Vulnerability analysis of the physical part of the internet

Ling Zhou

Laboratory for Safety Analysis,
Swiss Federal Institute of Technology Zurich (ETH),
Sonneggstrasse 3, 8092 Zurich, Switzerland
E-mail: zhou@mavt.ethz.ch

Abstract: The internet, one of the most important critical information infrastructures (CII), compasses both, the physical and cyber part. Industrial network operators and service providers often draw more attention to the protection of the cyber part. However, some recent incidents, especially those causing cascading failures and big disruptions in social functions, are reminders that physical part may be most vulnerable and deserve adequate protection efforts. Therefore, this paper put the vulnerability of the physical part of the internet in focus. Most important outage causes of the internet backbones, like cable cuts and (inter-)dependency among different networks, are presented. Analysis approaches based on network theory are introduced and applied in a case study to the Swiss research and educational network (SWITCH). Recommendations on good practices are offered finally.

Keywords: critical infrastructures; internet; critical information infrastructure protection; CIIP; internet physical part; vulnerability analysis; network theory; optical networks; internet backbones; outage causes; traffic simulation; cable cuts.

Reference to this paper should be made as follows: Zhou, L. (2010) 'Vulnerability analysis of the physical part of the internet', *Int. J. Critical Infrastructures*, Vol. 6, No. 4, pp.402–420.

Biographical notes: L. Zhou is a Research Associate and a Post Doctoral at the Laboratory for Safety Analysis of ETH. She received her PhD in Availability Analysis and Optimization in Optical Transport Networks from the Department of Information Technology and Electrical Engineering of ETH. Her research interests include reliability and vulnerability analyses for critical infrastructures, especially for information and communication systems.

1 Introduction

Effective functioning of today's societies is based on critical infrastructures (CI), i.e., large scale infrastructures whose degradation, disruption or destruction would have a serious impact on health, safety, security or well-being of citizens or the effective functioning of governments and/or economy. They are usually summed up in critical sectors, such as energy, finance and health sectors. All CI widely use information and communication technologies (ICT) and strongly depend on them. Information processes supported by ICT that are CI for themselves, or critical for the operation of other CI, are

called critical information infrastructures (CII) (Białas, 2006). Owners and operators play a major role in critical information infrastructure protection (CIIP). In general, CII stakeholder (e.g., network operators and infrastructure providers) involvement appears largely deficient. They show a vision strongly related to their own infrastructure and business framework, with a limited attention on inter- and cross-sector interface elements and trans-domain consequences (Bologna et al., 2006). CIIP is a new challenge for today's societies whose functioning is based on ICT.

The project of research is about 'CI protection', of which vulnerability analysis of CII in Switzerland is of major interest. In a first step, the information and telecommunications systems, especially the internet backbone, is in focus of attention.

The internet, as a global system of interconnected computer networks including millions of private, public, academic, business, and government networks, is one of the most important CII and in the focus of this paper. The internet consists of physical components (e.g., physical cables, servers, bridges and hubs, routers, personal computers), network components [such as domain name servers (DNS)] and information components (operating systems, network software, databases, web servers and browsers) (IRGC, 2006). In short, there are both physical and cyber part in the internet. However, more attention is often drawn to the protection of cyber part, e.g., ICT security or cyber-security, by industrial network operators and service providers. The private sector running the infrastructures perceives the risk mainly as a local, technical problem or in terms of economic costs, because the owners and operators are in a position to install technical safeguards for ICT security at the level of individual infrastructure (Brunner and Suter, 2008). To ensure the socially desirable functions of CII, the physical part also deserves adequate protection efforts. Protecting fibre cables from physical damage is critical for the dependability and reliability of the telecommunications network, because of ever increasing amounts of traffic over single fibre cables.

In 2003, PhD candidate Sean Gorman famously mapped US fibre-optic paths for his dissertation at George Mason University and found it was easy to locate critical choke points from public records and data. More and more of the nation's critical communications merge into fewer and fewer fibre-optic cables. Carriers do not want to spend money to run redundant fibre-optic lines. Such geographic limitations have spawned another dangerous trend. Different companies tend to install their cables alongside the same limited number of roads and railways, often unknowingly (Poulsen, 2006). On the other hand, city planning boards require operators to use the same cable duct in order to avoid construction works all over the city. All this brings critical communication backbones into a potentially unsafe condition.

For instance, a coincidence of weather damage and a construction accident knocked Sprint's Western US network out of service for three and a half hours on the afternoon of January 9, 2006. The outage occurred following a fibre cut due to a dig-up by a backhoe near Phoenix, Ariz., which took place hours after traffic from a flooded area of California had been routed through Phoenix to enable emergency maintenance (Wilson, 2006). Experts say, "Sprint outage is a reminder that with all the attention paid to computer viruses and the latest Windows security holes, the most vulnerable threads in US's critical infrastructures lie literally beneath our feet" (Poulsen, 2006). Deliberate saboteurs or terrorists may make big network outages with some rented backhoes and careful target selection. In practice, dual-span failures may occur accidentally more often than expected due to frequent cable cuts, span maintenance or upgrade operations, and shared risk span

groups. Dual-span failures may prevent traffic rerouting around a ring, which is a general topology of long haul networks, and cause network outages.

Network vulnerability receives attention from many fields, particularly those dealing with network design since it is required to embed vulnerability measurements in designs for reliable or survivable networks. Motter and Lai (2002) demonstrate that the heterogeneity of different complex networks makes them particularly vulnerable to attacks in that a large-scale cascade may be triggered by disabling a single key node. Gorman et al. (2004) focus on computer data networks and the spatial implications of their susceptibility to targeted attacks. Simulations are run to determine the repercussions of targeted attacks and what the relative merits of different methods of identifying critical nodes are. Grubesic and Murray (2006) explore the topological complexities associated with network interconnections. The loss of vital nodes, like telecommunication switching centres, is evaluated utilising a developed spatial optimisation model.

This paper focuses on the possibly underestimated vulnerability of the physical part of CII, especially on the internet backbones. Causes of network outages and performance problems of the internet backbones are given in Section 2. Topological methods based on network theory are introduced in Section 3. A case study of SWITCH is given in Section 4, followed by a summary and outlook.

2 Outage causes of the internet backbones

The internet backbone refers to the principal data routes between large, strategically interconnected networks and core routers in the internet. These data routes are hosted by commercial, government, academic and other high-capacity network centres, the internet exchange points and network access points that interchange internet traffic between the countries, continents and across the oceans of the world.

Today's internet backbones are provisioned to provide excellent performance. However, performance degradation and service disruption are likely in the case of failure, such as human errors (misconfiguration), power failures, cable cuts, hardware failures, congestion, attacks, software bugs, etc. From these causes, it is easy to find that, human errors, power failures and physical failures account for a big part of the outage problems.

