A Conceptual Model for Service Availability*

Judith E. Y. Rossebø[1,2], Mass Soldal Lund[3,4], Knut Eilif Husa[5], and Atle Refsdal[3]

[1] The Norwegian University of Science and Technology
[2] Telenor R&D, Norway, judith.rossebo@telenor.com
[3] University of Oslo, Norway, atler@ifi.uio.no
[4] SINTEF ICT, Norway, mass.s.lund@sintef.no
[5] Ericsson Applied Research Center, Norway, knut.eilif.husa@ericsson.com

Abstract Traditionally, availability has been seen as an atomic property asserting the average time a system is "up" or "down". In order to model and analyse the availability of computerized systems in a world where the dependency on and complexity of such systems are increasing, this notion of availability is no longer sufficient. This paper presents a conceptual model for availability designed to handle these challenges. The core of this model is a characterization of availability by means of accessibility properties and exclusivity properties, which is further specialized into measurable aspects of availability. We outline how this conceptual model may be refined to a framework for specifying and analysing availability requirements.

1 Introduction

Availability is an important aspect of today's society. Vital functions as e.g. air traffic control and telecom systems, especially emergency telecommunications services, are totally dependent on available computer systems. The consequences are serious if even parts of such systems are unavailable when their services are needed.

Traditionally, the notion of availability has been defined as the probability that a system is working at time t, and the availability metric has been given by the "uptime" ratio, representing the percentage of time that a system is "up" during its lifetime [1]. This system metric has been applied successfully worldwide for years in the PSTN/-ISDN telephony networks along with failure reporting methodologies [2].

With this traditional understanding, a web-based application such as a concert ticket sales service may have $99,999\%$ availability, however if it is down for the 5 minutes when concert tickets to a popular artist are put out for online sale while at the same tickets can be purchase via competing distributors, this means a considerable loss of profit for the adversely affected ticket sales website even though the service is considered to be highly available along traditional lines. Service availability needs a more enhanced metric in order to measure availability in a way that meets the demands of today's services which have been shown to have much more bursty patterns of use than traditional

* The research on which this paper reports has been funded by the Research Council of Norway project SARDAS (152952/431). Thanks to Manfred Broy, Rolv Bræk, Øystein Haugen, Terje Jensen, Fabio Massacci, Birger Møller-Pedersen, Ina Schieferdecker, Ketil Stølen and Thomas Weigert for commenting on earlier versions of this paper

PSTN/ISDN services [3]. Such burstiness in usage patterns also affects the ability of the service to provide to all users requiring the use of a service at a given moment.

Indeed, as the environment where services are deployed becomes more and more complex [4] a more "fine-grained" view on "what is availability" is needed. Several global virus attacks have recently showed that availability is indeed affected by security breaches, e.g., when e-mail servers are flooded by infected e-mails, the availability for "real" e-mails decreases. Another example is the so called denial of service attack, for which a service is overloaded with requests with the only purpose of making the service unavailable for other users.

In this paper we motivate and introduce an augmented notion of availability. In the heart of the resulting conceptual model lies a characterization of availability as aspects of accessibility and exclusivity. Further, we seek to preserve well-established definitions from our main sources of inspiration: security, dependability, real-time systems, and quality of service (QoS). The paper shows how the conceptual model may be used as a basis for specifying service availability requirements in a practical setting.

In Sect. 2 we provide the basis for our analysis of availability including our analysis of different viewpoints and approaches on availability and other aspects in the fields of security and dependability. Motivated by this discussion on related work in the fields of dependability and security research, we identify the requirements a conceptual model of availability should satisfy. In Sect. 3 the properties of availability are discussed, in Sect. 4 the means to achieve availability are classified, and in Sect. 5 we present some of the threats to availability. In Sect. 6 the overall conceptual model including an availability measure is presented. Summary and conclusions are provided in Sect. 7.

2 Requirements to a Refined Notion of Availability

The setting for our availability analysis is derived from the fields of dependability and security, and we therefore strive to conform to the well-established concepts and definitions from these fields where there is a consensus. We also look to different approaches and viewpoints in dependability and security research to motivate and derive a set of requirements for an availability concept model which enables an augmented treatment of availability that is more suited to securing availability in today's and future services.

