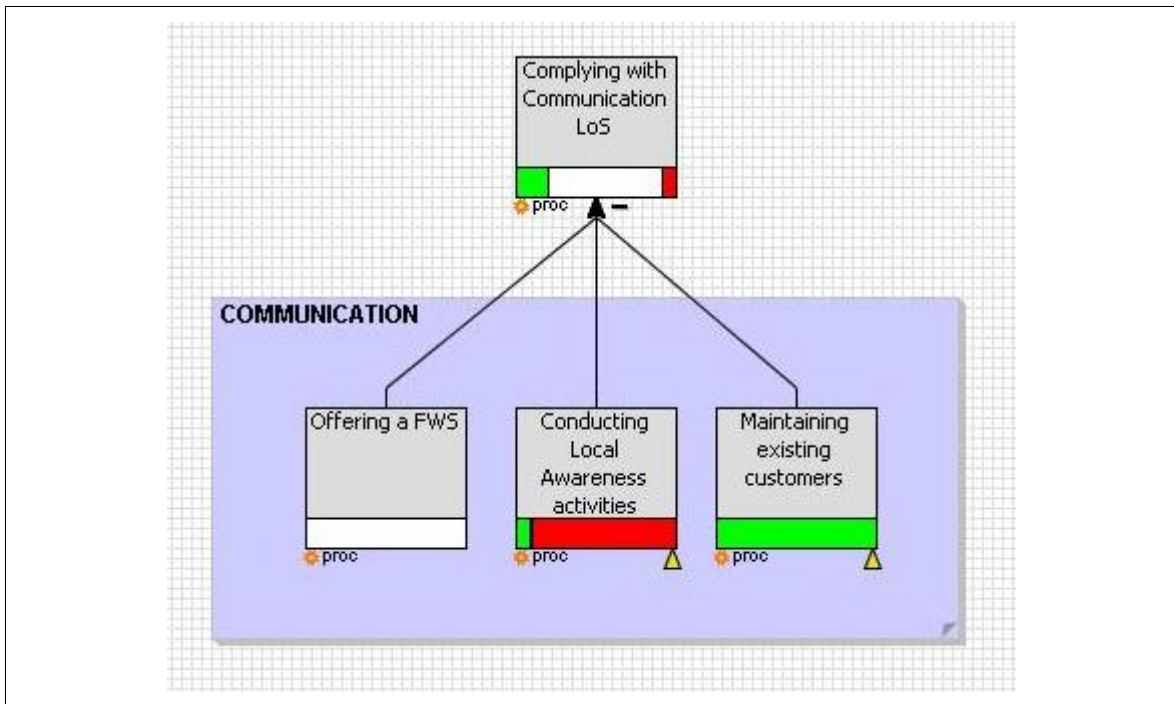


Figure 4.9 Coverage model: communication.

There is one further process in the coverage model which applies both to the three provision elements discussed above and the communication aspect. This is the identification of the appropriate LoS. Without this process it would be impossible to measure the level of compliance.

Building the *Perimeta* coverage model helped to hone our understanding of the interaction between, and the relevant importance of, the various FIM processes. The model was restructured several times as our understanding of the links and processes grew. *Perimeta* proved to be a useful tool that enabled us to break down the FIM service into its component parts and identify the most significant elements that could affect performance of this part of the FIM process. An example of the full coverage model is shown in Figure 4.15.

Structure of the 'service effectiveness' model

The Flood Warning Performance Measures work instruction document defines 'service effectiveness' as the proportion of flooded, serviced properties that were sent a flood warning, i.e.:

$$\frac{\text{(Number of flooded serviced properties sent a warning)}}{\text{(Number of flooded serviced properties)}}$$

The top (or 'root') process of the service effectiveness model indicates the strength of evidence in support of the Environment Agency's ability to measure and report with confidence the effectiveness of the FWS.

Measuring the effectiveness of the FIM service is significantly more complex and inherently more subjective than measuring the 'coverage'.

Our model for FIM service effectiveness began by reviewing the elements of the coverage model to assess whether they would also be relevant to a model of service effectiveness. This required us to expand and convert the elements and introduce a qualitative assessment. The aim of the service effectiveness model is to measure and

report the effectiveness of the FWS. The structure of the final model recognises that in order to achieve this aim:

- accurate, timely and reliable flood warnings must be sent;
- flooded, serviced properties must be identified.

These processes are independent of each other; their lack of interaction dictates their position in the model structure. Neither feeds into the other, but both are required if the overall aim of the model is to be achieved. This criterion has been applied throughout the model build to indicate where processes and sub-processes are required.

The major part of the service effectiveness model focuses on the sub-processes which contribute to the first of these two independent processes (i.e. sending accurate, timely and reliable flood warnings). By contrast, the only child processes which contribute to the identification of flooded, serviced properties are data collection processes.

The process of sending accurate, timely and reliable flood warnings is also split into two independent sub-processes, namely:

- ‘having effective technical systems’;
- ‘making good flood warning decisions’.

Again, the independent nature of these two sub-processes dictates their position in the model. The model structure described above is illustrated in Figure 4.10.

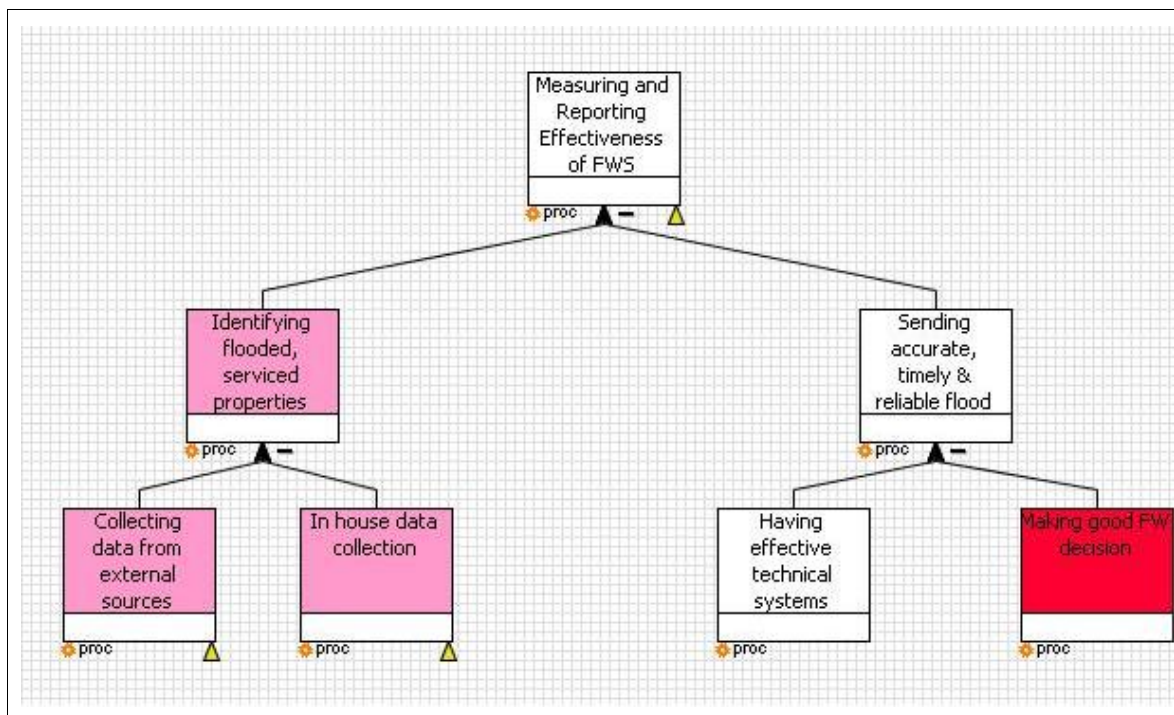


Figure 4.10 Service effectiveness model: aim and top processes.

Both processes – ‘having effective technical systems’ and ‘making good flood warning decisions’ – sit above a large number of sub-processes. These are described below.

The area of the model that investigates technical systems follows a similar pattern to the coverage model: three sub-processes cover the areas of detection, forecasting and dissemination. In this model, however, it is the ability of the systems to function effectively that important, not the simple existence of the systems. The model structure used to represent the three systems is shown in Figure 4.11 to Figure 4.13.

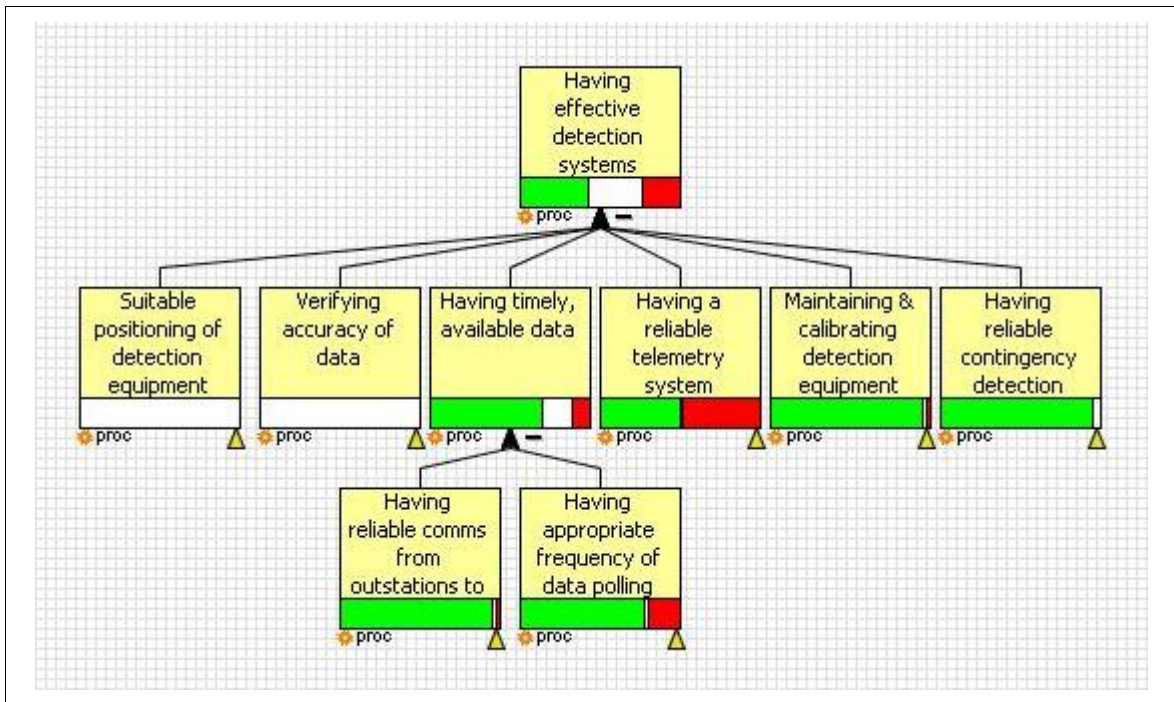


Figure 4.11 Service effectiveness model: effective detection system.

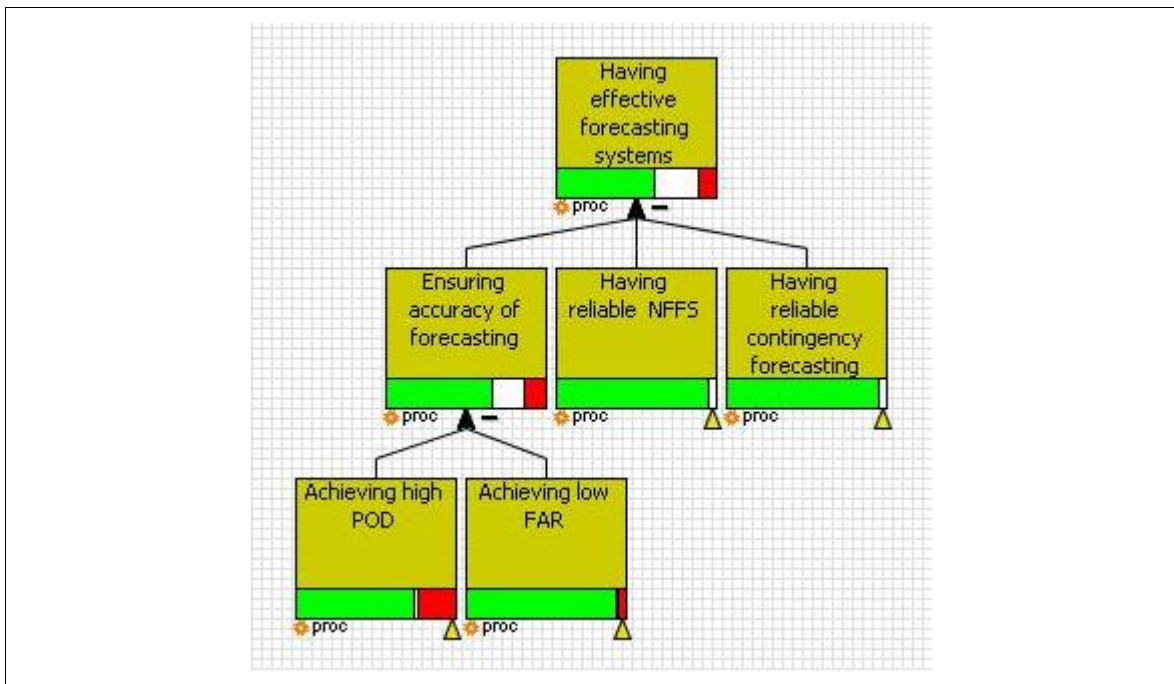


Figure 4.12 Service effectiveness model: effective forecasting system.

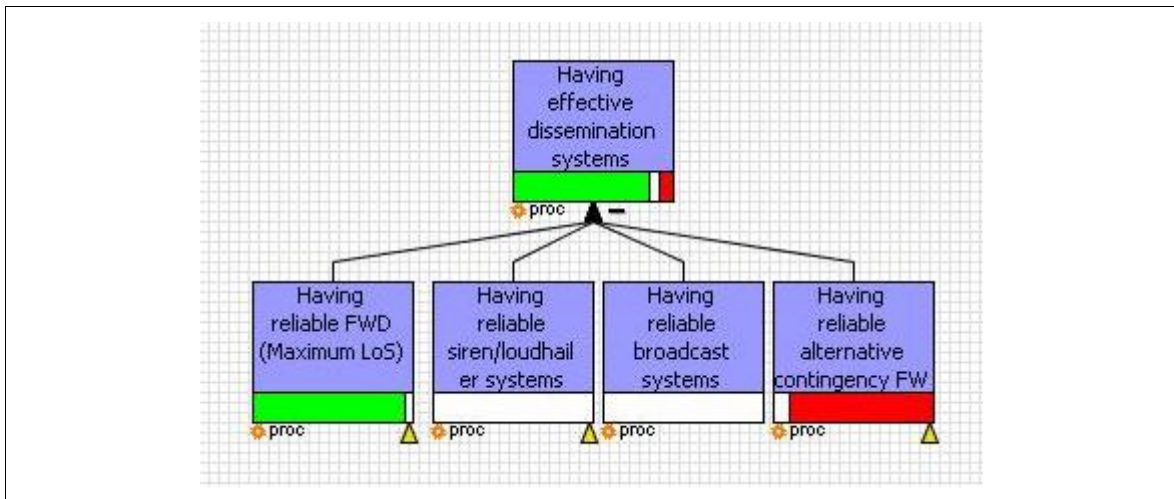


Figure 4.13 Service effectiveness model: effective dissemination system.

These figures also show the lowest level child processes for which direct, case study evidence is required.

The part of the model that examines flood warning decision-making has four sub-processes, specifically:

- having effective staff;
- having live effective procedures in place;
- validating flood warning thresholds;
- having good communication and information management.

The first of these sub-processes, 'having effective staff', has several sub-processes of its own, as detailed in Figure 4.14.

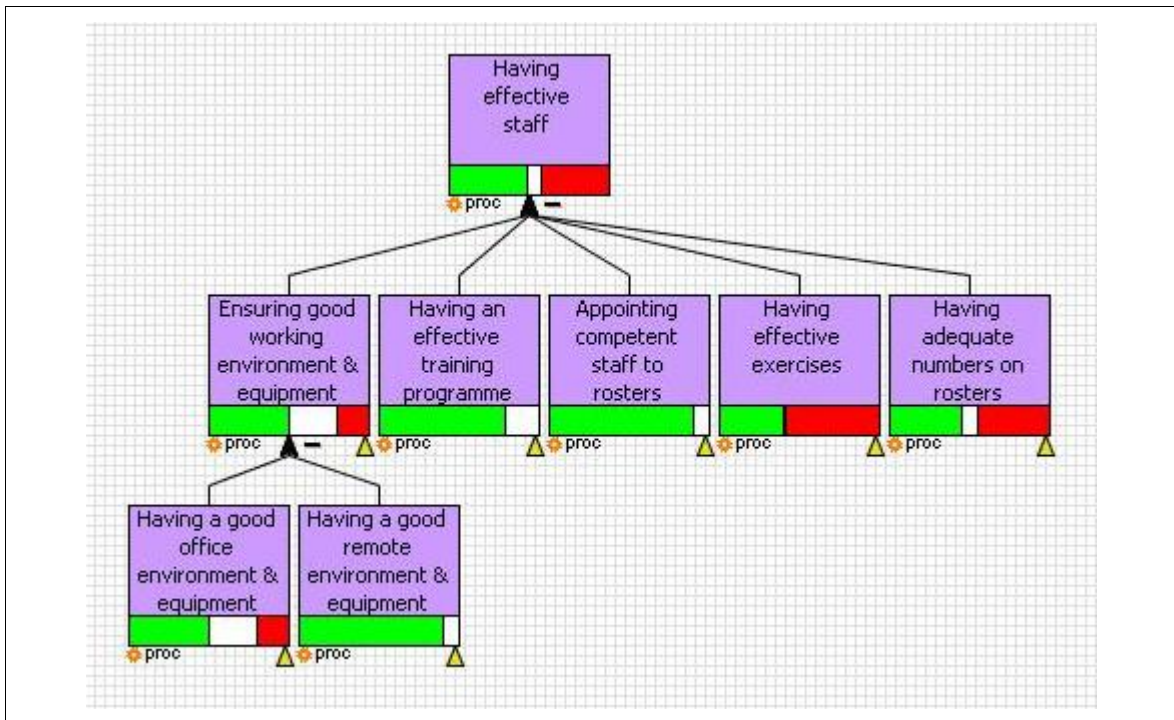


Figure 4.14 Service effectiveness model: effective staff.

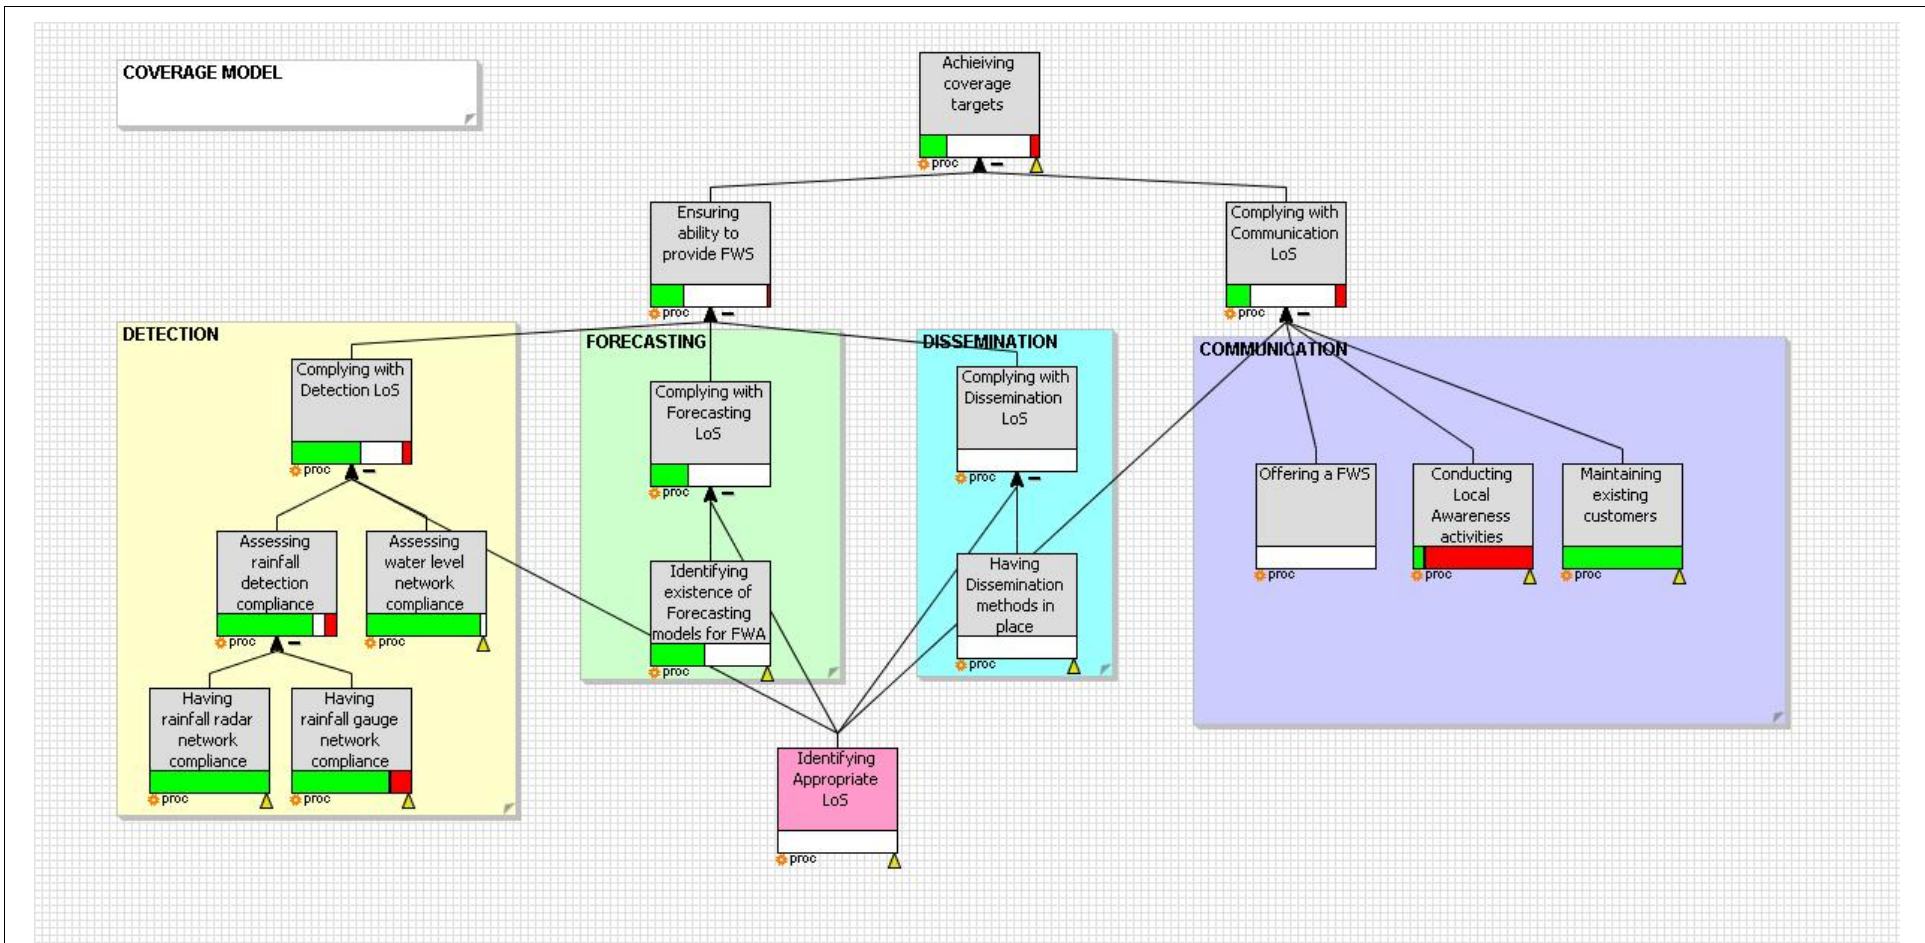
The decision-making sub-process 'having good communication and information management' also has its own sub-processes. The child processes associated with 'having good communication and information management' relate to 'communication between staff of all levels' and 'the systems and processes used to manage a flood incident'.