A report from the Alliance for Telecommunications Industry Solutions found that cable dig-ups were the single most common cause of telecom outages over a 12-year period ending in 2004 with the number of incidents dropping in recent years but with the severity and duration of the outages increasing (Poulsen, 2006). In 2002, the Federal Communications Commission of the USA (FCC) statistics showed that metro networks experience 13 cuts annually and long haul networks experience three cuts for 1,000 miles of fibre. Even the lower rate for long haul networks implies a cable cut every four days on average in a network with 30,000 route-miles of fibre (Grover, 2004; Vernon and Portier, 2002). Cable cuts may occur much more frequently than people expect. Most cable cuts occur accidentally as a result of construction and utility crew activity, making it nearly impossible to predict when and where cable cuts might occur (Hoffman, 2002).

From the beginning of 2005, the availability of outage records is prohibited to the public according to new FCC regulations. The lack of failure data from operational networks has further limited the investigation of failures in the internet backbones. However, considering the frequency and impacts of cable cuts, to identify most vulnerable cables and make an in-depth understanding becomes extremely valuable.

While construction crews are often to blame for cable cuts, accidents and acts of nature account for nearly all of the remaining causes of cable cuts. Cable cuts have reportedly been caused by vehicles running into aerial poles, fires, stray bullets, floods and fallen trees. Animals have been implicated in cable cuts. Hunters have also been reported to be responsible for cable outages. In September 2008, the fibre-optic lines delivering cable, internet and phone services in Southern Indiana were severely broken, due to dove hunters shooting at doves perching on the fibre-optic lines. Although a repair personnel was deployed immediately, it took much longer than originally expected before services were restored, because the hunters had shot the cable in more than one place (McNamara, 2008).

The other important cause of network outages can be rooted in a feature of CI, i.e., (inter-)dependency. Many infrastructures are interdependent or depend on a host technology, such as electric power or ICT. Initial disturbances or even destruction in parts of one CI, may lead to failure in the infrastructure itself or/and in the other (inter-)dependent CI, triggering a disruptive avalanche of cascading and escalating failures (Vespignani, 2010). While modern societies depend heavily on the proper functioning of CI, the societal disruption caused by infrastructure failures is therefore, disproportionately high in relation to actual physical damage (Chang, 2009). Therefore, the most dangerous vulnerability may hide in the (inter-)dependencies across different CI. Cascading failures are common in most of the complex electricity, communication and/or transportation networks that are the basic components of our lives and industry.

The 'mini telecommunication blackout' in 2004, Rome, typically shows the (inter-)dependency among different infrastructures. On January 2, 2004, a cooling plant pipe leak caused flooding of a telecommunication centre of Telecom Italia in Rome. Several boards and devices failed for short circuit and storage batteries dropped out of their capacities, too. Part of wired and wireless services tilted, causing problems and delays in different infrastructures, including Fiumicino Airport (stop of check-in, ticketing services and of luggage acceptance and switching), ANSI print agency, post offices and banks, ACEA power distribution and the communication network (GARR), connecting the main Italian research institutions (Ciancamerla et al., 2007; Bonanni et al., 2008).

Although most failures emerge and dissolve locally, largely unnoticed by the rest of the world, a few trigger avalanche mechanisms that can have large impacts on the entire networks. The presence of spare capacity, which can absorb unanticipated stresses or unexpected demands, can make a system more robust. However, greater pressures for economic efficiency and lower price of service have often squeezed out redundancy in systems, resulting in tighter operational margins. Vulnerability can also arise when separate smaller systems are integrated into a larger system, thus, creating greater complexity and the increased potential for trans-boundary propagation of disturbances, both of which have occurred in the case of the electric power grid, as in the 2003 blackout in the Northeastern USA and Southeastern Canada.

In many realistic situations, the flow of physical quantities in the network, as characterised by the loads on nodes, is important. For such networks where loads can redistribute among the nodes, intentional attacks can lead to a cascade of overload failures, which can in turn cause the entire or a substantial part of the network to collapse. This is relevant for real-world networks that possess a highly heterogeneous distribution of loads, such as the internet and power grids. In the case of the internet, traffic is

rerouted to bypass malfunctioning routers, eventually leading to an avalanche of overloads on other routers that are not equipped to handle extra traffic. The redistribution of the traffic can result in a congestion regime with a large drop in the performance (Crucitti et al., 2004).

3 Vulnerability analysis of optical networks

According to the approaches described by Eusgeld et al. (2009), statistics and simulation based on network theory are applied in the screening phase of vulnerability analysis. Both methods can identify vulnerable parts of a network directly from various aspects. In the case study, risk to the physical networks, like physical damage to the network backbone and the subsequent capacity limitations, are mainly considered. Credible results of these methods have to be based on complete and detailed information of the network. Therefore, it is suggested that network owners can also conduct this kind of vulnerability analysis for their own networks, considering the confidential information and results related to their network vulnerability. However, a more comprehensive risk assessment and vulnerability analysis is still necessary and needs to be conducted based on the analysis of individual networks by taking the interdependency of different networks into account.

3.1 Basic definitions

This section lists some terms (Grover, 2004; Zhou, 2007) which are often used in the analysis of optical networks.

- *Link*

  A link represents a basic unit of capacity between two adjacent nodes. Switching devices operate on the unit to interconnect capacity.

- *Span*

  A span is a physical entity, i.e., a fibre-optic cable, which collects all links between two adjacent nodes. The physical damage of a span, e.g., a cable cut, will affect its entire links simultaneously.

- *Route*

  A route is a concatenation of spans.

- *Path*

  A path is a specific cross-connected sequence of individual links on spans. A path carries a specific unit-capacity flow and cross-connected through a sequence of transmission spans.

- Service path or connection

  A service path or connection is a light path providing a specific traffic flow between two nodes with a built-in redundancy for protection or restoration.

3.2 Topological statistical data

In order to identify vulnerabilities from the physical topology of a network, at least the following two kinds of statistical data have to be collected:

- *Degree of each node:* Counting the degree of each node in a network is generally the first step of an analysis. By this, the most-connected node which plays an important hub-like role can be found. The importance of all nodes can be sorted according to their degrees. Then, the corresponding security and safety rules around these nodes can be deployed. On the other hand, some nodes with lower degrees may also play a key role if they are connecting cables with heavy traffic in a sparse network and the spare capacity of any alternative links is limited.
- *Active links in each cable duct or span:* When trying to identify vulnerable parts of a network from the aspect of physical cable routing, we have to investigate the number of lit (active) links in each part of a cable duct and identify the shared risk link group (SRLG), in which several logically distinct links are routed through a common duct. With these data at hand, we are able to find out how many links may be affected due to a cable-cut and how the traffic of the network may be affected.

Topology-driven analysis of vulnerabilities provides an essential support to a screening analysis aiming at identifying obvious vulnerabilities and guiding in-depth analysis, based on, e.g., object-oriented modelling (Eusgeld et al., 2009).

