

Case Study on Risk Analysis for Critical Systems with Reliability Block Diagrams

Kim Weyns

Department of Computer Science
Lund University
kim.weyns@cs.lth.se

Martin Höst

Department of Computer Science
Lund University
martin.host@cs.lth.se

ABSTRACT

This paper presents a practical risk analysis method for critical, large-scale IT systems in an organisation. The method is based on reliability block diagram modelling and was adapted to fit the requirements of governmental organisations and to reduce the effort required to capture complex failure behaviour. Through the use of different failure categories the risk analysis can be simplified, the input data becomes easier to estimate and the results are easier to use in an organisational risk and vulnerability analysis. The paper first explicitly describes the different steps of the method and then presents a case study in which the method was applied and evaluated in a real-life setting. The method is meant to help an organisation to communicate internally about the reliability of their critical IT systems and to prioritise proposed improvements to this reliability.

Keywords

Risk Analysis, Reliability Block Diagram, Government, Critical System, Availability, Case study

INTRODUCTION

Problems with critical IT systems can hinder critical services in our society. Therefore it is important that these systems are an integral part of the major risk and vulnerability analyses conducted by the organisations that administrate these systems. A thorough risk analysis approach (Swedish Emergency Management Agency, 2003) needs to combine information about the reliability and the availability of these systems with an assessment of how critically dependant the organisation is on its most critical IT systems.

In this paper we focus on the first part of this risk assessment, namely analysing the availability and reliability of critical IT systems. For many systems, an analysis based on historical failure data is not an option. Instead, a risk analysis technique based on the structure and components of the systems is necessary. A number of risk analysis techniques exist for the assessment of the reliability of technical systems and more specifically IT systems. Many of these techniques, such as Fault Tree Analysis (Ericson, 1999) or Failure Mode and Effect Analysis (Beauregard, 1996), require a detailed analysis of all components of the systems and are therefore more suited for small, embedded systems with a very specific function. To conduct a similar analysis for large, complex systems, such as a patient data system, would require too large an effort to be realistic for most organisations.

In an earlier paper (Weyns et al., 2012) we identified a need and the requirements for a risk analysis method for large-scale IT systems that requires fewer resources, but that can still capture the complex failure behaviour typical for these systems. In this paper we present a risk analysis method that can be used in practice by governmental organisations to analyse the reliability and availability of their most critical IT systems. The method is based on Reliability Block Diagrams (RBD), and was adapted to make it easier for the organization to collect or estimate the input data needed and to make the results more easily applicable in the organisation's business continuity management. The method was applied and evaluated at a Swedish municipality on a real-life medical journal system.

RELATED WORK

In an earlier study specifically concerning IT dependability problems at Swedish municipalities (Weyns et al., 2009) we identified a number of important issues, one of which was the need for practical methods and techniques for risk analysis. A large number of detailed risk analysis techniques for technical systems already exist (U.S. Department of Defense, 2007). For example, Fault Tree Analysis (Ericson, 1999) is a top down

technique used to analyse all possible conditions that can lead to a certain failure in a technical system. Failure Mode and Effect Analysis (Beauregard, 1996) is a risk analysis technique to assess the probability and effects of possible failures of the system. A third technique, called Reliability Block Diagrams (RBD) (Staley and Sutcliffe, 1974), forms the basis for the method proposed in this paper. Reliability Block Diagrams, have previously been used to analyse the reliability of specific systems such as uninterruptible power supply systems (Rahmat et al., 2011) and UMTS networks (Dharmaraja et al., 2008). Each of these techniques requires a very detailed system model. Therefore these techniques are most suited for systems with well-defined components for which the failure behaviour can be predicted. In contrast, in this paper we propose a method better suited for complex, distributed systems, where the other techniques would require too much effort to be of practical use.

Several authors (Office of Government Commerce, 2007, Weyns et al. 2010) have proposed a process oriented approach to IT dependability management, in which the risk management process is one of the most important aspects in assessing the reliability of IT systems, but no specific methods for this risk analysis are discussed.

