A dependability case approach to the assessment of IP networks

Ilkka Norros and Pirkko Kuusela
VTT, Technical Research Center of Finland
P.O.Box 1000
FI-020044 VTT, Finland

Pekka Savola
CSC – Scientific Computing Ltd.
P.O. Box 405
FI-02101 Espoo, Finland

***Abstract*—IP networks, composing the Internet, form a central part of the information infrastructure of the modern society. Integrated approaches to the assessment of their dependability are, however, only emerging. This paper presents three contributions for meeting this challenge. First, we propose an adaptation of the safety case methodology, a 'dependability case' approach, as a practical form of organizing heterogeneous information concerning the dependability of a large communication network. The idea is to build structured argumentation for the support of dependability claims, making use of various kinds of evidences. Second, we suggest a conceptual framework for considering the dependability of IP networks in an integrated way. Third, we propose to structure dependability cases of IP networks according to the main aspects of dependability, rather than structural units or layers of the network. The proposed methodology is tested on the Finnish University Network and found promising.**

I. Introduction

It is commonly believed that the three main sectors of electronic communication – voice, television, and data transfer – are being unified. They converge to a infrastructure, where the key role is played by one generic service, the global delivery system of Internet Protocol (IP) packets. However, Internet and IP were not designed for a mission-critical global infrastructure, rather they have grown to such due to tremendous flexibility and modifiability in the design.

When a large part of existing services and functions are moving to the Internet, an extensive loss of IP connectivity would in fact paralyse the society. It is more than natural to ask questions about the dependability of IP networks, ie., questions about their availability, reliability, controllability, vulnerability, security, etc. Can one rely on this new infrastructure as much as on the traditional technologies and the old ways of acting?

The question on overall dependability of IP networks as a whole is rather new, although robustness etc. have been important aspects in improving existing algorithms and protocols. We contribute to the task by proposing a methodology for addressing the overall dependability of IP networks. We borrow ideas from safety cases, which have gained support as efficient means of assessing safety, and aim at a tool that can be used not only to evaluate dependability but also to take care of it. The key emphasis in the methodology is to be able to handle complicated structures and dependencies in an understandable and traceable way. The extension of the notion of safety case into that of a dependability case was done by Despotou and Kelly [1]. Our work seems to be its first, tentative application to networking. It was created in the IPLU-II project (http://iplu.vtt.fi) and the case study part of it was first time internationally presented by the authors in the NORDUnet2008 Conference. Here the aim is to motivate and present the methodology and its application to an IP network in a general setting.

This paper is organized as follows: First the dependability case methodology is introduced and motivated in Section II. In Section III we analyse the aspects and actors of dependability of IP networks. A proposal and an example on how to actually implement a dependability case of a network is presented in Section IV, and an illustration of argumentation is given in Section V. Discussion and conclusions can be found in Section VI.

II. Dependability case

In safety critical industry, e.g. nuclear power plants, the safety of a system must usually be assessed and accepted by a governmental regulator. It is far from obvious how such assessments would be most effectively done for large and complex systems. One way of solving this task is a goal oriented approach called safety case [2], [3]. It has gained considerable support as an efficient means of identifying and addressing safety concerns at each step of the system's lifecycle. We propose the adoption of the basic ideas of safety case to the assessment of the dependability of IP networks, and call this transformed methodology "dependability case". The key difference between a safety case and a dependability case is the increased heterogeneity of aspects that must be included in addressing dependability [1].

By adjusting a definition of the safety case, we may define the dependability case of a network as

> a documented body of evidence that provides a convincing and valid argument that the network is adequately dependable, taking all aspects of dependability into account, for a given application in a given environment.

The motivation of dependability case is to develop methodology for addressing dependability of a complex system. The idea is to gather dependability-related information into one document (or document structure) that is usable later on to demonstrate the dependability of the system. Key features

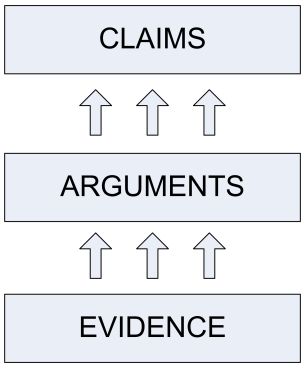

Fig. 1: Basic structure of a dependability case.

needed would be the ability to indicate complicated structures and dependencies in an understandable manner. Putting the separate details and various observations in their own context, while keeping the whole organized, has value as well. The effort to explicitly record evidence and argumentation helps in traceability.

