Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.

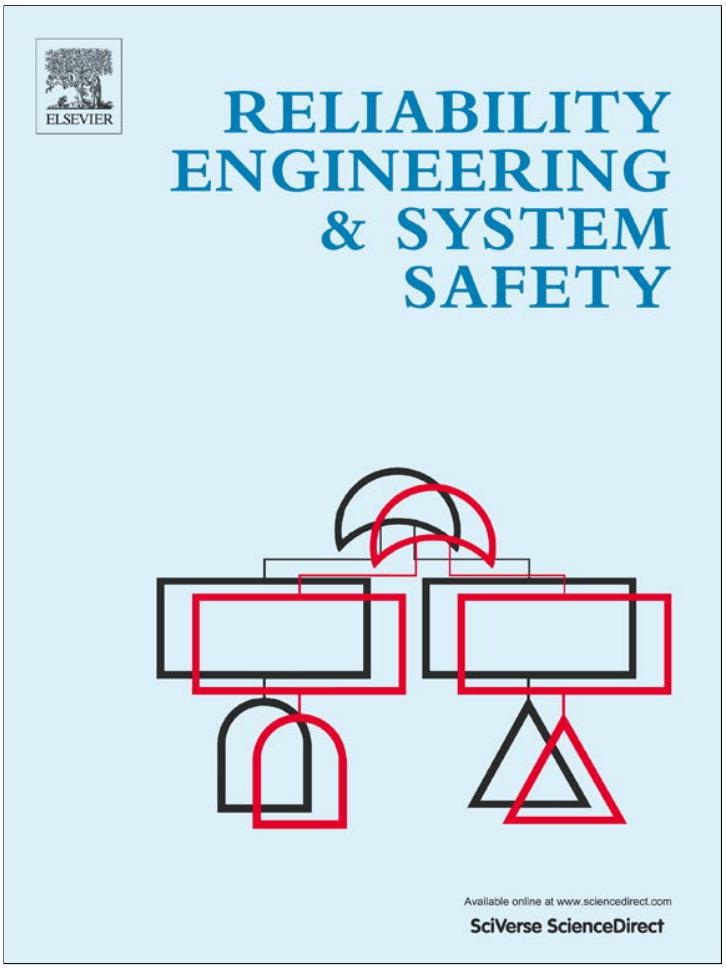


(This is a sample cover image for this issue. The actual cover is not yet available at this time.)

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Author's personal copy

Reliability Engineering and System Safety 109 (2013) 160–174

Contents lists available at SciVerse ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Review

# Towards risk-aware communications networking ☆

Piotr Chołda [a,1], Eirik L. Følstad [b,2], Bjarne E. Helvik [b,2], Pirkko Kuusela [c,3], Maurizio Naldi [d], Ilkka Norros [c,*,3]

[a] AGH University of Science and Technology, Kraków, Poland
[b] Q2S, Norwegian University of Science and Technology (NTNU), Trondheim, Norway
[c] VTT, Technical Research Centre of Finland, Helsinki, Finland
[d] Dipartimento di Informatica Sistemi Produzione (DISP), Università di Roma "Tor Vergata", Roma, Italy

ARTICLE INFO

*Article history:*
Received 13 May 2011
Received in revised form
16 July 2012
Accepted 15 August 2012
Available online 28 August 2012

*Keywords:*
Dependability
Network design
Recovery
Reliability
Risk
Service level agreement (SLA)
Survivable networks

ABSTRACT

We promote introduction of risk-awareness in the design and operation of communications networks and services. This means explicit and systematic consideration of uncertainties related to improper behavior of the web of interdependent networks and the resulting consequences for individuals, companies and a society as a whole. Central activities are the recognition of events challenging dependability together with the assessment of their probabilities and impacts. While recognizing the complex technical, business and societal issues, we employ an overall risk framing approach containing risk assessment, response and monitoring. Our paradigm gathers topics that are currently dispersed in various fields of network activities. We review the current state of risk-related activities in networks, identify deficiencies and challenges, and suggest techniques, procedures, and metrics towards higher risk-awareness.

© 2012 Elsevier Ltd. All rights reserved.

**Contents**



☆This work was done within the EU FP7 NoE Euro-NF (http://www.euronf.org) framework (Specific Joint Research Project RISKASIP).
* Corresponding author.
*E-mail addresses:* piotr.cholda@agh.edu.pl (P. Chołda), eirik.folstad@q2s.ntnu.no (E.L. Følstad), bjarne@q2s.ntnu.no (B.E. Helvik), Pirkko.Kuusela@vtt.fi (P. Kuusela), naldi@disp.uniroma2.it (M. Naldi), Ilkka.Norros@vtt.fi (I. Norros).
[1] The researcher has been also partially funded by the Polish Ministry of Science and Higher Education, Grant No. O R00 0119 12.
[2] "Q2S - Centre for Quantifiable Quality of Service in Communication Systems, Centre of Excellence" appointed by The Research Council of Norway, funded by The Research Council, NTNU and UNINETT.
[3] The researchers have been also partially supported by Project "GOODNET—Control and Care in the Risk-aware Networking" funded by Tekes, the Finnish Funding Agency for Technology and Innovation.

0951-8320/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ress.2012.08.009

Author's personal copy





## 1. Introduction

It is a generally recognized fact that people depend more and more deeply on a complex web of interdependent communications networks, in particular the Internet. Risks associated to the failures of electronic communications reverberate directly across society. We must be aware of those risks and decide how to face them.

The literature dealing with risk is extensive, ranging from Lowrance's classic [1] to the recent [2–5]. Looking only into a few papers recently published in RESS like [6–10], we can see that the topic is gaining momentum in various contexts, also in networking. Most of the risk literature is, however, either general or focused on industries in which life danger or societal consequences of failures are spectacularly high (e.g., safety of energy provisioning, aviation, railways, or oil drilling platforms). In contrast, the context of communications networks is underrepresented, although the inclusion of risk in reliability analyses has been advocated in some papers (e.g., in [11] for networks or [12] for dependable computing). With this paper we intend to promote the issue of networking risk in a comprehensive way, and illustrate possible approaches.

Lowrance [1] defined risk as 'a measure of the probability and severity of adverse effects.' Kaplan and Garrick [13] specified the notion of risk as a triplet consisting of a risk scenario (including the event sequence leading to the unfortunate event), the likelihood of that scenario, and the consequences of the scenario (damage created). Along these lines, we consider that also in networking there are three basic components that should be considered to properly address risk:

- clear recognition of *events that challenge network dependability*,
- assessment of their *probability* (conditioned on available information, thus often subjective/Bayesian) and *extent*, trying to take into account uncertainties involved, and
- assessment of their *impact*, an element mostly neglected in our context so far.

However, as we will discuss extensively in this paper, the dependability problems of communication networks and the associated risks are to a large extent different from those in the fields traditionally considered safety-critical or risk-prone. One of the main challenges is that the rhythm of innovation in communications is now so fast that the networks must be considered as continuously changing. We can still recognize the usage of technical innovations as proceeding through three phases: from being a technological *toy* to becoming an *alternative* to the devices and services in place, to finally representing the *dominant solution*. In the world of Internet applications, the progress from a toy to a dominant way of acting happens often at very fast pace. Moreover, networking paradigms are changing —consider, e.g., the current trend towards cloud computing.

So far, risk in network design and operation has been present mainly implicitly, while here we emphasize the need to deal with it explicitly. Nowadays, a typical approach to reliability is technology-oriented. The occurring failures are classified according to their roots, frequencies, and time durations. Granting full connectivity and survivability to those failures is the ultimate end in itself. For that purpose, recovery methods are introduced, with a basic focus on connectivity restoration, assessment of actual downtimes incurred before traffic can be sent again, and optimization of backup resources. However, in practice we must recognize that there are a number of higher-level issues that require a finer attention to the consequences of failures and the people or companies affected by them. We can identify both business and societal issues. The service provided by an operator is typically just a part in a longer service chain, so that the service provider has to cooperate with the other providers supplying other links of the chain. At the same time that service belongs to a service portfolio offered by the operator to customers of different relevance and profile, where each service has different technical needs. Service disruptions do not have all the same business consequences.

The liability of a network operator if the service provided to its customers degrades is defined in Service Level Agreements (SLA). They typically include performance bounds on some basic service parameters (e.g., the minimum guaranteed bitrate), but generally do not go much beyond granting a basic degree of dependability in terms of service features like Quality of Service (QoS), availability and security [14]. However, communications services are nowadays so pervasive that a number of human activities, sometimes vital, rely on them: the influence of critical infrastructures to people's lives (and the consequences of catastrophes) is so heavy that they are typically supervised by state authorities, and go quite beyond the simple business relationship between service provider and customer. Therefore, evaluation of services just through some technical dependability parameters leads to a very narrow viewpoint. Adopting a *risk-aware networking* approach allows us to consider and reconcile both technical and higher-level perspectives.

Although dependability can be affected by intentional attacks on the network, in this paper we limit our interest to risks related to service disruptions due to network failures that are not caused by malicious actions,[4] considering all unintentional events after which some network elements (hardware, software, protocols) cease to work properly. We structure our discussion adopting notions of NIST Special Publication 800-39 [16] on the management of information security risk in a company, which is the most relevant level of responsibility in this context. This recommendation divides the risk management process into the following four activities:

- *Risk framing*: The umbrella action that produces a whole risk management strategy for the organization, where assumptions on challenging events are enumerated and the three points below are accomplished.
- *Risk assessment*: To identify possible problems, and estimate their frequencies and impact, first qualitatively and then quantitatively.
- *Risk response*: To determine reactions to predicted risk, where basic options include acceptance, avoidance, mitigation, sharing, or transfer.

[4] This topic has been described quite well, see for instance [15].

Author's personal copy



- *Risk monitoring*: To check whether the selected responses have performed well, and provide feedback to re-evaluate and update response policies.

Unfortunately, the development of scientific methods to assess networking risks is still in progress. Although there are many notable activities contributing to this area, some already with a long history and a considerable maturity (see Sections 3–5), the entire view has not been grasped yet in a sufficiently comprehensive and organized way. Rather, a lot of work has been done concerning the risk-aware paradigm we are advocating, but it is dispersed in many various regulatory, standardization, research, planning, operational and maintenance as well as infrastructure resilience activities. They should be recognized and promoted as a whole, so that each part would really add to the meaning and significance of the others in the context of the risk-aware networking.