2.1 Classifying Availability

Availability has been treated by the field of dependability and the field of security. The definitions of availability commonly used in these fields are:

1. Readiness for correct service [5].
2. Ensuring that authorised users have access to information and associated assets when required [6].
3. The property of being accessible and usable on demand by an authorized entity [7,8].

We find the first of these definitions to be insufficiently constraining for practical application to design of systems and services with high availability requirements. An

integral part of securing availability is ensuring that that the service is provided to authorised users only; this is not addressed by the first definition. This aspect is addressed by the second, but neither of these two definitions captures the aspect of a service being *usable*. The third definition, however, does capture all of these aspects, and therefore is the basis for our analysis of availability in more detail.

We claim that there is a need to provide an enhanced classification of availability in order to thoroughly analyse and enable the rigorous treatment of availability throughout the design process depending on the requirements of the individual services. *Our availability model should therefore characterise the properties/attributes of availability.*

2.2 Classification of Threats and Means

The IFIP WG 10.4 view on dependability is elaborated by J. C. Laprie in [5]. This conceptual model of dependability consists of three parts: the *attributes* of, the *threats* to and the *means* by which dependability is attained [9]. This is a nice approach which motivates us to use a similar approach in our classification of availability. *Clearly, threats to availability such as denial of service, and means to availability such as applying redundancy dimensioning techniques, have an important place in our availability model.*

However, in order to classify threats to availability and means to achieve availability in a security setting, we are also motivated by the approach used in the security field of risk analysis and risk management as in [10,11].

This is because, incidents resulting in loss of availability do not necessarily transpire due to faults and therefore classification of means in terms of faults as in [5,9] is, in our view, insufficient for availability analysis. An example is the hijacking of user sessions by an attacker or group of attackers, preventing the authorised user or group of users from accessing the service. This incident results in loss of service availability for a set of users, without incurring a fault in the system. An *unwanted incident* is defined in [12] as an incident such as loss of confidentiality, integrity and/or availability. A fault is an example of an unwanted incident. *The availability model should therefore classify the means to achieve availability in terms of countering unwanted incidents.*

In [5,9], the *threats* to dependability are defined as faults, errors and failures, and these are seen as a causal chain of threats to dependability:

$$\text{fault} \longrightarrow \text{error} \longrightarrow \text{failure}$$

This understanding of threats serves nicely in the dependability model, however, we use the definition of threat, as defined in [8]: a *threat* is a potential cause of an unwanted event, which may result in harm to a system or organisation and its assets. Unlike [9], we do not consider such a causal chain alone as the sole threats to availability, as service availability may be reduced by e.g. a denial of service (DoS) attack which reduces the service availability without causing a fault, error, or failure to the actual service itself. *The conceptual model of availability should classify known threats to availability while conforming to existing literature on the classification of security threats.*

2.3 Viewpoints for Analysing Availability

For our availability analysis, it is appropriate to evaluate whether we should consider a system from a black box or white box perspective. In [13,14], Erland Jonsson provides a conceptual model for security/dependability with a black box view.

In this system model view, Jonsson considers availability to be a purely behavioural aspect related to the outputs of the system, solely with respect to the users. Availability is defined as the ability of a system to deliver its service to the authorised user [13]. This viewpoint is valid and useful for some aspects of availability analysis; however, we see the need for evaluating availability from other viewpoints as well. Availability aspects of the internal components of the system must also be analysed.

We claim that aspects of availability must indeed be observed from both the input and output sides as well as the internal components of the system. For example, denial of service attacks can be observed as malicious input to a system to either flood the system and render it unavailable, or in order to alter the integrity of the system, e.g., by deleting a group of users from the database of authorised users. In the latter case, the input messages of the intruder can be observed, and the changes to the internal database, resulting in a loss of availability for those users that were deleted, will also be registered.