The above processes were all integrated to form an overall model of service effectiveness, an example of which is shown in Figure 4.16.

4.3.2 Assessing the evidence requirements

The definition of appropriate evidence for each process was an iterative process of discussion and model refinement (in itself a fundamental part of developing a *Perimeta* model). Direct evidence was entered into the coverage and service effectiveness models via the child processes that sit at the bottom of the model structure. The evidence is propagated up through the model by the *Juniper* algorithm which combines the direct evidence from one or more child processes.

An important part of model construction is an assessment of the degree of uncertainty inherent in the direct evidence input into the models. This had to be estimated by the model builders. For each child process, evidence in favour of the process (green) and evidence against (red) is entered, along with an assessment of the degree of uncertainty. These values are represented in the *Perimeta* model as an 'Italian flag'. Tables A1.1 and A1.2 in Appendix 1 detail the evidence entered for each of the child processes in the *Perimeta* models.



Note: Since this flood event occurred before new operating instructions were implemented, there is complete uncertainty (i.e. white) in three of the lowest sub-processes).

Figure 4.15 Coverage model based on data for the 2007 flood event.

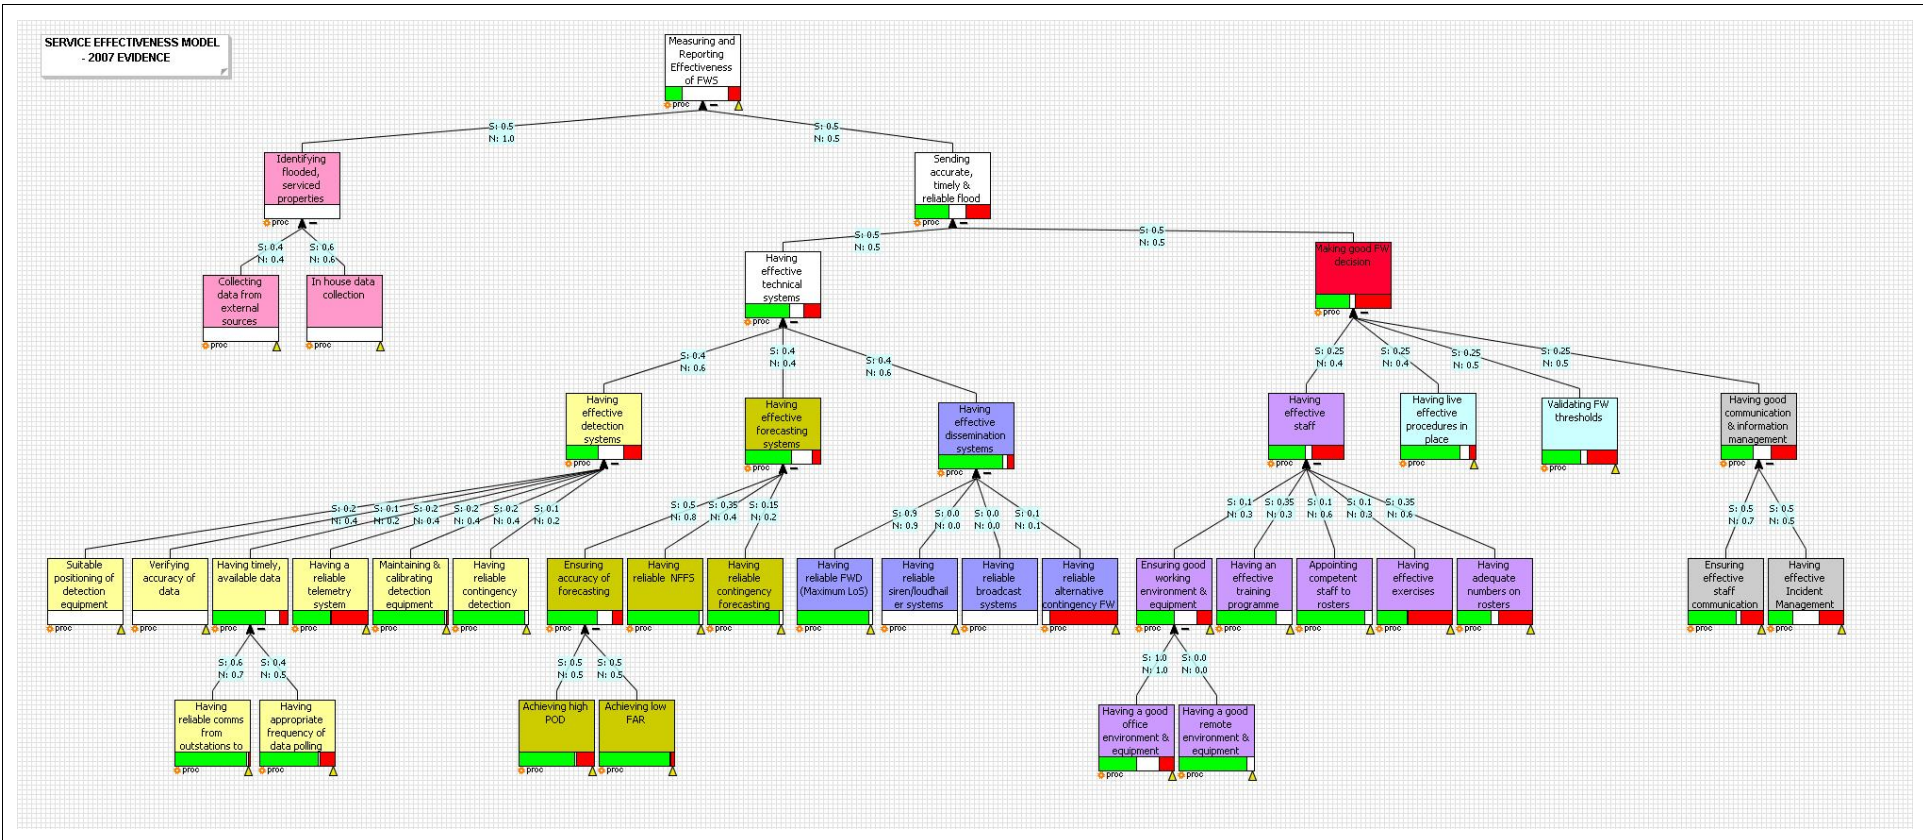


Figure 4.16 Service effectiveness model based on data for the 2007 flood event.

4.3.3 Setting the ‘necessity’ and ‘sufficiency’ values

Once the *Perimeta* model structures were established, it was then necessary to set the ‘necessity’ and ‘sufficiency’ values attached to the links between parent and child processes. These values within the model effectively determine the influence (or weight) that the evidence associated with a child process has on its parent process.

Setting the ‘necessity’ and ‘sufficiency’ values is an inherently subjective task. Appendix 2 contains a guidance document provided by Dr John Davis of the University of Bristol to assist with setting of these two values.

Necessity and sufficiency are defined as follows:

- i. **Sufficiency** is a measure of the strength of influence the success of a given child process has on the success of its parent. Sufficiency is related to the extent of the positive contribution that success of the child process makes towards ensuring success of the parent process.
- ii. **Necessity** is a measure of the extent to which failure of a child process will cause failure of its parent. If the parent process is likely to fail if the child process fails, then there is a case for increasing necessity.

Necessity and sufficiency both range from 0 to 1. As the value of each increases, the greater the influence of the evidence within the child process will be on the parent process. Looking at the extremes, a *sufficiency* of 1.0 means that the child process will on its own define the success of the parent. Conversely, a sufficiency of 0 (zero) implies that the success of the child will have no influence on the success of the parent. In the case of *necessity*, a value of 1.0 means that if the child process fails then so will the parent process, irrespective of the success or failure of any other child processes. A value of 0 (zero) implies that the failure of the child process will have no influence on the failure of the parent.

In practical terms, if the sum of the *sufficiencies* feeding into a parent process sum to 1 then this implies that the processes influencing the success of the parent are fully defined. If the sum of the *sufficiencies* is less than 1 this could indicate that the processes influencing the success of the parent are not fully defined (i.e. one or more processes is missing).

As the *Perimeta* model was developed and refined, conflicts were encountered within the model. Conflicts can be caused by a range of factors. One such case is where the evidence is assumed to be independent and the *sufficiencies* are set high – if the different sub-processes are presenting a different picture, then conflict between these forms of evidence will arise. This is analogous to having two experts, both of whom are trustworthy and believable, one of whom says ‘yes’ and the other says ‘no’ to the same question.

In practical terms, conflicts in the model are removed by reducing the influence of the individual sources of evidence (i.e. the *sufficiency* and/or *necessity* values) feeding in to the parent process where a conflict is present.

The process of setting the *necessity* and *sufficiency* values was (like much of the model building process) an iterative process. Once the model structure was more-or-less finalised, we defined an initial set of necessity and sufficiency values for each linked process within the model, based on our understanding of the FIM system. These values were then reviewed and discussed during the case study meetings with the Environment Agency. Subsequently, amendments were made to some of the values.

The *Perimeta* software has a facility to add comments to each process within the model structure; this was used to record the thinking and logic behind the choice of necessity and sufficiency values. For the majority of processes, the same *necessity* and *sufficiency* values were used for the two flood events (July 2007 and January 2008). There were, however, a small number of processes for which we could justify using different values. For example, within the service effectiveness model, the child processes of the process 'ensuring good working environment and equipment' were given different *necessity* and *sufficiency* values for the two flood events. We reasoned that the split between office-based duty work and home-based duty work was different during the two events.

Tables A1.3 and A1.4 within Appendix 1 contain the *necessity* and *sufficiency* values for the 2007 coverage and service effectiveness models respectively. The justification for each of the values is also given in the tables.

The process of setting the *necessity* and *sufficiency* values revealed *Perimeta's* high degree of sensitivity to these values. This highlighted the importance of reviewing the values as part of the model-building process and ensuring that the Environment Agency staff, who represented potential end-users, were involved throughout the process of model building and application.

4.3.4 Testing and validating the *Perimeta* models

Once the coverage and service effectiveness *Perimeta* models had been built and refined (see Section 4.3.1), each model was tested and validated through the case studies.

Evidence requirements and sources of evidence

To test and validate the *Perimeta* models through the case studies, we required specific evidence for the two flood events (summer 2007 and January 2008) in the West Area of the Thames Region. This evidence was obtained from several sources, specifically:

- Environment Agency performance measures;
- judgement based on Environment Agency staff experiences of the events;
- post-event reports (e.g. 2007 Technical Flood Report, Internal Debrief Report and Action Plan from the Multi Agency Forward Look Event etc.).

Where possible, measured evidence was used in favour of evidence based on judgement in order to minimise the subjectivity and uncertainty involved. Two sets of evidence were required for each of the processes: evidence *for* and evidence *against*.

There was no evidence available for a small number of processes; in these cases, there was no option but to leave the success or failure of the process as completely uncertain. An example of this is within the service effectiveness model, where the evidence associated with the two child processes of the parent process 'identifying flooded, serviced properties' is unknown. This lack of evidence introduced significant uncertainty into the model.

The amount of uncertainty present (i.e. the white portion of the 'Italian flag') reflects our degree of confidence about how well we have assessed the consequences of success or failure of a process (parent or child). In some cases the Environment Agency staff indicated that they were highly confident about the effect of success or failure of

individual processes, with no uncertainty about their assessment. In many cases however, there was some uncertainty inherent in evidence on performance; in the absence of an estimated value, we set an assumed level of uncertainty (typically 5 for 10 per cent) was assumed .

The Environment Agency staff indicated, on several occasions during the case study, that high levels of uncertainty or a lack of evidence associated with specific processes indicated the parts of the system that needed more investigation and possible investment in order to improve performance.

Model validation

Once all the available evidence had been inserted into the models, it was then possible to undertake limited validation testing on both the service effectiveness and coverage models for the 2007 and 2008 flood events. The purpose of the validation was to verify each model's representation of the system at both the top level and intermediate levels and establish our confidence in the structure of the model, its parameters and their input values.

Validation is ideally undertaken by comparing how the evidence is propagated up the model's hierarchy with actual measured evidence (derived from performance indicators) for the top and intermediate processes within the model. However, for the majority of processes, there was no measured evidence available. Thus the models were only partly validated, based on informed judgement rather than hard evidence.

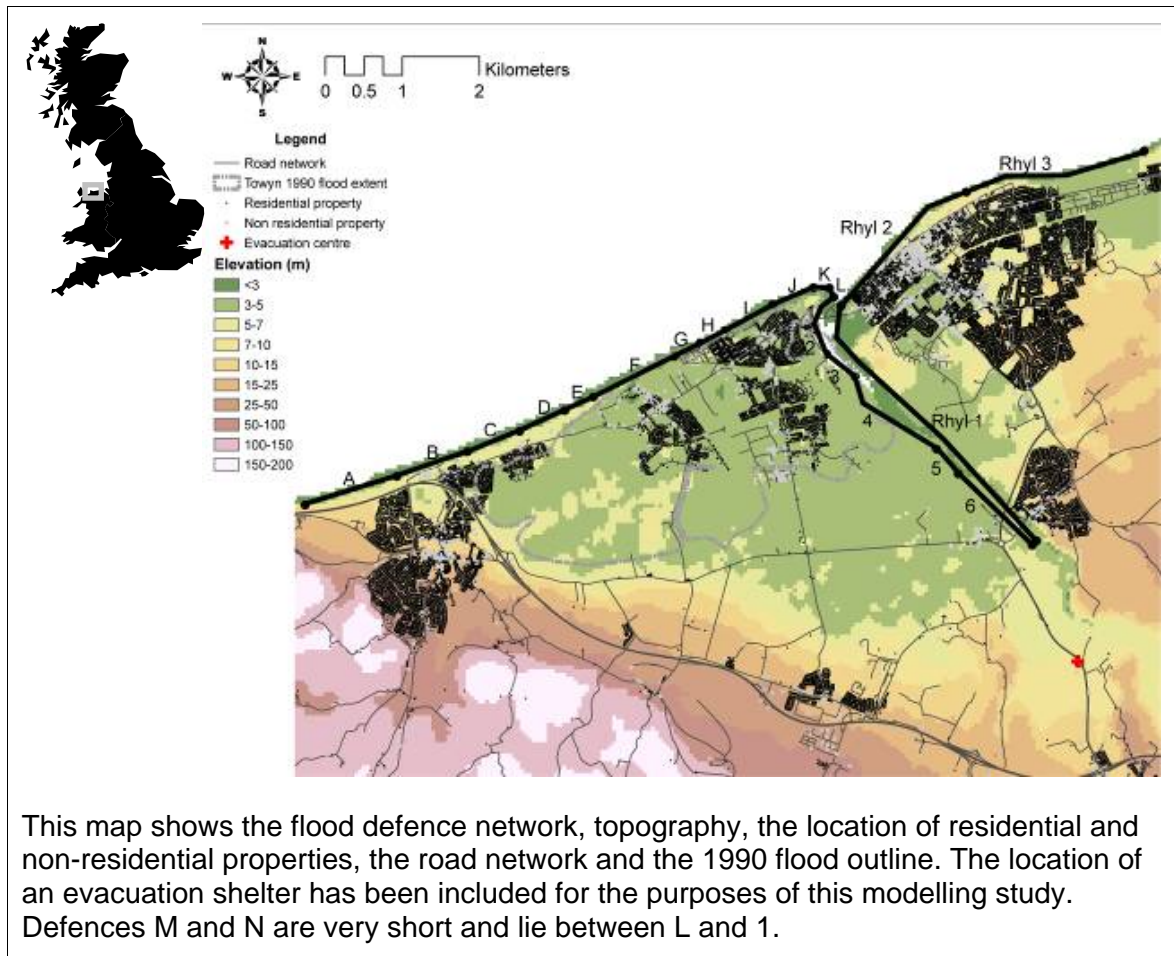
The processes at which validation was attempted within each of the models are summarised in Table 4.2.

Table 4.2 Processes validated within the coverage and service effectiveness models.

Coverage model	Service effectiveness model
Complying with Detection LoS	Having effective detection systems
Complying with Forecasting LoS	Having effective forecasting systems
Complying with Dissemination LoS	Having effective dissemination systems
Complying with Communication LoS	Having effective staff
Achieving coverage targets	Making good flood warning decision
	Sending accurate, timely & reliable flood warnings
	Identifying flooded, serviced properties
	Measuring and reporting effectiveness of FWS

The participants involved in the model building benefited considerably from this exercise: they were able to develop a shared picture of what sub-processes contribute to the parent process and gain a better understanding of how they interact. Agreement on model structure and parameterisation through active dialogue contributes to the validation process – the aim of which is to develop a model that users agree provides a reasonable representation of the system being modelled.

4.4 Testing the agent-based model



This map shows the flood defence network, topography, the location of residential and non-residential properties, the road network and the 1990 flood outline. The location of an evacuation shelter has been included for the purposes of this modelling study. Defences M and N are very short and lie between L and 1.

Figure 4.17 Map of Towyn, Wales.

This study has built from scratch an agent-based model of FIM that couples hydrodynamic simulations with human behaviour. The model has been coded in NetLogo² which is a freeware agent-based model development environment. NetLogo provides a flexible, free and convenient platform for such model development. Other freeware environments (e.g. RePast, Swarm) are likely to have been as convenient, whilst AnyLogic is a good commercial environment.

As it currently stands, the model built during the course of the case study work is not suitable for operational deployment within real-time situations, but it has been designed with this ultimate goal in mind; it is expandable and generic (i.e. not tied to a particular case study location). This version of the model can currently be accessed online at:

<http://www.staff.ncl.ac.uk/richard.dawson/FloodEventABMDemo/FloodEventABMDemo.html>

This version is a java applet compiled from the NetLogo code. It can be run directly from a webpage, although there are frequently java compatibility issues so the applet is not always able to run. A video of a simulation is also provided (at the same site) for download.

² A quick summary of the license is that use is unrestricted, including commercial use, but there are some restrictions on redistribution and/or modification (<http://ccl.northwestern.edu/netlogo/faq.html>)

It was necessary to construct an agent-based model from scratch because no open source flood event management agent-based model exists. The Life Safety Model is owned by BC Hydro and was developed with dam-break floods in mind, although this model has been tested against flooding in Canvey Island. However, it was necessary in this study for the model developers to have access to the code for model development and testing purposes. Furthermore, the model developed in this project has no proprietary IPR issues and can therefore be used flexibly and extended in further research.

Dataset processing and importing

The model uses nationally available information which can, where possible, be imported directly into the model using a series of scripts. These scripts were developed especially to ensure that the model could be transferred to other locations with minimal difficulty. These scripts load and import into the NetLogo model:

- topographic data from LiDAR and IfSAR (imported as a raster grid);
- the OS MasterMap ITN transport network (imported as a topological network that describes the road and paths, their connectivity, road type, and where applicable, direction);
- the National Property Database/AddressPoint point data (used to identify the location of buildings and building type);
- data on the location of flood defences from the National Flood and Coastal Defence Database (NFCDD).

In addition to the national datasets described above, the IfSAR data has been supplemented with local manhole surveys and re-sampled to a resolution of 50m for the inundation modelling. The flood outline from the 1990 flood was used to calibrate the flood model in the agent-based model. Although the tidal surge for the 1990 event is fully recorded, our current model implementation only describes a surge level, rather than a time series covering the full tidal cycle.

The coastal defence sections A–N and fluvial defence sections 1–6 have been extracted from the NFCDD. For this modelling analysis, only breaches of the coastal defences were considered. The Rhyl defence data was not readily available, so three defence sections have been assumed for the purposes of case study. These defences have been assigned the same properties as defence section F. We decided to expand the simulation beyond Towyn to demonstrate the interaction between two towns that were linked by limited transport connections (the bridges at defence section L and at the southern end of defence section 6 are shown in Figure 4.17). Floods that cut off these links can have enormous implications for any FIM response. The 2007 flooding highlighted the importance of being able to analyse the inter-regional impacts of flood; in the 2007 floods key transport corridors were blocked, thwarting attempts to install temporary barriers and stretching local flood fighting capacity.