Our emphasis is somewhat different from that of complex systems science, where new theoretical tools are developed for better understanding of emergent features and phenomena of large networks (see, e.g., [17]). Such insights can be valuable for understanding general relations like scaling laws, and perhaps even for rough evaluation of some risks. For example, it is a mathematical fact that if the distribution of node degrees is appropriately heavy-tailed, then even a random placement of links yields a 'softly hierarchical' network topology with short connections and high robustness [18]. We, however, are calling for higher risk-awareness of all players in actual, practical communications networking. Properly understood, it includes awareness of theoretical advances in complex network research.

In this paper, following the description of the risk management process recalled above, we investigate four realms of interest for risk-aware networking. In Section 2 about *methodology*, we claim that risk-aware networking needs a methodology (risk framing) that reflects the multi-faceted nature of the subject. In Section 3 about *techniques*, we show that there are several techniques of mostly mathematical character that are essential for assessing network dependability and risk. In Section 4 on *recovery methods*, we show that the main technical response to risk is its mitigation when we design the network aiming at its survivability, and risk sharing combined with service differentiation when the agreements between the service provider and its customer implicitly include a trade-off between cost and quality. Finally, in Section 5 on *practices*, we claim that network design and operation practices must be adapted to cater for proper risk monitoring, an indispensable part of risk-awareness. We conclude by expanding the view to the whole ecosystem of networking companies, and emphasizing the role of regulators in defining the responsibilities and accountability of network and service providers.

## 2. Risk framing: on the methodology of risk-aware networking

The current Internet is both large and extremely complex. It is based on an ecosystem of continuously evolving technical organizations and companies providing communications facilities and services. Its design is based on the work of practitioners and scientists, while various rules for its construction and use, as well as guidelines for its evolution, are set by national and international regulation authorities.[5] While discussing the risk due to failures of this communication infrastructure, we have at least four classes of actors playing a role, and having different points of view, as shown in Fig. 1. The four actor classes differ as to their interests and the threats they face, as well as for their range of action, see Table 1. The class represented by network and service *providers* is the only one that is able to directly influence the network and modify it, and therefore it is natural that the *provider*'s point of view slightly dominates the risk perspective. The aim of *researchers* is mostly to find algorithms and computational techniques for tasks like resilient design and efficient monitoring. In this paper we also seek to contribute to high-level understanding of the complexity of risk-related issues and support risk-aware networking. The achievement of this goal is valuable for the *providers* and *regulators*, and eventually for the *users*.

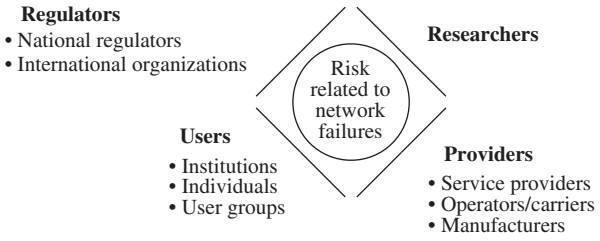


**Fig. 1.** Different classes of actors in the field of network risk.

### 2.1. Risk management cycle

Even the hardware-centric world of providers ought to study the network as a compound socio-technical system, functioning on one hand within human societies and on the other hand in a natural physical environment. A system approach is necessary, but detailed system-theoretic modeling is possible only for well-defined sub-systems. For more vaguely characterized systems, and in particular when human factors have a prominent role, one could prefer methods like Checkland's Soft Systems Methodology [19]. According to this approach, the researcher does not see the world as a set of well-distinguished systems, but as a realm of complex purposeful activities for which models are constructed within a learning process aimed also at improving those activities.

In this spirit, we consider risk-aware networking, seen principally from the *provider*'s viewpoint, as a control cycle presented in Fig. 2. An existing network is continuously *observed* (by monitoring, customer feedback, etc.), and some kind of a *picture* of its dependability status is maintained. However, when speaking of *risk* monitoring, special attention should be directed to failures with potentially severe consequences, and losses experienced by users should also be monitored as far as possible. Various degrees of analysis and processing of the raw observation data may be required here. On the basis of the observed status, decisions on network changes, including those of network support systems and relations to external networks, are taken. An important task here is the adequate assessment of risks related to the network. While risk factors related to 'normal' failures can, at least in principle, be estimated quantitatively on the basis of failure and loss statistics, thorough qualitative analyses may reveal system vulnerabilities with still a high level of risk. Many other factors than dependability-related ones influence the decisions, but we do not deal with them here.

Note that the inner cycle of Fig. 2 could basically apply to any network provider, whereas the tasks of analysis and risk assessment require additional expertise and are often neglected. The consequence of this is that the 'picture' and the risk response can be unfounded. Contrary to the current state, they should be mandatory in risk-aware networking. The most relevant

[5] For instance, re-shaping of the Internet top-level governance from its US-centric origin to a global entity has undergone a long discussion among Internet authorities, governments, international organizations and other stakeholders.

Author's personal copy



**Table 1**
Actor classes and their characteristics relevant for network risk.

| Actors | Interests | Threats | Actions |
|---|---|---|---|
| *Providers* | • Profit<br>• Customer satisfaction | • Penalties<br>• Loss of customers | • Care for dependability<br>• Enterprise risk management |
| *Users* | • Availability<br>• Quality and price | • Loss of connectivity or service<br>• Business or life consequences | • SLA adjustment<br>• Choice of a provider |
| *Regulators* | • Common benefit and societal needs<br>• Competition | • Anarchy and monopoly<br>• Breakdown of critical infrastructures | • Regulations<br>• Collection of statistics |
| *Researchers* | • Innovative solutions<br>• Understanding | • Lack of funding<br>• Lack of focus | • Public promotion of ideas<br>• Standardization |

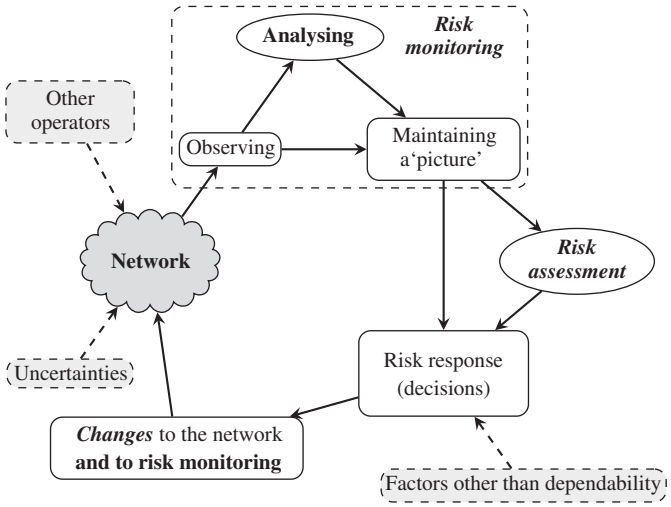


**Fig. 2.** The care-taking cycle of risk-aware networking of a provider.

techniques and practices related to those tasks will be discussed in detail in Sections 3 and 5, respectively.

### 2.2. Risk framing for communications networks

There is a tendency to see the network as a single entity, but this is far from reality. The network is composed by a number of autonomous subnetworks, which may provide services at different layers, and be owned and operated by a number of different parties. In this context an autonomous subnetwork is to be intended as capable of taking its decisions independently of other players, e.g., concerning its design, its operations, its interconnection, its purchase of services, its cooperation on a peer-to-peer basis. Those subnetworks interact in a complex manner to provide end-user services. It is more correct to regard the network as an ecosystem of networking companies, where each of them has its own business model and a place in the value chain. In this setting, each actor simultaneously competes and co-operates with other market actors. A typical example of the inter-relationships in this context is illustrated in Fig. 3. In that picture each cloud represents an autonomous entity, while the clouds depicted within each dotted box represent entities belonging to the same company (i.e., companies A, B, and D), with a common governance and coordination.

The end *service providers* deliver services (e.g., cellular telephony) to end users. These co-operate for roaming and interconnection, but concurrently they may compete for the same

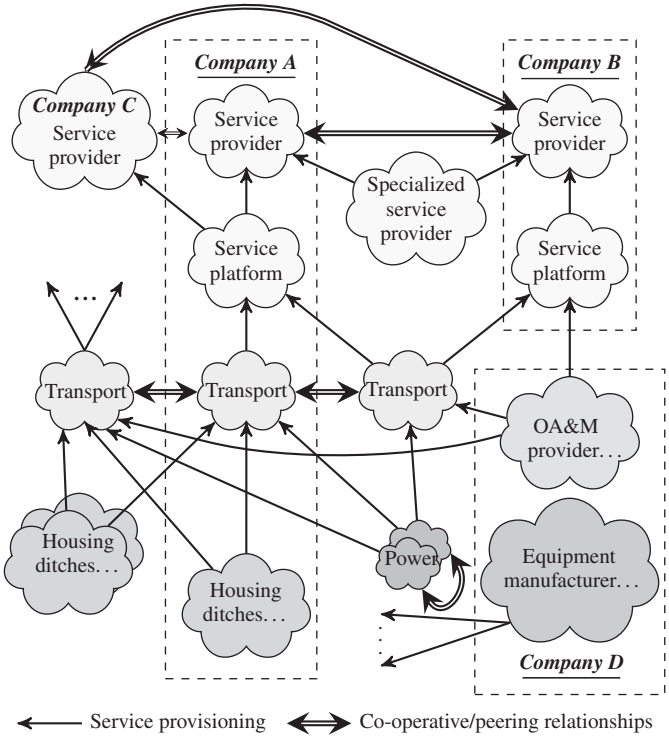


**Fig. 3.** Example of dependencies between autonomous market actors resulting from their co-operation.

subscribers. In Fig. 3, companies A and B have their own *service platforms* for delivery, while company C leases all its delivery services from company A. The latter is a vertical provider, operating all facets of the service delivery. However, also for this kind of providers, it is common to rent the housing and transport capacity from others. The transport capacity may be provided either as links/leased lines or through the use of virtual routers owned and operated by another party. The fact that the different technology layers may be owned and operated by different entities hinders a global view. A *specialized service provider*, which may play a role of a location provider or clearing house, is also indicated. It is foreseen that in the future there will be an increasing number of such providers, forming a complex value chain for a compound end user service. In this context, it should also be reminded that end users may be served by several access networks and a number of intermediate Autonomous Systems to reach the service platform of a provider they are looking for.