Vriezevolk et al. (2011) have proposed a risk analysis method for a specific type of distributed systems, namely telecommunication systems. They also explicitly discuss a list of requirements for such a method and their method is also based on step-wise refinement of the system model.

Concerning IT dependability management in the public sector, Santos et al. (2008) have previously investigated the relation between IT technology and cooperation in emergency response organisations and Zimmerman et al. (2006) discussed the interconnections between information technology and other critical infrastructure for emergency response. In Sweden, the Swedish Civil Contingencies Agency published Basic Level for IT Security (Swedish Emergency Management Agency, 2003), a collection of guidelines concerning IT safety for governmental organisations based on international standards.

BACKGROUND

Municipal Risk analysis

An earlier paper (Weyns et al., 2012) identified a number of important requirements for a method to be suited for analysing of critical systems at Swedish municipalities, an analysis they are required to conduct by law. These requirements pose restrictions on both the input and output data of the method, as well as on the system model used in the method and on the amount of effort required for the whole analysis.

Reliability Block Diagrams

The method described in this paper is based on a variation of Reliability Block Diagrams (Staley and Sutcliffe, 1974). A Reliability Block Diagram (RBD) models a system as a collection of blocks representing the system's logical or physical components as shown in Figure 2. The availability of the system is modelled as a path from left to right. The most common configurations are blocks in series (all subsystems must be available) and parallel (at least one subsystem must be available), but more complex relations (such as k-out-of-N) can be modelled. Given the availability functions of the individual blocks (often modelled with the help of exponential or Weibull distributions), the total availability of the system can be calculated directly or through Monte-Carlo simulation. Many commercial software packages are available to assist in the creation and analysis of RBDs.

PRACTICAL RISK ANALYSIS METHOD

The practical risk analysis method presented in this paper consist of 6 steps presented in the following sections:

1. Scoping
2. Definition of failure categories
3. Construction of the system model
4. Estimation of the failure frequency of each component in each category
5. Analysis
6. Presentation of the result

Scoping

The first step in the risk analysis is to clearly define the limits of the system being analysed. This includes defining what is considered a failure of the system for which the availability is being analysed. For example,

there is a difference in calculating the availability of a distributed system from the view of one specific computer, of one specific user with access to multiple ways to connect to the system or for all the users.

Definition of Failure Categories

The second step is to define categories for the duration of the unavailability. The categories are meant to represent both different severities of failures for the users depending on the system, and different efforts to repair the system. For an example, see below in the case study. The addition of the categories of failure is exactly what makes it easier to estimate the necessary input data for the analysis and to later interpret the results of the analysis (see below).

Construction of the System Model

Step three in the risk analysis is to construct a Reliability Block Model for the system. This can be done iteratively, by replacing subsystems by a more detailed RBD representing this subsystem. The refinement of each subsystem should be done by experts of the respective subsystems. The detailed RBD of the subsystems that are present in multiple systems can later be reused for constructing the RBD of other systems.

The iterative refinement of the model can be repeated until a level is reached where the reliability of each component can easily be described by failure frequencies defined in the categories defined in step 2. As will be illustrated by the case study below, the use of these categories makes that many types of failures can be modelled at high level components, keeping the system model simple and reducing the effort required for the analysis. The possibility to still capture complex failure behaviour without having to model low-level components is one of the strongest advantages of the proposed method.

The components modelled in the RBD do not have to be easily identifiable physical components. The most important in the model is that the failures of the components have to be independent. If two components could fail for the same cause, a common component in series with the two components needs to be added to the model to capture this failure. For example if two components are dependent on the same power supply, the power supply should be a separate component in the model in series with the two former components. A power failure is then no longer included in the estimated reliability of these two components but instead attributed to the new power supply component.