With a black box view only, as in [13], only the externally observable behavioural aspects of availability can be studied. However, it is also important to observe and analyze the internal behaviour in the system in order to analyze the availability aspects of components, in particular service components which collaborate to deliver the service. Motivated by a service-oriented system view, it is important to consider a whitebox view also, so that the internal means to achieve availability can be specified and internal causes that affect availability can be examined. *The conceptual model should therefore address internal and external concerns of availability.*

2.4 Requirements of Different Services

In the current and future telecommunications market, there are many different types of services each of which may have different requirements with respect to availability. Telephony services, and in particular, emergency services, are examples of services with stringent availability requirements. Internet-based services, however, have somewhat different requirements. Requirements for what may be tolerated of delays or timing out of services are rather lax currently for e.g., online newspaper services. Yet, a citizen who leaves the tax return to the last minute before the deadline for filing requires urgently that the online tax return submission service is available at that particular moment [15].

For traditional telecommunications services, the traditional availability requirement of $99,999\%$ availability is still valid, however, it does not sufficiently address all of the differentiated requirements with respect to service availability. More precisely, as advocated by the Service Availability Forum (SAF) [16], there is also a need for a customer centric approach to defining availability requirements. The availability concern of the Service Availability Forum is readiness for correct service and in particular continuity of service, with a focus on the demands of the customers.

We intend to incorporate the ideas of the SAF in our model, to enable customer oriented availability requirements, however, extending these to include the aspects of ensuring that unauthorised users cannot interrupt, hijack, or prevent the authorised users from accessing a service. *The model must address the availability requirements in a flexible manner, in order to address the different aspects of availability.*

2.5 Measuring Availability

As discussed in the introduction, we need a more fine grained measure of availability than pure "up" or "down". Services can exist in numerous degraded but operational/-usable/functional states between "up" and "down" or "correct" and "incorrect". For example, an online newspaper may behave erratically with slow response times for displaying articles browsed without going down or becoming completely unavailable. It should be possible to describe various states of availability in order to specify just how much a reduction of service quality may be tolerated.

While both the Common Criteria [17] and Johnson [14] define security measures and provide techniques for measuring security in general, there is a need for a more fine grained metric for measuring availability that takes into account, for example, measurement of how well user requirements are fulfilled, as well as a need for measuring the ability to adequately provision a service to all of the authorised users requiring the service at a given moment. Such a metric needs to take into account the appropriate set of parameters, not just the usual average based on the mean time to failure (MTTF) and the mean time to repair (MTTR). *Our aim is to incorporate techniques from the existing initiatives in the fields of security and dependability in order to arrive at a more complete composite measure of availability.*

3 Properties of Availability

Availability encompasses both exclusivity, the property of being able to ensure access to authorised users only, and accessibility, the property of being at hand and useable when needed. As such, contrary to, e.g., the IFIP WG10.4 [18], which treats availability as an atomic property, we see availability as a composite notion consisting of the following aspects:

- Exclusivity
- Accessibility

We elaborate on these two properties in Sect. 3.1 and Sect. 3.2.

3.1 Exclusivity

By *exclusivity* we mean the ability to ensure access for authorised users only. More specifically, this involves ensuring that unauthorised users cannot interrupt, hijack, or prevent the authorised users from accessing a service. This aspect is essential to prevent the denial of legitimate access to systems and services. That is, to focus on prohibiting unauthorised users from interrupting, or preventing authorised users from accessing

services. Our definition of exclusivity involves both users and non-users, i.e., ensuring access to users while keeping unauthorised users out. This is in order to properly address means to achieve exclusivity. Some of these will address ensuring access for authorised users and others will address techniques for preventing unauthorised users from accessing or interrupting services.

The goal with respect to exclusivity is to secure access to services for authorised users in the best possible way. Essentially this means:

- Secure access to services for the authorised users.
- Provide denial of service defence mechanisms. Here we focus on prohibiting unauthorised users from interrupting, or preventing users from accessing services.
- Ensure that unauthorised users do not gain access to services.

Note that attacks via covert channels or by eavesdropping can lead to loss of confidentiality without loss of exclusivity as the attacker is not accessing the service, but passively listening in on service activity. Confidentiality, however, consists of exclusivity and absence of unauthorised disclosure of information.