Author's personal copy



In the risk assessment context we should also take into account that networks that may appear as completely independent of each other, may in fact share strong commonalities in their environment. For example,

- power to networks or components that should act as backup of each other may be provided by the same supplier;
- cables owned by different carriers may be placed in the same cable duct or ditch;
- equipment operated by different providers may reside in the same housing facility.

For instance, it is quite common for the cellular access that the base stations of different operators, possibly using various technologies, are co-located.

The last but not least player that should be mentioned in this networking ecosystem is the equipment manufacturer. Usage of homogeneous equipment and software throughout the network is a recognized source of common mode logical and design failures and should be accounted for in a risk analysis. Apart from this fact, there is a trend among network operators to outsource the operation and maintenance of their networks to companies associated with the manufacturers, see for instance [20]. This is also a source of dependencies between services delivered by apparently independent providers and should be addressed in a risk analysis. See also a discussion on theoretical aspects of this issue below in Section 3.2.

Fig. 4 collects the major aspects of Internet networking that should be considered at the risk framing stage. The three dimensions of complexity are:

- horizontal sectioning;
- vertical layers of network technology/protocol;
- market elements related to the commercially optimized sharing of resources within the system of different network providers.

While developing the methodologies supporting risk-awareness, we may take guidance from networking systems, such as aviation and railways, where risk assessment methodologies are more mature. The field of communications shares some properties with those, since:

- there can be many competing providers sharing their resources,
- business solutions are networked,
- outsourcing is used, and
- there is a continuous adaptation to new technologies.

However, there are major differences. We point out that:

- communications networking does not have any central control unlike the ones used for common airspace,
- in communications networking the availability is rarely sacrificed if problems arise (to stop unsafe services), and
- the challenge in the communications risk assessment is the incomparably fast, heterogeneous development of technologies, components, and network usage patterns.

Thus, on the whole, communications networks are heterogeneous and fast moving objects for risk assessments and a large flexibility is needed in the methodology. On the other hand, when time scales and focus are narrowed and targets restricted, we become closer to a well-defined homogeneous (sub-)system.

A risk-aware network operator has also to take into consideration diverse techniques and practices related to causes of failures, network fault-resilience, failure statistics and the estimation of losses, together with their underlying assumptions and limitations. Moreover, the difference of the *user*, *provider* and *regulator* viewpoints highlights the need for various system methods and models according to the particular requirements of risk assessment, starting from a *user*'s interests up to the *regulator*'s concerns on the dependability of the communication infrastructure as a whole. Note that the network recovery with respect to 'normal failures' is not unrelated to the occurrence of big crashes and resilience against them. Serious disturbances may result from unexpected coincidences or from hidden vulnerabilities, but both cases can be counteracted (although never totally prevented) by continuous care for network dependability, e.g., through:

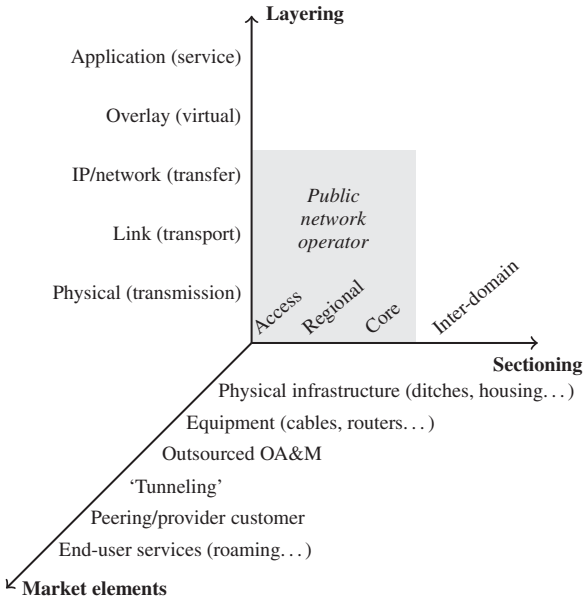


**Fig. 4.** Network aspects that influence the Internet risk.

- resilient network design,
- research on the functioning of algorithms and software,
- tracking of vulnerabilities, and
- high safety culture at core network facilities.

## 3. Risk assessment and modeling techniques

First, we elaborate on mechanisms used for dealing with failures as the main factors responsible for increasing risk related to telecom operations. Then, we present the extent of dangers related to failures occurrence along with problems of gathering meaningful statistics, giving a global view of this phenomenon. Afterwards, we sketch foundations of the theoretical aspects of risk modeling. All of the presented techniques can be used for quantification of the design and prediction of the resilient network behavior.

### 3.1. Network reliability

The reliability of a network is typically evaluated by the probability that two nodes are (physically) connected through a communications chain. A more effective approach would include the evaluation of the actual capacity provided to carry the customer's traffic. Two nodes may be connected but with a traffic capacity between them inadequate to provide the desired services. In short, it can be said that a satisfactory level of QoS or even Quality of Experience (QoE) are the sufficient conditions for successful networking, while connectivity is only a necessary one.

Author's personal copy



**Table 2**
Summary of models and methods on statistical network failures, basic results.

| Data | Model or analysis method, outputs | Reference |
|---|---|---|
| *Large IP backbone; links, routers, PoP* | Overlapping failures, failure classes; number of events: power-law; time between failures: typically Weibull | [27] |
| *IP-core and regional network, node and link failures* | Distance-dependent correlation between failures of network elements | [28] |
| *Small IP-core, node failures* | Router model with exponential uptime and Pareto downtime durations | [30] |
| *Small IP-core, node and link failures* | Nodes and local links, Weibull uptimes; long haul links, gamma uptimes | [31] |
| *Access network, node and link failures* | Spatial and temporal localities, impact of a link failure to a node outage, impact of network design on outages | [32] |
| *Wireless network, software/hardware failure data* | Typically Weibull distribution for *MTTR* or *MTTF* | [33] |
| *Large IP network* | Failure clustering | [34] |

When neither repairs nor recovery methods are taken into account,[6] the *reliability function* can however be used as the basic measure for a single network element (e.g., a node or a link). It determines the probability that an element operates properly beyond an assumed mission time $t$, starting at time 0. Any single element alternates therefore between up and down states. The uptime, i.e., the period of the uninterrupted work, can be taken as a measure of failure frequencies,[7] and is usually estimated by its mean, *MUT*. Similarly, we measure the average duration of downtimes, *MDT*. Since the downtime can be viewed as the time needed for the element to recover from its failed state, MDT and MUT are equivalent to the mean time to recovery, *MTTR*, and the mean time to (first) failure, *MTTF*, respectively.

Despite the fact that we have many theoretical and advanced models related to the network reliability theory, commonly applied network modeling of repairable systems has been limited to steady-state availability analysis, which combines the information on mean up- and downtimes. Turning from single elements to connections (or services), we may define the availability $A$ of a connection as the probability that the connection is up, expressed through the ratio $A = MUT/(MUT+MDT)$. An important property of this metric is that it might be directly perceived by an end user. From the viewpoint of risk assessment however, availability alone is not sufficient as a risk assessment basis. It can be misleading, since it is an average and does not account for the possibly wide dispersion around the average. Additionally, it does not allow to recover the frequency and impact of failures, since it provides just the ratio of the up- and downtimes. Therefore, a quantitative approach focused on service continuity is necessary [21].

Reliability metrics can be included in networking risk assessment using Kaplan and Garrick's [13] triplets $(s_i,p_i,c_i)$, where $s_i$ represents the event, $p_i$ is the probability of that event, and $c_i$ is the impact (consequence) of that event. In the networking context, we can say that for instance: $s_1 =$ fault of link 1; $p_1 =$ unavailability of link 1, or *MTTF* for this link; and $c_1$ is traffic loss related to the failure of link 1, provided an assumed recovery method is applied, etc. Alternatively, $c_i$ may be expressed in a way that is more relevant for the provider's business or the customer's satisfaction, e.g., as penalties due to SLA violation (in money units), as customer churn rates, or through quality deterioration metrics (increase of download times, decrease of streaming video QoE, etc.).

[6] Practically, repairs and recovery procedures are carried out as a response to reduce the risk, and they should be taken into account. See Section 4.1.

[7] In the case of contemporary multi-layer recovery, faults in some layers (e.g., IP) can be masked by fast survivability procedures triggered in layers beneath (e.g., optical), and then the higher-layer service perceived by a client is not broken, allowing for treatment of such a situation as an uninterrupted working time.

### 3.2. Failure statistics

A network operator is risk-aware if it is able to credibly predict failure probabilities and their impact. Measurements of up- and downtimes are needed not only for direct computing of performance indicators but also for forming adequate modeling assumptions. Better models would help reliability prediction, and consequently risk assessment, of future networks. Two issues are challenging in this context:

- dependences between failure events,
- the non-Markovian character of failure processes.

Moreover, we assume that many large operators assume risk-aware policies that require high-availability networks. Therefore, single points of failure should be avoided. This raises the crucial problem of a proper modeling of multiple failures. So far, the most commonly applied approach to reliability considered failures as independent events despite there are some notable scientific contributions going beyond it, for instance the hazard potential approach [22]. Other generic models for correlated failures include Markovian models [23,24], a martingale approach to failure dynamics [25], and a recent general framework based on normal copulas [26]. The independence assumption is very useful for mathematical convenience to decrease complexity, but it is generally false. In fact, measurements [27,28] indicate that failures in communications networks may often be correlated. Overlooking this would result in dangerous overestimation of the actual reliability [29], which leads to undervaluation of the actual risk incurred.

Measurements on the operational Sprint network [27] indicate that about 30% of failures take place either simultaneously or within a few seconds of one another on different links. Assuming independent link failures would suggest that joint ones be extremely rare. In addition to sheer chance, reasons behind a multiple failure may fall in the following three categories:

- *Structural*: Two systems share a common service or component;
- *Dynamic*: A failure of one component increases the stress on another;
- *Epistemic*: The first failure remains unobserved until the second occurs.