Estimation of the failure frequency of each component in each category

The main disadvantage of traditional RBDs is that they usually employ complicated models for the reliability of each of the components. This requires good estimates for the parameters of the distribution of the mean-time-to-failure (MTTF) and the mean-time-to-repair (MTTR) of each single component. Also the output of the analysis is then formulated as a distribution for the MTTF and MTTR of the complete system. To make it easier to estimate the reliability of the components and to make the results easier to interpret in the context of the organisation's business continuity management, this paper proposes to replace the input distributions for MTTF and MTTR of each component by the frequency of the occurrence of a limited number of failure categories, based on the duration of the unavailability of the component.

These parameters are easier to estimate for professionals involved in the risk analysis. These categories also correspond to the process of IT service at the organisation, where some problems can quickly be fixed by restarting or replacing the failed component by IT personnel of the organisation, while other types of component failures always take longer to repair because they require technical support of outside suppliers. With each of the estimates it is important to clearly document the underlying assumptions and the source of the estimates, together with a measure of the uncertainty on each number.

This step implicitly requires an identification of the risks and threats to the availability of each component. This risk identification should preferably involve brainstorming by experts on each of the components combined with study of the available documentation for the components, including public reports of failures that have been reported in the media or in official reports.

Analysis

The calculation of the overall availability of the system can be simplified significantly compared to the general reliability block diagrams. Although Monte-Carlo simulation can still be used for the calculation, the overall reliability of the system can now also be calculated with a few simple rules (Staley and Sutcliffe, 1974).

First, for each component the unavailability (expressed as a percentage of time) in each category can be calculated by multiplying the failure frequency with the average failure duration. The availability is then equal to 100% minus the unavailability.

For the combination of two components coupled in series (both have to be available for the system to be available) the combined availability is simply the product of the availability of both components.

For the combination of two components coupled in parallel (at least one has to be available for the system to be available) the combined unavailability is the product of the unavailability of both components.

The total availability of the system or of a component across the different failure categories can also be calculated by multiplying the availability in the different categories. This assumes that the failure rates in the different categories are completely independent. This is, in practice, not always the case (e.g. a power failure might still be equally likely to occur while the system is being restored from crash caused by a software bug, but a software bug is probably less likely to be triggered during a power failure). However, because the availability of most components will be very high the error introduced by this approximation will be much smaller than the uncertainty on the result attributed to the uncertainty of the estimated failure frequencies.

With these simple formulas, the total availability of the system can be calculated, together with the number of expected failures in each failure category and the contribution of each component to these results.

Analysis of RBDs also makes it possible to analyse the sensitivity of changes to the reliability of one of the components to the reliability of the complete system. This model also makes it easy to calculate the effect of improvements to the structure of the system (for example by adding extra parallel components to those components responsible for most of the failures) on the overall availability of the system.

The failure behaviour modelled by this system model is equivalent to a traditional reliability block diagram where each component is replaced by a set of components in series, one for each category of failures. In this RBD each of these subcomponents is then assigned a corresponding exponential failure time distribution and a constant repair time.

Presentation of the result

The final step of the risk analysis is to present the results. In this step it is important to adapt the presentation of the results to the different target groups. The IT personnel responsible for the maintenance of the systems is probably most interested in the technical details of the analysis, while managers who have to decide about the budget for the system are most interested in the cost-benefit analysis of possible improvements to the system. Safety managers might be especially interested in a clear overview of the source of the estimates of the reliability of the components, to increase the level of trust that can be placed on the results.

An advantage of the use of failure categories is that the output of the risk analysis will also be described with the help of these outage duration categories, which naturally correspond to the type of outages the organisation needs to prepare for in their emergency management. For many types of systems in a governmental organisation, short outages will cause only nominal disruption, while longer outages will cause considerable problems for the organisation and the services it provides. Therefore, these concrete definitions of different outage categories will make it easier for the results to be included in the overall risk management process than when the availability is expressed as a complex failure distribution.

The results of the analysis can best be presented graphically in pie diagrams showing the contribution of different components to the unavailability in the different failure categories.