Hydrodynamics

The inundation process is represented using a simple raster cell model (in the spirit of Lisflood and JFlow) where flow is predominantly governed by topography, water surface and friction (Figure 4.18). Although this approach is ultimately not as accurate as more sophisticated and computationally expensive codes such as TUFLOW, it is sufficient to describe the key dynamics of the inundation to enable FIM responses to be tested in this study. Later model developments may incorporate more sophisticated

inundation models. We couple the hydrodynamics into the agent-based model, rather than drive the model using the hydrodynamics as an external boundary condition. This allows us to test within the wider dynamics of the system any FIM responses that are designed to divert flow paths or that require actions that take place within the floodplain. Ultimately, this is a more flexible design than importing exogenously generated flow paths.




Figure 4.18 Representation of flow between raster cells (based on Bates and DeRoo, 2000).

Network model

The road and path network is imported from the OS MasterMap dataset. Frequently agent-based models use a raster grid which the agent is able to move through. However, the transport system is crucial for the movement of agents (vehicles and pedestrians) through a town. We capture the connectivity and directionality of the road and path system by importing the network directly into the model; agents are able to move along this network.

The use of a network model also has computational advantages when modelling over large domains because the raster cells can be larger, reducing the computational expense of the hydrodynamic model. The network is reported on the raster surface in the model viewer so that every raster cell which contains a road component is shaded.

Agent behaviour

This type of agent-based model must address two key issues, namely:

- the starting conditions for the agents;
- the reaction of the agents during a flood event.

The starting conditions reflect what a typical moment may look like prior to the flood event. These conditions have temporal variation, for example more people will be in their homes during the night, more at work during a weekday etc. However, it is impossible to truly anticipate the location of each and every person and vehicle at any given time, yet their location is likely to alter the model output (in terms of lives lost etc.) of any given simulation.

To address these challenges, each agent is described by:

- the possible states which it can take;
- the actions it can take;
- the transitions between states.

It is unrealistic to represent the precise behaviour of each individual, so the population is divided into groups based on representative behaviour patterns. These could include for example 'single professionals', 'married man with family – one employed family member', 'married woman – second employed family member', 'working single mum' etc. A 'typical' day for these types of population agent are constructed from the National Travel Survey (ONS and DfT, 2006), WICID (Stillwell and Duke-Williams, 2003) and other data sources where they are available. To represent the natural variation within these groups and in the National Travel Survey (NTS) data, movements are described in probabilistic terms. Figure 4.19 shows the format used to describe population type and movement; Figure 4.20 translates this text algorithm into a flow diagram.

A simple example that describes how the agent 'working single mum' might move around during a day has been expanded. Here the agent is assigned to 'home 1' until around 8am. With a normal distribution centred at 8am, with a standard deviation of 15 minutes, the agent has a probability of 1 that they will travel from home to school. Approximately five minutes after this, with a standard deviation of one minute, the agent has a probability of 0.9 that they will travel to work, and a probability of 0.1 that they will definitely travel to some shop before going to work two hours (with a standard deviation of one hour) later. The term 'some shop' randomly selects a shop in the modelling domain, whilst the term 'nearest shop' would send the agent to the nearest building categorised as a shop. Times can be specified in hours (h), minutes (m) or seconds (s). When given the order to evacuate, the agent (in this example) responds, with probability of 0.7 in this example, to the evacuation warning within five minutes (with a standard deviation of one minute) and heads to a random evacuation point. However, 30 percent of the time the agent will not evacuate and will resume its previous behaviour. The proportion of agents that receive and successfully act on this warning can be estimated from post-flood event reviews and other data where available. As there are often insufficient observations of flood event behaviour, a sensitivity test can explore how the lack of information on the proportion of people who evacuate affects the spread of results from the model.

There is no perfect method for representing human activity. However, the use of travel surveys and census data to identify types of behaviour is regularly used by transport modellers (e.g. Horowitz, 2006), although there are of course other possible approaches to classifying populations and generating behaviour rules.

The approach described above provides enormous flexibility for setting up the agent-based model. It reflects the uncertainties in the travel survey information and provides the inevitable randomness of traffic conditions and agent location that would be observed at any one instance. These issues are discussed in more detail in Chapter 5.

The agent-based model also incorporates non-residential activity by adding a certain amount of through traffic into the model domain along the major roads. Natural extensions to this framework would be to explore how the number of additional drivers on the road during adverse weather conditions (how many more people commute to work on a wet day?) that are likely to be associated with a flood event. Although this work has not been implemented here, it could be achieved simply by altering the proportion of agents travelling by road.


```

{"single young professional"}
{"married man – main earner"}
{"working single mum" "home 1" [
  ["home 1" "daily 8h 15m" "school" 1]
  ["school" "5m 1m" "work" 0.9]
  ["school" "5m 1m" "some shop" 0.1]
  ["some shop" "2h 1h" "work" 1]
  ["work" "daily 17h 15m" "home 1" 0.25]
  ["work" "daily 17h 15m" "school" 0.5]
  ["work" "daily 17h 15m" "supermarket" 0.25]
  ["supermarket" "45m 10m" "home 1" 1]
  ["school" "5m 1m" "home 1" 0.5]
  ["school" "5m 1m" "supermarket" 0.5]
  ["supermarket" "45m 10m" "home 1" 1]
  ["home 1" "daily 20h 2h" "nearest recreation" 0.1]
  ["some recreation" "1h 30m" "home 1" 1]
  ["evacuate" "5m 1m" "some evacuation point" 0.7]
  ["evacuate" "0s 0s" "resume" 0.1]]}
{"retired couple"}
etc.

```

Figure 4.19 Example of the text file used to describe population behaviour.

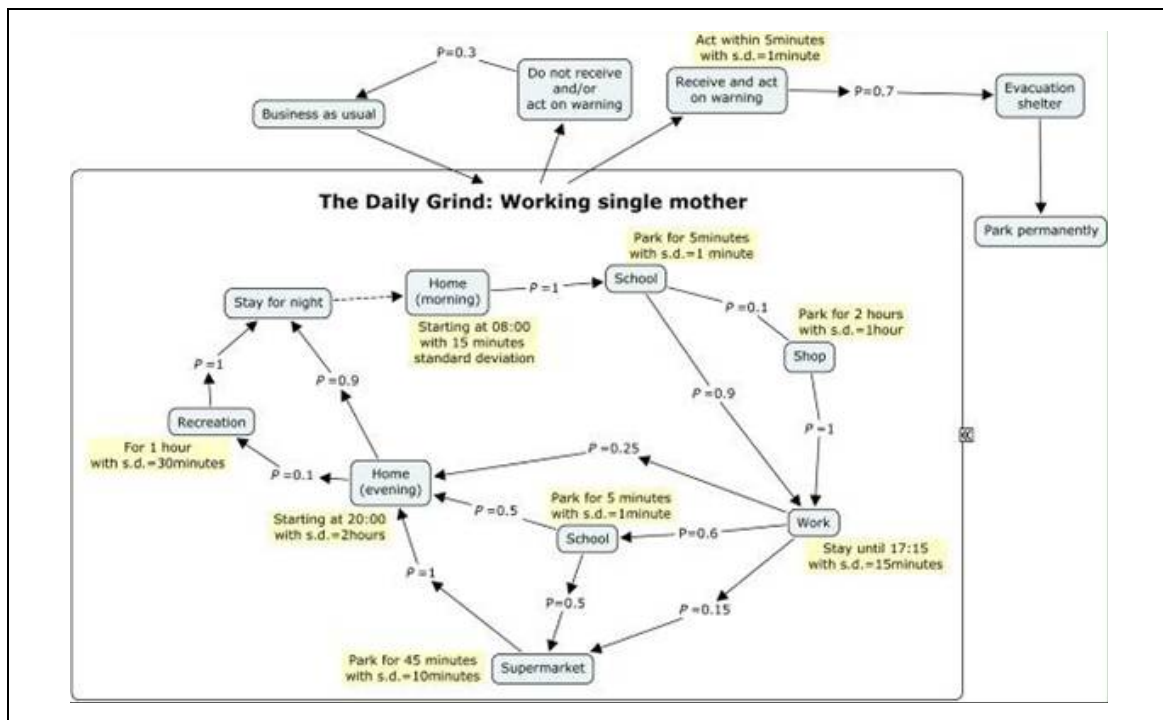
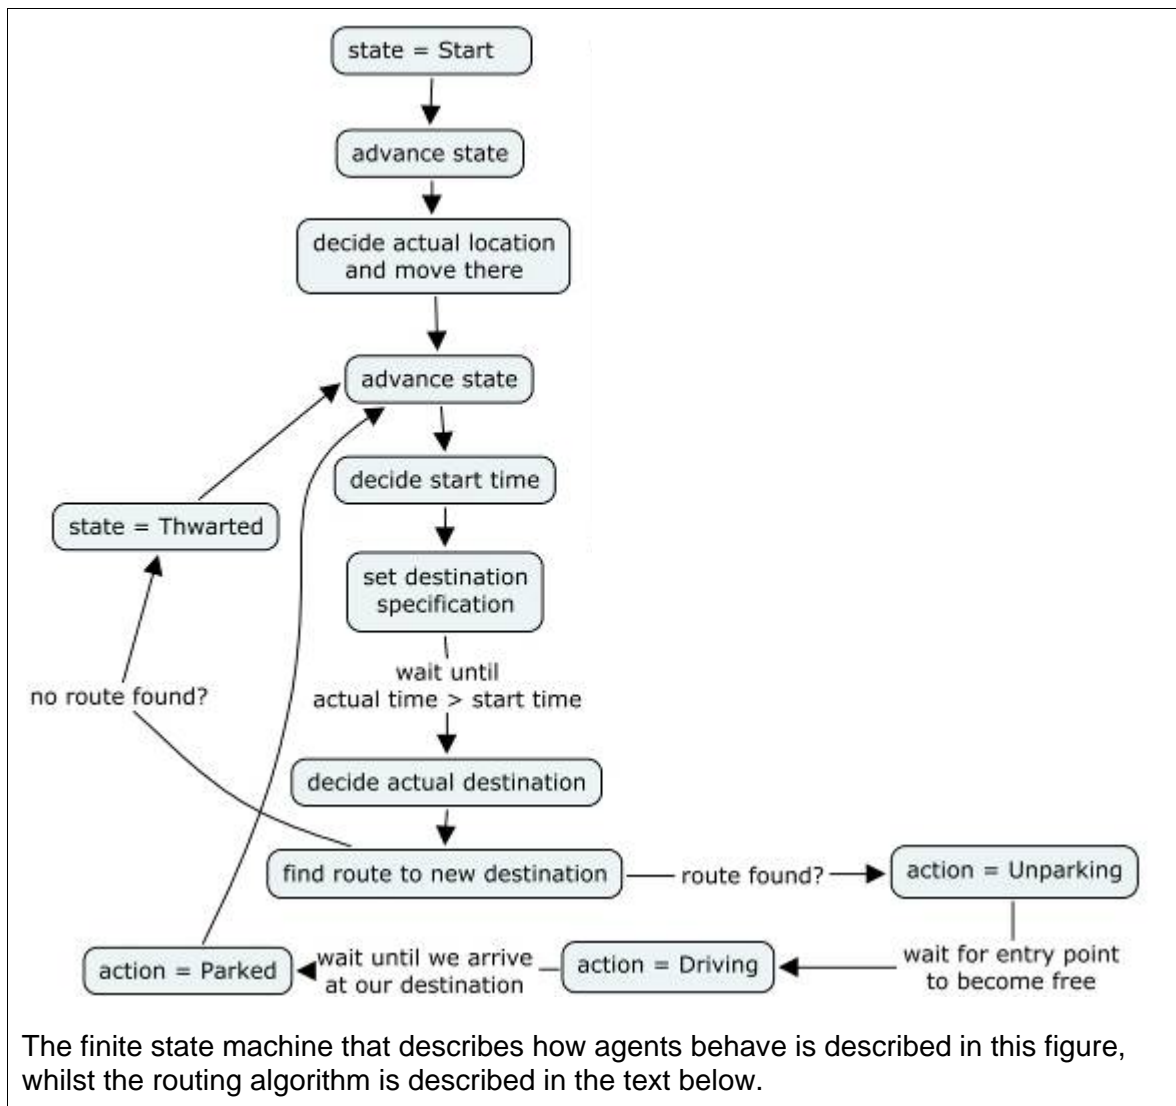


Figure 4.20 A flow diagram explaining the resulting agent behaviour described in the text file in Figure 4.19.

Traffic rules

Vehicles move through the network model according to a set of traffic rules. These have been selected because they are easy to implement globally within the model; rules can be generated automatically from the mapping datasets, but are sufficiently sophisticated to capture the main dynamics of traffic flow through a road network. Detailed micro-simulation rules could be included (e.g. traffic light regulation), but they

would add significantly to the model construction time and computational overhead of each simulation. The algorithm used to advance each agent from one state to another (e.g. from 'home' to 'school') is shown in Figure 4.21.



The finite state machine that describes how agents behave is described in this figure, whilst the routing algorithm is described in the text below.

Figure 4.21 Algorithm for routing agents.

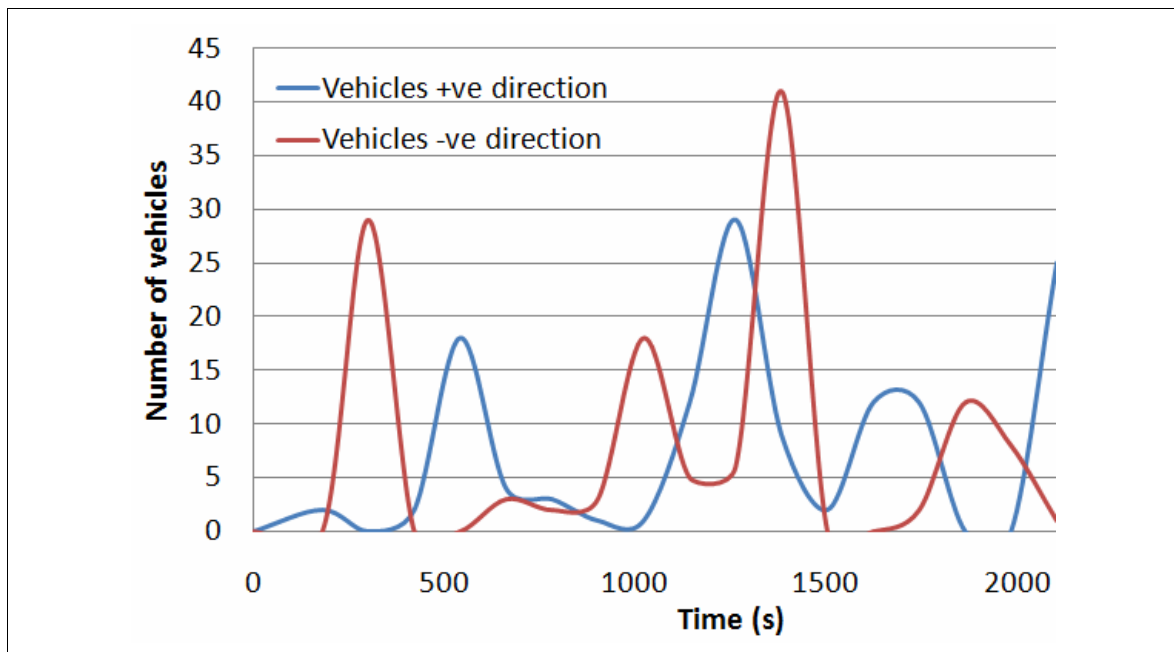


Figure 4.22 Vehicle numbers in both directions changing through time on a road link.

The main traffic rules are:

- i. **Connectivity**
Vehicles can move along the network where connections exist.
- ii. **Direction**
Where the OS MasterMap indicates a one way street, agents can move along this road in one direction only.
- iii. **Queuing**
Vehicles are not able to pass another vehicle, and slow down when close to another vehicle.
- iv. **Junctions**
Vehicles slow at a junction and cannot move on to another road until there is sufficient space in the traffic.
- v. **Road type**
The maximum speed of vehicles is limited by road type – 60mph for A roads and 30mph for other roads. Road type is extracted automatically from the OS MasterMap dataset.
- vi. **Floodwater**
Where floodwater deeper than 20cm is blocking a road (i.e. where a raster grid is wet and this intersects with the road network), if the vehicle is already on the link when it floods, it slows down and turns around and recalculates a route to its destination. If the vehicle is about to turn on to a road that it sees is already wet then it will also recalculate its route. Each vehicle remembers which of the roads it knows to be flooded. In the current model, vehicles do not communicate information to each other – although this rule could be added later.
- vii. **Drowning**
If a vehicle gets caught in floodwater deeper than 20cm then it is classified as "drowned" in the case study presented here (although more sophisticated vulnerability functions could be incorporated into the model at a later date).

Mortality function

The model incorporates a binary mortality function where agents in water depth less than 20cm are considered safe, and over 20cm are considered drowned. This function has predominantly been implemented to demonstrate that mortality functions more sophisticated than 'wet/dry' could be used at a later date.

This approach recognises that 'wet/dry' is unrealistic because vehicles can still drive through very shallow floodwater and very shallow flood depths rarely lead to instant drowning. However, this mortality function does not consider the velocity of floodwaters or the duration of exposure to floodwater; these factors can have a significant impact on mortality and longer term morbidity issues. Likewise, individuals in houses or buildings can frequently move to the second floor – evacuations in recent flood events have taken place in significantly deeper water (albeit usually with support of the Fire and Rescue Service who are not currently represented in this model). The recommendations section of this report (Chapter 7) describes a more complete approach to assessing individual vulnerability.

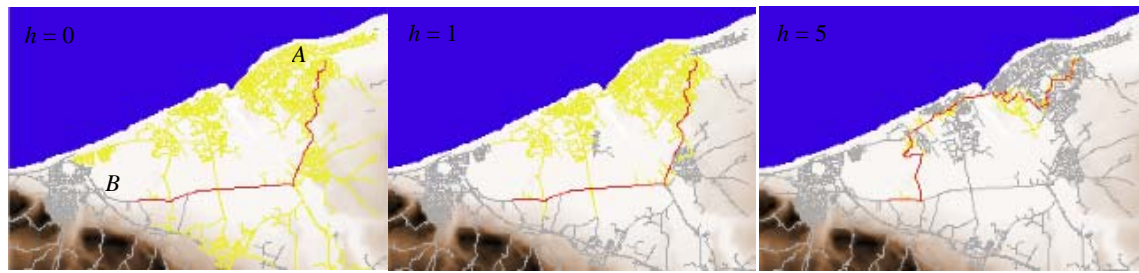
Computationally efficient routing of journeys

Describing journeys along a network requires a routing algorithm that defines how a vehicle travels from point *A* to point *B*. Two main algorithms are available: a random walk algorithm and a routing algorithm.

We selected to use a routing algorithm because most journeys are more purposeful than a random walk; the 'A*' algorithm has been implemented to identify the route. There are a number of network search algorithms, but the A* algorithm is computationally efficient and effective at finding an approximate solution of the most efficient route between *A* and *B*, even accounting for speed limits on different roads in its routing.

The A* algorithm, and others, are described in more detail in Dechter and Pearl (1985). The accuracy of the algorithm can be adjusted and traded off against the computational cost of the routing calculation. For the geographical area covered in our model, even an extremely rapid algorithm finds a route within five per cent of the shortest distance, but saves enormous computational expense.

We acknowledge that many drivers do not take a near optimal route from *A* to *B* (perhaps because they do not know it, or prefer a scenic route), but some assumption about driving behaviour is required. To represent this 'less efficient' behaviour in some drivers, or to capture local preferences, the optimisation algorithm could be made probabilistic, thus representing a range of behaviours. However, this possibility has not been implemented in the current model. Figure 4.23 shows how an agent, when it can no longer complete its route (due to flooding, or an evacuation warning), tries to re-route itself.



The yellow (lightest lines in greyscale) shaded roads are the routes searched by the algorithm where a heuristic factor of 0 denotes searching every possible route combination between *A* and *B* and therefore calculates the shortest route. The higher the heuristic factor, the smaller the search space and the faster the calculation (albeit at the expense of a small drop in accuracy) of the shortest route. The algorithm accounts for road type.

Figure 4.23 Route finding algorithm operating for three different heuristic factors (0, 1 and 5 respectively) shows the calculated route between *A* and *B* in red (darkest line in greyscale).

NetLogo interface

In the interface (Figure 4.24) the user can specify the:

- directory where the model data is stored;
- sea level;
- total number of agents in the modelling domain;
- the time of day in which the simulation starts running.

The user can also view the:

- time of day;
- number of residential and non-residential buildings flooded;
- number of agents moving through the road network;
- agents that have been diverted by floodwater, or isolated so that their route has been completely blocked by floodwater;
- counter and time series plots showing the number of agents that have been 'drowned' by floodwater.

During a simulation the user can interact with the model and choose to:

- change the sea level;
- randomly destroy one of the defence sections;
- issue an evacuation warning.