The last case means that if network monitoring is not implemented carefully and thoroughly, single faults may go unnoticed, hidden by automatic recovery, and only multiple faults will lead to a system failure. As regards structural reasons, one can point to common equipment among providers, common physical infrastructure among carriers and between seemingly diverse access networks, common operation and maintenance activities among providers, and other commonalities along the market elements in Fig. 4.

Author's personal copy



The analysis of network failure data suggests that assuming failure processes to be Markovian—memoryless, possessing exponentially distributed up- and downtimes—may not be justified. Downtime durations of network elements have been reported to be sub-exponential or heavy-tailed. Thus, long failure durations are more frequent than what simple Markovian models suggest, and this should be taken into account during risk assessment in relation to the events' impact. A summary of statistical network failure modeling and analysis methods appearing in recent literature is presented in Table 2.

Unfortunately, appropriate failure measurement data are rarely publicly accessible, with [35] as one of few positive exceptions. Still, any up and downtime data are valuable in developing reliability modeling as it prompts for more practical models and reveals areas in which the current data collection needs to be improved. Even with simple data and network information it is possible to provide estimates of the experienced availability for network customers [30] widening the scope of network planning and management.

### 3.3. Loss estimation

Loss is a major metric to evaluate the impact of risk-incurring events on the service layer, and easily perceived by customers. There is a whole chain of losses at different levels in a communications network. At the bottom (the Layering-axis of Fig. 4), there are losses of data at the physical layer. But, because of unsuccessful recovery or lack of retransmission mechanisms, such losses propagate to higher layers, resulting in application losses, the ones that generate real harm to a service. The proper context for risk assessment is therefore located at the application layer. In fact, the same amount of traffic loss at the physical layer can result in different service disruptions for the customer: compare for instance a loss of single private voice call vs. emergency call. At the design stage, we can set an upper bound for losses under the reference design conditions, and obtain therefore a conservative estimate. Two groups of losses can be identified as follows.

- *Direct losses*: Related to the accumulated unfinished work due to failure occurrence [36,37], that is traffic lost by being sent to NULL interfaces or traffic that would be carried if a connection were not broken.[8] Exact modeling of this type of losses has rarely been used. Instead, the loss is assumed to be proportional to downtime and transfer rate. This is not an extensively studied phenomenon; as far as the authors know, [38] is the only paper that discusses this group of losses. Sometimes, experimental data are analysed to fit known distributions, as in [30].
- *Indirect losses*: Generated by secondary effects, e.g., when traffic flows are re-routed from failed paths and can cause congestion and buffer overflows on the paths that have received the re-routed traffic. Though the customer is not affected directly, QoS/QoE is decreased (a negative externality). This category of losses is typically not taken into account. Its evaluation, however, has to consider current and dynamic network conditions such as topology, load, and routing.

### 3.4. Risk metrics and networking business

Previously, we dealt with the statistical characteristics of network failures. However, a new dimension should be added to the reliability issue, by moving beyond the operational view and considering the business risk for the service providers, society and consumer activity associated with downtimes. Here we delineate a mathematical risk theory for networks, drawing from non-networking contexts.

A mathematical risk theory, providing models for the economical losses associated with adverse events, is well developed in the contexts of the finance [39] and insurance business sectors [40]. In the former, the aim is to analyze the variation of the value assigned to a portfolio of securities consisting of stocks or bonds, as the market conditions change. Downturn events may result in the fall of the value of a stock (the market risk) or in the default of a money borrower (the credit risk). On the other hand, in the case of insurance, the best known method for *risk transfer*, the aim is to evaluate the economical losses associated with an insurance policy, should the insured-against event take place (e.g., the loss of an asset or the occurrence of damages), and correspondingly, to set the insurance premium.

The aim of a mathematical risk theory for networking would be to provide mathematical tools to associate a risk measure with the overall set of downturns that may affect the network service, for instance: failures, indirect traffic losses, or malfunctioning. The introduction of a mathematical risk theory may also provide a different view of the protection means against such downturns.

We first consider the issue of direct risk metrics. Above all, the translation of the risk concept to the networking context calls for the identification of risk in this case. In the insurance/banking context the risk is naturally expressed in monetary units. This is not the case in the networking context. Here the risk is related to the failure to provide services as embodied in the operator–customer relationship. But such a failure is typically expressed through reliability-related events, and namely in network-centric units. For example, we are used to consider the network connectivity, the failure occurrence rate and its duration, the degradation in the Quality of Service (loss rate, packet delay). Therefore, we need to convert those reliability or QoS metrics into economical losses. Examples of network downturns that may be converted into monetary expressions are:

- the amount of lost traffic when the Quality of Service is degraded,
- the amount of lost traffic when the customer's connection is broken,
- the penalty paid by the operator to its customer (individual or institutional) under SLA or a threat of judiciary actions;
- the effort that has to be spent to restore the service, both as capital expenditures for purchase of new equipment to replace the failed one and as operational expenditures on a maintenance team.

A proposal to assess the economic value of some of these downturns is reported in [41] in the context of an economics-based traffic management system, where the penalties stated in SLA and the market price of leased lines are used to evaluate the economical damage associated with traffic loss. We can note that the relationship between reliability and economical losses is not a straightforward one: for instance, in [42] it has been shown that a system with a larger reliability level is not necessarily characterized by smaller losses caused by failures.

In network operations, losses are not isolated cases concerning a single point of failure and a single customer. They occur quite continuously (luckily, on a small scale, most of the time) and involve a number of customers each time. In the finance context the risk for a security owner is determined by the aggregation of the securities it holds, that is, its portfolio of securities, some of which may lose money and others do not. An analogous situation

[8] This amount is by definition unobservable, since it is not produced by customers who notice failures and stop sending traffic. Nevertheless, this unsent amount of traffic is not charged and represents a loss of potential revenues to an operator, thus also being a failure impact.

Author's personal copy



takes place in the banking context for the credit risk of a money lender, and in the insurance context for the portfolio of insurance policies held by an insuring company. A similar approach holds for the networking context as well, where we may consider a portfolio composed of customers/services (an approach proposed in [26]) and the losses associated with such a portfolio.

We need anyway a method to map the adverse events occurring on the network into a quantity characterizing the economic risk incurred by the network operator. As hinted above, this is actually a two-step process. First, we convert the network-centric measure of service disruption into an economic measure representing the losses associated to the service disruption, and then we compute a single metric summing up the overall effect of those losses. If we indicate the random variable representing the losses by $X$, an overall measure of risk on $X$ is a functional $\rho(X)$ that maps the probability distribution of the random variable $X$ into a positive real number. The most common measure of risk developed in the contexts recalled above is the *Value-at-Risk* (*VaR*) [39]:

$$VaR(\alpha) = F_X^{-1}(\alpha), \tag{1}$$

where $F_X^{-1}(\cdot)$ is the inverse probability distribution function of $X$. $VaR(\alpha)$ represents the maximum loss incurred with probability $\alpha$ (the probability that the loss is greater than the $VaR(\alpha)$ value is $1-\alpha$). Despite being widely used as a reference metric in many risk management environments, *VaR* has received a considerable criticism (see [43], where alternative measures of risk are also surveyed, and [44] for a more general critical analysis of risk measures). A major problem with *VaR* is that it does not consider the potential extreme losses exceeding the *VaR* itself. At least this shortcoming can be avoided if we employ the alternative metric *Tail Value-at-Risk*, also known as *Conditional Value-at-Risk* (*TVaR* or *CVaR*), which measures the expected value of the losses larger than *VaR* (hence, it is larger than *VaR*):

$$TVaR(\alpha) = \mathbb{E}[X \,|\, X \geq VaR(\alpha)] = \frac{1}{1-\alpha}\int_{\alpha}^{1} VaR(\xi)\, \mathrm{d}\xi. \tag{2}$$

When $\alpha$ is very close to 1, the resulting $TVaR(\alpha)$ or $VaR(\alpha)$ can be taken as proxies for the maximum loss to be expected. Both measures are shown for comparison in Fig. 5, where a sample probability density function of $X$ (e.g., the losses incurred during a time period of one month) is reported. The *Volatility* is a measure of the dispersion of losses around the average value (in practice, the standard deviation of losses in the reference period of time).

Both *VaR* and *TVaR* represent *univariate* measures of risk, since they provide a single value. They are the metrics of choice when all the risk facets can be aggregated, possibly after a preliminary conversion into the same unit of measure, e.g., monetary value. However, it may happen that different sources of risk need to be considered separately, so that the joint value of different risk components is of interest rather than the aggregated risk. Another situation arises when one wishes to consider risk components that are not amenable to a monetary expression, for instance by considering at the same time QoS, security, and monetary loss metrics. In those situations we may resort to multivariate risk measures, such as those analyzed for instance in [45].

Though so far we have relied a lot on metrics derived from the finance context, some caution needs to be exercised when translating those metrics in the networking area. A major difference is that the events of interest in finance or insurance industries are point events, such as the default of a company or the disaster occurrence to the insured company. Instead, in the networking context most events of interest have an associated duration (e.g., the time needed to repair a failure), and losses grow with the duration of the event. Therefore, we need to associate a time dimension to the risk measure, so that we should refer the Value-at-Risk to a specified time period, e.g., *VaR* over a year. An example of the application of the *VaR* metric in the networking context is provided in [26].

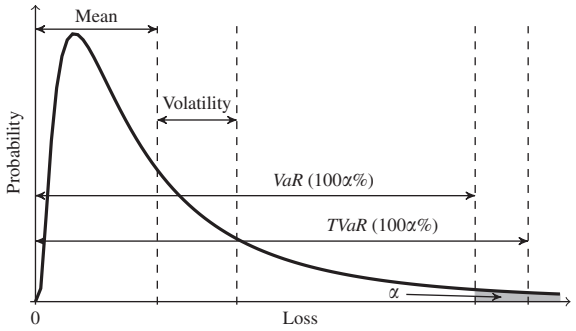


**Fig. 5.** Basic risk measures.

The evaluation of any risk metric in networks is, however, a difficult task for a number of reasons. In fact, as it was said in Section 3.3, a relevant component of the risk incurred during failures is the traffic loss when QoS or QoE is degraded or the connectivity is lost, with all associated problems. In addition, the computation of the risk measure often involves correlated variables. There are two reasons for this:

- many failures are correlated or depend on a common source of failures (see Section 3.2);
- even if correlated failures are neglected, risks associated to SLAs are correlated since they may refer to the same network region.