3.3 Simulation method

A simulation is an imitation and modelling of a real thing, state of affairs, or process in order to gain insight into their functioning and potential failure combinations. Key issues in simulation include acquisition of valid source information about the referenced item, selection of key characteristics and behaviours, the use of simplifying approximations and assumptions within the simulation, and fidelity and validity of the simulation outcomes. Here, we refer to a computer-based simulation in particular, which is an attempt to model a real-life or hypothetical situation on a computer so that it can be studied to see how the system works. By changing variables, predictions may be made about the behaviour of the system. In this study, a network traffic simulation is conducted.

Network traffic simulation is a process used in telecommunications engineering to measure the efficiency of a communications network. Network traffic simulation usually follows the following four steps:

- modelling the system as a dynamic stochastic (i.e., random) process
- generation of the realisations of this stochastic process
- measurement of simulation data
- analysis of output data.

The simulation model can be built with any programming language (e.g., C++, VB, and Java). In this study, the software AnyLogic, a Java-based hybrid simulation development environment, which can present a simulation easily by animation, is used for modelling.

However, accurate simulation model development requires extensive resources. It is necessary to check that the data is statistically valid by fitting a statistical-distribution and then testing the significance of such a fit. Therefore, simulations are also expensive to build and to make. The present model in the case study is built based on the basic network information obtained up to now. When more detailed information can be collected, more factors can be added into the model for further research.

4 Case study of SWITCH

4.1 Introduction of SWITCH

SWITCH, the Swiss national network for research and education, has been serving the Swiss science network since 1987. SWITCH, as a research centre, represents the interests of Switzerland in numerous bodies and its key role therefore makes an important contribution to the development and operation of the domestic internet.

Due to its relatively concise topology (see Figure 1) compared with other commercial networks, SWITCH is selected as reference for the first case study. SWITCHlan (the network of SWITCH) provides access to research networks worldwide as well as to the commodity internet for all Swiss universities, the two Federal Institutes of Technology and the major research institutes. The number of its internet PoPs (points of presence, access points to the internet) is comparatively small (about 40 PoPs), which greatly reduces the complexity of its simulation.

Figure 1 SWITCHlan topology (see online version for colours)

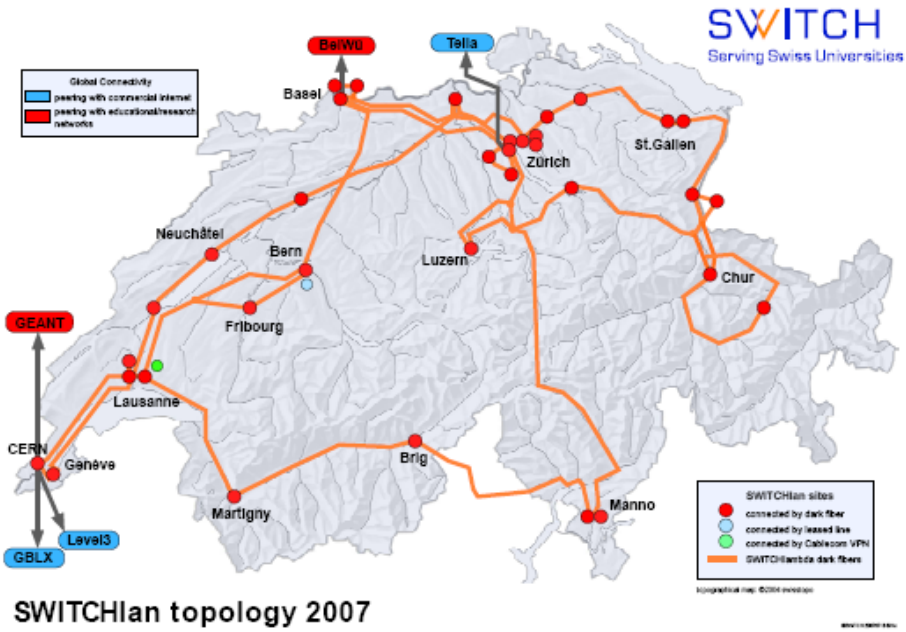

The SWITCHlan backbone entirely relies on dark fibres that SWITCH owns or rents from various dark fibre providers. SWITCH takes responsibility for all the equipment used to light the long-distance fibres. Two different types of technology are currently deployed: stand-alone wavelength division multiplexing (WDM) systems on the core rings, either dense WDM (DWDM) or coarse WDM (CWDM); and simple point-to-point Ethernet links everywhere else. SWITCHlan is actually a 'metro system' offering precisely the following range of functions:

- regeneration-free links up to 600 km, long enough to cover Switzerland
- multi-gigabit Ethernet and 10 gigabit Ethernet transponders
- scalability to 16 optical channels
- support for single fibre operation.

This combination of backbone connections on one fibre and daisy-chained point-to-point gigabit Ethernet links is now standard throughout SWITCHlan. This solution does not need an expensive DWDM node at each branch connection. Figure 2 shows the schematic structure of typical SWITCHlan connections.

Figure 2 Typical SWITCHlan connections (see online version for colours)

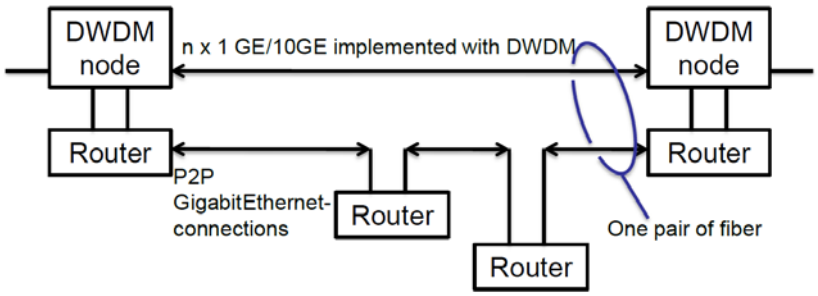

Source: SWITCHweb (2008)

4.2 Information collection from SWITCHlan and goals of case study

Because SWITCH itself is a research and educational network, the research project can be fully supported by SWITCH. Through several intensive interviews and frequent telephone and email contact with experts from SWITCH, the backbone of SWITCH network is given with information about physical routing of each cable, physical location and equipment of each node, average traffic statistics on each link as well as relevant technical details. Some information, like Figure 3 showing average traffic of SWITCHlan, can be obtained directly through SWITCH official website. However, considering vulnerability results of SWITCHlan may deal with some confidential problems, only limited details can be presented.

Most PoPs of SWITCHlan have two routers working in an active and standby manner. Each router represents a node except that co-located routers of a PoP are regarded as one node. The capacity of each path can be found from the traffic weather

map. Flow on each link can be simulated according to the method described in the following Subsection 4.4.