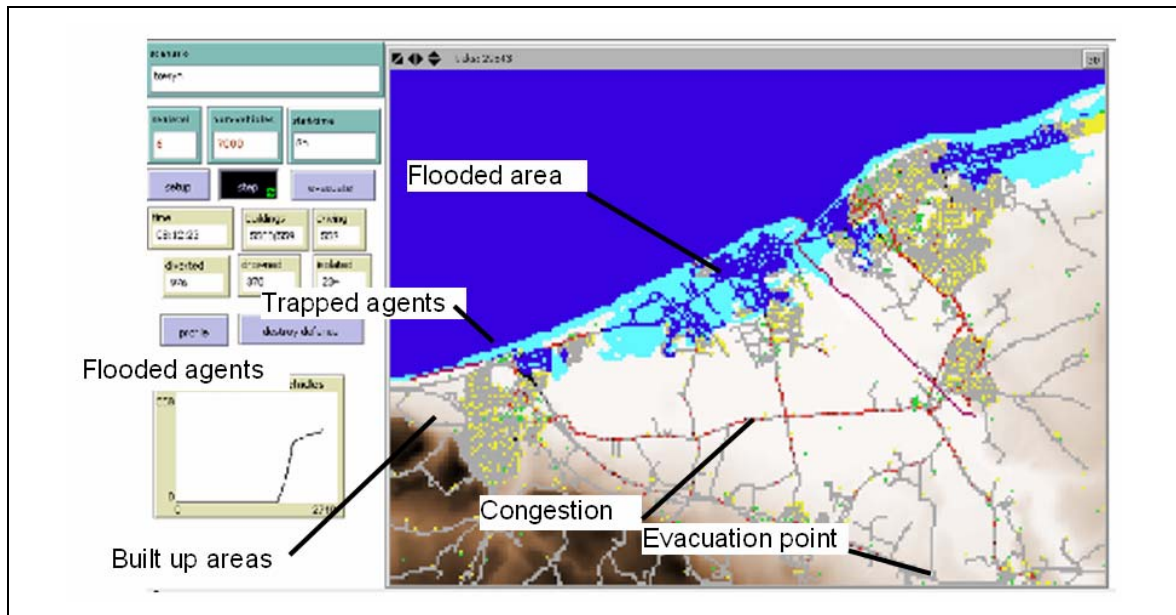


Figure 4.24 Screenshot of agent-based model showing how congestion can build up in certain routes and agents can be flooded.

5 Outputs from the case studies

5.1 Introduction

This chapter considers the results of examining and testing, through case studies, the ability of performance matrices, root cause analysis, *Perimeta*-based hierarchical modelling and agent-based modelling to help assess the reliability and evaluate the performance of FIM.

5.2 Performance matrices and root cause analysis

The potential for applying performance matrices to performance assessment, and applying root cause analysis via a fish-bone diagram to explore reasons for system failure, were explored with the case study team. Insights gained from applying these and the other tools tested (see below) are summarised in a series of text boxes.

Performance matrices and the fish-bone diagram – key insights

- Both tools provide a means of rapid post-event assessment.
- Both tools are easy to use, and require only qualitative data.
- Performance matrices can assess changes in performance over time in the same area, or compare performance in different areas during the same event.
- Fish-bone diagrams and a root cause analysis can help structure a more detailed *Perimeta* model.

5.2.1 Performance matrices

The case study application of the performance matrices provided evidence of performance during two successive flood events at a strategic or overview level.

A comparison of the model outputs for the July 2007 and the January 2008 flood events indicates better FIM performance during the more recent event. This was attributed to:

- improvements in FIM processes following the July 2007 event;
- the smaller scale and severity of the January 2008 flood event compared to the July 2007 event.

The assessment of FIM performance using performance matrices is a relatively quick and easy method for comparing performance between flood events; such matrices can also be used to demonstrate different levels of performance. Assessments can be undertaken either spatially in different Environment Agency regions, areas or catchments, or across time if applied to different flood events in the same location. The matrices clearly show performance levels but do not identify the reasons for any improvements or regression.

5.2.2 Fish-bone diagram

Root cause analysis was itself not tested during the case study work, although the potential for its application has been discussed earlier in this report. As described in Chapter 4, the 'fish-bone diagram' was found to be useful because it assisted with the structuring of *Perimeta* models. The diagrams helped to ensure that all the main factors that influence coverage and service effectiveness were represented in the models.

5.3 *Perimeta* models

The *Perimeta* models were constructed by the case study team to assess the current level of performance of the 'coverage' and 'service effectiveness' aspects of FIM. The team was not able, within the resources of this study, to extend the case studies to apply *Perimeta* models to help assess interventions to improve FIM performance (i.e. the models were not tested for their potential to help develop a business case for interventions to improve FIM performance).

***Perimeta* – key insights**

- The process of model building helped the case study team to better understand the processes involved in FIM systems and their interactions with one another.
- Application of *Perimeta* revealed to the case study team the influence of significant areas of uncertainty in some of the FIM processes.
- The models tested in the case studies, using actual data, were able to indicate to the case study team where process improvements could be considered.
- Once validated, the models are able to test the effects of improvements on FIM performance and can be used to help determine the benefits of different options.

5.3.1 The results from the *Perimeta* coverage model

The Flood Warning Performance Measures work instruction document defines the 'Coverage' term within the Flood Warning Service Limit (FWSL), as:

$$(Number\ of\ serviced\ properties) / Number\ of\ properties\ at\ risk$$

The 'Italian flag' (figure of merit) for the top process in the *Perimeta* model of 'coverage' represents the strength of evidence that supports coverage is in line with the flood warning LoS target. In other words, it represents the degree of confidence that we have in being able to meet the performance target we are aiming to satisfy in terms of flood warning 'coverage'.

This probabilistic form of evidence can be used to estimate a value of *expected cover* in the flood damage avoided (FDA) equation, in place of a deterministic value for 'coverage'. The model indicates the level of assurance that the coverage target will be met, not the actual percentage of coverage itself.

The results of the latest *Perimeta* model for the July 2007 event show the following evidence relating to coverage:

- 25 per cent (green) of the FWSL area is certainly covered in line with the FW LoS;
- up to 26 per cent (white) of the FWSL area may be covered in line with the FW LoS;
- 49 per cent (red) of the FWSL area is certainly not yet covered.

This evidence is based on input data provided by the Thames Region for the July 2007 event.

Using data for the 2008 event, the *Perimeta* output relating to coverage was:

- 40 per cent (green) of the FWSL area is certainly covered in line with the FW LoS;
- up to 13 per cent (white) of the FWSL area may be covered in line with the FW LoS;
- 47 per cent (red) of the FWSL area is certainly not yet covered.

The reduction in uncertainty and the increase in evidence in favour of good coverage between the 2007 and 2008 events is largely due to a significant decrease in the uncertainty associated with the sub-processes in 2007 (see Figure 4.15) relating to: 'identifying appropriate LoS' and 'having dissemination methods in place'.

The case study team decided that the propagated evidence is plausible, though the evidence for the 'complying with detection LoS' process seemed a bit low. However, the actual measured evidence lies within the area of uncertainty.

The model showed that measures addressing reliability or uncertainty could be considered to improve the performance of the following sub-processes:

- 'Offering a flood warning service' – improve reliability (thus reducing the red portion of the 'Italian flag');
- 'Identifying appropriate LoS' - reduce uncertainty;
- 'Identifying existence of forecasting models for FWA' - reduce uncertainty;
- 'Having dissemination methods in place' - reduce uncertainty;
- 'Having rain-gauge network compliance' - improve reliability;
- 'Assessing water level network compliance' – improve reliability;
- 'Conducting local awareness activities' - improve reliability.

The process of building and applying the above model revealed the following insights:

- The model is reasonably robust and could be used in a forward-looking planning**
The model can be run under various 'what-if' scenarios to identify areas where investment should be targeted in order to maintain or improve FIM performance. This may be where the real value of the models lies.
- Changes to individual sub-processes have only a small impact at the top of the model**
Amending the evidence for individual sub-processes in isolation was found to have a small impact on the propagated evidence at the very top of the

model (i.e. the influence of evidence from individual sub-processes becomes more diluted the higher up the system one looks). This suggests that there are no quick-wins for improving FIM, which the Environment Agency confirmed reflects reality.

iii. **Additional performance measures at an intermediate level within the *Perimeta* model would be useful**

These intermediate processes include 'complying with detection LoS' and 'complying with forecasting LoS'. However, assessing the performance of these processes in practical terms would not be straight-forward.

iv. **Identification of uncertain model processes**

The model made it possible to identify processes that have significant uncertainty associated with the evidence (e.g. 'Identifying appropriate LoS' and 'Offering a FWS'). This is useful because it highlights parts of the overall process that need more investigation and/or investment.

To prioritise investment within the different areas of the FIM system, the results from the *Perimeta* models would need to be used as input to a value function that could indicate the benefits arising from these investments. These benefits can then be compared with the costs of investments. *Perimeta*, used in a forward-looking planning mode, could thus be used to check that the plans are making best use of resources and investment.

A significant amount of the uncertainty associated with the input data (evidence) arises from the fact that the Environment Agency does not routinely assess the performance of the lowest sub-processes within the *Perimeta* model. Reducing the uncertainty at the lowest levels would significantly reduce the uncertainty in the evidence that is propagated through the model.

5.3.2 The results from the *Perimeta* service effectiveness model

The Flood Warning Performance Measures work instruction document defines 'service effectiveness' as the proportion of flooded, serviced properties that were sent a flood warning:

$$\frac{\text{(Number of flooded, serviced properties sent a warning)}}{\text{(Number of flooded, serviced properties)}}$$

The 'Italian flag' (figure of merit) for the top process in the *Perimeta* model of 'service effectiveness' is the strength of evidence in support of the fact the Environment Agency can measure and report the effectiveness of the flood warning service. The sub-process 'sending accurate, timely and reliable flood warnings' is the real measure of service effectiveness, but this is not at the top of the model because the above definition of service effectiveness requires knowledge of the number of flooded, serviced properties. This is brought into the model through the process 'identifying flooded serviced properties' and its sub-processes.

The results of the latest *Perimeta* model for the July 2007 event show the following evidence relating to service effectiveness:

- 23 per cent (green) evidence in support of the Environment Agency was able to measure and report the effectiveness of the flood warning service;
- 61 per cent (white) evidence to suggest that it is unknown whether the Environment Agency was able to measure and report the effectiveness of the flood warning service;

- 16 per cent (red) evidence suggesting that the Environment Agency was not able to measure and report the effectiveness of the flood warning service.

This evidence is based on input data provided by the Thames Region for the July 2007 event.

The results of the latest *Perimeta* model for the January 2008 event show the following evidence relating to 'service effectiveness':

- 32 per cent (green) evidence in support of the Environment Agency being able to measure and report the effectiveness of the flood warning service;
- 55 per cent (white) evidence to suggest that it is unknown whether the Environment Agency was able to measure and report the effectiveness of the flood warning service;
- 13 per cent (red) evidence suggesting that the Environment Agency was not able to measure and report the effectiveness of the flood warning service.

This evidence is based on input data provided by the Thames Region for the January 2008 event. By comparing the 'Italian flag' for the data for 2008 against the flag for 2007 indicates that the evidence in support of good performance increased and the amount of uncertainty reduced between 2007 and 2008. This suggests that improvements were made over the intervening period.

An unfortunate feature of the service effectiveness model is the lack of evidence available for some of the sub-processes; these are set as '100% uncertain' within the model. For the July 2007 and January 2008 events there were six and five processes, respectively, where no evidence was available. This lack of evidence increases the overall uncertainty of the top process of the model, particularly because some of the sub-processes where direct evidence is unknown are relatively high up the model (i.e. underneath 'identifying flooded, serviced properties') meaning that this uncertainty has more influence than if it were further down the model structure.

The modelling team found the propagated evidence to be plausible, particularly given that the model shows an improvement between the 2007 and 2008 events which is in line with Environment Agency expectations.

The model showed that measures addressing reliability could be considered to improve the performance of the following sub-processes:

- 'Having a reliable telemetry system' – improve reliability;
- 'Having a good office environment & equipment' – improve reliability;
- 'Having effective exercises' – improve reliability;
- 'Having adequate numbers on rosters' – improve reliability;
- 'Ensuring effective staff communication at all levels' – improve reliability;
- 'Having effective Incident Management System' – improving reliability.

The evidence for the following sub-processes is completely uncertain; testing measures to improve the measurement of the performance of these processes would allow the relative effect each has on overall process performance to be determined:

- 'Suitable positioning of detection equipment' (2007 and 2008 event)
- 'Verifying accuracy of data' (2007 and 2008 event)

- 'Collecting data from external sources' (2007 event only)
- 'In-house data collection' (2007 and 2008 event)

The evidence associated with the two sub-processes shown below is shown as being completely uncertain within the models because neither of these dissemination methods is used in the Thames Region.

- 'Having reliable siren/loudhailer systems' (Intermediate LoS)
- 'Having reliable broadcast systems' (Minimum LoS)

However, these two processes were included in the model because they are defined within the Flood Warning LoS work instruction document and will be relevant to other regions of the Environment Agency.

The process of building the above model also revealed the following insights:

- There is currently very little measured evidence regarding service effectiveness**
The evidence associated with the top process of the model reveals that there is currently very little measured evidence (and therefore a lot of uncertainty) that can confirm that service effectiveness is being achieved. The Environment Agency confirmed that this is a reasonable reflection of the situation in the West Area in the Thames Region.
- There is significant uncertainty in the identification of flooded properties**
This suggests that further investment is required in this area, though the evidence provided for the two events does show that data collection improved between the two events.
- Additional information on the performance of specific sub-processes would be useful**
This additional evidence would help to better identify the most influential parts of the FIM system.

As with the coverage model, the service effectiveness model can be used to help the Environment Agency identify areas where further investment is required to report on and/or measure performance. The caveats applying to the model structure, parameterisation and input data need to be fully understood before the models can be used appropriately.

5.4 Comparing results with Environment Agency performance targets

The Environment Agency's Flood Risk Management Strategy sets performance targets related to the FWS over the next five years. These performance targets are:

- 80 per cent Levels of Service by 2011**
This target can be measured at the planning stage and relates to the percentage of at-risk properties within the FWSL³ that are 'serviced' in accordance with the Flood Warning LoS. The appropriate dissemination and communication of flood risk must be established. It does not state that appropriate detection and forecasting need to be established.

³ Defined by the Extreme Flood Outline = 1000 year undefended

- ii. **75 per cent Appropriate action maintaining this measure to 2011**
This target relates to the percentage of flood warning recipients that take action to reduce damage to property following receipt of a flood warning message. This is measured during the post-event stage, usually by mail-shot or market research.

Targets have been assigned by the Environment Agency to each of the six performance factors in the Flood Damage Avoided (FDA) equation in the Flood Warning Investment Strategy covering the period 2003 to 2014. Table 5.1 summarises these performance targets (taken from Table 1 within the Flood Warning Performance Measures work instruction document).

Table 5.1 Flood Warning Investment Strategy 2003/4 to 2012/3. Benefit targets for English Regions and Environment Agency Wales for Declared Targets Option 3.

	English Regions				
	Year 03/04	Year 06/07	Year 07/08	Year 09/10	Year 12/13
Damage Reduction	30%	35%	37%	40%	40%
Coverage	70%	78%	78%	80%	80%
Service Effectiveness	65%	75%	77%	80%	80%
Availability	63%	75%	77%	80%	80%
Ability	80%	85%	85%	85%	85%
Effective Action	50%	75%	78%	85%	85%

	Wales				
	Year 03/04	Year 06/07	Year 07/08	Year 09/10	Year 12/13
Damage Reduction	30%	35%	37%	40%	40%
Coverage	50%	68%	72%	80%	80%
Service Effectiveness	49%	61%	65%	80%	80%
Availability	61%	72%	70%	75%	75%
Ability	75%	80%	75%	85%	85%
Effective Action	61%	75%	78%	85%	85%

Notes: Table reproduced from Table 1 within the Flood Warning Performance Measures work instruction document.

5.5 The agent-based model

The agent-based model is able to produce multiple simulations to explore flood event and emergency response characteristics and associated uncertainties; this gives a broad understanding of their influence on the performance of FIM. The ability of an agent-based model to provide this information was explored in the Towyn case study. Towyn Council has shown interest in taking this work further to assess the usefulness of the agent-based model in helping assess measures to improve public safety and community resilience in the face of potential flood hazards.

Agent-based modelling – key insights

- Able to represent interactive and dynamic processes within emergency response systems.
- Can indicate the relative influence of flood event characteristics and FIM responses on a range of public safety outcomes.
- Indicates where resilience could be improved and can be used to test the effects of such improvements.
- Can be used to explore the sensitivity of public safety indicators to FIM actions and emergency response measures.
- Towyn Town Council is considering applying agent-based modelling to help improve evacuation contingency planning and increase community resilience.

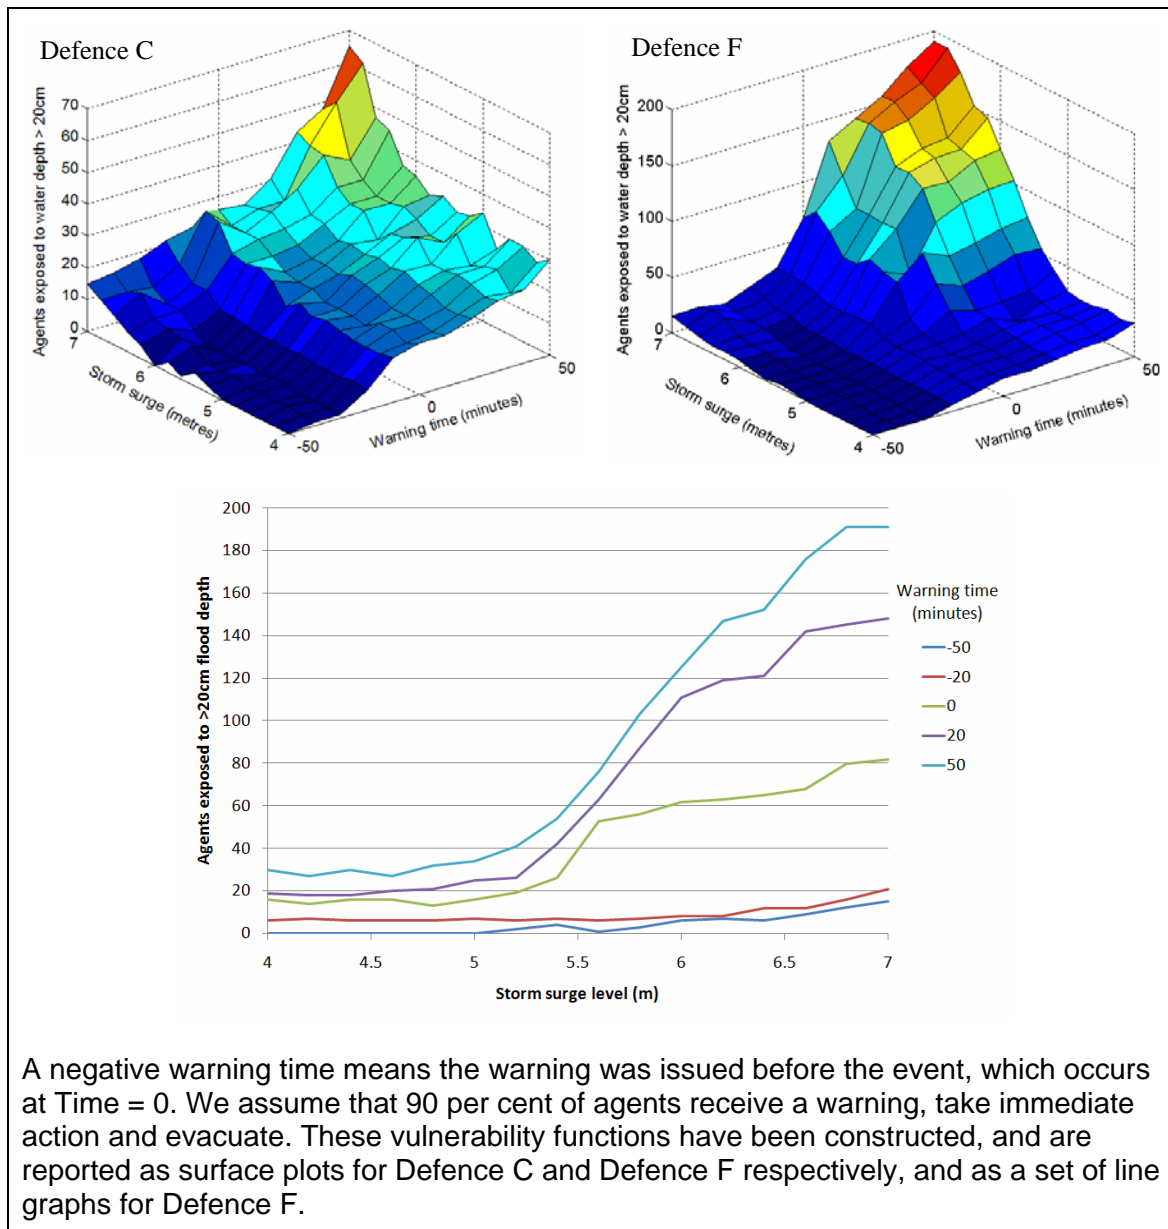
5.5.1 Results

The agent-based model can be used to produce multiple simulations that explore uncertainties and give a broad understanding of the performance of the FIM system. A number of results are presented within this section.

Figure 5.1 plots of the number of agents exposed to water depths above 20cm as a function of storm surge level and warning time (these plots can be considered like the output of a vulnerability function). Figure 5.2 shows the same simulations, plotting the same exposure for all defences, but only against storm surge level.

The model outputs are reported here in terms of the number of agents exposed to water depths greater than 20cm. This description is used to reflect the fact that a ‘wet’ agent is not necessarily ‘drowned’ for the reasons described in Section 4.4. Whilst the functions appear to be generally monotonic (i.e. the number of exposed agents increases with storm surge height and shorter warning time) they are quite ‘jagged’, largely because:

- Probabilistic description of agent behaviour**
As described in earlier sections, the time at which agents move between states is described probabilistically to capture uncertainties in agent behaviour. This can lead to a large degree of variation in the number of cars on the road at any time between different simulations.
- Initial model conditions.**
The behaviour of each agent is assigned randomly at the start of each simulation. This reflects the fact that although we have an approximate range of the types of behaviour from surveys, these do not describe each individual in the model domain precisely. Therefore, in this model, the agents’ home, shop, business and school choices are allocated at random.
- Model sequencing.**
The order in which the behaviour commands of individual agents are processed by the computer. For example, if two agents are stopped at a crossroads, each on a road perpendicular to the other, the order in which one agent crosses the crossroads before the other may have a knock-on effect in the way those two agents interact with other agents later in a simulation (e.g. through resultant congestion, or subsequently ending up on a road as it floods where no delay at the crossroads might have allowed an agent to miss the flood). This issue is common to all agent-based simulations.