Finally, downturns are typically rare and their evaluation in a complex environment is likely to call for a simulation-based evaluation where we have to resort to variance reduction techniques.

### 3.5. Techniques for assessment integration

In addition to the mathematical techniques recalled in the previous sections, which allow us to evaluate the impact of risky events under precise modelling assumptions, there are semi-formal methods that may be useful in risk assessment even when the network system as a whole is not modeled formally. An example is represented by *dependability cases*, which are defined in [46] 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.' A basic structure of the dependability case consists of three key elements, whose chain of relationships can be illustrated as follows:

*Evidence* → *Arguments* → *Claims*

Those elements are characterized as:

- the available *evidence*, i.e., all kinds of documented facts about the network, including maintenance procedures, failure data etc.; and
- the stated goals, or *claims* about different aspects of dependability; they are usually subdivided into a hierarchy of sub-claims; as well as
- explicitly formulated *arguments*, which provide support to the claims on the basis of evidence.

The above scheme is an adaptation of the definition of so-called *safety cases* used in the safety assessment of large systems

Author's personal copy



like nuclear power plants [47,48]. A dependability case gathers dependability-related information on a complex system in one document (or document structure), organized according to the claim-argument-evidence logic.

Although an experimental dependability case of the Finnish University Network was reported in [49], network risk assessment has not been approached yet with the 'case' methodology, as far as we know. However, we would like to point out at this possibility for the future.

## 4. Risk response: recovery methods and service differentiation

We divide approaches to risk response into two groups. The first of them has a long tradition in designing a network so that it applies mechanisms that provide survivability to failures. What is new, is the emphasis laid on necessity for introduction of differentiated mechanisms suited for various types of services in order to properly address many levels of risks related to them. The second group has an economical character and paradoxically can be in many cases a quick win option for an operator that does not have sufficient resources to effectively use technical means or assumes additional methods of protecting its value chain.

### 4.1. Technical means

Unintentional outages due to failures represent the main risk in communications networks. Thus, *risk mitigation* is achieved mainly by making networks resilient (fault-tolerant, survivable) in the face of failures, through recovery mechanisms that automatically adopt redundancy (spare resources, backup) to switch the traffic affected by failures. Various recovery mechanisms can also be employed to introduce service differentiation. It represents a form of *risk sharing*, since the customer and the network operator may use the SLA to agree on the respective levels of responsibility in the presence of service degradation.

We can adopt either of two approaches to network resilience against failures: the *engineering* approach focused on the implementation of technically available mechanisms, and the *operations research* approach emphasizing optimization goals or the pursuit of mechanisms meeting requirements related to selected network services. The first approach can be treated as a bottom-up one: the choice of network technology limits the set of available recovery methods that can be used; then the resulting costs and quality are influenced by that choice. On the other hand, the application of the operations research approach, i.e., an approach not restricted by what is present in the equipment or standards, is top-down: first a designer assumes some constraints, for instance related to the quality, and then tries to select an optimized recovery method from a broad spectrum of possibilities. A typical goal is the minimization of cost. In this case, the set of recovery methods is usually loosely limited, and sometimes quite unrealistic options are treated as available (e.g., optical layer re-routing). Now, with the advent and implementation of sophisticated recovery methods, e.g., supported by the automatic control plane under Generalized Multiprotocol Label Switching in fixed networks, the engineering and operations research perspectives converge and enable the adoption of applicable cost- and risk-optimized recovery methods, making risk response feasible.

In order to facilitate the design process, several classifications of possible recovery methods have been introduced. They have a twofold meaning: (1) a presentation of the spectrum of possible methods to mitigate risk, i.e., an educational role; and (2) an indication of the degrees of freedom associated to each method, so as to allow for a preliminary rough design selection. The classifications typically adopt some rules of thumb that include cost and quality features, and allow for service differentiation (and the subsequent risk sharing). Here we report a quite comprehensive classification, which includes five dimensions [50]. That classification is also sketched in Fig. 6:

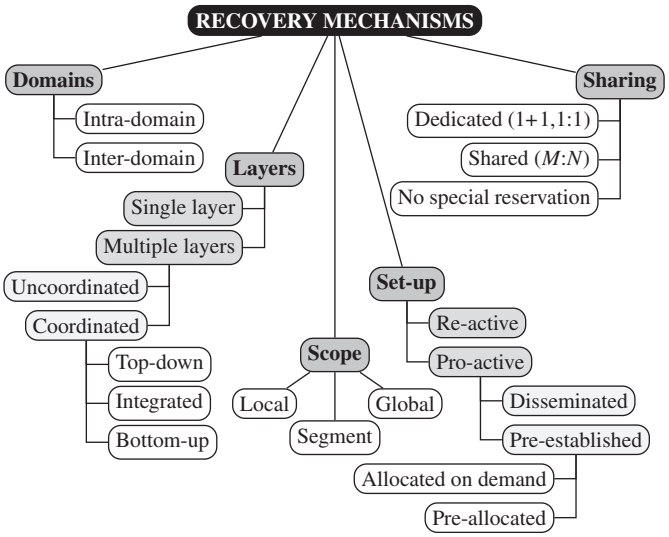


**Fig. 6.** Classification of network recovery mechanisms. The prefix *pre-* is used to mean that the process is performed before a failure occurs, i.e., with a pro-active approach.

(I) *Domains*, involved in recovery operation (see the Sectioning-axis in Fig. 4): (1) intra-domain (single domain): fast (i.e., decreasing impact), but enabling only local optimization, and (2) inter-domain (multiple domains): assuring global optimization, but slower, hardly enabling quality control, and prone to operator's resistance to reveal sensitive information.

(II) *Layers* in which the recovery operates (see the Layering-axis in Fig. 4): (1) *single layer* based, where the recovery may take place either at lower layers—fast but usually expensive—or at higher layers—slower but potentially cheaper and more flexible (enabling better service differentiation due to the involved connection granularity)—and (2) *multiple layers* based, with the actions on multiple layers being either uncoordinated—simple but costly—or coordinated—potentially cheaper, but complex.

(III) *Scope* of recovery, defining which part(s) of the end-to-end connection are to be recovered (the scope is typically limited to intra-domain recovery): (1) *local* (single link, node): fast but involving complex optimization and potentially expensive, (2) *global* (path, end-to-end): slower, easier to optimize, and (3) *segment*: intermediate between the two above.

(IV) *Set-up* method of recovery resources, which defines the timing relationship between the failure and the determination of the recovery action, with the two basic cases being: (1) *re-active*, computed on demand (known as *restoration* or *re-routing*), coming from the IP world, being flexible but slow, and (2) *pro-active*, pre-computed (called *protection*), which is robust and fast, but rather costly, typical for connection-oriented fixed telecommunications networks (SONET/SDH, MPLS), having many sub-types dependent on the technology applied.

(V) *Sharing* of recovery resources, which defines the degree of sharing of the redundant resources adopted for recovery: (1) *dedicated*: very costly but fast, (2) *shared*: quite robust

Author's personal copy



and with a reasonable cost, and (3) with *no special reservation* of resources, relevant for all types of re-active methods, flexible and cost-efficient but slow.

The options associated to each dimension employed in the classification allow for a very wide range of risk-aware choices, which affect both network design and management, with implications on several time-scales. Most typically, we have three levels of actions, with their own time horizon. We have *dimensioning* activities, which are accomplished with years-long perspectives. We have *routing* that allocates existing resources for time periods lasting months or days. And on the smallest time-scale we have *traffic engineering*, where decisions are taken hour-by-hour.

Since the engineering and the operations research approaches have somewhat different goals, their outcomes are expected to be diverse. While the former, directed at operational goals, aims at reducing disruption, the latter tries to minimize costs. The optimization solutions that consider thresholds on quality parameters agreed in an SLA are quite rare, although some works have recently been published [51–55]. However, the disruption is usually considered statically, and the aim is the minimization of the mean downtime or steady-state probability of unsuccessful recovery. But the perceived disruption, impacting QoE, is not the recovery time reduction itself, but either the simulated loss of traffic or the delay incurred due to failures. The divergence of results obtained through the two approaches stems also from the fact that cost-optimized methods tend to involve sharing and a larger scope of recovery (global or segment), while fast recovery, minimizing disruption, is achieved through dedication of resources and acts on a local level. Methods combining cost-effectiveness and limitation of failure impact should rearrange a connection after a performance threshold is exceeded (e.g., concerning the maximum number of failures), but they intrinsically involve difficult dynamic optimization algorithms. Additionally, when we consider the layer at which the recovery action takes place, higher layer methods are better from the cost viewpoint due to their finer granularity, while coarse lower layer methods are fast and expensive as they operate on bulky data but do not allow for differentiation. All those aspects should be taken into account by a risk-aware network designer.

In order to introduce a risk-aware approach and allow for risk sharing, it should be avoided to take into account just the operator's perspective and neglect the customer's perception of the service. In fact, the customer is directly affected by failures and is the party mostly interested in the correct evaluation of their impact. Risk assessment without taking into account the customer's role leads to a very partial view. Reliability, risk, and costs are linked by the service features agreed on in the SLA, which represents a relevant tool to enforce a fully risk-aware view by the network operator. For example, the customer's perspective can be included when a broad portfolio of service classes is considered, with different levels of resilience, and the associated various levels of risk sharing [56]. Although the existing SLAs are quite general from this viewpoint (they typically take into account just the steady-state availability), even the existing standardization and recommendation solutions envisage a very large set of metrics that can constitute a base for service differentiation. Aside from measures inspired by the reliability theory, we can find metrics as various as QoE or the number of concurrent failures [57]. With the adoption of a risk-aware approach, it is time to use such a broad set of metrics as a basis to set performance thresholds and define penalties for their violation in SLA.