As dependability case is a tool for assessing dependability it is efficient in dependability communication between operator and regulating authority or operator and business customer. In addition, when incorporated in the design early on, it serves the operators' needs in evaluating, developing and managing dependability. The structure of dependability case allows the evaluation procedure to be repeated and results to be updated in a manner where the development in time can be easily seen.

Structure of a dependability case

The key elements of a dependability case, illustrated in Figure 1, are

- the stated goals, or *claims* about different aspects of dependability; the claims are usually subdivided into a hierarchy of subclaims;
- the available *evidence*, and
- explicitly formulated *arguments*, which provide support to the claims on the basis of evidence.

The claims are statements about properties of a system or a subsystem. They may originate from the regulatory, user, design or operation requirements, and typically they have different viewpoints such as goal, vulnerability or standard/requirement orientation.

The evidence consists of any kind of information relevant to the assessment of the network's dependability. This can be general information about the network's structure, "hard" monitoring or measurement data collected from the network, or "soft" information gathered by interviewing network operating personnel, providing both facts and opinions. Evidence could also be a result from tests specially performed for the assessment. Further, if the dependability case has several levels with nested structure, evidence can also be a sub-claim in a lower hierarchy level. Supporting a claim by using several and preferably independent evidence increases robustness to tolerate flaws in a single argument. In any case, it is crucial that the evidence is explicitly registered and available for study.

The heart of the dependability case is the argumentation. Arguments give the meaning of the data (evidence) in the context of claims and specific targets of dependability evaluation. The arguments, providing a link to the claims, can be deterministic, probabilistic or qualitative in character. Deterministic arguments can be formal proofs or demonstrations of the fulfillment of dependability requirements. Probabilistic arguments involve statistical reasoning to establish a numerical value of some statistical dependability level of the system, such as the availability of the system. Qualitative arguments are interpreted as compliance with rules that have an indirect link to the desired attributes of the system, like utilizing good maintenance practices etc.

The claim – argument – evidence structure allows the dependability case to be visualized in a logical manner. As with safety case, a dependability case is not only a vehicle for assessment and approval, but also for *taking care* of dependability. The operator of the system can maintain awareness of the dependability status and judge where to allocate efforts for its improvement.

In UK the regulators of safety critical industries and infrastructures demand safety case studies to be conducted [1]. Will the regulators of communication infrastructures require assessments on dependability in the future?

III. Dependability of IP networks: a conceptual analysis

Dependability can be defined as a system's ability to avoid service failures that are more frequent or more severe than what is acceptable [4]. After such a general orientation, dependability is usually defined as a collective concept that combines several aspects or attributes. The choice of the most relevant aspects varies, however, according to the type of the system. We present here a conceptual analysis of the dependability of IP networks that yields a selection of six main aspects: *robustness* (of basic protocols), *invulnerability* (in the sense of a continuous struggle against vulnerabilities), *controllability* (in the sense of the choice and implementation of various control mechanisms), *maintainability*, *reliability* and *availability*. These aspects are (i) necessary for understanding the various origins of failures in IP networks and the strategies for their avoidance, (ii) more or less sufficient for the same purposes, and (iii) rather orthogonal, as well logically and also by referring mostly to different relevant actors and to different types of activity as regards improvements.

Before discussing the six aspects in detail, we distinguish four kinds of generic actors with different viewpoints to IP networks. All together are collected in Figure 2.

A. Generic actors

User: The generic user represents both individual and corporate users. When the IP convergence proceeds, the availability of IP connectivity, with certain quality, becomes the central service demand that the user directs to the network. The actual services run end-to-end over the network.

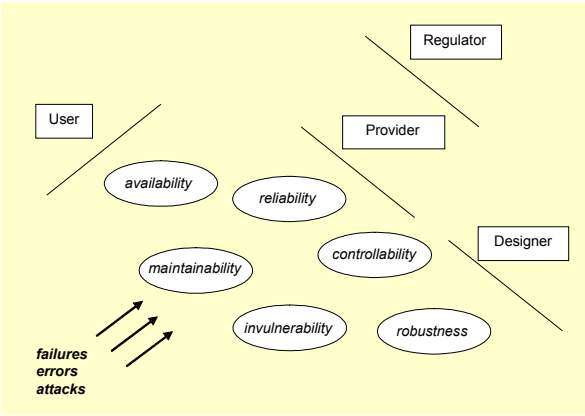

Fig. 2: Main aspects of dependability of IP networks (ovals) and related actors (rectangles).