3.2 Accessibility

We define *accessibility* as the quality of being at hand and usable when needed. The notion of "service" is rather general, and what defines the correctness of a service may differ widely between different kinds of services. Accessibility is related to quality of service (QoS) [19,20,21], but what is considered relevant qualities vary from one domain to another. Furthermore, QoS parameters tend to be technology dependent. An example of this is properties like video resolution and frame rates [20], which are clearly relevant for IP-based multimedia services and clearly not relevant in other service domains, such as SMS or instant messaging services.

What all services do seem to have in common is the requirement of being timely; for a service to be accessible it must give the required response within reasonable time. In addition to being timely, a service will be required to perform with some quality to be usable. Hence, we divide accessibility properties into two major classes of properties: *timeliness* properties and *quality* properties. Timeliness is the ability of a service to perform its required functions and provide its required responses within specified time limits. A service's quality is a measure of its correctness and/or how usable it is.

Consider an online booking service. From the viewpoint of a user at a given point in time, we could say that the quality of the service is either 1 or 0 depending on whether the user gets a useful reply (e.g. confirmation) or unuseful reply (e.g. timeout). (Over time this can be aggregated to percentages expressing how often one of the two kinds of responses will be given.)

In a multimedia service like video streaming, the frame rate may be seen as a timeliness property (each frame should be timely) while the resolution of each frame and the colour depth are quality properties.

In both these examples we may see a dependency between timeliness and quality. In the first example (Fig. 1) we may assume a deadline t_2 for the response to the user for the service to be accessible. However, we must also assume some processing time t_1

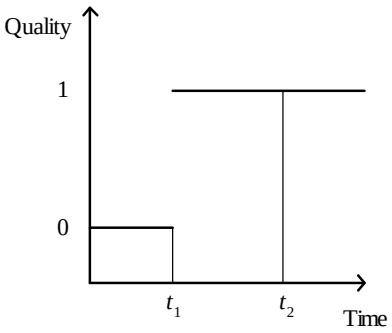

Figure 1. Quality vs. timeliness

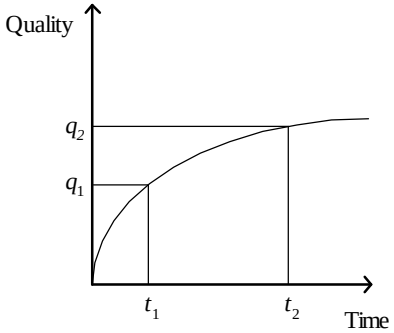

Figure 2. Quality vs. timeliness

for the service to be able to produce an answer. This means that the quality requirement enforces a lower bound on the timeliness; if the deadline is too short the user will always receive the timeout message. In other words we must have that $t_1 < t_2$ for the service to be accessible.

In the other example (Fig. 2) we may assume that higher quality requires more processing time per frame. This means that a required quality q_1 provides a lower limit t_1 on the processing time of each frame. Further, to get the required frame rate there must be a deadline t_2 for each frame, which provide an upper bound q_2 on the quality. This means the service must stay between this lower and upper bound to be accessible. This approach may be seen as an elaboration of Meyer's concept of *performability evaluation* [22].

These considerations motivates a notion of *service degradation*. We define service degradation to be reduction of service accessibility. Analogous to accessibility we decompose service degradation into timeliness degradation and quality degradation, and see that these are quantities mutually dependent on each other. For example, graceful degradation in timeliness may be a way of avoiding quality degradation if resources are limited, or the other way around. A combination of graceful degradation in timeliness and graceful degradation in quality may also be applied. Related to QoS, accessibility may actually be considered a QoS tolerance cut-off, i.e., the point at which the QoS deteriorates to a level where the service is deemed no longer usable, so that the service is considered unavailable.

4 Means to Ensure Availability

Traditionally, the approach to meeting availability requirements has primarily focused on ensuring accessibility aspects of availability such as by introducing redundancy, and by service replication. This is a valid approach to availability, but it does not ensure, e.g., that the service is accessible to authorised users only. There are costs involved in introducing redundancy and replication, which need to be justified. The goal should be to obtain more comprehensive, more cost-effective means to achieve availability, and to

specify, design, and implement a set of measures that enable delivery of services and/or systems according to availability requirements.