Figure 3 SWITCH traffic weather map (see online version for colours)

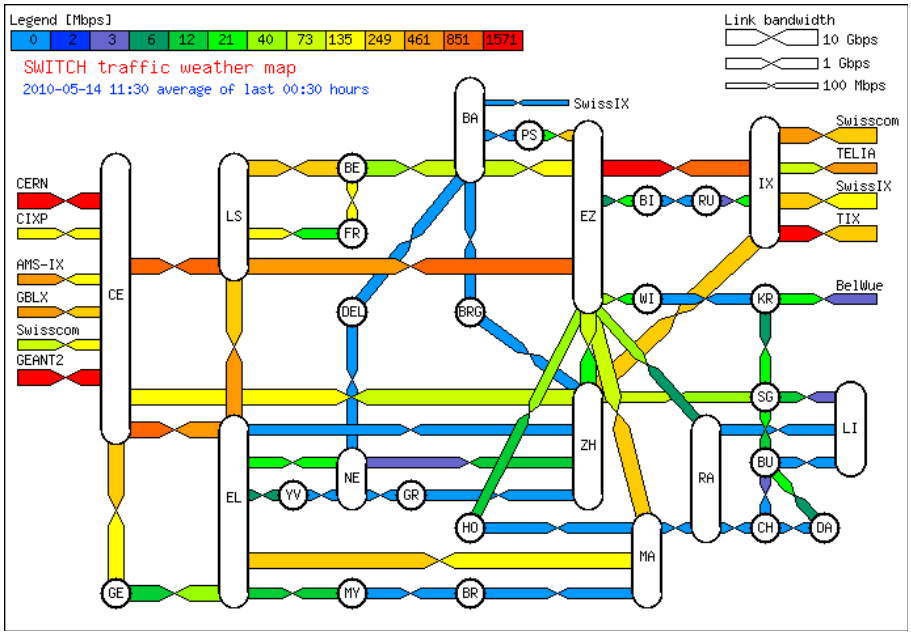

As mentioned in Section 2, targeted attacks on critical nodes and spans can degrade network performance, cause outage and lead to a cascade of overload failures for networks where flows redistribute among the nodes. The goal of the case study is to identify the critical nodes and spans of SWITCHlan. The analysis only takes into account of data level of obtained information (e.g., topological statistics) instead of dealing with telecommunication services.

Since restoration or dedicated protection schemes can protect a network against single-span failure, performance may be affected when multiple-spans or nodes fail in this network. SWITCHlan, however, with no protection mechanisms in the underlying optical fibre infrastructure, relies on internet protocol (IP) routing protocols (open shortest path first, OSPF) to perform dynamic rerouting when optical spans fail due to fibre breaks, power outage, etc. Therefore, one cable cut in the SWITCH network can degrade performance if spare capacity is limited and lost traffic cannot be fully restored on the IP level. Since single-span failures occur much more frequently than dual-span failures, we only consider the situation of one cable cut in the case study of SWITCHlan. Node failures seldom occur as compared with cable damages in practice (Willebeek-LeMair and Shahabuddin, 1997; Zhou et al., 2007). The node equipment usually comprises a redundancy, and its repair time is short. Therefore, node failures are not treated in the simulation. Critical nodes are identified based on topology analysis only.

4.3 Statistical data of SWITCH

Figure 4 shows the percentages of nodes of various degrees in SWITCHlan. The chart shows that over 50% of nodes have at least two links, while only about 2.6% of nodes connect to 12 other nodes. These few high-degree hub-like nodes, however, play a crucial role, which hold network together but are fragile to targeted attacks. Full redundancy is the most fundamental design requirement for a core backbone, and a topology of connected rings was favoured (Kugler, 2006). According to Figure 4, SWITCHlan can be regarded as a well-designed biconnected network.

Figure 4 Distribution of nodes' degrees in SWITCHlan (see online version for colours)

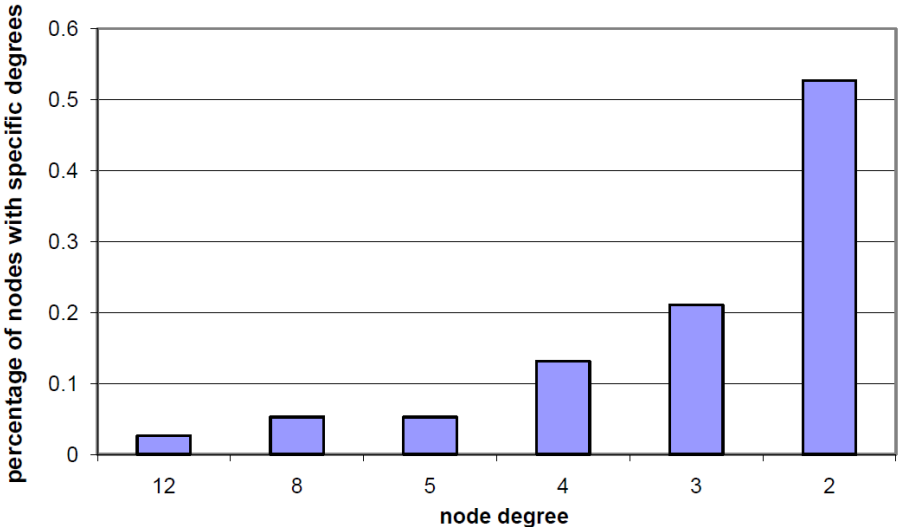

Figure 5 Percentage of spans with various active links

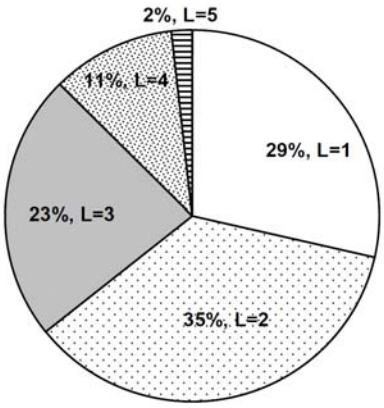

Notes: *L* indicates the number of active links in a span. The percentage value shows the ratio of a span with an *L* in the whole SWITCHlan.

Figure 5 displays percentages of spans with various active links in a pie chart. *L* indicates the number of active links in a span. Apart from a few exceptions, SWITCH is deploying a mix of different technologies (DWDM, CWDM and simple point-to-point gigabit Ethernet links) in bidirectional operating mode, i.e., only a single fibre is used for traffic in both directions. Therefore, here, each service path includes two directions between its two end nodes (Kugler, 2003). The percentage value shows the ratio of a span with an *L* in the whole SWITCHlan. 64% of spans in SWITCHlan contain one or two active links and 36% of them have three or more active links. The spans containing more active links are crucial. A cut of this kind of spans may greatly affect the traffic rerouting of the network. Special safety mechanisms, like installing more spare capacity in alternative spans and applying span restoration scheme (Grover, 2004), have to be considered on these cables.

4.4 Statistical distribution of traffic

In probability theory and statistics, the triangular distribution is a continuous probability distribution with lower limit *a*, mode *c* and upper limit *b*. The triangular distribution is typically used as a subjective description of a population for which there is only limited sample data, and especially in cases where the relationship between variables is known but data is scarce (possibly because of the high cost of collection). It is based on a knowledge of the minimum and maximum and an 'inspired guess' as to the modal value (WikipediaTD, 2008).