Provider: The generic provider includes both network operators and network equipment manufacturers. The provider is the actor who is directly responsible for the network and its dependability. In the present de-regulated networking business, the generic provider splits into a mosaic of specialized companies, all driven by the imperative of profitability.

Designer: By the generic designer we denote all those instances which have created the IP architecture and protocols and develop them further. Thus, the designers include scientists, engineers, standardization organizations among others. Network equipment manufacturers are strong players also in the role of designer.

Regulator: The generic regulator represents the society's general will regarding the IP network infrastructure. Regulators are regulatory agencies and legislative bodies. The IP networking technology and business develop so fast that standardized certification actions are not common. Although IP-specific regulations are emerging, most of the existing regulations still concern telephone networks only.

B. Main aspects of the dependability of IP networks

We now discuss the six above-mentioned aspects in a unified format. The first three aspects — robustness, invulnerability and controllability — concern general design features and choices. The logic of our presentation of them is a kind of 'Hegelian negations' of the vision/dream on an ideally robust packet network. The other three — maintainability, reliability and availability — are standard aspects of dependability of any technical system, but we need to point out their concrete content in the context of IP networking.

1) Robustness of basic protocols: The basic service offered by an IP network is to deliver a packet from any interface A to any other interface B. However, it is seldom sufficient to consider an IP network as a closed system — rather, they should be considered as parts of the global Internet.

The successful functioning of the Internet as a global and highly open information infrastructure relies on the robust performance of its basic protocols, in particular (i) routing of the datagrams (BGP for inter-domain routing, a few alternatives for intra-domain routing) and (ii) traffic congestion avoidance (so far mainly TCP, but in near future this picture expected to become more complicated).

Definition. By the robustness of protocols we mean their ability to maintain the fundamental functionality of the system despite of unpredictable variability of inputs (traffic), and despite of frequent changes of the system's detailed structure (both intentional and unintentional, like component failures).

Metrics and criteria. There should not be heavy oscillations in routing (so called route flapping, sometimes observed by BGP), nor in traffic goodput, and packet losses should remain moderate even by heavy overload. The goal should be that the algorithms used in the basic protocols can be mathematically proven to behave in a robust way.

Related activities. The provider (operator) of an IP network can and must to some extent rely on the robustness of the basic protocols, so they can be considered as given when the dependability of a particular network is assessed. However, even then it is important that the network provider understands their functioning sufficiently deeply. Moreover, protocols develop and are sometimes changed to better alternatives. This happens at a slow timescale and the generic actor responsible for it is the Designer. This evolutionary nature is fundamentally different from a system like a nuclear power plant, where the basic design is made once and for all.

2) Struggle for invulnerability: However robust the basic protocols are, no protocol design can be absolutely robust. Rather, there are always vulnerabilities against unexpected inputs (both 'on the scene' like targeted DoS attacks, and 'behind the scene' like errors in routing configuration), and struggling with vulnerabilities is a continuing activity in IP network provisioning.

Definition. Vulnerabilities of a system are possibilities of behaviour that lead to serious degradation of performance by causes that come from outside of the system aspects taken into account in the robust basic design.

Metrics and criteria. Statistics on detected vulnerabilities, intrusions (attempted and successful), configuration errors etc. should be collected.

Related activities. The struggle against vulnerabilities happens on several fronts. Whereas the Designer studies and proposes improvements to the inherent robustness of the system (say, to the problem that forging an origin address is too easy), the software providers are busy with identifying and correcting errors and security holes, and the operating personnel develops good work practices to avoid configuration errors. The Regulator helps to coordinate the struggle on national and international levels, mainly as regards security problems. Any vulnerability can be considered at least as a potential security problem.

3) Controllability by network operator: Despite the robustness of the basic protocols and the openness it allows, an IP network provider should not be void of tools to exert control on the network and its traffic. However, whereas the basic

protocols are essentially identical in every network, it depends on choices of the operator what additional control mechanisms are implemented and how quickly they can react. This is why we like to raise the controllability to an independent aspect of an IP network.