CASE STUDY

Case description

The risk analysis method described in the previous section was evaluated in a case study in the city of Helsingborg, where the availability was analysed of a large patient data management system used in elderly

care. This system was chosen because it was identified as the municipality's most critical system while at the same time it was a very mature system that had been in use for over 10 years (although in different configurations) and experts of the system were available to provide the necessary input data for the risk analysis.

Helsingborg is relatively large Swedish city, with about 130,000 inhabitants at the end of 2011. Elderly care in Sweden is a responsibility of the municipal governments. The municipality of Helsingborg operates 20 elderly homes, and provides home care to a few thousand elderly people. The system studied in this case study is used for all planning and record keeping of all medical and non-medical elderly care provided by the city. The system is critical for the organisation as any longer unavailability of the system will disrupt the information flow concerning elderly care, which might lead to patients receiving incorrect care because care givers don't have access to their most recent medical history. The system has about 1000 active users, most of them nursing staff and has been in use since 2000.

The system is a large commercial system with some specific adaptations for this specific installation. The daily maintenance of the system is done by an internal IT organisation, administrating the IT systems for the whole city government.

Risk Analysis Results

For the risk analysis we strictly followed the six steps described above.

Scoping

After discussions with the system administrators from the elderly care department, the scope of this risk analysis was set to only consider unplanned unavailability of the system affecting many users at different locations. The system is deployed at many locations and in case the system is unavailable at one of them the users there can easily access their data from other computers at other locations without any critical consequences for the organisation or its customers. This would only cause a minor inconvenience compared a large scale outage of the system. Therefore it was decided the results for the availability will explicitly not include:

- the unavailability of the system from a single PC (for example caused by hardware failure in this PC),
- the unavailability of the system to a single user (for example authentication problems for one user),
- the unavailability of the system at one specific elderly care facility (for example caused by power failure at this facility),
- the unavailability of parts of the system that are not used in elderly care and
- the unavailability of the system because of planned maintenance of the system, which is done on a regular basis at clearly specified times to cause minimal inconvenience to the users.

Definition of failure categories

In this study we used the following four categories of failures:

- Failures of very short duration lasting under one hour (with an average duration of 30 minutes)
- Failures of short duration lasting between 1 and 5 hours (with an average duration of 3 hours)
- Failures of long duration lasting between 5 and 24 hours (with an average duration of 14.5 hours)
- Failures of very long duration lasting longer than 24 hours (with an average duration of 3 days)

These categories are chosen both because of their significance for the users and for the IT-staff maintaining the system. For the organisation depending on the system, very short and short failures present no more than a minor inconvenience, as the system is not used for emergency care. The most important information is printed out on paper on a daily basis to decrease the dependence on the system. After a short service interruption the same staff who administered the care to the patients will still be able to enter the information into the system, only with a slight delay, possibly causing some limited overtime work for some staff. Longer failures will be more disruptive as they will impact personnel working in multiple shifts, and the failure of the system disturbs the information transfer between consecutive caregivers for the patients. This information transfer will then have to be done manually, increasing the likelihood that necessary information is lost or delayed. The organisation has a strong safety focus and manual routines have been prepared for this type of situations. Very long failures are much more disruptive as they necessitate information to be shared on paper between multiple caregivers over a longer period of time. During the service interruption the printed information will become more and more outdated and when the system has been restored a large effort will be required to enter all the manually recorded

information into the system again. In these situations the risk is the greatest that critical information, such as for example information on previously administered insulin doses for diabetes patients, is unavailable.

From the point of view of the IT-personnel, very short failure represent those failures that can be solved by a restart of the system or one of the components of the system. The IT department provides round-the-clock support for the system. Short failures are those failures that can be solved by the organisation's own personnel by replacing failed components such as routers or hard drives or by moving parts of the system to a different server. The servers of the system are in one server room, but backup servers are available in cold standby at second server room located in another part of the city. Long failures are those for which the organisation's internal IT support lacks the skills or the replacement parts to restore the system. For most of the components of the system service level agreements with the suppliers guarantee a maximum response time. Very long failures are these failures for which the cause lies outside the organisation and which will affect more than just this one system and even more than the organisation, such as a large power blackout affecting a large part of the city or a virus infecting the operating system on one of the systems components.