The capacity, average and maximum traffic of each path in both directions on a monthly basis are available on the SWITCH website. If the minimum traffic is assumed to be zero, the triangular distribution can be exploited in a simulation to generate a series of traffic data for each path in one direction, which would be a good guess of traffic at a point in time.

4.5 Internet traffic pattern and growth ratio

The internet is rapidly growing in number of users, traffic levels, and topological complexity. At the same time, it is increasingly driven by technical developments and economic competition. These developments render the characterisation of network usage and workloads more difficult, and yet more critical (Thompson et al., 1997). Traffic patterns on internet links can look quite chaotic compared with telephony networks, where many low-bandwidth calls of relatively long duration and relatively well known arrival processes add up to nice bell-shaped curves. According to some studies, domestic link traffic, especially its low-frequency components, follows a stable predictable 24-hour pattern that repeats daily (Figure 6). Figure 7 shows the simulated traffic based on the triangular distribution. Although no clear bell-shape can be found from the simulated traffic, it does give a good guess when minimum, maximum and average data have been determined. In the simulation (as shown in Figure 8), the traffic usage is denoted by the ratio of an absolute volume to the capacity of the studied link and displayed in different colours. Cold colours (e.g., dark blue) represent low utilisation while warm colours (e.g., yellow or red) represent high utilisation. Red colour indicates 'overflow' particularly.

Figure 6 Traffic usage on one SWITCH link (see online version for colours)

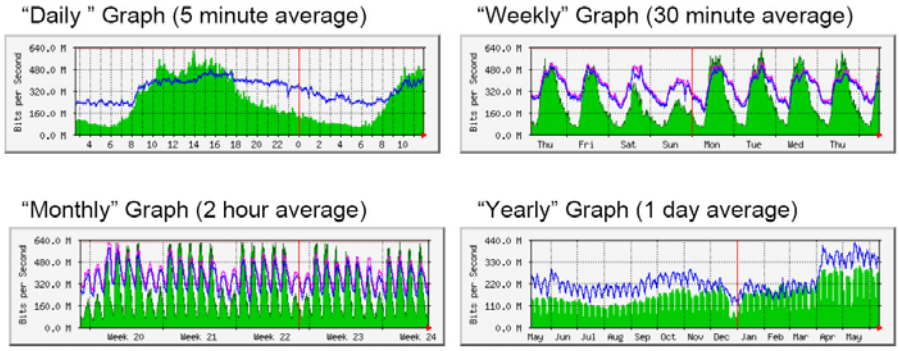

Notes: GigabitEthernet3/2 to Telia. Green: incoming traffic in bits per second; blue: outgoing traffic in bits per second; dark green: maximal five minute incoming traffic; magenta: maximal five minute outgoing traffic.

Source: SWITCHweb (2008)

Figure 7 Simulated traffic based on the triangular distribution (see online version for colours)

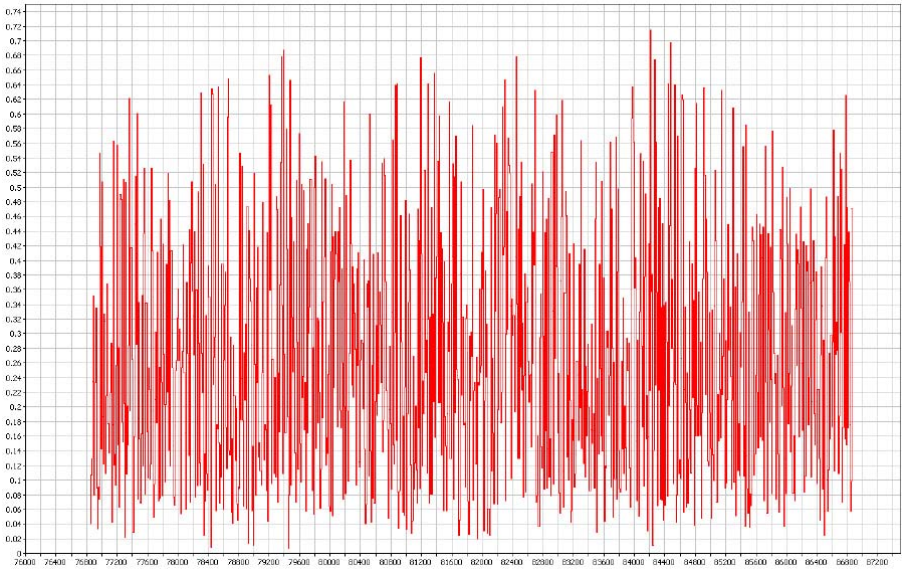

Internet traffic has been increasing very rapidly. In 1995 and 1996, there were periods in which internet traffic could double in as little as a hundred days, a rate of growth that looked almost terrifying. But it was not long before it began to slacken. By the late 1990s, traffic was doubling each year; that is, it was growing at 100% a year. But from 2002–2007, the growth rate has dropped, and it now hovers at 50% to 60% a year. Figure 9 shows a global internet traffic growth factor graph between 2002 and 2008 (InternetTraffic, 2008). The statistics were collected by University of Minnesota from

112 sites. The data may not be sufficient, but approximately indicate that the mean annual growth rate is 1.732, meaning the internet traffic grew at 70% a year.

Figure 8 Simulation interface (see online version for colours)