Definition. By controllability of an IP network we indicate all implemented techniques that assure that the the network is under the hold of its operator. It reflects the possibilities of the network operator to accept or route traffic offered to the network and to open or close individual services or the whole network.

Metrics and criteria. Reaction times to serious anomalies like attack and overload, both observation and action.

Related activities. Monitoring capabilities are a necessary part of controllability, and their adequacy depends on what the operator has chosen to implement. Note that sheer collection of data does not make the monitoring effective without efficient techniques for making inference from the raw data. Control actions include traffic engineering techniques (e.g., traffic differentiation, MPLS routing) as well as traffic filtering (e.g., of malicious traffic).

4) Maintainability: By controllability we emphasized the choices the network operator has made when implementing various mechanisms. Once these have been specified, we turn to the maintainability of the network.

Definition. Maintainability means the system's ability to undergo modifications and repairs [4]. In communication networks it is mainly characterized by the possibilities to extend, renew, and update the network.

Metrics and criteria. Network load statistics provide the basis of dimensioning. For many other activities related to maintainability, metrics and formal criteria are difficult to set despite of quality control and assurance procedures. The kind of relevant evidence depends on the question — for example, a certificate of a training course can be used in an argument for a claim on the qualification of the operating staff.

Related activities. Proper network dimensioning is crucial for maintaining an IP network, as long as the traffic grows steadily as it has done so far. Hardware upgrading usually introduces new software into the system, and good safety practices are important by software installations and updates. Administrative and economic activities like subcontracting, making agreements on repair, and organization of personnel for network operation and various support functions are also crucial for maintainability.

5) Reliability: The question on reliability is posed for a well-defined system, in our case a network with specified constitution — not for the network seen from a wider perspective, where its structure is seen as evolving and various principal choices are made in its build-up. Thus, we are basically in the realm of classical reliability theory, except that the difficult area of software reliability cannot be avoided. According to this paradigm, a system is broken hierarchically into subsystems and finally into components that either work or not. Networking software can, however, usually not be localized to individual hardware components, nor does the performance of software have a clear on/off character according to the presence of errors. Reliability is built into networks through component redundancy and protection mechanisms. Thus, we subsume the automated resilience features, like path protection, under the reliability aspect, whereas part of the resilience of IP networks comes directly from the robustness of the basic protocols.

Definition. Reliability means the continuity of correct service [4].

Metrics and criteria. Reliability is measured by the probability that the system works, as a function of the probabilities of each component to work. As regards software reliability, formal checks of correctness should be aimed at. Estimates of the amount and kind of remaining errors in a piece of software provider an easier but weaker alternative. Assessments of the reliability of automatic protection mechanisms may require sophisticated analyses.

Related activities. Taking care of reliability includes collecting and maintaining component reliability statistics and certificates, analyses, and computing the optimal choices where to build redundancy.

6) Availability: At a particular moment, the only aspect of dependability that the user experiences is the availability of the service. Note, however, that the service provided by an IP network is not an on/off matter but has various qualities like throughput and delay.

Definition. Availability means readiness for correct service [4].

Metrics and criteria. The proportion of time during which the datagram transfer service is delivered, with some specified qualities.

Related activities. Availability should be measured and estimated. Explicit goals should be set.

IV. Building a dependability case of an IP network

We now turn to the following question: what is the most suitable architecture of a comprehensive dependability case of a given IP network? We approached this problem constructively by building an experimental dependability case of the Finnish University and Research Network (Funet, http://www.csc.fi/funet), which is the national core network connecting universities and research institutions. Our 'soft' evidence consisted of three interviews with personnel operating the network. The 'hard' evidence used in the case consisted of the network topology and its routing rules, ping-based core router downtime data and link traffic data, both from the period 2000–2007. Adelard's ASCE tool [5] was used for keeping the case together. It should be clearly noted that our exercise was not a real dependability assessment, but possibly points a way how to make one.

Having got some touch with the material, we decided to build the case according to the main aspects of dependability considered in the previous section. Thus, the top claim 'Funet is highly dependable' was split into subclaims of type 'the service provided by Funet has high availability' etc. (Note

that Funet has not set any explicit target value on availability.), which were in turn divided into 1-3 further layers of subclaims. Here is a more detailed account on some matters encountered in this exercise.

Robustness of basic protocols: This aspect is mainly a concern of the protocol Designer, and we did not formulate any particular claims here. Funet offers pure IP connectivity.