By means to ensure availability we address *protection* of the service from incidents leading to a loss of availability. Therefore, in our model, we categorise the means into the following three groups: *incident prevention*: how to prevent incidents causing loss of availability; *incident detection*: how to detect incidents leading to loss of availability; and *recovery from incident*: the means to recover after an incident has lead to a loss of availability. We do not attempt to create an exhaustive list of all such measures, but do provide examples that illustrate the different aspects of securing availability.

4.1 Incident Prevention

Preventative means are defined as the internal aspects of a system that are designed to prevent, stop or mitigate intrusions, faults, errors, or other incidents which have a negative effect on the availability of a system.

Access control is an important preventative means for achieving the exclusivity aspect of availability. Access control is the prevention of unauthorised use of a resource, including the prevention of use of a resource in an unauthorised manner [7].

Providing *integrity protection* mechanisms is important for example, in order to protect against manipulation and redirection of messages resulting in denial of service for the authorised user.

It is also important to ensure that the required resources e.g. in the network that an authorised user has permission to use during a session are indeed allocated to the user to ensure that the service is delivered according to the user availability requirements.

Another example of a means for avoiding loss of availabilty is *graceful degradation* [23], that is degradation of a system in such a manner that it continues to operate, but provides a reduced level of service rather than failing completely. By applying graceful degradation schemes a complete loss of availability can be prevented.

4.2 Incident Detection

Incident detection consists of means to discover incidents such as denial of service attacks, faults, errors or failures, which lead to a loss or reduction of availability.

Detective measures will commonly be coordinated with recovery aspects of the system in order to adapt and restore system availability. Fault detection, traffic flow monitoring, intrusion detection systems (IDS), and accounting audits are all examples of detective measures.

For an efficient approach to unwanted incident detection, it is wise to combine monitoring, fault detection and IDS techniques along with audit logs generated and process the information and data collected in real time or close to real time in order to detect and thwart attacks or incidents that have the potential to result in loss or reduction of availability.

4.3 Recovery from Incident

Recovery from incident consists of the means to recover from incidents leading to loss or reduction of availability. This includes techniques for adapting the service, e.g. in the

case that anomalies are detected by the IDS so that major unwanted incidents of loss of availability are avoided. Recovery means may entail, e.g., making changes to the internal aspects of the system, such as correction of faults or removal of system vulnerabilities. Additionally, external filters may be implemented to filter away the discovered cause of the incident such as malicious traffic or traffic from unauthorised users. Recovery addresses the *adaptability*, *robustness*, *maintainability* and *redundancy* aspects of the system.

5 Threats to Availability

The most explicit threat to availability is *denial of service* (DoS) attacks. *Replay*, *masquerade*, *modification of messages*, *man-in-the-middle* and *misuse of service* are examples of other kind of active threats that may affect availability. Threats may originate on the inside (inside attackers) or the outside (outside attackers) of the system. The impact of threats varies with the nature of the threats; some threats may result in degradation of the service, others in complete loss of service. Going into detail on this issue is outside the scope of this paper, but below we give some examples on how some of these threats may affect availability.

Denial of service attacks may lead to loss of use due to unauthorised use of the service preventing authorised users from accessing the service. Unauthorised use may also create over-usage problems having an overload effect and in this way degrading the quality of the service for the authorised users.

In a masquerade, an attacker steals the identity of a real user and obtains fraudulent access by masquerading as the real user while preventing the valid user from accessing services. Or, the other way around, an attacker replaying or masquerading as a service may deceive the user, and the service the user intended to access is then not available.

6 Conceptual Model for Service Availability

Based on the requirements from Sect. 2 and our discussion above we propose the overall model presented in Fig. 3 (represented in UML 2.0) and further explained in the following text.

In the figure the relationships between availability, threats and means are shown. Availability is affected by means and threats. Means ensures availability and protects against threats. Threats may cause reduction of availability.

There are many different types of services, and they may have different requirements with respect to availability. Availability requirements should be flexible enough to address the different services consistently. We propose that availability is specified by the means of availability policies and predicates over measurable properties of services, and that these policies and predicates are decomposed in accordance with the decomposition of availability in the conceptual model. An availability policy consists of an accessibility policy (e.g., required resources) and an exclusivity policy (e.g., which entities have permissions to use the service or system).

