# Internet under Threat: Simulation of Survivability with INESS

Wolfgang Fischer, Niels Lepperhoff, Angelika Volst

Programme Group Systems Analysis and Technology Evaluation (STE)
Forschungszentrum Jülich GmbH
52425 Juelich, Germany
wo.fischer@fz-juelich.de
n.lepperhoff@fz-juelich.de
a.volst@fz-juelich.de

**Abstract:** There are several studies analysing the Internet's topology to obtain insights into the consequences of disturbances. But concentrating on topology misses the point that the Internet's survivability depends on the relationship of topology, capacity and traffic. Using the INESS simulation model we compare several disturbances of Internet exchange points to reveal the consequences for the survivability of the Internet. We classify disturbances in the aggregated Internet Exchange Points (IX) Frankfurt, Los Angeles, London, Paris, Miami, San Francisco, San Francisco & Los Angeles together, and Tokyo along two dimensions: packet loss and the number other of significantly hampered Internet exchange points. Two important phenomena of disturbances in networks are revealed: the "remote effect" ("Fernwirkung") and the "Janus face".

## 1 Introduction

The "information structure", in particular the Internet, is one of the critical infrastructures. In politics and society it was a long-held conviction that the Internet is tolerant of disturbances *because* it is a packet-oriented network. Nevertheless, there are some incidences having an impact on the Internet. There are three principal categories of incidents: random (technical) disturbances; major natural events (e.g. earthquakes); intentionally attacks, focusing on the most vulnerable parts of the infrastructure. Our work is focused on the last two disturbances. They have a high damage potential that is much greater than the damage caused by "normal" technical occurrences.

Concentrating on topology only reveals the existence of alternative paths and single point of failures like backbones or IX without alternative paths. Even if an alternative path exists, topology alone does not indicate whether this path is able to cope with the additional amount of traffic caused by the disturbance, but capacity and traffic does. Therefore studying the survivability of the Internet, one has to look at *topology*, *capacity* and *traffic* at the same time. A promising technique is *simulation*. With our “Internet Security Simulation” (*INESS*) model we systematically compare the effects of disturbances in several central IX to get a better understanding of how vulnerable locations of the Internet are. Because the Internet’s topology and use is constantly changing and data on topology are sometimes elusive, simulation results should be treated in a qualitative manner only.

## 2 Short description of the INESS model

At the start of a simulation experiment, the simulation control “generates” the network infrastructure based on the scenario description, in which the infrastructure used, the quantity of data packets to be generated, the disturbance and other control information is defined. After the infrastructure has been set up, the IXs “learn” this infrastructure for a period laid down in the scenario description. The routing tables are built up using a procedure based on the BGP routing protocol. In each time step, it is first checked whether disturbances have occurred or whether existing disturbances have been remedied. Each “source”[7] then generates new packets. The IXs and the backbones then redirect the packets that have arrived.

Backbones are connected to each other in the IXs. For each packet arriving, the IX must therefore select the backbone via which this packet is to be forwarded. Packets destined for the IX are passed on to the first source in the IX so that they can be “processed” there. All other packets are passed on to the backbone that is free and that has the lowest number of IXs to the packets' destination. “Routing” is the mechanism which guides this packet traffic. The backbone, which is appropriate for a certain destination, can be taken from the routing table. This describes the central idea of the Internet protocol, IP. The BGP is used as a routing protocol between autonomous systems. Only the central idea of this protocol is required for our simulation.

At periodic intervals, subject to slight random fluctuations, each IX dispatches new path updates to its neighbours. If an IX has not received any routing updates from a neighbour for a certain period, this IX is removed from all pathways in its routing tables and the removal is communicated to its neighbours (“hello protocol”).

---

[7] Within an IX one or more sources produce or receive packets. The local Internet links, to which providers and users of the Internet are connected, are combined in the respective source. Each source works for one IX only. An IX may have several sources each representing, for example, one type of Internet use.

At each time step, every source unloads its “queue”, in which those packets have been deposited that are addressed to the IX assigned. A successful delivery of the packets is recorded and the packets are deleted. New packets are generated. Packets can either be distributed unsecured as described by the UDP or secured (TCP). They do not represent an application since the model is on the packet level.

## 3 Backbone network and Internet use

We compiled a map based on the backbone data published by Telegeography [Te02a; Te02c; Te02d]. But the data omit the backbones within countries with the exception of the USA, and Telegeography concentrates on backbones with the largest capacity. Therefore we have compiled the European backbones from the European fibre optical cable networks [Te01; Te02b]. The resulting aggregated backbone network consists of 209 IX and 558 backbones with a capacity of more than 240 Mbit/s each.[8] There is some evidence that IX have less capacity than the sum of all adjacent backbones. We therefore estimated the IX capacity as twice as high as the maximum simulated throughput measured without any disturbances [for details see LFV02].