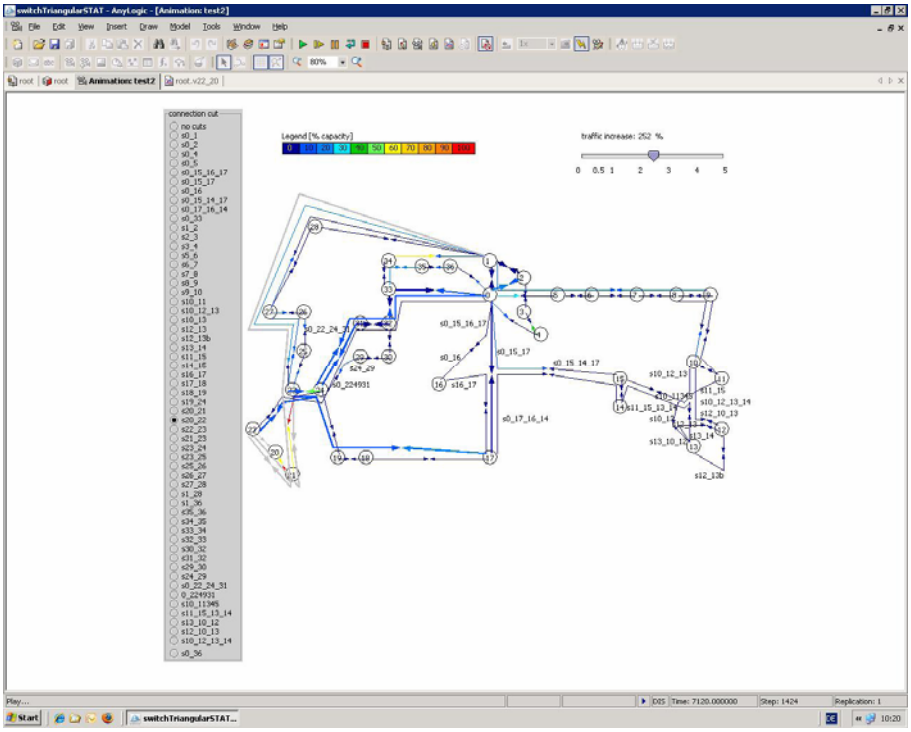

The SWITCH traffic follows a similar rule with an average annual growth rate of 1.55 (Figure 10). In the simulation, a traffic growth factor is designed to range from 0 to 5. If 0.5 is assumed to be the average growth factor per year, the simulation can imitate the traffic variation of SWITCH within ten years. According to the specialists from SWITCH, the service paths of SWITCHlan are generally upgraded gradually and the present routers can still serve for about five years. Therefore, the designed factor in the simulation may make sense when varying between 0 and 2.5. However, any vulnerable paths can be identified easily by adjusting the factor to its maximum value.

4.6 Traffic simulation of SWITCH

Based on the information obtained from SWITCH, the first step of the vulnerability analysis is to build a network model and imitate the traffic of each path.

When one cable cut occurs, the traffic carried by this cable will be redistributed among the remaining survival network. Because some cables in SWITCHlan contain several active links carrying high load, rerouting traffic of this kind of cable may cause overflow in some parts of the network, which would also become vulnerable parts.

Figure 9 Global internet traffic growth factor graph (see online version for colours)

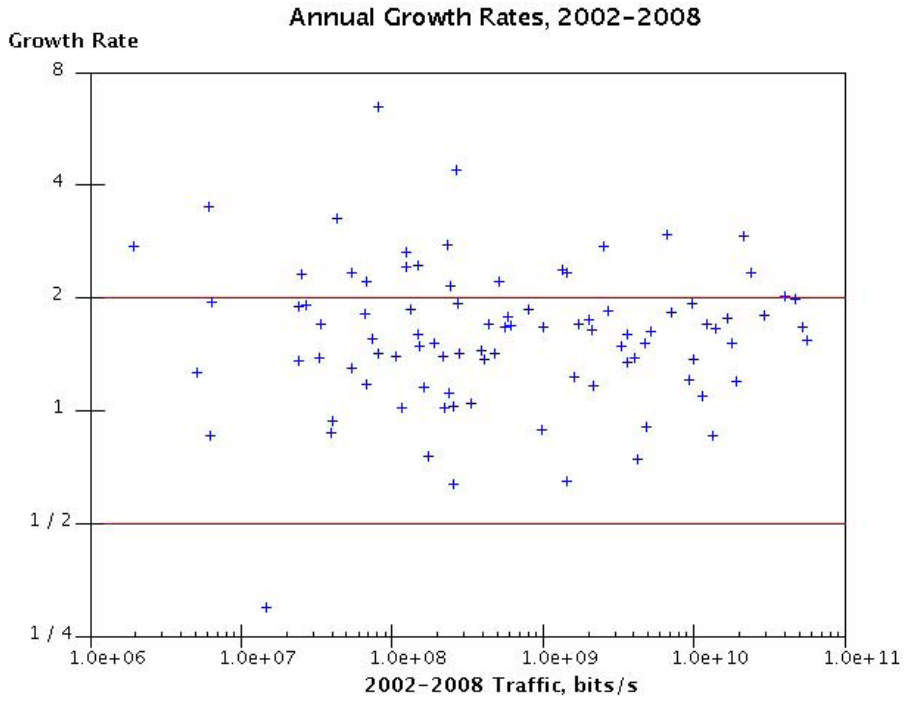

Notes: 2002–2008 statistics from 112 sites. Volume-weighted mean annual growth rate is 1.732.

Source: InternetTraffic (2008)

SWITCHlan is a network with a lot of spare capacities. The simulation is based on the published data of February 2008 on the SWITCH website. If the border links connecting to the other international networks are not considered, only seven one-directional paths' utilisations out of 120 one-directional paths are bigger than 10%. The remaining paths carry little traffic which generally occupies less than 5% of their respective capacities. Because of the huge spare capacity existing in SWITCHlan to provide enough space for possible rerouting traffic, one cable cut will not cause any overflow problem in the analysed year of 2008. However, as the internet traffic grows, this may change in the years to come. Network traffic simulation gives us a chance to take a look at a future situation and prepare for future risk.

According to Section 2, cable cuts are the single most common cause of telecom outages. Therefore, only this cause is considered for the simplification of the simulation. When a cable fails, traffic is rerouted following the rule of shortest path first.

According to the SWITCH specialists, a path will be upgraded when its utilisation frequently exceeds 30%–40%. Figure 11 shows the maximum numbers of paths with their utilisation exceeding 30% in five years in two cases, i.e., no cable cuts and one cable cut, respectively. When no cable fails at all, there are two paths of which the utilisations have exceeded 30% (in fact about 60%) in the analysed year of 2008. These two paths have to be upgraded immediately; otherwise they may soon become bottlenecks of the

network. Within two years of the analysed year of 2008, five more paths will have increased traffic which can be over 30% of their individual capacities. These paths have to be monitored specially and upgraded in time. Otherwise, there will be 20 paths with the problem of high utilisation in five years. The case of one cable cut follows a similar trend as the case of no cable cuts, certainly with more problematic paths. Figure 12 presents the maximum overflow paths in five years in the two cases, i.e., no cable cuts and one cable cut, respectively.

Figure 10 SWITCH internet traffic growth factor graph (see online version for colours)

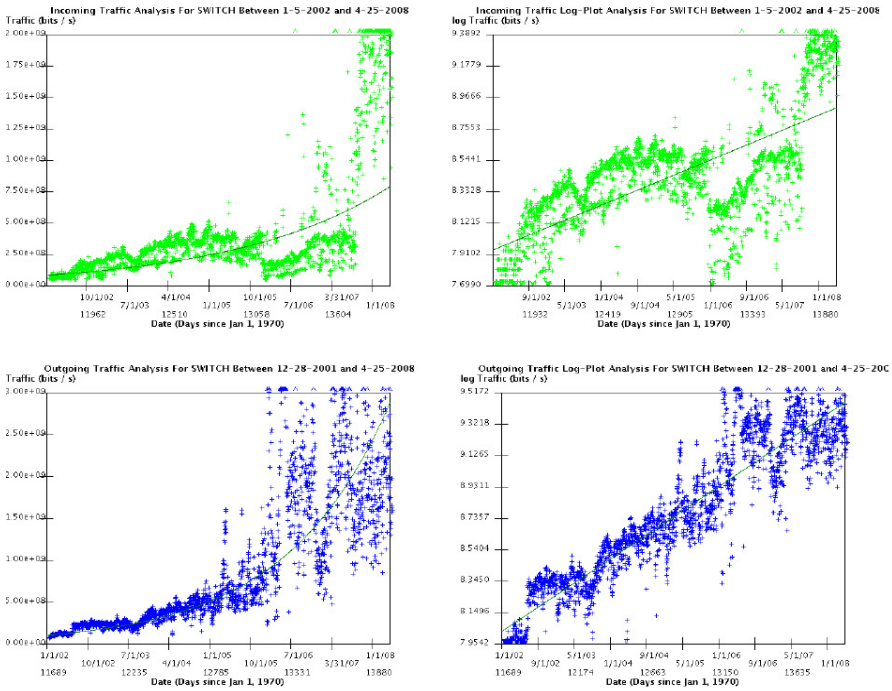

Notes: 2002–2008 statistics. The right-hand charts are the log-plot style of their corresponding left-hand charts. Incoming annual growth rate is 1.42. Outgoing annual growth rate is 1.68.