### 4.2. Market means

A risk management approach to reliability allows us to consider risk mitigation and hedging strategies as accompanying the network-related ones, such as redundancy or fast recovery provisioning. In fact, we can adopt protection measures at a management level, not relying on the operator's own network. We can classify such *risk mitigation* measures under the following two categories:

1. *expansive strategies*, where we aim to preserve the revenue stream;
2. *protective strategies*, where we aim to recover the damage due to failures.

The former category includes those measures by which the operator keeps the service going for its customers, though relying on the networks of other operators rather than its own. If the operator keeps the service going, it keeps cashing on the traffic delivered for its customers, though it will have to pay the alternative operator for that. In order to have this alternative available at the sudden and unpredictable time of failure, the operator must have bought rights of usage on other networks in advance. Buying such rights on demand could not be possible, because the alternative network is not available, or could prove to be too expensive when the offering party uses its dominant position, taking advantage of the urgent need of the provider having outage problems in its network. A preventive purchase of usage rights is represented by the option for leasing spectrum examined in [58].

On the other hand, the operator may choose to accept the loss resulting from the failure of its network, but recovering at least part of its loss by subscribing to an insurance policy. In that case the operator pays a premium against network-related disasters, and receives the compensation on the occurrence of failures. An example of this strategy is the one against security risks on the Internet, proposed in [59]. An alternative form of protection is that guaranteed by the so-called CAT bonds (CAT is for 'catastrophe'), issued for natural disasters as well as for man-made and malicious attacks [60]. Here the time sequence of money exchanges is reversed with respect to classical insurance policy. In an insurance policy the network operator pays a premium upfront and receives a compensation if and when the insured-against event takes place. In CAT bond the operator issues a bond, which is bought by investors wishing to take on the risk faced by the operator. Though the operator receives the price paid by the investors upfront, it is then compelled to pay a periodic coupon to compensate them for their risk-taking. However, if the event that is insured against by the bond takes place before the bond's expiry time, the provider is not obliged to pay the principal back to the investors, who then suffer the risk related to the failure.

## 5. Risk monitoring and related practices

In this section we discuss some issues related to the operation of networks and provisioning of services that will have great impact on the risk associated with networking. Practical models should cope when confronted with some of the real problems:

- absence of single entity to control the network;
- limitation of the insight into underlaying failures and fault handling processes; or
- the network is under constant change and evolution, in fact times between equipment and topology changes in communications networks have the same order of magnitude as times between severe failures.

The most important issues are discussed below. Firstly, we emphasize the relationships of Internet market players. Secondly, we position the role of regulators in this market. Thirdly, practical aspects of network operation are discussed, with the emphasis

Author's personal copy



put on human factors. And at the end of the section, we give an overview of current practices and identify the challenges in collecting and using data for active risk monitoring.

### 5.1. Communications networking marketplace and regulators

Fig. 3 illustrates just a few of the inter-relationships for a tiny fraction of the network. The overall set of relationships is large and immensely complex, and has not been mapped yet, as far as the authors know. If we limit ourselves to view the Internet as an interconnection of Autonomous Systems (AS), a substantial effort has been spent to map the network of ASs, as for instance in [61–63]. However, even for this limited case, no one claims to have a complete map.

Hence, performing an adequate risk treatment of communications networks is extremely challenging, also because of missing information about:

1. the propagation of the consequences of network failures through value chains and peering relationships between market actors;
2. an extensive number of commonalities among market actors due to mutual provisioning of underlaying services, common infrastructure, common OA&M (Operations, Administrations, and Maintenance) procedures, etc.

An example of an attempt to help the end user manage the risk associated with these commonalities in a multi-provider, multi-technology mobile access setting is represented by the suggestion to extend the media independent handover databases of the IEEE 802.21 standard with information about equipment supporting the individual access points and their dependability characteristics [64]. Such an approach is, however, far from sufficient to allow an overall risk assessment of services provided in a complex market.

In fact, regulators enforce international and domestic laws to determine requirements on operators providing public service. Two examples of the regulators' activities are the EU Directive on universal service [65] at the European level, and the body of telecommunications regulations defined by the FCC (Federal Communications Commission) in the US [66]. The aim of such regulations is to ensure the availability to the end users of affordable, good quality, and future-oriented services in a competitive framework. We can note that enabling risk assessment and risk management should also be one of the aims. We expect such regulation to define the operators' obligations relevant to risk management. Note also that the requirements set by the regulators have significant impact on the ecosystem of networking companies, and thereby on the risk associated with the services provided. Such impact is not always a risk-reducing one. For instance, the regulation on operators with a significant market power reduces the duplication of infrastructure, contributing to the overall efficiency of the telecommunication system. But such reduction makes the infrastructure more liable to failures and less available, with a subsequently increased risk. On the other hand, the regulators might use means to manage the social impact through special requirements for services related to national safety and provided to prioritized users. For example, in the FCC's National Broadband Plan [67], the reliability and resiliency of communications networks are treated as important issues and will be addressed.

A number of ongoing activities in national and international bodies will have a significant impact on the risk associated with communications networks and services, e.g., by the FCC on network neutrality [68], JAIPA (Japan Internet Providers Association) on packet shaping [69], and ICANN (Internet Corporation for Assigned Names and Numbers) on Internet governance and addressing. Unfortunately, the risk issues are often not explicitly dealt with.

We must finally note that addressing the risks associated with communication networking in a societal context requires information about the network structure, about operational and commercial cooperation among service providers, and operational statistics to be available outside the individual market actors that originally own them. This is contrary to their current practice, which keeps much of this information confidential. Its diffusion will also cause additional costs, so that a resistance may be expected from the service providers themselves. Hence, if the public aim is to control those risks, a firm and decisive approach from national and international bodies is required. This must be followed up by the standardization bodies in defining which information shall be disseminated, how, and by which format.

### 5.2. Network design and operation practices

The usage of planning, operational and surveillance tools is quite often reflected by the maturity of the operator and, of course, by his economy. The tools depend on the technologies used and the services provided, as well as the operator's organization. The varieties of tools support all parts of the network, like infrastructure (e.g., buildings), the logical and physical network elements and structures, as well as service management and provision. In order to assess and monitor risks, we need the total view of the network structure coming from all the tools, and all the information obtained from the working equipment and operations support systems. Dimensioning, scalability and effective dependability of a network are the result of the balance among conflicting requirements, due for example to the expected market structure, the user behavior, economical aspects, or risk issues (both investment and operational expenses). Most operators use quite simple rules of thumb for dimensioning and design to get the desired level of robustness in their networks (recall for instance the classification of recovery mechanisms given in Section 4.1).

For risk-aware networking the insight into the operation and maintenance processes used by the network operator is important.

### 5.3. The human factor in operations

It is a common view in the networking domain that a high percentage of network failures is caused by human errors [70]. Therefore, it is highly relevant for risk-aware networking to understand, how human operators keep the 'invisible' infrastructure functioning. As one of very few studies in this area, we cite results of [71]. Twenty representatives of the staff operating the networks of a large national operator were interviewed and the answers were analyzed with the methodology described in [72]. The special work demands in high-tech environments that are intrinsically implied by features of the work domain can be generally classified as being related either to its (1) dynamics, (2) complexity, or (3) uncertainty. These three dimensions covered and structured well also the features of the present domain.

*Dynamics* appears due to the network nature itself, that is: (i) high frequency of faults and disturbances, implied by the size of the network; (ii) continuous renewal of technology; (iii) network growth caused by new services, and as requirements for fast action of the staff. A reason for urgency may be the scale of disturbance, the criticality of the failure, and a strict SLA. The criticality of a fault must be found out quickly, and actions are often needed before the root cause of the disturbance is identified.

*Complexity* appears first as related to technologies: the staff of a large operator has to master a large set of technical concepts,

Author's personal copy



products and versions. In addition to this, the operators must manage the historically produced complexity of the existing networks, including ownership and management relations.

*Uncertainty* is encountered on one hand as a technical constraint related for instance to the wear-out of hardware, hidden errors in software, differences in the implementation of standards and impossibility of exhaustive testing. On the other hand, the operators must often act on the basis of incomplete or flawed information, where additionally all effects of change at the network cannot be known beforehand. The uncertainty is escalated by activities changing the network where modifications are carried out simultaneously at several sites.

Erroneous acts were found at two levels: individual and organizational processes. As regards the work of an individual network staff member, haste, stress, handling several tasks simultaneously and night work were identified as factors increasing the vulnerability of work performance. Two error types identified in work performances were lapses and confusions, in particular during configuration and subcontractors works. The vulnerability of performance at organizational level seems to be caused by factors related to habit and culture, for example: meaningless repetitions, slackening of attention during work, neglecting of knowledge or instructions, as well as weakening of interest and motivation. In summary, the human factor in OA&M is an important risk factor and should be taken into account at the risk monitoring stage.

### 5.4. Data collection for risk monitoring

Collection of dependability-related information for risk monitoring is of utmost importance for network operators. Though all operators collect huge amounts of data, the collection is aimed neither at reliability prediction nor at risk monitoring and assessment. Rather, this data is used for network and service provisioning, management, and OA&M tasks. In addition, there is no common approach in failure data collection among network operators and service providers. And, naturally, this data is not made available in the public domain. Hereafter, we provide a general description of the current state of data collection, based on our insight, and some guidelines on its requirements under a risk-aware approach.

Fig. 7 is a simplified sketch of the risk-related data flow between networking market players and the users. The interrelationships between the network operators and service providers illustrated in Fig. 3 are depicted in Fig. 7 with the cloud crossing several network providers. The services are internal (delivered by the operator to its own end users) as well as external (provided to other operators).

As said in Section 3.3, service failures are a major source of risk in communications networks. The impact of failures for the service provider is determined by their multi-faceted consequences: lost traffic and corresponding revenues, penalties incurred by not meeting the SLA, or other indirect costs related to the operator's reputation. However, in a societal context, the impact of communication failures may be far larger than that suffered by the provider of the service. In Fig. 7 the impact is shown spanning several network providers and the society. Consequences beyond the scope of the operator are very difficult to assess, since they are tightly linked to individual users and the context in which the services are used. The expenses associated with impact, and its measurement, may differ significantly among the users. Furthermore, not all consequences can be measured by monetary terms.