Construction of the system model

The main structure of the system is based on a client-server architecture. In a few iterations a sufficiently detailed model of the system, can be derived from the basic client-server system model, as seen in Figure 1. The components of the system and their main failure modes are described in Table 1.

Table 1. Description of the main components of the studied system and their failure modes

Component	Description and failure modes
Client computer hardware	Although the hardware in each singly computer is not very reliable, many hundreds of available personal computers are located at many different elderly care facilities around the city. A simultaneous failure in the hardware of sufficiently many of these PCs to disturb the necessary access to the system is extremely unlikely and therefore the reliability of this component is so close to 100% that we can ignore this component in the rest of the analysis.
Operating System (OS)	All PCs run on the same commercial operating system and are therefore vulnerable to a virus or a previously unknown bug affecting this operating system.
Client software	The client software is a commercial product with custom adaptations for the organisation. Software faults in the client could negatively affect the availability of the system.
Citrix client	In this component we model failures in the virtualized workspace client installed on each computer.
Power supply	Power blackouts that affect a large part of the city or more would also limit access to the system as most locations don't have backup power.
Helsingborg's City Data Network (HSDN)	The city has its own internal network. Network failure or overload could disturb the communication between the clients and the server. The network architecture has sufficient redundancy to prevent failures in single routers or network cables from affecting traffic in the entire network. Because we are not modeling the availability of the system at specific locations, more detailed components of the network do not need to be modeled in detail. The most critical parts of the network are the connections to the server rooms. These are included in the description of the server rooms below
Application Server	This includes the specific server software of the system as well as the server cluster's hardware and operating system.
Database Cluster (DB)	The data in the system is stored in a separate database server cluster, which could experience both hardware and software failures.
Authentication Server	Because of the sensitive medical information in the system, the system uses a separate authentication server. Also this server can be affected by both hardware and software failures
Citrix Server	Also for the workspace virtualisation server hardware and software failures need to be considered.
Server rooms	All servers are located in one of two server rooms. The server rooms are in geographically separate locations and are located in buildings with a high level of physical security. For all servers, backup machines are available in cold standby. Because of this setup, any physical damage (e.g. because of fire or flooding) in one of the server rooms will cause a short failure in the systems located in this server room. For this reason the two server rooms cannot be considered as true parallel components in the reliability block diagram. Both server rooms have access to multiple connections to the network and to the power supply and also have access to backup power generators. The probability of a failure affecting both server rooms simultaneously is sufficiently small to be negligible in comparison to the other failure probabilities considered.

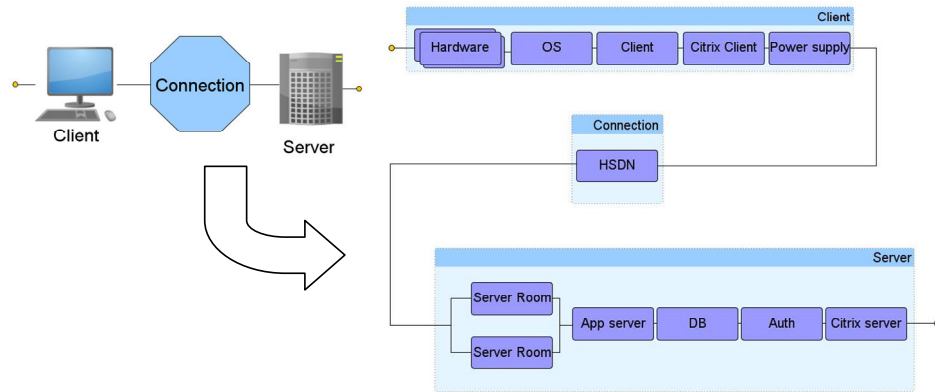


Figure 1. First and last iteration of the system model used in the analysis

Estimation of the failure frequency of each component in each category

To estimate the failure frequencies of the different components, we used a combination of historical data, service guarantees in service level agreements with the suppliers and expert opinion. For some failures like power outages and some components that have been in use for a longer time, historic incident data was available, but for most components, experts with long experience working at the city IT department provided their best estimates for the failure rates in each category. The results are shown in Table 2. It can be seen that about half the estimates are 0, this is because some of the components only experience short term failures (because they are easily replaced) and some others experience only long term outages (because they always require a large effort to restore the system). This reduces the complexity of the analysis later.