Source: InternetTraffic (2008)

Both Figure 11 and Figure 12 only show the maximum number of specific paths and do not indicate the exact location of the problematic paths, which can be different when a cable cut occurs in different places. No overflow paths occur in both cases in the analysed year of 2008. However, in the following year two overflow paths will appear even when no cables fail at all.

In the simulation, nine key spans have been identified out of 56 spans. One cut of this key span may cause 2 to 5 other service paths to overflow in one to five years.

The simulation indicates that SWITCHlan works in a good condition at present and shows high robustness. Suitable upgrade actions have to be performed in time on the

identified vulnerable paths in order to mitigate the possible overflow situation in the near future.

Figure 11 Trends of maximum numbers of paths with their utilisation exceeding 30% in five years when no cable cuts occur and when one cable cut occurs, respectively (see online version for colours)

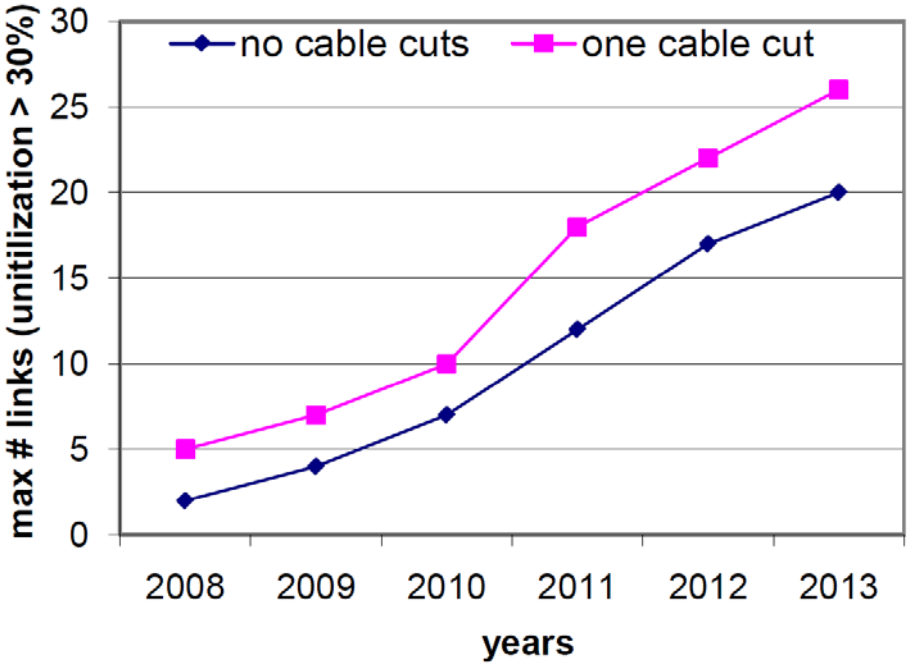

Figure 12 The maximum overflow paths in five years when no cable cuts occur and when one cable cut occurs, respectively (see online version for colours)

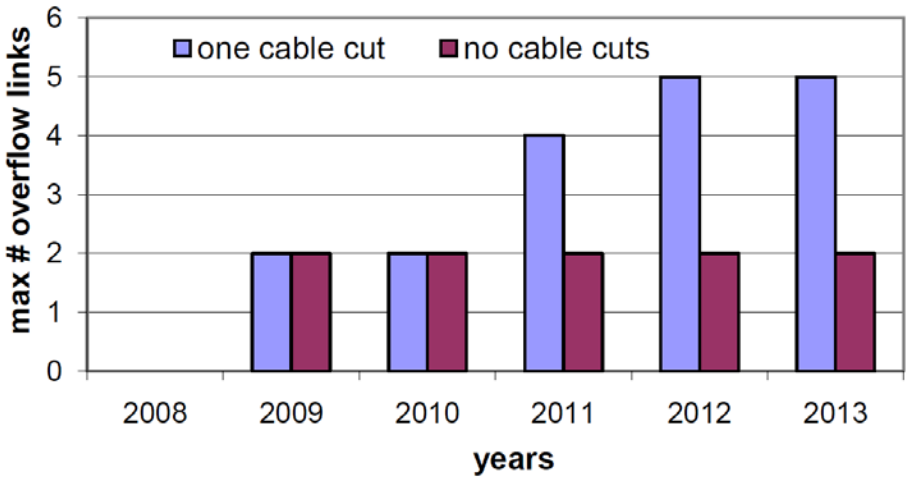

5 Summary and outlook

The physical part of networks handles the 'bit transport' for the transmission services. In this instance, vulnerability is linked specifically to its availability. Physical disturbances attributable to the intentional or unintentional disabling of cables are the most acute manifestations of this vulnerability (Luiijf and Klaver, 2000). In this paper, the SWITCH network was studied on its physical cable routing, network topology and traffic graph. In contrast to commercial networks, which often operate closer to their limits, SWITCH has very low traffic load and relatively high spare capacity in each link. Even if SWITCH has no problem in case of one cable cut in the analysed year of 2008, it may experience problems in next few years if no systematic and effective upgrade actions are executed at all. Thus, we can imagine the possible risk that commercial networks are facing.

A commercial network operator may argue, "SWITCHlan is simple and almost only a basic bi-connected network. We have a much more complex meshed network to reroute any failed traffic". That may be true to some extent. For some cable cuts, it seems that only one route is available in SWITCHlan for restoration. However, this situation can exist in more meshed networks, too, not to mention their possible or even more serious 'overflow' problem. Two nodes seem to be connected by several cables, where in fact only one or two active paths can work for the communication between the two nodes. These paths may even go through a common cable duct. In this case, the failure of the common duct will break all possible restoration routes between these two nodes. Therefore, we have to study both logical and physical topologies and traffic on each link of a network at the same time. Studying on one part of a network cannot help identify vulnerabilities correctly. All in all, a thorough investigation has to be conducted to identify vulnerability and reduce potential risk.

We recommend that the following points about the protection of physical information infrastructure should be specially considered by a country:

- regard common connection points as key protected infrastructure, e.g., the node and the place where several cables are gathering
- persuade and support backbone/service providers to arrange their backbone along diverse physical routing
- cooperate with service operators to define a series of basic safety standards for the completely unregulated ICT sector.