A negative warning time means the warning was issued before the event, which occurs at Time = 0. We assume that 90 per cent of agents receive a warning, take immediate action and evacuate. These vulnerability functions have been constructed, and are reported as surface plots for Defence C and Defence F respectively, and as a set of line graphs for Defence F.

Figure 5.1 The relationship between agents exposed to water depths greater than 20cm, the storm surge level and the warning time.




Figure 5.2 Vulnerability relationships for each defence (agents exposed to water depths greater than 20cm for different storm surge levels) when no flood warning is given.

Figure 5.3 shows how a particular simulation can be run many times to quantify the model's variability. This figure also quantifies the benefits of providing a flood warning. In this case, the expected number of agents exposed to a water depth greater than 20cm is 21.14 'with a flood warning' and 219.14 'without warning' with a standard deviation of 4.53 and 14.24 for the two scenarios, respectively. The flood warning provides a tenfold reduction in exposure of the population.

The increased variance for simulations with no flood warning occurs because more vehicles are located in the flood risk area at the time the flood starts, which means there is greater potential that the number of agents exposed to floodwater is more sensitive to the initial conditions and model sequencing.

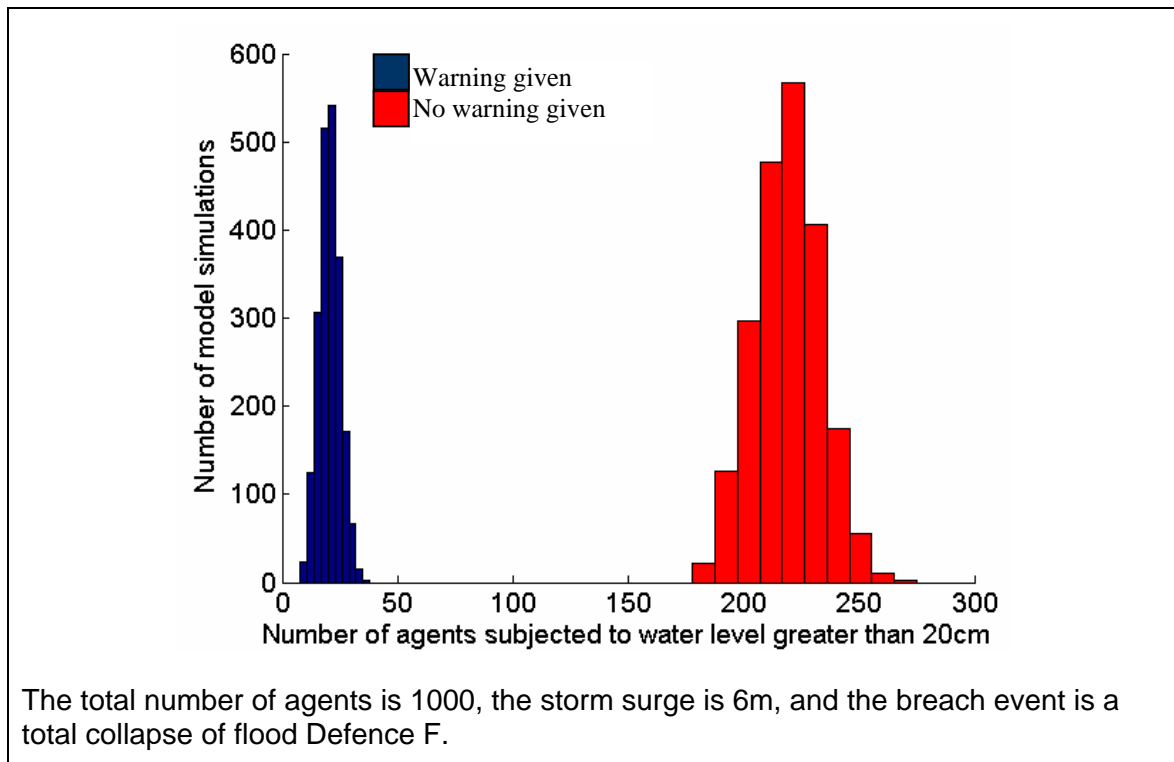


Figure 5.3 Histogram of number of drowned agents for a sample of 2173 simulations where a flood warning is given approximately one hour before the defence breaches and where no warning is given at all.

The results presented in Figure 5.1 to Figure 5.3 provide only aggregated measures of vulnerability. Figure 5.4 and Figure 5.5 extend these results to show how the model can provide spatial information on vulnerability and risk to humans.

The maps present the outputs of the same simulations shown in Figure 5.3 for a single defence failure. However, they show the locations where agents, whether driving or parked at some location, are likely to be exposed to water depths greater than 20cm. It is worth noting that many of the agents that are parked have been caught in commercial and industrial districts – a result that might be expected at this time in the morning (~8am) during the school run and start of the working day.

Figure 5.5 presents results from the same simulations again, but shows where in the model domain roads are more prone to congestion during a flood event evacuation. There is a general increase in road traffic activity because vehicles that would have been 'parked' at work or elsewhere evacuate by road instead. Perhaps unsurprisingly, congestion is most likely to be observed on the roads leading to the evacuation point and in locations of high residential and commercial density. However, if the emergency services were based in Rhyl and they were needed in Towyn, then it is clear that congestion on the two main routes into Towyn may reduce the effectiveness of their response. Testing the location of evacuation centres, emergency response bases and storage facilities (e.g. for temporary defences) can help identify more robust locations.

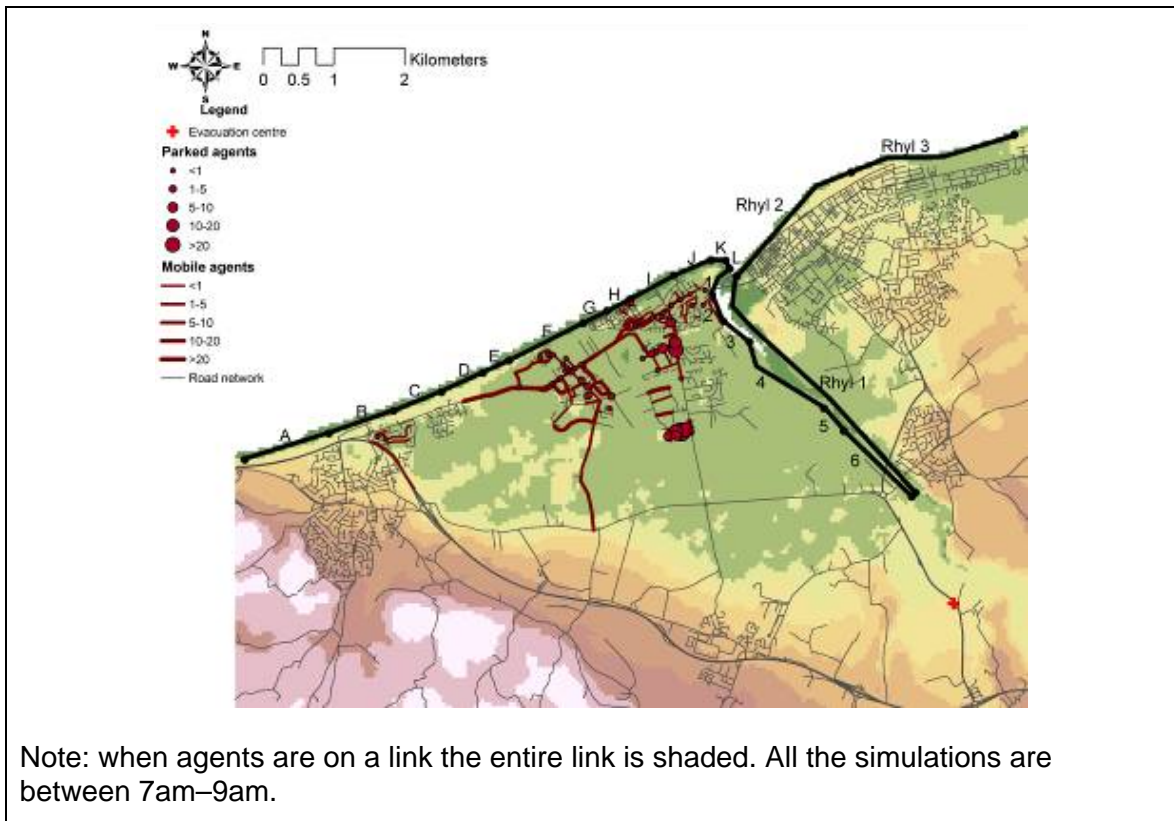


Figure 5.4 The location, and expected number of agents exposed to water depths over 20cm from 2173 simulations of a 6m storm surge and the failure of flood defence F.

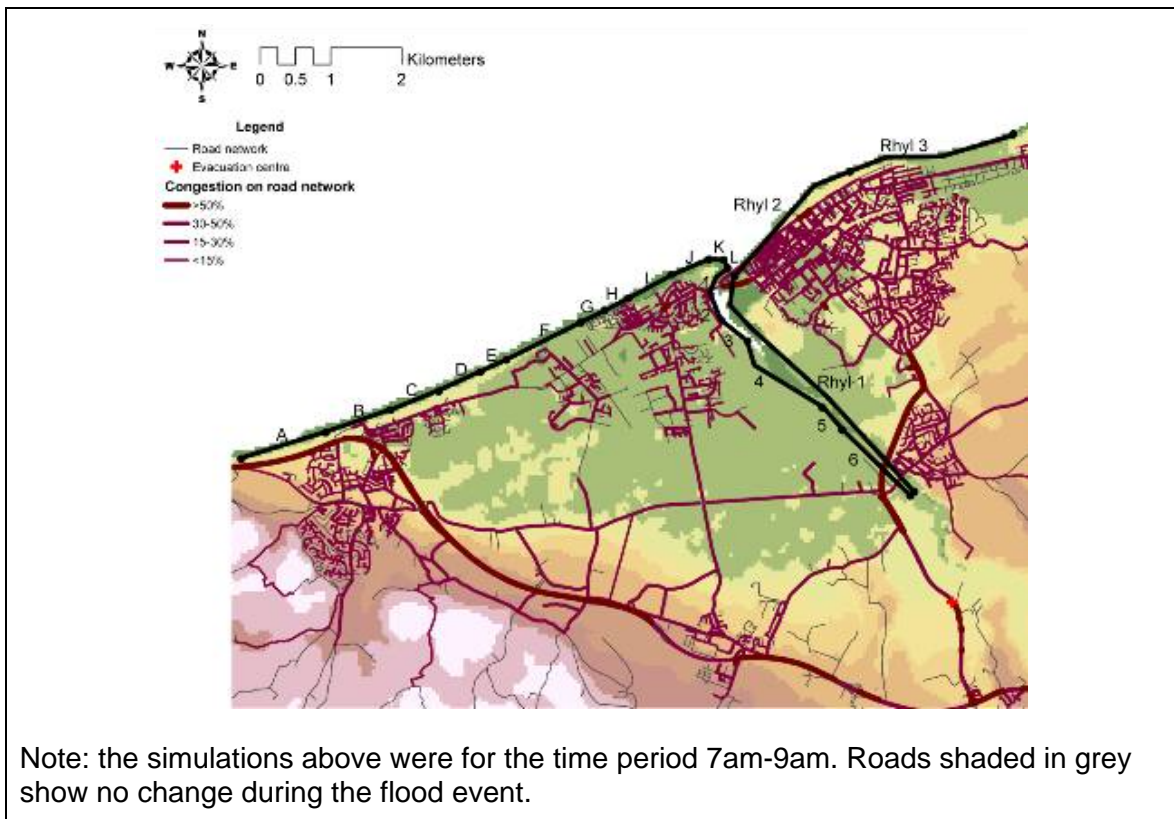


Figure 5.5 Roads prone to become congested, relative to normal traffic, during a flood event where congestion is measured in terms of the drop in average speed along a road compared to normal driving conditions.

5.5.2 Risk analysis

Flood risk, r , is calculated as a function of the probability of an event, ρ , and the consequences (usually economic damages, but any damage function can be used) of that event, d , for a set of input conditions, defined by the vector $\mathbf{x}:=\{x_1, x_2, \dots, x_n\}$, where each variable x_i , represents a particular property (e.g. fragility of flood defences) of the flood system:

$$r = \int_{\mathcal{X}} \rho(\mathbf{x})d(\mathbf{x})d\mathbf{x} \quad (1)$$

This calculation has been successfully applied to systems of flood defences in the RASP, NaFRA and Foresight: Future Flooding projects (Hall *et al.*, 2003; Dawson and Hall, 2006) where the variables in \mathbf{x} represent those associated with the reliability of flood defences, the location and type of property in the floodplain, the hydrodynamic boundary conditions and their frequency.

The inclusion of FIM policies essentially acts as a modifier to either the consequences of flooding (e.g. where FIM action leads to reduced damages because a flood warning has been successfully announced, received, understood and acted upon) or the probability of flooding (e.g. where FIM action reduces the probability of a flood defence failing because a flood warning has been successfully announced, received, understood and acted upon by shoring defences or putting up temporary barriers). In Equation 1 (above), the existence of a FIM system leads to the need for more variables in \mathbf{x} to describe the system. For example, these parameters will include factors such as the number of people subscribed to the FWS and the proportion of those who receive the message.

The benefits of FIM policies in terms of risk reduction, B_{FIM} , can be calculated as:

$$B_{FIM} = \int_{\mathcal{X}} \rho(\mathbf{x}')d(\mathbf{x}')d\mathbf{x} - \int_{\mathcal{X}} \rho(\mathbf{x}_{FIM})d(\mathbf{x}_{FIM})d\mathbf{x} \quad (2)$$

where \mathbf{x}' is the state of the system with no FIM policy, and \mathbf{x}_{FIM} is the state for a single or portfolio of FIM policies. These values are obtained from the simulation model; Figure 5.3 shows how different policies can lead to different outputs (i.e. \mathbf{x}' and \mathbf{x}_{FIM} according to whether a flood warning is issued, measured in terms of human impacts). These outputs can form the basis of FIM appraisal.

An example of the benefit of a FIM policy on flood risk is summarised in Table 5.2, where the benefits of an effective flood warning system can be estimated in terms of expected number of agents exposed to a 20cm depth or more of floodwater. These results suggest that the risk to people is reduced almost four-fold, or an expected reduction of 0.076 agents exposed per year. The total risk to people, expressed in terms of the expected number of agents exposed to 20cm depth of flooding, R_p , is calculate using Equation 3:

$$R_p = \sum_{i=1}^N \int f(H_s, W)P(B_i | H_s, W)A(W_{FP} > 0.2m)dH_s dW \quad (3)$$

where:

- N is the number of flood defences
- $f(H_s, W)$ is the joint probability density function describing the loading space of significant wave height, H_s , and water level, W
- $P(B_i | H_s, W)$ is the fragility function describing the probability of a breach in defence i occurring as a result of the joint loading conditions
- $A(W_{FP} > 0.2m)$ is the number of agents exposed to floodplain water depths, W_{FP} above the threshold of 20cm.

This calculation is repeated for the warning and the no warning simulations. The joint probability density function and fragility functions were the same as those used in the RASP Intermediate Level Method (Dawson and Hall, 2006) application in Towyn (some of the inputs are shown in Figure 5.6 and Figure 5.7).

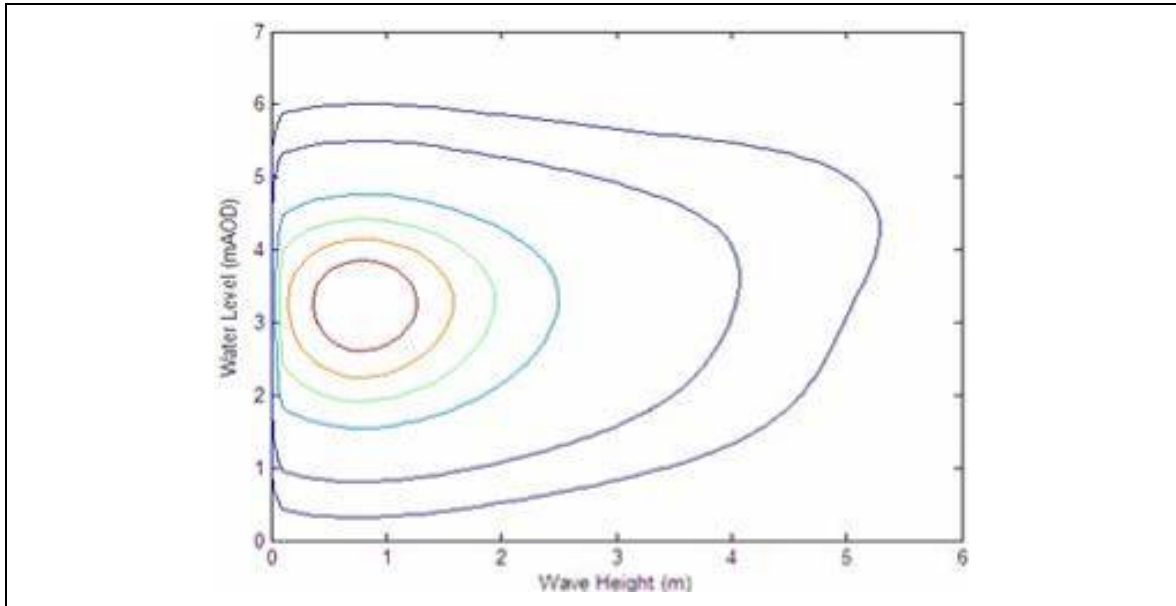


Figure 5.6 Contour plot of the joint probability density function for wave height and water level at Towyn.




Figure 5.7 Fragility functions for defences D and K respectively.

Table 5.2 Expected number of agents exposed to water depths greater than 20cm for each defence for simulations where a warning is issued one hour in advance and when no warning is given (as reported in Figure 5.3).

Defence	A	B	C	D	E
No warning	2.34E-03	9.89E-03	4.62E-04	1.77E-05	5.72E-03
Warning	2.68E-04	3.15E-03	1.58E-04	1.07E-06	1.83E-03

Defence	F	G	H	I
No warning	1.05E-03	1.83E-03	1.42E-02	4.23E-02
Warning	4.12E-04	4.59E-04	5.62E-03	1.03E-02

Defence	J	K	L	M	N
No warning	1.28E-02	1.91E-03	2.11E-03	1.28E-03	3.33E-03
Warning	3.55E-03	9.63E-05	2.31E-05	1.73E-04	3.37E-04

Defence	Rhyl1	Rhyl2	Rhyl3	Total
No warning	1.11E-03	1.65E-04	2.17E-03	0.103
Warning	4.10E-04	8.51E-05	1.17E-04	0.027

A risk-based calculation might also include benefits expressed in terms of damages avoided. For certain types of appraisal, this calculation can be readily achieved using existing equations (see Appendix 2 within Part A: Guidance). Whilst these risk-based calculation do not explicitly require the use of an agent-based model, they may benefit from such modelling that can identify congestion and the ability of agents to return to their houses to take suitable and timely action to avoid damages. This would help to refine the ‘response effectiveness’ parameter in that calculation.

Given successful implementation of the recommendations in Chapter 7, further analysis with the agent-based model will help to quantify how many more measures (e.g. deployment of temporary flood defences, shoring of defences during a flood, individual flood management actions such as domestic floodguards etc.) generate ‘damage savings’. Evidently, the benefits of flood event management decisions can be compared with the costs of implementing the policy, C_{FIM} , to provide a basis for prioritising investment that is consistent with current Project Appraisal Guidance.

5.5.3 Sensitivity analysis

Diagnostic (Newcastle University and Halcrow, 2006) and risk-based techniques (Hall *et al.*, 2003; Dawson and Hall, 2006) can be used to identify the risks and uncertainties within FIM components, processes and systems.

The results above provide evidence that probabilistic analysis provides a useful impression of the model results and uncertainties. It does not, however, provide a direct means of diagnosing the contribution that variables, individually or in combination, make to the overall uncertainty.

The conventional way to address this difficulty in analysing sensitivity is through ‘one at a time’ sensitivity analysis, where each individual variable is perturbed from its nominal value, while other variables are kept constant (several examples of this have been

demonstrated earlier in this report). This approach, applied above, has provided useful insights into the agent-based model and the FIM system.

However, this method provides only a cursory impression of sensitivity because it does not test the range of potential variability and it does not deal with variations in combinations of variables. Variance-based sensitivity analysis overcomes both of these problems by testing the sensitivity of the model's output over the range of variability of each variable individually and in combination. Variance-based sensitivity analysis provides a rational basis for identifying where investment in a particular data acquisition strategy, process or component (as long as it is represented in the model) is most likely to reduce flood risk.

5.5.4 Scenario analysis and training

In addition to risk-based appraisal of FIM measures, the agent-based model could be used to support training and scenario testing events. For example, an event with representatives of flood risk management organisations, blue light services and local authority planners could be used to test how planners respond interactively, using the model with pre-defined scenarios as part of a training exercise. Likewise, different types of flood event can be tested easily without the overheads of a desktop exercise. These workshops are also an opportunity for stakeholders to feed back into the development process of the model as part of an iterative ongoing development programme.

6 Discussion

6.1 Introduction

This chapter considers the insights gained during the course of our case study tests into the usefulness of performance matrices and root cause analysis, hierarchical modelling using the *Perimeta* software and agent-based modelling for assessing the reliability and evaluating the performance of FIM.

Although the techniques on which the tools are based are not new (some were developed in the 1990s), this report describes their first recorded application to FIM. This represents a significant contribution to the science base relating to FIM, in particular, and to non-structural flood risk management in general.

6.2 Evaluation of methods

The following sections summarise the strengths and weaknesses of the methods developed and tested in the case studies during the course of this study.