Under a risk-aware approach, statistics of performance and faults in the network serve as the basis for service monitoring and strategic decisions on company operations, like the re-assessment of the risk framework. In many cases the collection of those statistics is defined by regulations, contracts, or agreements. Several sources

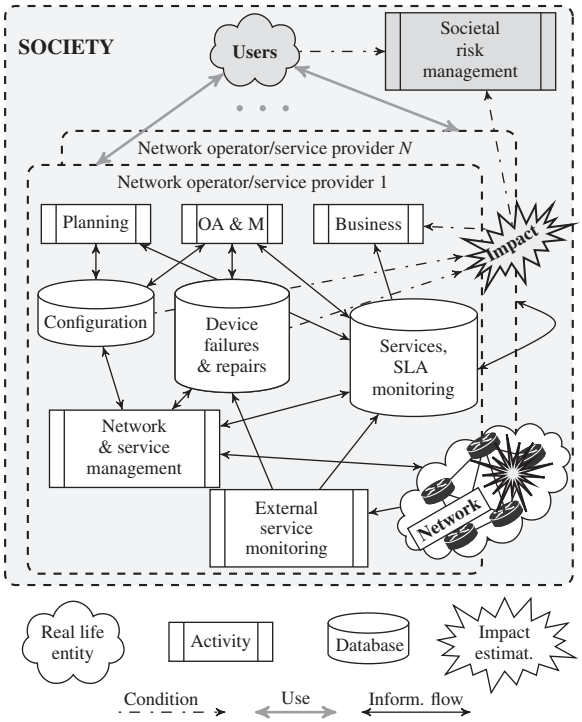


**Fig. 7.** Network risk related data flow.

collectively separated into internal and external ones are used for data collection. The configuration databases represent the physical and logical network and service provisioning of the planned and current state of the network. They must be maintained to make the information on network topology accessible and to reflect the interdependencies among resources in the network. The performance and fault statistics are stored in provisioning databases and provide information on the behavior of the network over time. Not only the network internal observations are sources for such data but also interconnected networks and customers are relevant data sources. Careful attention must be paid to which data should be collected and stored, to balance the needed information against the amount of data to be stored and analyzed.

Using the collected data for risk monitoring bears considerable challenges. As noted in Section 3.2, dependencies between failure events are likely and the failure processes have a non-Markovian character. If the correlations between failures are to be analyzed, a lot of information is needed, including network topology, traffic handling or recovery mechanisms. Unfortunately, such information is fragmented among different data collection systems stemming from various operational requirements. In addition, data from different market actors must be considered jointly.

A minimal *failure event record* contains the following information pieces about the event:

- equipment identifier,
- start time of the failure, and
- repair time of the failure.

In order to analyze the failures and possible correlations, the event data must be put into the context in which the event takes place. The context (network topology, etc.) should also be stored. The role of meta-data, that is the data crucial for the correct interpretation of the fault data, is underestimated and it is seldom stored in a comprehensive way. To properly address risk, one needs to have information about the process that gives rise to the failure statistics.

Author's personal copy



**Table 3**
Goals on the way to risk-aware networking.

| Objective | Features, indicators | Baseline | Target | Intermediary steps |
|---|---|---|---|---|
| *Design, planning & assessment taking into account risk-awareness* | Assumed level of risk | Qualitative treatment of risk at best | Risk (event-frequency-impact) in a goal function | Extension of a set of parameters involved in network design and SLA construction, risk as a constraint |
| *Proper risk assessment* | Used reliability metrics | Availability as an implicit measure of risk/loss (in the services context) | A set of explicit measures of risk (frequency of events and their severity, along with the assessment of their uncertainty) | Definition of the relation between network reliability and risk assessment |
| | | Mainly connectivity assessment<br>Loss measured at the traffic level | Quantification of the network ability to provide services with required QoS levels<br>Loss assessed at the service level | Inclusion of QoS/QoE measures in reliability assessment<br>Development of proper loss models taking into account layering and indirect impact |
| | Used risk metrics | Risk expressed implicitly as selected reliability metrics | Rich set of explicit risk metrics actively applied and induced by business and societal conditions | Definition of adequate risk measures for communications networking |
| *Risk-aware data collection* | Analysis-friendly collection of failure data | Detailed failure data utilized almost only re-actively, for repair purposes | Analysis of failure data as a continuous activity; results are used in risk assessment and network design | Working on existing data to improve monitoring and to reveal improper assumptions related to risk assessment |
| | Modeling of dependences | Basing on the independence assumption | Correlations between failures taken into consideration | Collection of more detailed data |

If network operators provide *incident information*, such as classification, root cause identification or priority, this data needs to be linked to the failure events. Such incident information is obtained under the most complete knowledge[9] and it cannot be reconstructed afterwards.

In data collection, one needs also to be aware of certain challenges that may be encountered in the analysis of the reliability data. One issue is that the most interesting data may be difficult to distinguish from erroneous information. For example, if downtime durations indicate heavy-tailed distribution (see Section 3.2), then the longest observed downtimes are very important observations. However, a long downtime duration may also result from an error in parsing the event information, or just from missing data. A fundamental problem is that network monitoring is typically implemented by using the very same network it is monitoring: its performance degrades if the underlying network degrades.

Physical equipment, software modules and system configurations are replaced, modified or changed quite frequently to meet traffic demands and functional requirements. This results in short unaltered periods of the systems compared to the time constants in their failure processes. Similarly, for cost efficiency and competitiveness, operation and maintenance have to adapt to new equipment/technology and the services provided. These continuous changes add challenges to data analysis and risk management.

For the assessment of correlation and interdependencies, so important in risk analysis, the physical network topology can be extracted from the various databases, but the logical topology might be difficult to obtain. This makes it difficult to maintain a correct network view to correlate failures triggered by error propagation of common conditions/events. There are several commercial systems on the market to perform tasks in data collection and in failure correlation, such as for example Hewlett–Packard OpenView and TeMIP or IBM Tivoli. However, none of them makes data easily accessible for risk assessment or enables (semi)automated risk management.

## 6. Conclusions

Throughout this work, we have characterized, with varying levels of maturity, risk-related approaches for the design and operation of dependable networks. We claim that the basic needs to be addressed by the research community to serve the industry, business, and society are the following:

*Establishment of studies on risk-aware communications networks*. Table 3 presents different networking aspects from this standpoint.

*Multi-level approach to risk management* (*e.g.*, *multi-carrier*, *multi-technology*, *multi-service*, *multi-metric*). Risk-awareness is not limited only to the design of the operator's own infrastructure. It must take into account different building blocks of the service, where various dimensions are represented. For example, operational aspects to be covered involve: roaming agreements, multi-homing issues, CAT bonds usage, or even maintenance of the user equipment as a part of the network.[10]

*Complex value chain*. The proper risk-aware design of networks involves a very large set of interacting partners. The complexity of those interactions makes it difficult to determine the details of their cooperation, especially as it is desirable that the risk-aware attitude is adopted not only by an individual player, but by the whole community of operators. The issue is hindered by the fact that risk-awareness is related to the steadily ongoing cyclic process of risk framing, assessment, response, and monitoring. Such difficulties are reflected, for example, in the definition of SLAs, where we have a variety of approaches to service provisioning guarantees against an abstract umbrella that is determined by regulatory or standardization bodies. Introduction of risk-

[9] No pre-programmed intelligence can fully detect anomalies or root causes as this may require information about abnormal events in the outside world at that very moment.

[10] A paradoxical example of the last issue is the massive provision of software patches to users of mobile phones of a popular brand, which had to be accomplished by the operator, rather than by the equipment manufacturer, to avoid risk of misconfiguration and loss of services to the customers.

Author's personal copy



awareness would involve emphasis also on the development of more sophisticated SLAs when necessary.

## Acknowledgments

The authors would like to thank Andrzej Jajszczyk, Przemysław Pawełczak, and Rafał Stankiewicz for their help while working on this paper.