To model all these failure modes with a traditional reliability block diagram with reliability distributions for each component, each of the components in our system model would have to be broken down in many subcomponents

Analysis

Table 2 also contains the results of the analysis. The total availability is estimated to an impressive 99.9%. This corresponds well to the experienced availability during the last years, where the system has generally been very reliable. The total downtime of the system (not including planned system updates) is estimated to only around 8 hours per year. From the data we can also see that very short term failures are expected once every other year, short failures once every eight months on average. Long term failures are estimated to occur only once in a 10-year period, and very long term failures once in a fifty year period on average, which is much longer than the expected life-time of the system. The availability of the system is thereby found to be more than adequate with relation to how critical the system is.

		< 1 hour		1 < x < 5 hours		5 < x < 24 hours		> 24 hours		Total	Downtime (min)
Client	PC	0	100,000%	0	100,000%	0	100,000%	0	100,000%	100,000%	-
	OS	0	100,000%	0	100,000%	0,02	99,997%	0,01	99,992%	99,988%	61
	Client	0	100,000%	0	100,000%	0,01	99,998%	0	100,000%	99,998%	9
	Citrix	0	100,000%	0	100,000%	0,01	99,998%	0	100,000%	99,998%	9
	Power	0,05	100,000%	0,1	99,997%	0,02	99,997%	0,01	99,992%	99,985%	80
Connection	HSDN	0	100,000%	0,05	99,998%	0	100,000%	0	100,000%	99,998%	9
Server	App Server	0	100,000%	0,5	99,983%	0,01	99,998%	0	100,000%	99,981%	99
	Database	0	100,000%	0,5	99,983%	0,01	99,998%	0	100,000%	99,981%	99
	Auth	0	100,000%	0,5	99,983%	0,01	99,998%	0	100,000%	99,981%	99
	Citrix Server	0,5	99,997%	0,01	100,000%	0	100,000%	0	100,000%	99,997%	17
	Server room	0,01	100,000%	0,01	100,000%	0	100,000%	0	100,000%	100,000%	2
System		0,56	99,997%	1,67	99,943%	0,09	99,985%	0,02	99,984%	99,908%	482

Table 2. Estimated failure frequencies (occurrences / year) and calculated reliability for all components and categories. Cells containing zeroes are filled in black and represent failures categories that are considered extremely unlikely to occur (< 1 per 1000 years). The final column contains the expected downtime of the components (min / year).

The uncertainty on each of the estimates is quite large, and especially on the smaller values in the table the true values might be off by a factor 2. Although the final result of the analysis is less sensitive to the uncertainty in the smaller values, the uncertainty should be taken into account when interpreting the results. However, the estimates are still the best estimates available, and the final result still gives a good indication of the order of magnitude of the unavailability of the system and the main components that are the cause of the most critical system failures.

Presentation of the result

The results of the analysis can be summarized in Figure 2. They offer a clear graphical overview of which components are responsible for most of the downtime of the system in each category. Similar graphs can easily be produced to show the breakdown of the total expected downtime by component or by failure category. The analysis shows that only a third of the expected 8 hours per year downtime can be attributed to critical failures (long or very long). The very long term failures are only expected from power blackouts or from failures affecting the operating system on the client computers, which are expected to be about equally unlikely. The long failures can be attributed to a large set of components, although the same two components still dominate here. This means that in case the availability concerning critical failures of the system needs to be improved, the main options are to invest in backup power at the different elderly care locations or to migrate to a more reliable operating system for the client computers. With this model, the improvements in availability from these measures can be easily calculated.