Invulnerability: We collected all security-related issues here. One could additionally identify here potential critical human errors, like misconfigurations of BGP, and procedures and tools for their avoidance. We, however, grouped them together with other operation practices under maintainability.

Controllability: Funet does not use MPLS, nor QoS differentiation, and mainly the traffic filtering capabilities to counter DoS attacks were counted as available control techniques. Monitoring capabilities could well be considered here, although we grouped them together with monitoring activities under maintainability.

Maintainability: The adequacy of network dimensioning was addressed here, using traffic data as evidence. Other items included monitoring the network, detection and correction of possible problems, software updating procedures as well as normal operation and maintenance practices. Respective arguments were formulated, referring to the interviews as evidence. A serious assessment would naturally need stronger evidence.

Reliability: We divided the consideration into components (routers, links, power supply), structure, and protocol software. No direct quantitative data on component reliability was available. Ping-based core router downtime data was used to get a rough idea of their failure probability. The approach of classical reliability theory was used for the hardware, whereas the software's reliability was assessed only by asking for the personnel's experiences. An example of structuring a part of a reliability argumentation is given in the next section.

Availability: The ping data provided information concerning the availability. The main lesson from this aspect was, however, that the monitoring data considered adequate for network management was not very well-suited for computing network availability estimates. In a well-made dependability case, the argumentation on availability needs to be based on hard evidence on component downtimes. However, the data collected from the network for monitoring and failure detection purposes may not, as such, be best suited for infering availability. Ideally, there should be dedicated monitoring and data-handling functions to measure the specific type of availability addressed in the case.

All together, the chosen structuring principle worked well in the Funet case. All our information concerning the network's dependability could be included and organized without forcing. Still more important, it was obvious how sharpening the argumentation by more evidence and more extensive analyses could happen at various points without changing the general structure. With the selected approach, similar types of issues are addressed in the same logical location, whereas qualitatively different types of issues are clearly separated from each other. Thus, specialists of different fields can contribute to a dependability case, each having a specified and well defined task in the case. Different viewpoints and argumentation results are summarized at higher and more abstract levels of the dependability case.

V. Illustration of argumentation in a dependability case

We illustrate the structure of the dependability case by sketching a part of an argumentation for evaluating reliability of the IP network structure. Visualization similar to one used in safety case tools is given in Figure 3. The approach is top-down: general claims are given first and then divided into sub-claims that are more context dependent. Evidence is indicated at the bottom. Argumentation is between the evidence and (sub)claims. Different types of arrows illustrate the possibility to indicate the degree or confidence in the assessment. Hence it is easy to see weak and strong parts of the dependability case or compare cases on a time-line.

Our highest level claim is that the *Network structure is reliable.* This is divided directy into sub-claims that there are *No single point of failure* and *No high risk 2 component failures*.

Main argumentation for failures is the cutset and risk analysis. Cutset analysis is based on the network topology and routing rules. The idea is to remove links and/or nodes, and then to check if network remains connected - using routing rules - and it has its targeted functionality (critical connections to outside still working). This is a purely combinatorial procedure that can be done unless the network is very large. When two or more components are removed, the network is likely to fail in functionality. The associated risk can be estimated, for example, as a product of the amount of lost traffic times the probability of the component failure pattern. The risk assessment requires estimates for failure probabilities and an estimate of the traffic matrix, computed from the link-level traffic data.

Supporting argumentation regarding failures is the analysis of the link traffic data to check that link loads are not so high that the traffic of a failed connection cannot be rerouted without causing overload.

The above illustrates mathematical argumentation incorporated into a dependability case. However, it may need to be supported by qualitative argumentation. Namely, a challenging feature of classical reliability analysis is the fact that layer 3 links that one would assume independent by a look at the corresponding topology graph may turn out to be dependent, because of sharing resources at layer 2. For this reason, we indicate the argumentation on the absence of hidden dependencies in components. This involves an examination of the network structure together with the system design and implementation.