## References

[1] Lowrance WW. Of acceptable risk: science and the determination of safety Los Altos, CA: William Kaufmann Inc.; 1976.
[2] Gheorghe AV, Masera M, Weijnen M, DeVries LJ. Critical infrastructures at risk. Securing the European Electric Power System, vol. 9 of topics in safety, risk, reliability and quality. Dordrecht, The Netherlands: Springer; 2006.
[3] Haimes YY. Risk modeling, assessment, and management. Hoboken, NJ: John Wiley & Sons, Inc.; 2009.
[4] Aven T. Quantitative risk assessment the scientific platform. Cambridge, UK: Cambridge University Press; 2011.
[5] Rausand M. Risk assessment. Theory, methods, and applications. Hoboken, NJ: John Wiley & Sons, Inc.; 2011.
[6] Aven T. On how to define, understand and describe risk. Reliability Engineering and System Safety 2010;95(6):623–31.
[7] Grøtan TO, Størseth F, Albrechtsen E. Scientific foundations of addressing risk in complex and dynamic environments. Reliability Engineering and System Safety 2011;96(6):706–12.
[8] Utne IB, Hokstad P, Vatn J. A method for risk modeling of interdependencies in critical infrastructures. Reliability Engineering and System Safety 2011;96(6): 671–8.
[9] Tømmerår J, Aven T. A framework for reliability and risk centered maintenance. Reliability Engineering and System Safety 2011;96(2):324–31.
[10] Levitin G, Hausken K, Taboada HA, Coit DW. Data survivability vs. security in information systems. Reliability Engineering and System Safety 2012;100: 19–27.
[11] Brush G, Marlow N. Assuring the dependability of telecommunications networks and services. IEEE Network 1990;4(1):29–34.
[12] Siewiorek DP, Chillarege R, Kalbarczyk ZT. Reflections on industry trends and experimental research in dependability. IEEE Transactions on Dependable and Secure Computing 2004;1(2):109–27.
[13] Kaplan S, Garrick BJ. On the quantitative definition of risk. Risk Analysis 1981;1(1):11–28.
[14] Helvik BE. Perspectives on the dependability of networks and services. Telektronikk 2004;100(3):27–44.
[15] Wheeler E. Security risk management. Waltham, MA: Syngress; 2011.
[16] Managing Information Security Risk. Organization, Mission, and Information System View. NIST Special Publication 800-39; 2011.
[17] Newman M, Barabási AL, Watts DJ. The structure and dynamics of networks. Princeton, NJ: Princeton University Press; 2006.
[18] Norros I, Reittu H. Network models with a 'soft hierarchy': a random graph construction with loglog scalability. IEEE Network 2008;22(2):40–6.
[19] Checkland P. Systems thinking systems practice: includes a 30-year retrospective Inc., New York, NY: John Wiley & Sons; 1999.
[20] O'Brien KJ. Ericsson and Nokia Siemens are managing just fine. The New York Times 2009; 〈http://www.nytimes.com/2009/04/13/technology/companies/13iht-network.ready.html?_r=2〉.
[21] Chołda P, Mykkeltveit A, Helvik BE, Jajszczyk A. Continuity-based resilient communication. In: Proceedings of the 7th international workshop on the design of reliable communication networks, DRCN 2009. Washington, D.C.; October 25–28, 2009.
[22] Brady K, Chandra J, Cui Y, Singpurwalla ND. Hazard potentials and dependent network failures. In: Proceedings of the 33rd Hawaii international conference on system sciences HICSS-33. Wailea Maui, HI; January 4–7, 2000.
[23] Spragins J. Dependent failures in data communication systems. IEEE Transactions on Communications 1977;COM-25(12):1494–9.
[24] Hecht M, Tang D, Hecht H, Brill RW. Quantitative reliability and availability assessment for critical systems including software. In: Proceedings of the 12th annual conference on computer assurance COMPASS'97. Gaithersburg, MD; June 16–19, 1997.
[25] Arjas E, Norros I. Stochastic order and martingale dynamics in multivariate life length models: a reviewIn: Mosler K, Scarsini M, editors. Stochastic orders and decision under risk. IMS lecture notes—monograph series, vol. 19. Hayward, CA: Institute of Mathematical Statistics; 1991. p. 7–24.
[26] Naldi M, D'Acquisto G. A normal copula model for the economic risk analysis of correlated failures in communications networks. Journal of Universal Computer Science 2008;14(5):786–99.
[27] Markopoulou A, Iannaccone G, Bhattacharyya S, Chuah CN, Ganjali Y, Diot C. Characterization of failures in an operational IP backbone network. IEEE/ACM Transactions on Networking 2008;16(4):749–62.
[28] Gonzalez AJ, Helvik BE, Hellan JK, Kuusela P. Analysis of dependencies between failures in the UNINETT IP backbone network. In: Proceedings of the 16th pacific rim international symposium on dependable computing, PRDC 2010. Tokyo, Japan; December 13–15, 2010.
[29] Johnson DM. QoS control versus generous dimensioning. BT Technology Journal 2005;23(2):81–96.
[30] Kuusela P, Norros I. On/off process modeling of IP network failures. In: Proceedings of the 40th annual IEEE/IFIP international conference on dependable systems and networks, DSN 2010. Chicago, IL; June 28–July 1, 2010.
[31] Gonzalez AJ, Helvik BE. Analysis of failures characteristics in the UNINETT IP backbone network. International Journal of Space-Based and Situated Computing 2012;2(1):3–11.
[32] Choi BY, Song S, Koffler G, Medhi D. Outage analysis of a university campus network. In: Proceedings of the 16th international conference on computer communications and networks, ICCCN 2007. Honolulu, HI; August 13–16, 2007.
[33] Matz SM, Votta LG, Malkawi M. Analysis of failure and recovery rates in a wireless telecommunications system. In: Proceedings of the 2002 international conference on dependable systems and networks, DSN 2002. Bethesda, MD; June 23–26, 2002.
[34] Lepropre J, Leduc G. Inferring groups of correlated failures. In: Proceedings of the 2nd conference on future networking technologies CoNEXT'06. Lisbon, Portugal; December 4–7, 2006.
[35] The Norwegian Research Network UNINETT: Downtime Statistics; 2011. 〈http://drift.uninett.no/downs/〉.
[36] Jæger B, Tipper D. Prioritized traffic restoration in connection oriented QoS based networks. Computer Communications 2003;26(18):2025–36.
[37] Heegaard PE, Trivedi KS. Network survivability modeling. Computer Networks 2009;53(8):1215–34.
[38] Gan Q, Helvik BE, Wittner O. Refined classification of service unavailability for comparison: shared path protection vs. rerouting. In: Proceedings of the 2006 international conference on communication technology, ICCT 2006. Guilin, China; November 27–30, 2006.
[39] McNeil AJ, Frey R, Embrechts P. Quantitative risk management: concepts, techniques and tools. Princeton, NJ: Princeton University Press; 2005.
[40] Denuit M, Charpentier A, editors. Mathématiques de l'Assurance Non-Vie; vol. 1, Principes Fondamentaux de Théorie du Risque. Paris, France: Economica; 2004.
[41] Iovanna P, Naldi M, Sabella R, Zema C. Economics-driven short-term traffic management in MPLS-based self-adaptive networks. In: Proceedings of the IFIP 4th international workshop on self-organizing systems, IWSOS 2009. Zurich, Switzerland; December 9–11, 2009.
[42] Todinov MT. Reliability analysis based on the losses from failures. Risk Analysis 2006;26(2):311–35.
[43] Szegö G. Measures of risk. European Journal of Operational Research 2005;163(1):5–19.
[44] Aven T. Misconceptions of risk, statistics in practice. Chichester, UK: John Wiley & Sons, Inc.; 2010.
[45] Cascos I, Molchanov I. Multivariate risks and depth-trimmed regions. Finance and Stochastics 2007;11(3):373–97.
[46] Despotou G, Kelly T. Extending the safety case concept to address dependability. In: Proceedings of the 22nd international system safety conference, ISSC 2004. Providence, RI; August 2–6, 2004.
[47] Bishop P, Bloomfield R. A methodology for safety case development. In: Redmill F, Anderson T, editors. Industrial perspectives of safety-critical systems: proceedings of the 6th safety-critical systems symposium. Springer-Verlag, Birmingham; 1998. p. 194–203.
[48] Kelly TP. Arguing safety. A systematic approach to safety case management. PhD thesis. York, UK: Department of Computer Science, University of York; 1998.
[49] Norros I, Kuusela P, Savola P. A dependability case approach to the assessment of IP networks. In: Proceedings of the 2nd international conference on emerging security information, systems and technologies , SECUWARE 2008. Cap Esterel, France; August 25–31, 2008.
[50] Chołda P, Jajszczyk A. Recovery and its quality in multilayer networks. IEEE/OSA Journal of Lightwave Technology 2010;28(4):372–89.
[51] Mykkeltveit A, Helvik BE. Adaptive management of connections to meet availability guarantees in SLAs. In: Proceedings of the IFIP/IEEE 11th international symposium on integrated network management, IM 2009. Long Island, NY; June 1–5, 2009.
[52] Bailey SR. Disaster preparedness and resiliency. In: Kalmanek CR, Misra S, Yang YR, editors. Guide to reliable internet services and applications. London, UK: Springer-Verlag Ltd; 2010. p. 517–43 [chapter 14].
[53] Xia M, Tornatore M, Martel CU, Mukherjee B. Risk-aware provisioning for optical WDM mesh networks. IEEE/ACM Transactions on Networking 2011;19(3):921–31.
[54] Ogino N, Nakamura H. Telecommunications network planning method based on probabilistic risk assessment. IEICE Transactions on Communications 2011;E94-B(12):3459–70.
[55] Vajanapoom K, Tipper D, Akavipat S. Risk based resilient network design. Telecommunication Systems, http://dx.doi.org/10.1007/s11235-011-9578-1, in press.
[56] Chołda P, Mykkeltveit A, Helvik BE, Wittner OJ, Jajszczyk A. A survey of resilience differentiation frameworks in communication networks. IEEE Communications Surveys & Tutorials 2007;9(4):32–55.
[57] Chołda P, Tapolcai J, Cinkler T, Wajda K, Jajszczyk A. Quality of resilience as a network reliability characterization tool. IEEE Network 2009;23(2):11–9.
[58] Mastroeni L, Naldi M. Options and overbooking strategy in the management of wireless spectrum. Telecommunication Systems 2011;48(1–2):31–42.

Author's personal copy



[59] Bolot JC, Lelarge M. A new perspective on internet security using insurance. In: Proceedings of the 27th IEEE conference on computer communications, INFOCOM 2008. Phoenix, AZ; April 15–17, 2008.
[60] Kunreuther H. The role of insurance in managing extreme events: implications for terrorism coverage. Business Economics 2002;37(2):6–16.
[61] Merit Network, Inc.: Internet Routing Registry; 2011. 〈http://www.irr.net/〉.
[62] The Cooperative Association for Internet Data Analysis CAIDA: Macroscopic Internet Topology Data Kit (ITDK); 2011. URL 〈http://www.caida.org/data/active/internet-topology-data-kit/〉.
[63] Kernen T. Public Route Server and Looking Glass List; 2011. URL 〈http://www.traceroute.org/〉.
[64] Følstad EL, Helvik BE. Managing availability in wireless inter domain access. In: Proceedings of the first international workshop on reliable networks design and modeling, RNDM 2009. St. Petersburg, Russia; October 12–14, 2009.
[65] The European Parliament Council: Directive on Universal Service and Users' Rights Relating to Electronic Communications Networks and Services; 2002.
[66] Text of Code of Federal Regulations/e-CFR: Title 47 Telecommunication; 2011.
[67] Federal Communications Commission: Connecting America: The National Broadband Plan. Directive 2002/22/EC; 2010.
[68] Federal Communications Commission: Internet Policy Statement; 2005.
[69] Japan Internet Providers Association (JAIPA), Telecommunications Carriers Association (TCA), Telecom Services Association (TELESA) and Japan Cable and Telecommunications Association (JCTA): Guideline for Packet Shaping; 2008.
[70] Kuhn DR. Sources of failure in the public switched telephone network. IEEE Computer 1997;30(4):31–6.
[71] Norros L, Norros I, Liinasuo M, Seppänen K. Impact of human operators on communication network dependability. Cognition, Technology & Work, http://dx.doi.org/10.1007/s10111-012-0225-8, in press.
[72] Norros L. Acting under uncertainty. the core-task analysis in ecological study of work. Espoo, Finland: VTT Technical Research Centre of Finland; 2004. VTT Publications: 546.