Figure 2. Calculated downtime in each failure category by source component (minutes per year)

Evaluation

To evaluate the usefulness of this method, the results were presented to representatives of the city's IT department, elderly care department and safety department at a joint meeting. The meeting started with a detailed walkthrough of the method and the results of the risk analysis. Next, the authors of this paper lead a structured discussion of both the conducted risk analysis and the method in general. In this discussion we focussed especially on the usefulness, the usability and the value of the method and its results. Also a number of advantages and disadvantages of this approach were identified and discussed. The following paragraphs present a summary of the minutes of this meeting.

First of all, this case study showed that this method can be applied to this kind of system and with only limited effort provide a quantitative overview of the main threats to the availability of the system. This type of overview of the main components of the system and the main measures being taken to ensure the systems reliability was something that was not available before this case study. This method provides a means for the IT department to communicate this type of information to the rest of the city government, which is something they had been

actively seeking before. It improves the understanding between the different departments and provides valuable input for further risk and vulnerability analysis at the elderly care department.

Secondly, one of the main advantages of the method is that the method is relatively easy to use. The modelling of the system is the most complex task requiring a detailed knowledge of the complete system and its components. However, because of the use of the categories, many different types of failures can be modelled on a very high level of abstraction. The method by itself does not guarantee that the risk analysis is complete and that all possible sources of failures have been identified. To be able to trust the results of the analysis, it is very important that the construction of the system model is done systematically and involves all available experts on the system.

It is also expected that parts of the system model can be reused when analysing other systems within the same organisation. Many of the organisations systems share some infrastructure such as the server rooms and the network. These components then don't need to be reanalysed in subsequent risk analyses.

The pie charts presenting the results are easy to understand and illustrate well which components present the largest threat to the system's availability. The participants in the evaluation meeting stressed the importance of adapting the presentation of the results to the different audiences for the results. For example, in case investments in the improvement to the system need to be approved by the political leadership of the city, a way has to be found to present the results without going into technical details of the system.

A clear disadvantage of the method is the uncertainty on the results, which is a direct result of the uncertainty on the estimates. At the same time, these estimates are the best available estimates and even if only the order of magnitude of the result can be relied upon, not the exact numbers, it is still better than having no estimate at all.

Except for the uncertainty on the estimates, there is also a risk of bias on the estimates of the reliability of the components as they are provided by experts who are often part of the maintenance personnel or the supplier of the system and therefore have an interest in the estimated reliability of the whole system. This is an issue often present in risk analysis and, whenever possible, data from multiple sources should be used. To deal with both uncertainty and bias it is important to clearly mark the source of all data to improve the traceability of the conclusions. The value of the result for the target audience is directly dependant on the trust they have in the data sources used in the analysis.

Another disadvantage of this method is that it does not take into account the unavailability associated with planned system updates. They cannot easily be included in the model in the same way as other sources of unavailability because their occurrence doesn't follow a random distribution. If the duration of these service updates could be carefully planned, their part in the unavailability could easily be added to the model, but experience has shown that these updates regularly take longer time than planned and cause longer than expected unavailability of the system. Further study is needed to examine how this can be accounted for in the model.

Concerning external validity, the system used during this case study was considered a mature system that had been in use for a long time and was well documented. The system was experienced as very reliable by the users and system administrators. Further research is necessary to evaluate how applicable the method is to systems that are newly installed, for which less documentation is available and for which the reliability of its components is unclear. Also, more study is necessary to determine how useful this method is to different types of organisations.

CONCLUSION AND FUTURE WORK

In this paper we have described the details of a risk analysis method to analyse the availability of large-IT systems. The method has been applied in a case study showing the practical applicability of the method. The method is based on reliability block diagrams, extended with failure categories. The addition of the categories makes it easier to collect or estimate the necessary input data for the model, simplifies the analysis and makes the result easier to incorporate in organisational risk analysis. The categories also make it possible to model many different types of failures based on a relatively high-level system model. These changes make it possible to also apply reliability block diagrams to much larger distributed systems instead of just relatively small embedded systems.