6.2.1 Overview methods

The overview methods were tested in a 'post-event' mode using information from recent flood events. They effectively indicated, in broad terms, levels of FIM performance. They provided a rapid means of comparing performance in different flood events. Table 6.1 and Table 6.2 summarise the strengths and weaknesses of applying the performance matrices and fish-bone diagrams to assessing FIM reliability and performance.

Table 6.1 Strengths and weaknesses of the performance matrices.

Strengths	Weaknesses
Offers a relatively quick and easy method of assessing performance	Performance descriptors are general rather than specific to a particular type and scale of flood incident
Matrices can readily be assessed at different spatial scales	Difficult to account for variable performance within individual elements of the matrices, particularly given the broad nature of many of the elements
Performance for different events can readily be assessed and compared	Matrices do not identify the reasons for poor / good performance
	Difficult to deal with uncertainty explicitly – however, uncertainty can be indicated by drawing an ellipse that covers more than one performance category

Table 6.2 Strengths and weaknesses of the ‘Fish-bone’ Diagram.

Strengths	Weaknesses
Easy and rapid to use, especially where structured assessment is required of reasons for failure or severe under-performance	Does not provide a measure of performance – it simply identifies the key factors that influence the failure or success
Serves as a useful precursor to developing a more detailed method of assessing performance (i.e. assisting the development of the <i>Perimeta</i> models in this particular case)	Does not (except in a general sense) indicate the degree of influence that a cause or situation has had on process performance

6.2.2 Perimeta

The *Perimeta* models were used in a ‘post-event’ mode to verify the model structure and its parameterisation. The models can also be used as a forward-looking planning tool to identify areas where investment should be planned in order to maintain or improve FIM performance.

Environment Agency staff indicated that used like this (i.e. ‘forward looking’) *Perimeta* could potentially be used to help justify the business case regarding plans to improve FIM. Before using the *Perimeta* models for such tasks, further sensitivity testing is required to demonstrate that the model structure and parameterisation are robust. It is important that this is done before using the models in a ‘forward-looking’ mode so that users will have confidence in the validity of model output when planning improvements to FIM.

Table 6.3 summarises the strengths and weaknesses of applying *Perimeta* to FIM.

Table 6.3 Strengths and weaknesses of *Perimeta* for use in FIM.

Strengths	Weaknesses
<p><i>Perimeta</i>'s hierarchical structure allows evidence of performance to be propagated through the model, allowing assessment of its influence on overall performance</p>	<p><i>Perimeta</i> is unable to represent dynamic simulation (i.e. feedback loops within systems)</p>
<p>The 'Italian flag' representation of performance is a powerful visual tool for engaging and communicating with stakeholders</p>	<p><i>Perimeta</i> is not good at representing situations where performance could vary significantly as system influences and other key variables change</p>
<p>Uncertainty in the evidence on performance can be represented by using interval probabilities – evidence can be based on expert judgement</p>	<p>Understanding and applying values to the weights (necessity and sufficiency) between respective system processes requires training of the model builder</p>
<p>The strength of influence of 'child' processes on the parent process can be represented</p>	<p>The need to resolve conflicts in evidence during 'what-if' scenarios may affect the model's use for forward-looking planning</p>
<p><i>Perimeta</i> is able to sit above a 'systems dynamics' model to show outputs from the simulation on performance</p>	
<p>The same model structure can be used with different evidence and parameter values for different scenarios or locations</p>	
<p>Building a <i>Perimeta</i> model generates a good understanding of the processes and linkages within FIM system</p>	

6.2.3 Agent-based modelling

The work undertaken within this study has gone a long way towards developing an agent-based model that can be used to provide risk-based analysis of the effectiveness of flood event management systems. This agent-based model simulation approach has demonstrated that it can provide insights that are not obtainable using other methods.

However, it is important to recognise the potential pitfalls of using an agent-based model. These are not predictive tools; they do not simulate the behaviour of individuals precisely, but they are useful for studying the aggregate response of the whole system and all its agents.

As with any model, agent-based models are not, on their own, recommended as a basis to justify policy and strategy decisions, but they can provide information that is unavailable from other existing techniques to support and enrich the evidence basis for certain FIM decisions. Sensitivity analysis can also be used to diagnose model behaviour and ensure the model and its limitations are well understood.

Agent-based models can inform higher level methods (e.g. through quantifying the impacts, in terms of exposure to floodwater, as a function of warning system effectiveness). Conversely they can have their parameterisation informed by high level measures of performance (e.g. the proportion of people likely to act on a flood warning if they receive one).

The focus in this study was to implement a proof of concept demonstration and explore whether an agent-based model can provide insights into FIM. We had to simplify a number of aspects of the flooding system in order to achieve this objective in the time and with the resources available. Therefore, the agent-based model produced by this study predominantly focuses on evacuation and flood warning effectiveness; it was not possible to extend the model beyond these processes given the constraints of the available project resources.

To implement a working model within the project's constraints, we limited the scope of the study. We did not represent all the possible agents involved in a flood event (e.g. no Fire and Rescue Service), nor changes in behaviour resulting from adverse weather conditions (e.g. the model does not allow for people being more likely to drive to work on a rainy day). The model does not include a wide range of behaviours for each agent; each agent could only carry on as normal or evacuate the floodplain immediately in response to a flood warning, whereas one might expect some people to help their neighbours, move property upstairs, rush to pick up their children from school etc.).

To facilitate the planning of future flood event management modelling activities, we have identified further steps for future work in the longer term (see Section 7.1).

Table 6.4 summarises the strengths and weaknesses of FIM agent-based modelling.

Table 6.4 Strengths and weaknesses of FIM agent-based modelling.

Strengths	Weaknesses
Complex system behaviour can be modelled relatively simply to provide valuable insights	Data-hungry, i.e. requires a significant amount of data to set up a working agent-based model
Can be used to test the impact of FIM policy and response measures, and also the sequential failure of different components of the FIM system	Requires prior knowledge of agent-based modelling and an experienced user to set-up the models
Specific rule-sets can be defined for each type of agent to allow agents to respond to external factors as realistically as possible	The availability of data for rule sets is often limited, particularly with regard to human behaviour under flood conditions
Agent-based models have a strong track record in other disciplines and show promise in FIM.	Very few FIM applications to date and the governing model equations (e.g. behaviour rules) will, by the very nature of modelling human agency, always be open to debate
The dynamics of the flood event can be simulated enabling a wide range of FIM options to be tested	Validation methods need to be improved
Can be used to test a wide range of decisions and FIM policies, and provide outputs for use in quantified, risk-based analysis	
Could be integrated with <i>Perimeta</i> modelling within the FIM decision framework	

7 Recommendations

7.1 Introduction

The decision-support tools (performance matrices and root cause analysis, *Perimeta* based hierarchical modelling and agent-based modelling) tested in this study have all been tested to a proof-of-concept level. In their current forms, they are emerging prototypes that can continue to be applied, in an exploratory way, to help assess the reliability of FIM and evaluate FIM performance in a more structured and systematic way than has to date been possible.

This study has provided an opportunity to test the application of these tools to FIM – in most cases this has been their first documented application to FIM. This represents a significant contribution to the science base relating to FIM, in particular, and to assessing the performance and reliability of non-structural flood risk management measures in general.

Further testing and progressive refinement of each of the tools is still necessary before they can be issued as general tools for use by Environment Agency staff. However, they are at a stage where their further development could be undertaken by selected Environment Agency staff for specific applications.

7.2 General recommendations

The following general recommendations arise from the technical and science work carried out in this study. They are presented under two timescale headings (short-term, and medium-term) and are listed in order of priority.

7.2.1 Suggested developments in the short term

1. Wider trialling of these models on a range of case study areas and flood types to assess more extensively:
 - a. their effectiveness in improving performance measurement;
 - b. the development of generic prototype forms (or templates) of each model for use within the Environment Agency;
 - c. to apply these prototypes to actual FIM planning issues and to business decision making within the Environment Agency.
2. Testing the benefits of a broader application of the methods to strategic assessments (such as the Long Term Investment Strategy (LTIS)) and developments in flood forecasting processes and reliability.

7.2.2 Suggested developments in the medium term

3. The development of more refined forms of model validation that check the ability of the models to represent human (behaviour) factors and systems as well as the physical, technical, and information systems involved in FIM.

4. The need to devise new performance measures within parts of the FIM system, for example, at an intermediate level within the *Perimeta* model to provide clearer evidence on the effectiveness of key processes.
5. The integration of the tools within a decision support framework, incorporating improved means of benefit assessment, to evaluate the efficacy of improvements to FIM performance and its influence on flood risk management (via future developments in the MDSF and RASP methods).

7.3 Specific recommendations

Further science R&D is required to refine each of the methods developed and tested during the course of this study. These are outlined below.

7.3.1 Overview-level methods

6. Further development of performance matrices to:
 - a. refine their structure through wider testing to ensure the adequacy and robustness of their representation of key aspects of the FIM system;
 - b. refine the wording of the performance descriptions, and supplement with suitable metrics, to reduce the level of subjectivity in applying the matrices;
 - c. relate the performance categories to levels of expected performance that take into account the nature, size, severity and scale of a flood event.
7. Further development of root cause analysis and fish-bone diagrams to:
 - a. test the application of root cause analysis in situations where serious system failures have occurred;
 - b. refine the structure of, and descriptors within, the fish-bone diagrams through wider testing;
 - c. develop generic templates for different aspects of the FIM system to facilitate their use and practical application.

7.3.2 *Perimeta* hierarchical process modelling

8. Further work on the *Perimeta* modelling to:
 - a. refine the structure and parameterisation of the *Perimeta* models through wider testing to ensure their adequacy to represent key aspects of the FIM system;
 - b. further test the sensitivity of the *Perimeta* models to determine the robustness of their structure and parameterisation;
 - c. assess the influence of the quality of input data and validation data on the quality of the models over a range of flood areas and flood events;
 - d. explore further the application of the model in forward-looking planning applications;

- e. develop a simple user guide and training course on *Perimeta* to allow the Environment Agency to build, test and use its own *Perimeta* models.

7.3.3 Agent-based modelling

Key recommendations with respect to agent-based modelling are outlined below. It is important to note that many of the recommendations below involve extending the flood event management model which has been built from scratch. By contrast, *Perimeta* is a more established modelling software system so the type of recommendations made below are quite different to those listed above.

Validation of agent-based models

A key issue is the development of improved validation methods. Currently each component of the model is parameterised as rigorously as possible using the best (easily) available information. Previous attempts to validate flood evacuation models focus on calibrating against an observed death toll (e.g. using MassVac (Hobeika and Jamei, 1985) or the Life Safety Model (Johnstone *et al.*, 2005)). The inadequacies in this method should be obvious from the results described in the previous sections. Figure 5.3 particularly shows how the variation in impacts under just one scenario are sensitive to the randomness in starting conditions. This leads to the following recommendation:

9. To develop an approach to calibration of agent-based models based on fusing data from multiple sources including:
 - a. spatial location of mortalities and injuries;
 - b. traffic monitoring data;
 - c. functional descriptions of organisational processes enacted during a flood;
 - d. post flood reports from organisations and eye witnesses of behaviour, flood responses and their effectiveness;
 - e. expert judgement from stakeholders on how closely the model replicates their understanding and experiences of flood events.

This approach would fuse the best available data to obtain the most representative simulation of flood events.

Other recommendations in the area of validating agent-based models are:

10. Further work to:
 - a. establish more rigorously the appropriateness of the agent behaviour types;
 - b. establish the sensitivity of the model to different behaviours;
 - c. identify whether 'national' standard classifications could be developed to facilitate the application of these methods in other locations.

Inclusion of multiple agents

Currently only members of the public are included in the agent-based model. We therefore recommend

11. A stakeholder mapping exercise should be held to identify other key agencies for inclusion in the agent-based model.

These additional agents could include vehicles and pedestrians; businesses; blue light services (police, fire and rescue service and ambulance); agents who support warning dissemination and rescue services; flood risk managers (Environment Agency, engineers, contractors responsible for deployment of temporary defences); and local authority emergency planners.

Representing these agents would greatly improve the capability of the agent-based model and the processes that it can capture. It is likely that the inclusion of these agents will require the model to be extended to include a number of additional transport modes, for example, boat equipment or helicopters to reach trapped agents.

Traffic modelling

Traffic simulation could be improved to represent driver movement more accurately such as the slowing effects of bends and hills, traffic calming measures, traffic light rules and other aspects of driving. We therefore recommend

12. Further exploration of the benefits of improved traffic modelling and micro-measures for FIM.

Micro-measure could include the location of warning signs; the adjustment of road travel directions (e.g. convert dual carriageways into one way to maximise evacuation capacity); or the use of traffic light systems to manage traffic evacuation.

Representation of organisational and individual responses

The reaction of individuals and organisations after they are informed of a potential flood event varies according to factors such as their previous experience of flooding events, their experience of receiving 'false alarms' and actions taken by the government to increase awareness of flood risk. We therefore recommend:

13. The use of structured interviews, role play or action research events and public surveys to obtain information on:
 - a. flood event management procedures;
 - b. the behaviour of individuals and organisations;
 - c. communication and interaction of individuals and organisations;
 - d. options open to flood event managers for mitigating the impacts of flooding;
 - e. how agents adapt their response to an emerging scenario, trying to find the most effective ways to avoid the hazard and limit the damage (to homes, transport links and key infrastructure) as the flood event unfolds.

This information can be used to construct a fuller process model than is currently possible. This information could help to represent the behaviour of individuals and organisations and their interactions in more detail.

The decisions involved will inevitably involve trade-offs and the proposed methodology is an effective way to identify the processes enacted, the trade-offs considered and the resulting choices and behaviour. Implementing this superior understanding would greatly improve the agent-based modelling approach.

Further FIM measures

Our current agent-based model considers a limited range of FIM measures, namely evacuation and flood warning. The current response to a flood warning also assumes the agent either responds immediately or resumes their previous task. Within the current model all agents receive the flood warning, regardless of whether they are in a car, at home or work. We recommend

14. Improvement to the rules in the agent-based model so that:
 - a. the dissemination efficiency of a flood warning reflects post-flood event survey data and the locality of the agent at the time the warning is issued;
 - b. the model allows other agent response to a flood warning, for example 'collect family from school or work' or 'move belongings upstairs' or 'remain in the house or place of work'.
15. Consider additional flood event management measures.

Other FIM measures could include:

- deployment of warning and diversion signs deployed near the flooded area;
- use of loud halers, public address systems and automated messaging systems;
- deployment of temporary defences;
- media reporting of flooding;
- public education;
- shoreing of embankments with sandbags;
- patrolling and inspecting embankments during a flood event;
- use of individual property protection measures;
- movement of valuables upstairs; and
- the role of emergency responders.

Extension of the agent-based model to include these measures would provide significant benefits from an operational and appraisal perspective.

Vulnerability functions

For an evaluation of flood risk and effectiveness of measures it is important to have insight in the human impact of floods. There are a number of existing models (e.g. Tapsell *et al.*, 2002; Defra and Environment Agency, 2006; Jonkman and Kelman, 2005), and we recommend

16. The evaluation of existing models of the human impact of floods and the coupling of these models with post-flood records of data on loss of life and injury (Jonkman and Vrijling, 2008) to develop a series of stochastic functions that parameterise vulnerability of individuals in terms of vulnerabilities (such as mortality, injury, trauma, and hypothermia).

These vulnerabilities will be conditional on flood characteristics such as inundation speed, water depth, rate of water rise, time of day etc. in addition to information on population characteristics such as age, health, mobility and other vulnerability measures. Integration of these types of functions into an agent-based model would greatly improve the impacts assessment.

References

- AXELROD, R., 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Princeton: Princeton University Press.
- AXELROD, R., and TEFATISON, L., 2006. A Guide for Newcomers to Agent-Based Modelling in the Social, In Tesfatison L and Judd K L (eds.) *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, Appendix A, *Handbook in Economics Series*, Amsterdam: Elsevier/North Holland.
- BATES, P. D. and DE ROO, A. P. J., 2000. A simple raster-based model for flood inundation simulation, *Journal of Hydrology*, 236: 54–77.
- BOSSOMAIER, T., JARRRATT, D., ANVER, M. M., SCOTT, T. and THOMPSON, J., 2005. Data integration in agent-based modelling, *Complexity International*, Vol. 11.
- BROUWERS, L. and VERHAGEN, H., 2003. Applying the Consumat Model to Flood Management Policies, *Agent-Based Simulation*, 4, 28–30 April, Montpellier, France.
- BROUWERS, L., HANSSON, K., VERGAHEN, H. and BOMAN, M., 2009. Agent Models of Catastrophic Events. Found on 27 March 2009 at people.dsv.su.se/~lisa/maamaw01.pdf.
- DAVIS, J. P. and HALL, J. W., 2003. A software supported process for assembling evidence and handling uncertainty in decision making, *Decision Support Systems*, vol 35, pp. 415–433.
- DAWSON, R., HALL, J. W. and DAVIS, J. P., 2004. Performance based management of flood defence systems, *Proc Institution of Civil Engineers: Water Management*, vol 157, pp.35–44 .
- DAWSON, R. J. and HALL, J. W., 2006. Adaptive importance sampling for risk analysis of complex infrastructure systems, *Proc. R. Soc. A*, 462(2075): 3343–3362.
- DECHTER, R. and PEARL, J., 1985. Generalized best-first search strategies and the optimality of A*. *Journal of the ACM* 32(3):505–536.
- Defra and Environment Agency, 2006. *Flood risks to people – Phase 2*, Defra, London.
- DI MAURO, M., and LAMBRUSO, D., 2009. Hyrdodynamic and loss of life modelling for the 1953 Canvey Island flood, In Samuels, P et al., (eds) *Flood Risk Management: Research and Practice*, London: Taylor and Francis.
- GALEA, E., OWEN, E. and LAWRENCE, P., 1996. The EXODUS model, *Fire Engineers Journal*, July 1996: 26–30.
- GILBERT, N. and TROITZSCH, K., 2005. *Simulation for the Social Scientist*, Milton Keynes: Open University Press.
- HALL, J. W., DAWSON, R. J., SAYERS, P. B., ROSU, C., CHATTERTON, J. B. and DEAKIN, R., 2003. A methodology for national-scale flood risk assessment, *Water and Maritime Engineering*, ICE, 156(3) 235–247. .
- HALL, J. W., LE MASURIER, J. W., BAKER-LANGMAN, E. A., DAVIS, J. P. and TAYLOR, C. A., 2004. A decision support methodology for performance-based asset management, *Civil Engineering and Environmental Systems*, vol 21, pp. 51–74.
- HARDING, N., GODFREY, P. and BENJAMIN, D., 2003. Delivering Best Value Highway Performance, presented at Transportation Research Board 83 Annual Conference, Washington DC.

- HELBING, D., FARKAS, I. and VICSEK, T., 2000. Simulating dynamical features of escape panic, *Nature* 407, 487–490 (28th September 2000).
- HOBEIKA, A. G. and JAMEI, B., 1985. MASSVAC: A model for calculating evacuation times under natural disaster, *Emergency Planning, Simulation Series* 15: 23–28.
- HOROWITZ, A. J., 2006. Statewide travel forecasting models, National Cooperative Highway Research Program, National Research Council (U.S.) Transportation Research Board.
- ISHIKAWA, K., 1990. *Introduction to Quality Control* (translated by Loftus J H), 448p, ISBN 4-906224-61-X.
- JOHNSTONE, W. M., SAKAMOTO, D., ASSAF, H., BOURBAN, S., 2005. Architecture, modelling framework and validation of BC Hydro's virtual reality life safety model. In: Vrijling et al. (eds) *Proceedings of the international symposium on stochastic hydraulics*, Nijmegen, 23–24 May 2005.
- JONKMAN, S. N., KELMAN, I., 2005. An analysis of causes and circumstances of flood disaster deaths, *Disasters*, Vol. 29 No. 1 pp. 75–97.
- JONKMAN, S. N. and VRIJLING, J. K., 2008. Loss of life due to floods, *J. Flood Risk Management*, 1: 43–56.
- KAPLAN, R. S. and NORTON, D. P., 1992. The balanced scorecard: measures that drive performance, *Harvard Business Review* Jan – Feb pp. 71–80.
- KAPLAN, R. S. and NORTON, D. P., 1993. Putting the Balanced Scorecard to Work, *Harvard Business Review* Sep – Oct pp2–16.
- KAPLAN, R. S. and NORTON, D. P., 1996. Using the balanced scorecard as a strategic management system, *Harvard Business Review* Jan – Feb pp. 75–85.
- LAMBRUSO, D., GAUME, E., LOGTMEIJER, C., MENS, M. and VAN DER VAT, M., 2008. *Evacuation and Traffic Management*, Report T17-0702, FLOODsite, WL/Delft Hydraulics.
- Newcastle University and Halcrow, 2006. *Thames Estuary 2100 Early Conceptual Options: A framework for Uncertainty Analysis in TE2100*, ECO Report R10X, Environment Agency.
- ONS and DfT, 2006. *Transport statistics bulletin: The National Travel Survey 2006*, Department for Transport, London.
- PARKER, D. J., 1991. *The Damage Reducing Effects of Flood Warnings*. Report Prepared for Halcrow National and for the National Rivers Authority (Anglian Region) Regional Telemetry Scheme Appraisal, FHRC, Enfield, UK.
- PARKER, D. J., TUNSTALL, S. M., McCARTHY, S. M., 2007. New insights into the benefits of flood warnings: results from a household survey in England and Wales, *Environmental Hazards*, 7, 193–210.
- SALLACH, D. and MACAL, C., 2001. The simulation of social agents: an introduction. *Special Issue of Social Science Computer Review*, 19, (3), 245–248.
- SAMUELSON, D. A., 2000. *Designing Organisations*, *OR/MS Today*, December.
- STILL, K. G., 1993. New Computer system can predict human behaviour response to building fires, *Fire*, 84(January 1993):40–41.
- STILLWELL, J., DUKE-WILLIAMS, O., 2003. A new web-based interface to British census of population origin – destination statistics, *Environment and Planning A*, 35(1):113-132.