The predicates place conditions on the allowed behaviour of the service. In order to express these predicates, there is a need to describe rules for allowed or prohibited

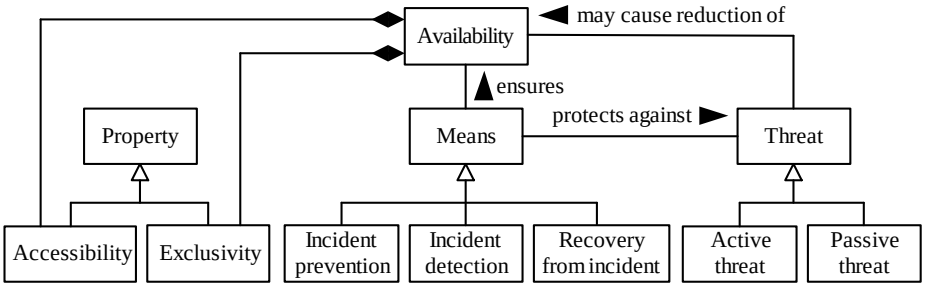

Figure 3. The overall picture

behaviour and to provide a means for measuring the availability properties of a service. Figure 4 illustrates how availability properties are related to services, i.e., as part of the relation between the service and the user entity using the service.

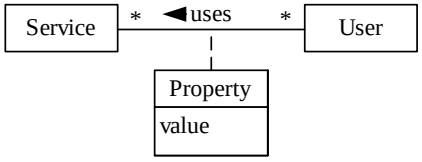

Figure 4. Service availability

Our conceptual model provides the foundation for an availability metric in that it provides decomposition of availability properties that may be mapped to measurable quantities. This metric includes behavioural measures, preventative measures, and correctness measures such as the measurement of degree of degradation.

The following is the mathematical representation of the availability metric for a service. Let A denote a service with an availability property for a user group U, and let X denote the availability metric for service A. We represent X as an n-tuple $X = (x_1, \ldots, x_n)$ where x_i is a measure of an aspect of availability. By this we mean that x_i describes requirements for a particular availability aspect. The minimum requirement for each x_i must be satisfied in order to fulfil the total availability requirement X.

Using our conceptual model this idea can be refined as follows: We represent X as a tuple $X = (X_1, X_2)$ where X_1 measures the exclusivity properties, and X_2 measures the accessibility properties.

Essentially, the requirement aims to describe the degree of accessibility and exclusivity that is sufficient for the user to be able to activate and use the service. Examples of measures of aspects of exclusivity may be illustrated by the following: For a measurement of exclusivity we need to be able to answer questions such as “how well does the system keep out unauthorised users while still granting access to authorised users?” This leads to the following examples of exclusivity requirements:

- The probability that an authorised user is denied access to the service at a given time t should be less than x.
- The probability that an unauthorised user obtains access to the service at a given time t should be less than y
- User u should be prohibited from accessing service s when user v is using the service.
- The number of intrusions at a given time t (e.g. during a critical moment) should be less than z.

Similar measures may be defined for accessibility. These may be defined with basis in measures for service degradation, timeliness, performance, and quality.

In order to apply the model, the availability requirements must be determined. Threats must then be analysed to understand what affects availability and means for ensuring availability need to be identified to meet requirements and counter threats. Measurements of the different aspects are then used to evaluate how well the availability requirements are met. A more in depth discussion of how to apply the model is the subject of further work. We are currently applying the model to our work on ensuring availability in service composition.

7 Conclusions

The contribution of this paper is a conceptual model for availability that takes into account a much broader spectrum of aspects that influence availability than previously addressed by work in this area. We have argued that exclusivity is an aspect of availability that has been generally neglected in the literature, and shown where it fits in an enhanced notion of availability. Further we have shown how QoS, real time and dependability considerations may be integrated in the model and treated as accessibility properties.

We have established that there is a need for a more fine grained metric for measuring availability and have provided a representation of the availability metric for a service that allows specification of the measurable requirements for exclusivity and accessibility properties.