The main advantages of the method are that it is sufficiently simple to be applied with limited resources and that part of the analysis of one system can be reused when analysing similar systems. The method was positively evaluated in a case study as an important improvement in the communication between then IT personnel and the rest of the organisation on the availability of critical systems.

Future work will focus on evaluating the method on different types of systems and on finding a way to incorporate the risk associated with system unavailability from planned updates of the system into the method. Further work is also necessary to evaluate this method against other risk analysis methods for example through controlled experiments.

ACKNOWLEDGEMENT

This research was conducted within the PRIVAD (Programme for Risk and Vulnerability Analysis Development) programme financed by the Swedish Civil Contingencies Agency. The authors would also like to thank the personnel of the city of Helsingborg for their participation and contribution to this study.

REFERENCES

1. M. R. Beauregard. (1996). *The Basics of FMEA*. Productivity Press
2. S. Dharmaraja, V. Jindal and U. Varshney. (2008). Reliability and Survivability Analysis for UMTS Networks: An Analytical Approach. In *IEEE Transactions on Network and Service Management*. Volume 5, Issue 3, pp. 132 – 142, 2008.
3. C. A. Ericson. (1999). Fault Tree analysis – A History. In *Proceedings of The 17th International System Safety Conference*.
4. Office of Government Commerce. (2007) *Information Technology Infrastructure Library*, Version 3.
5. M. K. Rahmat, and S. Jovanovic. (2011). Reliability Estimation of Uninterruptible Power Supply Using Reliability Block Diagram Method. In *International Review Of Electrical Engineering*. Volume 6, Issue 3, pp. 1109 – 1117, May 2011.
6. P. Reason and H. Bradbury-Huang. (2007). *The SAGE Handbook of Action Research: Participative Inquiry and Practice*, 2nd ed. Sage Publications Ltd
7. R.S. Santos, M.R.S. Borges, J.O. Gomes and J.H. Canós. (2008) Maturity levels of information technologies in emergency response organizations. In *Collaboration Researchers' International Workshop on Groupware (CRIWG). Volume 5411 of LNCS., Springer, Heidelberg*
8. J.E. Staley and P.S. Sutcliffe. (1974). Reliability block diagram analysis. *Microelectronics Reliability*. Volume 13, Issue 1, February 1974, Pages 33–47
9. Swedish Emergency Management Agency, SEMA. (2003) Basic Level for IT Security. *SEMA recommends 2003:2*.
10. U.S. Department of Defense. (2007). *Electronic Reliability Design Handbook. Department of Defense Handbook. MIL-HDBK-338*
11. E. Vriezokolk, R. Wieringa and S. Etalle (2011). A New Method to Assess Telecom Service Availability Risks. In *Proceedings of the 8th International Information Systems for Crisis Response and Management Conference, ISCRAM 2011*
12. K. Weyns and M. Höst. (2009). Dependability of IT Systems in Municipal Emergency Management. In *Proceedings of the 6th International Information Systems for Crisis Response and Management Conference, ISCRAM 2009*
13. K. Weyns, M. Höst, and Y. Li Helgesson. (2010). A Maturity Model for IT Dependability in Emergency Management. In *Proceedings of the 11th International Conference on Product-Focused Software Process Improvement, PROFES 2010*.
14. K. Weyns and M. Höst. (2012). Risk Analysis for Critical Systems with Reliability Block Diagrams. In *Proceedings of the 9th International Information Systems for Crisis Response and Management Conference, ISCRAM 2012*
15. R. Zimmerman and C. Restrepo. (2006) Information technology (IT) and critical infrastructure interdependencies for emergency response. In *Proceedings of the 3rd International Information Systems for Crisis Response and Management Conference, ISCRAM 2006*