TAPSELL, S. M., PENNING-ROWSELL, E. C., TUNSTALL, S. M. and WILSON, T. L., 2002. Vulnerability to flooding: health and social dimensions, Phil. Trans. R. Soc. A, 360: 1511–1525.

WALSH, C. L., HALL, J. W., STREET, R. B., BLANKSBY, J., CASSAR, M., EKINS, P., GLENDINNING, S., GOODESS, C. M., HANDLEY, J., NOLAND, R. and WATSON, S. J., 2007. Building Knowledge for a Changing Climate: collaborative research to understand and adapt to the impacts of climate change on infrastructure, the built environment and utilities. Newcastle University.

WONG, K. H. L. and LUO, M., 2005. Computational tool in infrastructure emergency total evacuation analysis, in P. Kantor *et al.* (eds.) Intelligence and Security Informatics, Springer, Berlin.

ZARBOUTIS, N. and MARMARAS, N., 2005. Investigating Crowd behaviour during Emergency Evacuations using Agent-Based Modelling, in Proc. EAM 2005, Oct. 17th–19th, Athens, Greece.

Appendix 1 *Perimeta* model data

The data used in the *Perimeta* models is detailed in Table A1.1 to Table A1.4.

Table A1.1 Evidence requirements for the coverage model.

Child process	Parent process	Evidence requirements
Having rainfall radar network compliance	Assessing rainfall detection compliance	Green = Rainfall radar coverage complies with LoS WI White = There is some rainfall radar coverage but it does not comply with requirements of LoS Red = No rainfall radar coverage
Having rainfall gauge network compliance		Green = Rainfall gauge coverage complies with LoS WI White = There is some rainfall gauge coverage but it does not comply with requirements of LoS Red = No rainfall gauge coverage
Assessing rainfall detection compliance	Complying with detection LoS	Propagated evidence
Assessing water level network compliance		Green = Water level gauge coverage complies with LoS WI White = There is some water level gauge coverage but it does not comply with requirements of LoS Red = No water level gauge coverage
Identifying appropriate LoS	Complying with detection LoS	This process is about full compliance with the FW LoS and identifying the correct service levels. Service level can be equal to or higher than required by the WI but not lower.
	Complying with forecasting LoS	Green = Correctly identified LoS i.e. in line with LoS WI White = Accuracy of LoS unknown Red = Known to have incorrect LoS assigned a – based upon Flood Zones rather than detailed modelling b – based on out of date assessment method
	Complying with dissemination LoS	
	Complying with communication LoS	
Identifying existence of forecasting models for FWA	Complying with Forecasting LoS	Green = Forecasting technique has been configured on NFFS White = Forecasting technique has NOT been configured on NFFS and / or is not appropriate. Red = No forecasting technique exists The target FAR and POD must be met before a forecasting technique can be used as the primary method
Having dissemination methods in place	Complying with dissemination LoS	Green = Appropriate Dissemination methods, as required by the LoS, are available White = The available Dissemination methods do NOT fully meet the requirements of the LoS. Red = No Dissemination methods exist
Offering a FWS	Complying with communication LoS	Green = Proportion of properties within the EFO that have been offered a FW Service White = Proportion of properties within the EFO for which it is not certain whether a FW Service has been offered Red = Proportion of properties within the EFO that have not been offered a FW Service

Table A1.1 continued overleaf

Table A1.1 continued

Child process	Parent process	Evidence requirements
Conducting local awareness activities	Complying with communication LoS	Green = Activities are in line with the LoS WI White = Some activities are carried out but NOT in line with the LoS WI Red = No activities carried out
Maintaining existing customers	Complying with communication LoS	Green = Contact is in line with the LoS WI White = Some contact but NOT in line with the LoS WI Red = No contact
Complying with detection LoS	Ensuring ability to provide FWS	Propagated evidence
Complying with forecasting LoS	Ensuring ability to provide FWS	Propagated evidence
Complying with dissemination LoS	Ensuring ability to provide FWS	Propagated evidence
Ensuring ability to provide FWS	Achieving coverage targets	Propagated evidence
Complying with communication LoS	Achieving coverage targets	Propagated evidence

Table A1.2 Evidence requirements for the service effectiveness model.

Child process	Parent process	Evidence requirements
Having reliable communications from outstations to RTS	Having timely, available data	Green = proportion of communication links (BT, mobile network etc) fully functioning. Red = proportion of communication links not functioning White = proportion of communication links where reliability unknown
Having appropriate frequency of data polling		Data has been polled from outstations in a timely manner appropriate to the catchment and forecasting requirements. Green = evidence to support polling is appropriate Red = evidence to support inappropriate polling White = no evidence
Suitable positioning of detection equipment	Having effective detection systems	Ensuring gauges are located effectively within their immediate environment, i.e. radar/rain gauges away from building, trees, sprinklers etc. River gauges in relation to bridges, structures and whether gauges are bypassed. Green = suitably located, unaffected by environs Red = not suitably located, known to be affected by environs White = suitability of location unknown
Verifying accuracy of data		Green = Verification is carried out and data is accurate Red = Verification is carried out and data is not accurate White = Uncertainty due to no verification Evidence = this process is a manual check as the data comes in, faith is placed that telemetry is showing correct data, unless it looks odd.
Having timely, available data		Propagated evidence
Having a reliable telemetry system		Green = proportion of system fully operational Red = proportion of system not operating White = measure of uncertainty of system performance.
Maintaining & calibrating detection equipment		Ensuring gauges are maintained and calibrated to AOD or local datum (whichever is appropriate), routine spot gauging for range of flows, rating reviews. Green = accurately calibrated (as far as technically feasible) Red = not calibrated White = inaccuracies/uncertainties in calibration

Table A1.2 continued overleaf

Table A1.2 continued

Child process	Parent process	Evidence requirements
Having reliable contingency detection systems	Having effective detection systems	Green = proportion of contingency system fully operational Red = proportion of contingency system not operating White = measure of uncertainty of contingency system performance.
Achieving high POD	Ensuring accuracy of forecasting	Green = proportion of forecasts that meet the FWLOS Red = proportion of forecasts that do NOT meet the FWLOS White = uncertainty as to whether forecast model is able to achieve POD rate
Achieving low FAR		Green = proportion of forecasts that meet the FWLOS Red = proportion of forecasts that do NOT meet the FWLOS White = uncertainty as to whether forecast model is able to achieve FAR rate
Ensuring accuracy of forecasting	Having effective forecasting systems	Propagated evidence
Having reliable NFFS		Green = proportion of NFFS functioning as required, i.e. all data required for forecasts to run is available and system working Red = proportion of NFFS not functioning as required White = measure of uncertainty of system performance.
Having reliable contingency forecasting systems		Green = proportion of contingency forecasting methods functioning as required, i.e. all data required for forecasts to run is available and system working Red = proportion of contingency forecasting methods not functioning as required White = measure of uncertainty of system performance.

Table A1.2 continued overleaf

Table A1.2 continued

Child process	Parent process	Evidence requirements
Having reliable FWD (Maximum LoS)	Having effective dissemination systems	Green = Proportion of FW recipients that were sent a warning from FWD. Red = Proportion of FW recipients that were NOT sent a warning from FWD. White = uncertainty as to whether FWD is performing / not measured.
Having reliable siren/loudhailer systems (Intermediate LoS)		Green = Proportion of FW recipients that were sent a timely warning as a result of loudhailers, sirens systems. Red = Proportion of FW recipients that were NOT sent a timely warning as a result of loudhailers, sirens systems. White = uncertainty as to whether loudhailers/ sirens performing as per service levels / not measured.
Having reliable broadcast systems (Minimum LoS)		Green = Proportion of FW recipients that were sent a timely warning as a result of broadcast systems (media/partners). Red = Proportion of FW recipients that were NOT sent a timely warning as a result of broadcast systems (media/partners). White = uncertainty as to whether broadcast system performed as per service levels / not measured.
Having reliable alternative contingency FW systems		Green = Proportion of FW recipients that were sent a timely warning as a result of alternative contingency systems. Red = Proportion of FW recipients that were NOT sent a timely warning as a result of alternative contingency systems. White = uncertainty as to whether alternative contingency system
Having a good office environment & equipment	Ensuring good working environment & equipment	Green = measure of adequacy of facilities in incident rooms and user satisfaction. Red = measure of inadequacy of facilities in incident rooms and user dissatisfaction. White = uncertainty of adequacy, either not measured or mixed views of satisfaction.
Having a good remote environment & equipment		Green = measure of adequacy of facilities at home/remote and user satisfaction. Red = measure of inadequacy of facilities at home/remote and user dissatisfaction. White = uncertainty of adequacy, either not measured or mixed views of satisfaction.

Table A1.2 continued overleaf

Table A1.2 continued

Child process	Parent process	Evidence requirements
Ensuring good working environment & equipment	Having effective staff	Propagated evidence
Having an effective training programme		<p>Effective training means having a programme in place that is of good quality, at correct frequency and valued by staff.</p> <p>Green = Effective training is implemented Red = Training is not implemented for some or all duty officers White = Training takes place but is not effective</p>
Appointing competent staff to rosters		<p>This is about having the right staff allocated to the right duty roles.</p> <p>Green = Duty officers are competent, appropriate and/or experienced enough to perform the role they are allocated Red = Duty officers not competent or experienced to perform the role they are allocated White = Uncertainty as to whether duty officer is appropriate, (i.e. duty officers undergoing training/ experience).</p>
Having effective exercises		<p>Effective exercising means exercises are carried out that are of good quality, at correct frequency and valued by staff.</p> <p>Green = Effective exercising is implemented Red = Exercising is not implemented for some or all duty officers White = Exercising takes place but is not effective / not measured</p>
Having adequate numbers on rosters		<p>This assumes the number of staff required for each role has been established.</p> <p>Green = established numbers of staff are on all rosters Red = shortfall of established numbers on rosters White = uncertain as to how many staff required for each role</p>

Table A1.2 continued overleaf

Table A1.2 continued

Child process	Parent process	Evidence requirements
Ensuring effective staff communication at all levels	Having good communication & information management	<p>This relates to all staff communications within Incident Management, e.g. effectiveness of communication between FWDO and MFDO.</p> <p>Green = a measure of effective (successful) communications during an incident. Red = a measure of ineffective (unsuccessful) communications during an incident, poor understanding of forecasts by users. White = Uncertainty as to reliability of communications / not assessed</p> <p>Source – debrief reports, subjective</p>
Having effective Incident Management System		<p>This relates to the effectiveness of Incident Management Systems and Processes used by the EA, including data and information management, resource management and tools to manage the incident.</p> <p>Green = a measure of the effectiveness or success of Incident Management Systems and Processes Red = a measure of ineffective or unsuccessful Incident Management Systems and Processes / non-existent White = Uncertainty as to how well Incident Management Systems and Process perform / not measured</p> <p>Source – Flood report, lessons identified etc</p>
Having effective detection systems	Having effective technical systems	Propagated evidence
Having effective forecasting systems		Propagated evidence
Having effective dissemination systems		Propagated evidence

Table A1.2 continued overleaf

Table A1.2 continued

Child process	Parent process	Evidence requirements
Having effective staff	Making good FW decision	Propagated evidence
Having live effective procedures in place		Includes procedures for all processes in detection, forecasting, dissemination and contingency. Must be in place and tested. Green = Effective procedures are in place and tested Red = Procedures are not in place White = Uncertainty relating to the effectiveness of procedures / not tested
Validating FW thresholds		Green = proportion of thresholds that are validated and correct. Red = proportion of thresholds that are validated and incorrect. White = no validation undertaken.
Having good communication & information management		Propagated evidence
Having effective technical systems		Sending accurate, timely & reliable flood warnings
Making good FW decision		Propagated evidence
Collecting data from external sources	Identifying flooded, serviced properties	Collecting numbers of flooded "serviced" properties (as defined in WI "FW Performance Measures") from professional partners. Where serviced = covered under the FWLOS (see Coverage model). Green = Evidence to support successful data collection, validated. Red = No data recorded White = Uncertainty of data recorded, unconfirmed/not validated.
In house data collection		Collecting numbers of flooded "serviced" properties (as defined in WI "FW Performance Measures") from Flood Data Recorders, Flood Ambassadors, Call Handlers database and Floodline, other EA staff. Where serviced = covered under the FWLOS (see Coverage model). Green = Evidence to support successful data collection, validated. Red = No data recorded White = Uncertainty of data recorded, unconfirmed/not validated.

Table A1.2 continued overleaf

Table A1.2 continued

Child process	Parent process	Evidence requirements
Identifying flooded, serviced properties	Measuring and Reporting Effectiveness of FWS	Propagated evidence
Sending accurate, timely & reliable flood warnings		Propagated evidence

Table A1.3 Necessity and sufficiency values for the 2007 coverage model.

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having rainfall radar network compliance	Assessing rainfall detection compliance	<p>S = 0.55 Adequate radar coverage alone will satisfy the FWLoS WI but sufficiency set slightly lower than that associated with the rain-gauge process since Agency possibly rely more heavily upon rain-gauges than radar. Sufficiency of both child processes increased to reduce uncertainty in parent process.</p> <p>N = 0.40 Failure of radar is not a 'show-stopper' as long as rain-gauge network is functioning ok</p>
Having rainfall gauge network compliance		<p>S = 0.70 Adequate rain-gauge network alone will satisfy the FWLoS WI but sufficiency set slightly higher than that associated with the radar network since Agency possibly rely more heavily upon rain-gauges than radar. Sufficiency of both child processes increased to reduce uncertainty in parent process.</p> <p>N = 0.60 Failure of rain-gauge network is not a 'show-stopper' as long as radar network is functioning ok, but is perhaps more of a problem than the radar network failing.</p>
Assessing rainfall detection compliance	Complying with Detection LoS	<p>S = 0.40 Same as for "Assessing water level network compliance" as this is equally as important as the water level network compliance but on its own is not sufficient for the parent process to succeed.</p> <p>N = 0.80 Non-compliance of the rainfall detection has a big influence on the non-compliance with the Detection LoS.</p>
Assessing water level network compliance		<p>S = 0.40 Same as for "Assessing rainfall detection compliance" as this is equally as important as the rainfall detection compliance but on its own is not sufficient for the parent process to succeed</p> <p>N = 0.80 Non-compliance of the rainfall detection has a big influence on the non-compliance with the Detection LoS.</p>

Table A1.3 continued overleaf

Table A1.3 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Identifying Appropriate LoS	Complying with Detection LoS	<p>S = 0.20 Set relatively low, since knowing the appropriate LoS does not mean that the relevant system will be compliant - the appropriate system still has to be put in place, which is where the 'real' work is</p> <p>N = 0.80 Set high since failure to establish the correct LoS makes it impossible to establish compliance of detection, forecasting, and dissemination. However, compliance may be achieved unknowingly from the historical work to set up detection, forecasting and warning systems.</p>
	Complying with Forecasting LoS	<p>S = 0.30 Set relatively low, since knowing the appropriate LoS does not mean that the relevant system will be compliant – the appropriate system still has to be put in place, which is where the 'real' work is</p> <p>N = 0.80 Set high since failure to establish the correct LoS makes it impossible to establish compliance of detection, forecasting, and dissemination. However, compliance may be achieved unknowingly from the historical work to set up detection, forecasting and warning systems.</p>
	Complying with Dissemination LoS	<p>S = 0.30 Set relatively low, since knowing the appropriate LoS does not mean that the relevant system will be compliant – the appropriate system still has to be put in place, which is where the 'real' work is</p> <p>N = 0.80 Set high since failure to establish the correct LoS makes it impossible to establish compliance of detection, forecasting, and dissemination. However, compliance may be achieved unknowingly from the historical work to set up detection, forecasting and warning systems.</p>
	Complying with Communication LoS	<p>S = 0.40 Set relatively low, since knowing the appropriate LoS does not mean that the relevant system will be compliant – the appropriate system still has to be put in place, which is where the 'real' work is</p> <p>N = 0.80 Set high since failure to establish the correct LoS makes it impossible to establish compliance of detection, forecasting, and dissemination. However, compliance may be achieved unknowingly from the historical work to set up detection, forecasting and warning systems.</p>
Identifying existence of Forecasting models for FWA	Complying with Forecasting LoS	<p>S = 0.70 Set relatively high as existence of forecasting models plays a big role in the success of complying with forecasting levels of service.</p> <p>N = 0.50 The absence of forecasting models does not necessarily mean parent process will fail, as there could be other means of forecasting available (e.g. peak-to-peak correlations, rate-of-rise extrapolations, rainfall-triggers, etc).</p>

Table A1.3 continued overleaf

Table A1.3 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having Dissemination methods in place	Complying with Dissemination LoS	S = 0.70 Dissemination methods are required to ensure compliance with the Dissemination LoS but if they are not appropriate then non-compliance will occur. N = 0.80 If there are no dissemination methods available then there is no compliance with the parent process.
Offering a FWS	Complying with Communication LoS	S = 0.40 Success in offering a FWS does on its own not guarantee success in complying with Communication LoS N = 0.60 Failure to offer a FWS implies failure to comply with Communication LoS
Conducting Local Awareness activities		S = 0.10 The success of this process will have a very minor impact on the success of failure of the parent process N = 0.10 The failure of this process will have a very minor impact on the success of failure of the parent process
Maintaining existing Customers		S = 0.20 The maintenance of existing customers is essential to complying with the Communication LoS but less so than some of the other child processes feeding into the parent process N = 0.20 Failure to maintain existing customers will cause the ability to communicate to deteriorate, but over a long period of time, but doesn't have as much weight to the above process as the other two child processes
Complying with Detection LoS	Ensuring ability to provide FWS	S = 0.33 Complying with any one of the LoS (Detection, Forecasting, Dissemination) does not guarantee the success of the parent process. N = 0.40 Failure to comply with the relevant LoS (Detection, Forecasting, Dissemination) does not mean that a FWS cannot be provided, it just means that it may not be compliant.
Complying with Forecasting LoS		S = 0.33 Complying with any one of the LoS (Detection, Forecasting, Dissemination) does not guarantee the success of the parent process. N = 0.40 Failure to comply with the relevant LoS (Detection, Forecasting, Dissemination) does not mean that a FWS cannot be provided, it just means that it may not be compliant.
Complying with Dissemination LoS		S = 0.33 Complying with any one of the LoS (Detection, Forecasting, Dissemination) does not guarantee the success of the parent process. N = 0.40 Failure to comply with the relevant LoS (Detection, Forecasting, Dissemination) does not mean that a FWS cannot be provided, it just means that it may not be compliant.

Table A1.3 continued overleaf

Table A1.3 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Ensuring ability to provide FWS	Achieving coverage targets	<p>S = 0.60 Relative to the Communication aspect of achieving the coverage targets, being able to provide a FWS is (perhaps?) more important, but on its own is not sufficient to achieve the Coverage targets.</p> <p>N = 1.00 If it is not possible to provide a FWS then the Coverage targets will not be met, hence the reason why necessity is set to 1.0</p>
Complying with Communication LoS	Achieving coverage targets	<p>S = 0.40 Complying with Communication LoS is not sufficient in itself to ensure success in achieving coverage target</p> <p>N = 0.50 Failure to comply with Communication LoS may not imply failure to achieve coverage as flood warnings are in practice disseminated by a range of methods some of which may not (strictly speaking) be compliant</p>

Table A1.4 Necessity and sufficiency values for the 2007 service effectiveness model.