Several more networks may be studied in the future, not only because of their different individual features, but because of the problem of (inter-)dependency among them. What will happen if the backbones of different networks go through the same cable duct and are subject to geographic common cause failures? How will the network operators react when several networks fail simultaneously? How to describe their behaviours?

This paper presents analysis approaches mainly based on network theory, which can only give a screening analysis according to the topological features and capacity limitations of a network. Due to CII complexity, heterogeneity, multidimensional interdependencies, significant risk dealing with the natural disasters and catastrophes, technical disasters and failures, espionage, international crime, physical and cyber terrorism, etc., these infrastructures require a new, holistic approach to their protection, considering not only information security methods and techniques but also achievements

of the safety domain. They also need cooperation on an international level. Improving dependability and survivability of CII is the key challenge (Białas, 2006).

Acknowledgements

This research has been funded as part of a project for the Swiss Federal Office for Civil Protection (FOCP). It does, however, not necessarily represent the official views of the Office or another Swiss government body. The author thanks Simon Leinen and Felix Kugler from SWITCH for their great support to our study.

References

Bia las, A. (2006) 'Information security systems vs. critical information infrastructure protection systems – similarities and differences', in *Proceedings of the International Conference on Dependability of Computer Systems (DEPCOS-RELCOMEX'06)*.

Bologna, S., Costanzo, G.D., Luiijf, E. and Setola, R. (2006) 'An overview of R&D activities in Europe on critical information infrastructure protection (CIIP)', in Lopez, J. (Ed.): *Proceedings of the 1st International Workshop on Critical Information Infrastructures Security*, LNCS 4347, pp.91–102.

Bonanni, E.C.G., Minichino, M., Clemente, R., Iacomni, A., Scarlatti, A., Zendri, E. and Terruggia, R. (2008) 'Exploiting stochastic indicators of interdependent infrastructures: the service availability of interconnected networks', in *Proceedings of the European Safety and Reliability Conference for Risk Analysis, Safety and Reliability*.

Brunner, E.M. and Suter, M. (2008) *International CIIP Handbook 2008/2009 – An Inventory of 25 National and 7 International Critical Information Infrastructure Protection Policies*, Center for Security Studies, ETH Zurich, chapter Conclusion, pp. 536–537.

Chang, S.E. (2009) 'Infrastructure resilience to disasters', *The Bridge*, Vol. 39, No. 4.

Ciancamerla, E., Minichino, M., Schmitz, W., Uusitalo, T., Linnemann, R., Wright, D., Khnert, C., Issacharoff, L. and Buzna, L. (2007) 'Deliverable D 2.2.2, 'tools and techniques for interdependency analysis'', Technical report, Integrated Risk Reduction of Information-based Infrastructure Systems (IRRIIS).

Crucitti, P., Latora, V. and Marchiori, M. (2004) 'Model for cascading failures in complex networks', *Physical Review, The American Physical Society*, Vol. 69, No. 4, p.045104(4).

Eusgeld, I., Kröger, W., Sansavini, G., Schläpfer, M. and Zio, E. (2009) 'The role of network theory and object-oriented modeling within a framework for the vulnerability analysis of critical infrastructures', *Reliability Engineering and System Safety*, Vol. 94, No. 5, pp.954–963.

Gorman, S.P., Schintler, L., Kulkarni, R. and Stough, R. (2004) 'The revenge of distance: vulnerability analysis of critical information infrastructure', *Journal of Contingencies and Crisis Management*, Vol. 12, No. 2, pp.48–63.

Grover, W.D. (2004) *Mesh-Based Survivable Networks, Options and Strategies for Optical, MPLS, SONET, and ATM Networking*, Prentice Hall PTR, New Jersey.

Grubesic, T.H. and Murray, A.T. (2006) 'Vital nodes, interconnected infrastructures, and the geographies of network survivability', *Annals of the Association of American Geographers*, Vol. 96, No. 1, pp.64–83.

Hoffman, M. (2002) 'Cable cuts', available at http://all.net/CID/Attack/papers/CableCuts.html (accessed on 25 September 2008).

InternetTraffic (2008) 'Minnesota internet traffic studies', available at http://www.dtc.umn.edu/mints/2002-2008/analysis-2002-2008.html (accessed on 25 September).

IRGC (2006) 'Managing and reducing social vulnerabilities from coupled critical infrastructures', White paper 3, International Risk Governance Council (IRGC).

Kugler, F. (2003) 'SWITCHlambda – a vision is taking shape', *SWITCH Journal*, pp.6–8.

Kugler, F. (2006) 'SWITCHlambda – rings closed', *SWITCH Journal*, pp.15–18.

Luiijf, H. and Klaver, M. (2000) 'In bits and pieces – vulnerability of the Netherlands ICT-infrastructure and consequences for the information society', Technical report, INFODROME.

McNamara, M. (2008) 'Dove hunters shoot up fiber optic lines, cause Indiana outage', available at http://www.multichannel.com/blog/1300000330/post/780032678.html (accessed on 01 October).

Motter, A.E. and Lai, Y-C. (2002) 'Cascade-based attacks on complex networks', *Physical Review E*, Vol. 66, No. 6, p.065102(R).

Poulsen, K. (2006) 'The backhoe: a real cyberthreat', available at http://www.wired.com/print/science/discoveries/news/2006/01/70040 (accessed on 01 October 2008).

SWITCHweb (2008) 'Switch website', available at http://www.switch.ch (accessed on 25n September).

Thompson, K., Miller, G.J. and Wilder, R. (1997) 'Wide-area internet traffic patterns and characteristics', *IEEE Network*, pp.10–23.

Vernon, A.J. and Portier, J.D. (2002) 'Protection of optical channels in all-optical networks', in *Proceedings 18th Annual National Fiber Optic Engineers Conference (NFOEC 2002)*, pp.1695–1706, Dallas, TX.

Vespignani, A. (2010) 'The fragility of interdependency', *Nature*, Vol. 464, pp.984–985.

WikipediaTD (2008) 'Triangular distribution', available at http://en.wikipedia.org/wiki/Triangular_distribution (accessed on 25 September).

Willebeek-LeMair, M. and Shahabuddin, P. (1997) 'Approximating dependability measures of computer networks: An FDDI case study', *IEEE/ACM Transactions on Networking*, Vol. 5, No. 2, pp.311–327.

Wilson, C. (2006) 'Dual fiber cut causes sprint outage', available at http://telephonyonline.com/access/news/Sprint_service_outage_011006/ (accessed on 01 October 2008).

Zhou, L. (2007) 'Availability analysis and optimization in optical transport networks', PhD thesis, Swiss Federal Institute of Technology Zurich (ETH).

Zhou, L., Held, M. and Sennhauser, U. (2007) 'Connection availability analysis of shared backup path-protected mesh networks', *Journal of Lightwave Technology*, Vol. 25, No. 5, pp.1111–1119.