Our conceptual model for availability embraces both a white box view as well as a black box view of availability and, hence, addresses both internal and external concerns of availability. The need for this is apparent in our current work on ensuring availability in service composition that encompasses a collaboration of roles, which are slices of behaviour across distributed systems. These must be composed correctly in order to achieve a service with the required availability.

The model also contains a classification of threats to availability and means to ensure availability, and establishes the relationship between threats, means and availability properties. Together these elements provide a framework in which all relevant views and considerations of availability may be integrated, and a complete picture of service availability may be drawn.

References

1. Ross, S.M.: Inroduction to probability models. 6th edn. Academic Press (1997)
2. Enriquez, P., Brown, A.B., Patterson, D.A.: Lessons from the PSTN for dependable computing. In: Workshop on Self-Healing, Adaptive and self-MANaged Systems (SHAMAN 2002). (2002)
3. Clark, D., Lehr, W., Liu, I.: Provisioning for bursty internet traffic: Implications for industry and internet structure. In: MIT ITC Workshop on Internet Quality of Service. (1999)
4. Arbaugh, W.A., Fithen, W.L., McHugh, J.: Windows of vulnerability: A case study analysis. IEEE Computer **33** (2000) 52–59
5. Laprie, J.C., ed.: Dependability: Basic Concepts and Terminology. Springer-Verlag (1992)
6. International Standards Organization: ISO/IEC 17799, Information technology – Code of practice for information security management. (2000)
7. International Standards Organization: ISO 7498-2, Information Processing Systems – Interconnection Rrefence Model – Part 2: Security Architecture. (1989)
8. International Standards Organization: ISO/IEC 13335, Information technology – Security techniques – Guidelines for the management of IT security. (2001)
9. Avižienis, A., Laprie, J.C., Randell, B.: Fundamental concepts of dependability. In: Third Information Survivability Workshop (ISW-2000). (2000)
10. den Braber, F., Lund, M.S., Stølen, K., Vraalsen, F.: Integrating security in the development process with UML. In: Encyclopedia of Information Science and Technology. Idea Group, 2005 (2005) 1560–1566
11. Lund, M.S., den Braber, F., Stølen, K.: Maintaining results from security assessments. In: Proc. Seventh European Conference on Software Maintenance and Reengineering (CSMR 2003), IEEE Computer Society (2003) 341–350
12. Standards Australia: AS/NZS 4360:1999, Risk Management. (1999)
13. Jonsson, E.: An integrated framework for security and dependability. In: The New Security Paradigms Workshop (NSPW'98). (1998) 22–29
14. Jonsson, E., Strömberg, L., Lindskog, S.: On the functional relation between security and dependability impairments. In: The New Security Paradigms Workshop (NSPW'99). (1999) 104–111
15. Ryvarden, E.: Skatte-servere tålte ikke trykket. digi.no, April 30 (2005) (In Norwegian).
16. Service Availability Forum: Backgrounder. (http://www.saforum.org/home, accessed March, 2004)
17. International Standards Organization: ISO/IEC 15408, Information technology – Security techniques – Evaluation criteria for IT security. (1999)
18. IFIP WG10.4: IFIP WG10.4 on dependable computing and fault tolerance. http://www.dependability.org/wg10.4/ (2005)
19. Barbacci, M., Klein, M.H., Longstaff, T.A., Weinstock, C.B.: Quality attributes. Technical report CMU/SEI-95TR-021, Software Engineering Institute, Carnegie Mellon University (1995)
20. Vogel, A., Kerherve, B., von Bochmann, G., Gecsei, J.: Distributed multimedia and QoS: A survey. IEEE Multimedia **2** (1995) 10–18
21. Group, O.M.: UML profile for modeling quality of service and fault tolerance characteristics and mechanisms. OMG Adopted Specification ptc/2005-05-02 (2005)
22. Meyer, J.F.: Performability evaluations: Where it is and what lies ahead. In: Proc. International Computer Performance and Dependability Symposium, IEEE (1995) 334–343
23. Shin, K.G., Meissner, C.L.: Adaption and graceful degradation of control system performance by task reallocation and period adjustment. In: Proc. 11th Euromicro Conference on Real-Time Systems, IEEE (1999) 29–36