VI. Discussion and conclusions

The Internet has become a critical infrastructure, and assessments of the dependability of IP networks will become

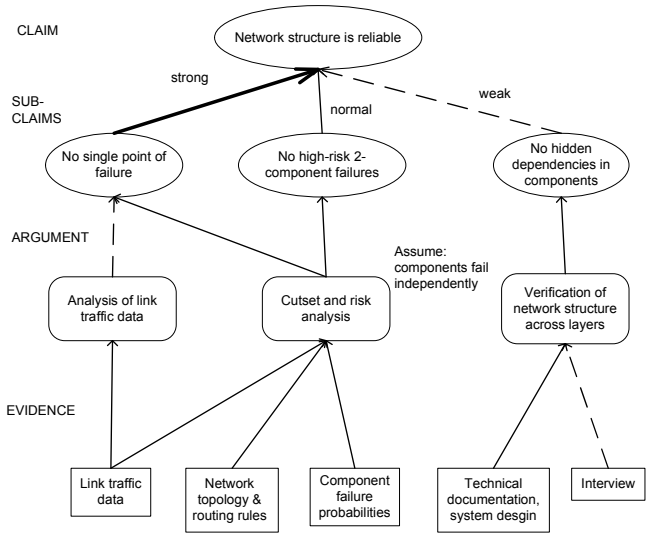

Fig. 3: Illustration of the claim-argument-evidence structure.

an indispensable part of taking care of it. Thus, means to address the overall dependability of IP networks are needed. This development work has to start today, as the tools will be demanded in the future.

We have proposed for this task a dependability case methodology following the general ideas of the safety case approach. Like the safety case methodology has become a practical tool in addressing safety issues, we hope that a dependability case can have a similar role in evaluating dependability and communicating dependability issues between various actors. In IP networking it would be fruitful both as regards Service Level Agreements (user–provider relation), as regards laws and regulations (provider – regulator relation), and as an internal tool of the provider alone.

We have found the dependability case methodology powerful in visualizing overall IP dependability, and it also indicates the structure and dependencies. The strength of this format is that it forces to make argumentation explicit, even when it happens to be weak, and technical arguments can be integrated smoothly into general arguments. Also, keeping the logic of the inquiry clear has high value in itself, together with the emphasis on indicating the evidence clearly. Moreover, the dependability case approach allows the structuring of research and development work in the network dependability area: it gives narrowly focused projects a reference frame and significance in the wider context.

There are some issues that distinquish a dependability case from a safety case. Despotou and Kelly [1] have addressed two important issues in a dependability case: the fact that the attributes of dependability are interrelated, and that they might even be in conflict with each other. They propose modularity and trade-off arguments to overcome these "horizontal" relations in the case.

An important issue that remains for further work is the scalability of the approach to very large networks. Safety cases have shown their power in large complex systems and methods to manage large cases have been developed [6]. Although a dependability case is richer than a safety case in terms of aspects to consider, we believe that, with a proper use of modularity and good architecture, dependability cases of larger systems are possible too. On the other hand, this methodology is in the development phase and needs to be applied to smaller networks first.

One interesting perspective of a dependability case would be to make it a living document. Many aspects like assuring power supply etc., remain relatively stable in time, but the network usage evolves in time. However, data on the network is collected all the time. One should use this data to monitor the dependability also in some long term manner. Less stable aspects, like availability, would thus be monitored, and the dependability case would be updated accordingly.

REFERENCES

[1] G. Despotou and T. Kelly, "Extending the safety case concept to address dependability," in *Proceedings of the 22nd international system safety conference*, 2004, pp. 645–653.

[2] P. Bishop and R. Bloomfield, "A methodology for safety case development," in *Industrial Perspectives of Safety-critical Systems: Proceedings of the Sixth Safety-critical Systems Symposium, Birmingham 1998*, F. Redmill and T. Anderson, Eds. Springer, 1998, pp. 194–203. [Online]. Available: http://citeseer.ist.psu.edu/bishop98methodology.html

[3] T. Kelly, "Arguing safety a systematic approach to safety case management," Ph.D. dissertation, University of York, Department of Computer Science, 1999. [Online]. Available: http://www-users.cs.york.ac.uk/~tpk/tpkthesis.pdf

[4] A. Avizienis, J.-C. Caprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Trans. Dependable and Secure Computing*, vol. 1, no. 1, pp. 11–33, 2004.

[5] Adelard LLP, "Assurance and safety case environment, ASCE software." [Online]. Available: http://www.adelard.com/web/hnav/ASCE/index.html

[6] T. Kelly, "Managing complex safety cases," in *In Proc. 11th Safety Critical Systems. Symposium, Heidelberg*. Springer Verlag, 2003. [Online]. Available: http://www-users.cs.york.ac.uk/~tpk/sss03.pdf