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having reliable comms from outstations to RTS	Having timely, available data	S = 0.60 Reliable comms are a vital part of having available data but are not on their own sufficient to ensure that data are also timely (this comes from the other child process) N = 0.70 If the comms are not reliable (i.e. fail), there will be no data available from the outstations via the preferred means, but the parent process will not fail completely as alternative means of obtaining data will be found by the Environment Agency (e.g. use of contingency system, on-site observations from staff etc)
Having appropriate frequency of data polling		S = 0.40 Polling data at the appropriate frequency is not on its own sufficient for the parent to succeed, but is required to ensure that the available data are timely N = 0.50 If data are not polled at the appropriate frequency then the available data may not be timely but may still be of some use (i.e. will not result in complete failure of parent). Data may also be available from elsewhere if the primary system fails
Suitable positioning of detection equipment	Having effective detection systems	S = 0.20 The positioning of detection equipment is one of several aspects that will influence the effectiveness of the detection system but is not on its own sufficient to ensure its effectiveness. It is deemed to be as influential over the parent process as several of the other child processes N = 0.40 Poor positioning of the detection equipment will have an impact on the effectiveness of the detection system but will not result in it being completely ineffective (assuming equipment is still reasonably positioned).
Verifying accuracy of data		S = 0.10 Verifying the accuracy of the data is one of several aspects that will influence the effectiveness of the detection system but is not on its own sufficient to ensure its effectiveness. It is considered one of the least important of the child processes, as even if the accuracy of the data is not verified it will still be useful N = 0.20 Inaccurate data will have an impact on the effectiveness of the detection system, but will not result in complete failure of the effectiveness of the detection system (assuming the inaccuracies do not mean that the data is completely useless). The accuracy of data is unlikely to deteriorate significantly over time
Having timely, available data		S = 0.20 Having timely and available data is one of several aspects that will influence the effectiveness of the detection system but is not on its own sufficient to ensure its effectiveness. It is deemed as influential over the parent process as several of the other child processes N = 0.40 Even if data are not available and timely then the detection system will not completely fail as there are other means of detecting what is happening (e.g. on-site observations, alternative detection systems, etc)

Table A1.4 continued overleaf

Table A1.4 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having a reliable telemetry system	Having effective detection systems	<p>S = 0.20 The proportion of the telemetry system that is working is one of several aspects that will influence the effectiveness of the detection system but is not on its own sufficient to ensure its effectiveness. It is deemed as influential over the parent process as several of the other child processes</p> <p>N = 0.40 For the proportion of the telemetry system that is not working no data will be available from the telemetry system but may be from elsewhere. This will reduce the effectiveness of the detection system but not result in its complete failure.</p>
Maintaining & calibrating detection equipment		<p>S = 0.20 Having accurately calibrated gauges working (i.e. maintained) is one of several aspects that will influence the effectiveness of the detection system but is not on its own sufficient to ensure its effectiveness. It is deemed as influential over the parent process as several of the other child processes</p> <p>N = 0.40 Gauges that are not accurately calibrated will result in less accurate data but will not result in it the detection system being completely ineffective.</p>
Having reliable contingency detection systems		<p>S = 0.10 Having a contingency system will not be sufficient on its own to ensure that the detection system is effective since this will largely be determined by the effectiveness of the main system (given that the contingency system will most likely be used very little, if at all)</p> <p>N = 0.20 As referred to above, since the contingency system is only likely to be used for a very small proportion of the time then overall if a proportion of it is not operational this will only have a small effect on the effectiveness of the detection system</p>
Achieving high POD	Ensuring accuracy of forecasting	<p>S = 0.50 Achieving a POD that complies with the FWLoS requirements will not on its own be sufficient to ensure the accuracy of forecasting but it will have a significant influence. Deemed as important as the other child process at determining the success of the parent</p> <p>N = 0.50 If the POD does not comply with the FWLoS requirements then this will have a significant impact on the accuracy of forecasts but will not necessarily mean that the forecasts are entirely inaccurate.</p>
Achieving low FAR		<p>S = 0.50 Achieving a FAR that complies with the FWLoS requirements will not on its own be sufficient to ensure the accuracy of forecasting but it will have a significant influence. Deemed as important as the other child process at determining the success of the parent</p> <p>N = 0.50 If the FAR does not comply with the FWLoS requirements then this will have a significant impact on the accuracy of forecasts but will not necessarily mean that the forecasts are entirely inaccurate.</p>

Table A1.4 continued overleaf

Table A1.4 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Ensuring accuracy of forecasting	Having effective forecasting systems	<p>S = 0.50 Being able to ensure the accuracy of forecasting is not sufficient on its own to guarantee the effectiveness of the forecasting system (e.g. other factors such of the reliability also affect the effectiveness). However this is considered to be the most important factor in determining the success of the parent process.</p> <p>N = 0.80 If the forecasts are inaccurate then this will strongly influence the effectiveness of the forecasting systems, but not mean that it is completely ineffective. Deemed to be the most influential of the child processes at determining the failure of the parent process</p>
Having reliable NFFS		<p>S = 0.35 Having a reliable NFFS is not sufficient on its own to ensure that the forecasting systems are effective (e.g. accuracy of the forecasts generated on the NFFS will also determine the effectiveness)</p> <p>N = 0.40 If certain parts of the NFFS are unreliable this will not necessarily mean that the effectiveness of the forecasting systems is entirely compromised as these parts may not be required or there may be contingency systems that can be used instead</p>
Having reliable contingency forecasting systems		<p>S = 0.15 Having a contingency forecasting system will not be sufficient on its own to ensure that the forecasting system is effective since this will largely be determined by the effectiveness of the main system (given that the contingency system will most likely be used very little, if at all)</p> <p>N = 0.20 As referred to above, since the contingency system is only likely to be used for a very small proportion of the time then overall if a proportion of it is not operational this will only have a small effect on the effectiveness of the forecasting system</p>
Having reliable FWD Maximum LoS)	Having effective dissemination systems	<p>S = 0.90 Since FWD is the primary means of disseminating warnings the reliability of FWD will largely determine the effectiveness of the dissemination systems. However, in the event of failure of FWD, alternative (contingency) means of disseminating the warnings will be used by the Environment Agency (e.g. dissemination by emergency services, local authorities etc)</p> <p>N = 0.90 Following on from the above, if FWD is not reliable then this will have a large influence on the effectiveness (or ineffectiveness) of the dissemination systems. However, given that alternative (ad-hoc?) means of dissemination are likely to be used by the Environment Agency if the situation requires it, failure of the FWD will not mean complete failure of the parent process</p>
Having reliable siren/loudhailer systems (Intermediate LoS)		<p>S = 0.00 Not used in Thames West Area</p> <p>N = 0.00 Not used in Thames West Area</p>

Table A1.4 continued overleaf

Table A1.4 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having reliable broadcast systems (Minimum LoS)	Having effective dissemination systems	S = 0.00 N/A in Thames West Area, but will be in the future N = 0.00 N/A in Thames West Area, but will be in the future
Having reliable alternative contingency FW systems		S = 0.10 Not necessary to use any contingency techniques in 2007, but small amount of weight placed on sufficiency as FWLoS specifies that a contingency method of dissemination should be in place N = 0.10 Not necessary to use any contingency techniques in 2007, but small amount of weight placed on necessity as FWLoS specifies that a contingency method of dissemination should be in place
Having a good office environment & equipment	Ensuring good working environment & equipment	S = 1.00 All work during the 2007 floods was undertaken from the office due to the scale of the event so this child process entirely defines the success of the parent N = 1.00 Following on from the text above, since all the work during the 2007 floods was undertaken from the office the state of the office environment and equipment determined the success or failure of the parent
Having a good remote environment & equipment		S = 0.00 Remote working not undertaken during 2007 event, hence value of 0 N = 0.00 Remote working not undertaken during 2007 event, hence value of 0
Ensuring good working environment & equipment	Having effective staff	S = 0.10 Having a good working environment and equipment will influence the effectiveness of staff, but to a lesser extent than some of the other influencing factors. N = 0.30 A poor working environment and equipment will reduce the effectiveness of staff but not to the extent that staff are completely ineffective
Having an effective training programme		S = 0.35 An effective training programme will help ensure the competence of duty officers, but is not sufficient on its own to ensure that duty officers are effective N = 0.30 If a proportion of duty officers are not trained then this will reduce the effectiveness of the duty officers but not result in complete ineffectiveness as this is determined by a range of factors (see other child processes)
Appointing competent staff to rosters		S = 0.10 Having appropriate staff for the duty roles will have an impact on the effectiveness of the staff in their duty roles but is not on its own sufficient to ensure the effectiveness of staff. Other factors are deemed more important. Staff who are not initially competent can become competent over time through training, exercises, experience etc N = 0.60 If the wrong staff are appointed to the duty roles then the effectiveness of these staff to undertake the duty roles will be much reduced

Table A1.4 continued overleaf

Table A1.4 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having effective exercises	Having effective staff	<p>S = 0.10 Training is seen as more important than exercising for ensuring that duty officers are competent in their roles, hence the reason why the sufficiency is lower for the exercise programme. Exercises tend to be undertaken relatively infrequently and are therefore considered to be less important at influencing the effectiveness of staff</p> <p>N = 0.30 If some of the duty officers have not taken part in an effective exercise programme this will reduce the competence of duty officers but to a lesser extent than not having the training (hence the lower necessity value here).</p>
Having adequate numbers on rosters		<p>S = 0.35 Adequate numbers of staff on the duty rosters will influence the effectiveness of staff but is not sufficient on its own to ensure that staff are effective. However this is considered one of the more important child processes</p> <p>N = 0.60 If there are not enough staff available for the duty roles then duty officers will be stretched during an event and their effectiveness will be reduced, but not to the extent that they are completely ineffective. However, it is likely that if there are insufficient numbers of staff on the rosters then staff from other areas will be called into help during an event.</p>
Ensuring effective staff communication at all levels	Having good communication & information management	<p>S = 0.50 Effective staff communication will help ensure good communication and information management, but is not sufficient on its own to ensure the success of the parent process</p> <p>N = 0.70 If staff do not communicate effectively during events then the communication of information will be poor. However, this does not necessarily mean that the management of information will be poor too</p>
Having effective Incident Management System		<p>S = 0.50 The effectiveness of incident management systems will help ensure that good communication and information management occurs during an event, but these are not sufficient on their own</p> <p>N = 0.50 If the incident management systems are not effective then this will affect the ability to maintain good communication and information management during events</p>

Table A1.4 continued overleaf

Table A1.4 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having effective detection systems	Having effective technical systems	<p>S = 0.40 Having an effective detection system is one of three processes that are required to ensure that the technical systems are effective, so on its own is not sufficient to ensure the effectiveness of the technical systems. Considered equally as important as the other child processes at determining the effectiveness of the technical systems</p> <p>N = 0.60 If the detection system is not effective, this will mean that the overall effectiveness of the technical systems is compromised as the actual situation 'on the ground' will not be reliably known which will in turn influence the effectiveness of the forecasts and ultimately the warnings that are disseminated.</p>
Having effective forecasting systems		<p>S = 0.40 Having effective forecasting systems is one of three processes that are required to ensure that the technical systems are effective, so on its own is not sufficient to ensure the effectiveness of the technical systems. Considered equally as important as the other child processes at determining the effectiveness of the technical systems</p> <p>N = 0.40 If the forecasting systems are not effective, this will mean that the overall effectiveness of the technical systems is compromised as the warnings will be based on poor forecasts. However, flood warnings can still be sent on the basis of the detection and dissemination systems, so arguably this is the least important of the three child processes in terms of influencing the ineffectiveness of the technical systems</p>
Having effective dissemination systems		<p>S = 0.40 Having effective dissemination systems is one of three processes that are required to ensure that the technical systems are effective, so on its own is not sufficient to ensure the effectiveness of the technical systems. Considered equally as important as the other child processes at determining the effectiveness of the technical systems</p> <p>N = 0.60 If the dissemination systems are not effective, this will mean that the overall effectiveness of the technical systems is compromised as it will not be possible to disseminate the warnings effectively.</p>

Table A1.4 continued overleaf

Table A1.4 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having effective staff	Making good FW decision	<p>S = 0.25 The effectiveness of staff will have a relatively strong influence on the ability to make a good flood warning decision, but is not sufficient on its own to ensure a good decision is made. Considered equally important as the other child processes</p> <p>N = 0.40 If staff are not effective then the chances of a good flood warning decision being made will be reduced</p>
Having live effective procedures in place		<p>S = 0.25 Having live effective and tested procedures in place will go a long way to ensuring that a good FW decision is made, but is not sufficient on its own to ensure that a good FW decision is made. Considered equally important as the other child processes</p> <p>N = 0.40 If there are no live effective and tested procedures in place this will make it more difficult to make good FW decisions but will not completely compromise the decisions. Reduced from original value following discussions with Agency during case study meetings.</p>
Validating FW thresholds		<p>S = 0.25 The greater the proportion of thresholds that have been validated then the more effective the FW decision will be, but on its own this is not sufficient to ensure that a good decision will be made. Considered equally important as the other child processes</p> <p>N = 0.50 The more thresholds that are not validated then the less effective the FW decision will be, but assuming that the initial thresholds are reasonable this will not have a big impact on the FW decision that is made. Reduced from original value on the basis of discussions with the Environment Agency. Considered the least important of the child process in determining the likelihood of a poor FW decision being made.</p>
Having good communication & information management		<p>S = 0.25 Having good communication and information management systems is not sufficient on its own to ensure that a good flood warning decision is made, but they will certainly assist with making good decisions. Considered equally important as the other child processes</p> <p>N = 0.50 Poor communication and information management systems will affect the ability to make good flood warning decisions, but not to the extent that good FW decisions cannot be made. Considered the most influential of the child processes in making a poor FW decision</p>

Table A1.4 continued overleaf

Table A1.4 continued

Child process	Parent process	Sufficiency (S) and Necessity (N) value, including justification of values
Having effective technical systems	Sending accurate, timely & reliable flood warnings	<p>S = 0.50 This is one of the two processes required to allow accurate, timely and reliable flood warnings to be sent. On its own it is not sufficient, but it is considered as important as making a good flood warning decision.</p> <p>N = 0.50 If the technical systems are not effective then this will affect the success of the parent process but will not result in its complete failure</p>
Making good FW decision		<p>S = 0.50 Making a good flood warning decision is part of the process of ensuring the accuracy, timeliness and reliability of flood warnings, but on its own is not sufficient. Considered as important as the other child process</p> <p>N = 0.50 If a good flood warning decision is not made then it will be more difficult to send accurate, timely and reliable flood warnings, but not make it impossible. Reduced from original value of 1.0 following discussions with Agency</p>
Collecting data from external sources	Identifying flooded, serviced properties	<p>S = 0.40 Assuming that more details of flooded, serviced properties are typically obtained from in-house data collection (as opposed to external sources), then this process will contribute to the success of the parent process but to a lesser extent than the in-house data collection process</p> <p>N = 0.40 If no data are collected from external sources then it is likely that not all of the flooded serviced properties will be identified</p>
In house data collection		<p>S = 0.60 Assuming that a larger proportion of the flooded serviced properties are identified from in-house data collection techniques than from external sources, in-house data collection will have more of an influence on the identification of flooded serviced properties than the collection of data from external sources but is not sufficient on its own</p> <p>N = 0.60 Not collecting data in-house will mean that a larger proportion of the flooded serviced properties are not identified than if the collection of data from external sources is not undertaken</p>
Identifying flooded, serviced properties	Measuring and Reporting Effectiveness of FWS	<p>S = 0.50 This is one of the two processes required to allow the measurement and reporting of the effectiveness of the FW System. On its own it is not sufficient.</p> <p>N = 1.00 If the flooded serviced properties are not identified then the measuring and reporting of the effectiveness of the FWS will fail.</p>
Sending accurate, timely & reliable flood warnings		<p>S = 0.50 This is one of the two processes required to allow the measurement and reporting of the effectiveness of the FW System. On its own it is not sufficient.</p> <p>N = 0.50 If accurate, timely and reliable flood warnings are not sent then this will not altogether prevent the measuring and reporting of the effectiveness of the FWS.</p>

Appendix 2 Rules of thumb for setting sufficiency and necessity

A2.1 Introduction

Sufficiency is a measure of how much the success of a child process influences the success of its parent. Although the term comes from a logical origin, some find it easier to think of the term as a measure of 'importance' or 'significance'. Sufficiency is related to the positive contribution of the child process on the parent.

Necessity is a measure of how much the failure of a child process will cause failure of its parent. Therefore necessity is related to failure or poor performance. If the parent can't succeed without the child succeeding then there is a case for increasing the value of the necessity term.

In extreme cases the sufficiency and/or necessity terms will equal 1.0.

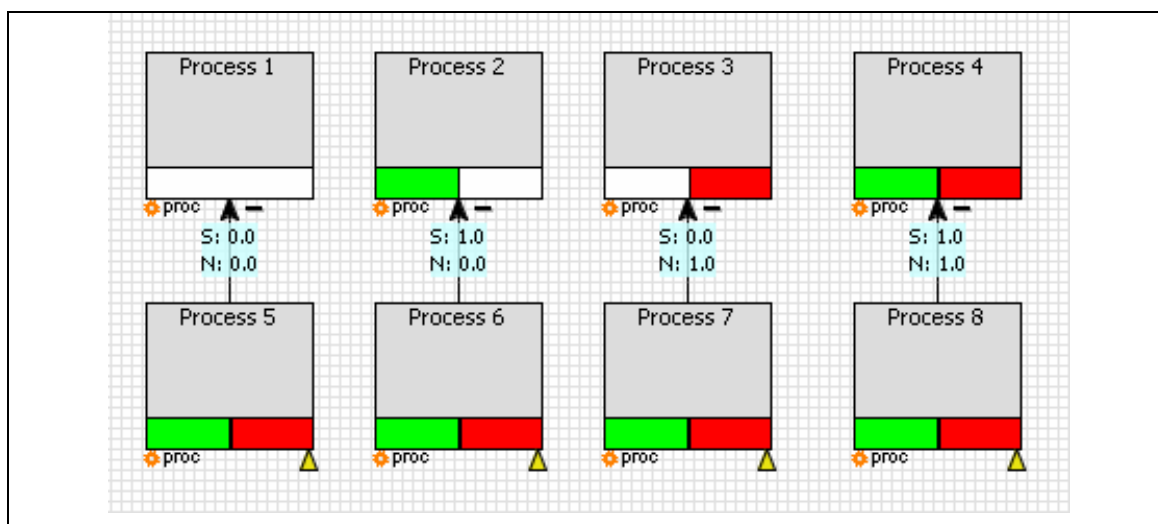


Figure A2.1 Diagram of the influence of child processes when sufficiency and/or necessity are set at either zero or 1.0.

In the first case in Figure A2.1 (i.e. process 1/5) necessity and sufficiency are both zero; the sub-process is not required and makes no contribution to the parent.

In the second case (process 2/6) sufficiency is set to 1.0 necessity to 0.0; the sub-process fully defines the success of the parent.

In the third case (process 3/7) sufficiency is 0.0 and necessity is 1.0. Failure of the parent (evidence against) is determined completely by the sub-process.

In the fourth case (process 4/8) necessity and sufficiency are both set to 1.0; the sub-process fully defines the parent for both success and failure.

In general it is rare to meet these extremes. We will consider the partial cases in the following sections.

A2.2 Sufficiency

A sufficiency of 1.0 means that nothing else is required. This sub-process alone will produce success. Hence it is very unusual to find this situation; it would mean that you only need one sub-process and hence you have a one-to-one mapping. But this implies the parent process is effectively the same as the sub-process.

However, it is important to understand a special case where a sufficiency of 1.0 would occur: when the logical connection between the parent and sub-processes is an OR statement. That is, the parent will succeed if we have either sub-process A working OR sub-process B. The current *Juniper* algorithm does not model this situation. In this case you have to make a note on the diagram and disconnect one or the other sub-process, depending on which is most effective.

The sufficiencies of sub-processes are not normalised; they don't have to add up to 1. The reasoning behind this can be seen in the Venn diagram in Figure A2.2.




Figure A2.2 Venn diagram of the degree of influence of three sub-processes (thin circles) on their parent process (thick circles).

In diagram A of Figure A2.2, there are three sub-processes with low importance/relevance (i.e. they have low sufficiency). The sufficiencies are not normalised because they don't contribute enough between them to fully satisfy the parent process.

In diagram B the three sub-processes have high importance/relevance (i.e. they have high sufficiency). Between them they fully satisfy the evidence requirement for the parent process but in doing so they overlap each other – there is dependency between them. So to properly account for this situation we increase dependency or more simply reduce the presumed sufficiency.

The practical implications of diagram B in Figure A2.2 are that it is unusual for the sum of the sufficiencies to add up to more than 1.0 (however, the way the heuristic works

means that the more sub-processes there are the more leeway there is on the sum of sufficiencies coming to more than 1.0). A sum of 1.25 is generally the upper limit.

A2.3 Necessity

The value of the necessity term can be argued in a similar way to the sufficiency term although the terms are not symmetrical. A necessity of 1.0 means that there is no way the parent can succeed if the child fails. In real life and man-made systems, it is rare that the necessity would ever equal 1.0 – it is usually possible that a parent may still succeed, for reasons that are not included within the model. So necessities may be set very high, but almost never at 1.0.

It is important to beware in the model that the combination of necessity set at 1.0 and totally negative evidence causes a fault in the logic; all other processes are swamped, causing a propagation error further up the tree.

We have the same logical modelling problem with necessity that we have with sufficiency. We may want to set several sub-processes with high necessities so that the parent will fail if any of the sub-processes fail. Again, this must be specifically marked in the modelling software, then, if one of the sub-processes does fail, the other sub-processes should be disconnected.

Naturally, sometimes we set sufficiency and necessity the same when we want there to be a balanced response.

A2.4 Dependency

The dependency looks at the relationship between sub-processes. It is rare that evidence from different sources is completely independent, although often it is assumed to be so. If we treat sources of evidence as independent we could be introducing bias or 'double dipping'. So we need a way to express any dependency between the sub-processes. In the *Juniper* algorithm, only the area between independent and fully dependent is active. 'Fully dependent' means the sub-processes are really the same thing (i.e. there is nothing new here).

A2.5 Too much evidence – conflict, incoherence and confusion

In some circumstances the evidence may be assumed to be independent and the sufficiencies are set high. But if the different sub-processes presenting a different picture, there is evidently some conflict (see Figure A2.3).

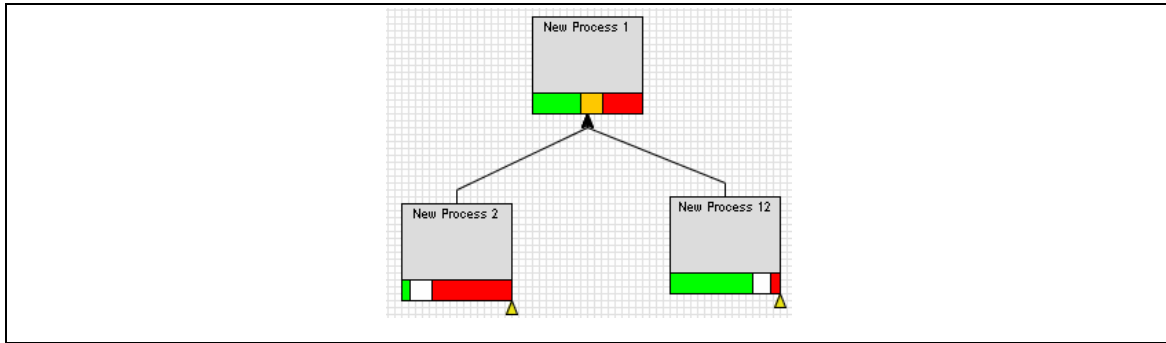


Figure A2.3 Diagram of conflicting evidence.

Figure A2.3 illustrates how conflicting evidence is revealed where the red and green stripes cross over, producing an amber band in the middle. This is a warning sign that you have to make a decision that no maths can help you with. It is like having two experts in front you, both of whom you would normally believe, one of whom says yes and the other says no. You either have to decide to discount one, reduce the overall sufficiencies or let the indecision propagate and deal with it later. This is a very important issue and one which many systems of this kind try to normalise away.

**Would you like to find out more about us,
or about your environment?**

Then call us on

08708 506 506* (Mon-Fri 8-6)

email

enquiries@environment-agency.gov.uk

or visit our website

www.environment-agency.gov.uk

incident hotline 0800 80 70 60 (24hrs)

floodline 0845 988 1188

*** Approximate call costs: 8p plus 6p per minute (standard landline).
Please note charges will vary across telephone providers**



Environment first: This publication is printed on recycled paper.