Unfortunately, there is no worldwide data about the traffic generated from the different users. Thus, the number of packets generated for each type of service has to be estimated. Our estimation is based on empirical studies of Internet use [details in VLF02].

## 4 Description of scenarios and measurement

The goal of our simulation study is to compare the effects of IX[9] disturbances at different locations of the backbone network. We compare the effects in two ways: We compare the disturbance effects with the base-line scenario, which has no disturbance; we compare the disturbance effects with each other.

For each location of a disturbance we vary the traffic to analyse the impact of different levels of network utilization on the effects of a disturbance: The normal package production (100%) is increased by 50% and doubled.[10] This duplication reduces the IX reserve capacities to zero. Therefore, a doubling is the extreme case giving information about the maximal possible effects, whereas an increase by 50% gives an indication of whether the effects of disturbances constantly increase with the traffic increase or rise abruptly.

---

[8] For some European telecommunications providers like the German Telecom data are not available in public. Therfore some backbones are missing in our simulation.

[9] An IX in our simulation is aggregated. In reality it consists of several distributed IX in the city. We assume that the failure of such an aggregated IX could not happen due to technical reasons. The probability of such a major disturbance and, in the case of an intentional attack, the motives behind, are not discussed here.

[10] The utilization can be doubled either by doubling the packet production or halving the capacities of IX and backbones (e.g. by insolvencies). INESS can use both variants.

To select the location of a disturbance we have identified the most important IX for a world-region. Telegeography [Te02c] measures the topological importance of an IX for by counting the international backbone capacity which connects the region with the IX (Figure 1). Therefore we analyse in this paper the effects of disturbances in these IXs.

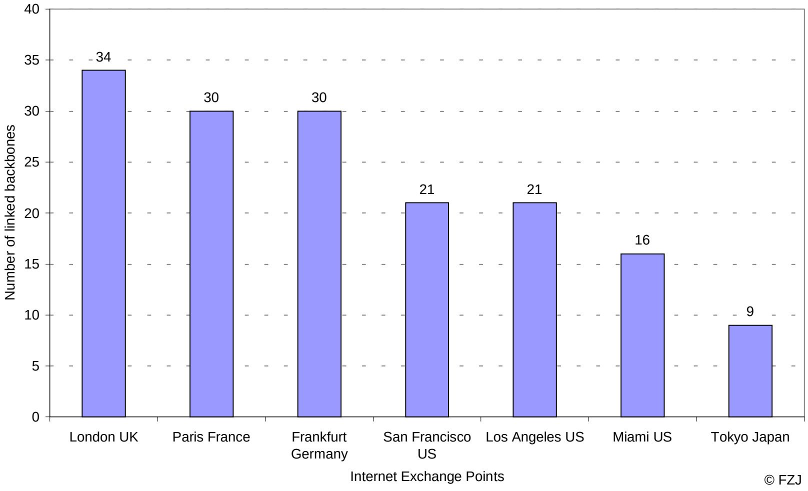

Figure 1: Number of backbones linked to selected Internet Exchange Points.

The Internet is a medium for transporting data onto which various services such as e-mail, file transfer and WWW have been grafted. The "packet loss" can be used to measure the quality of this service. If a packet is lost the transport service has not been provided, thus the reliability of the Internet decreases.

We distinguish two kinds of packet losses: A packet is "*undeliverable*" if the IX does not know a path with suitable capacity to the destination IX. A "*transport packet loss*" occurs if more packets arrive than the IX can handle – the surplus packets are dropped. *Total packet loss* is the sum of both kinds of packet losses and indicates the networks status on the *system level*.

Another important level of analysis is the *connection level*. A connection is a sender-receiver relationship: IX A sends packets to IX B. From an user's perspective the quality of connections is more important, because if a desired connection is hampered the user's activity is disturbed even if the network status on the system level is rather healthy.

To monitor the connection quality we plug in the network eleven “probing sources”. They stand for a “virtual” agent, e.g. a company, an international organization, making use of the network. For methodological reasons we have located theses sources at the border of the network to ensure that the probing packets have to travel through a large part of the network[11]. The probing packets are marked to monitor their path and the quality of probing connections. At every time step, each probing source sends one packet to each other probing source. Thus, the number of packets sent for each connection is well defined. A hampered connection will receive fewer packets. Therefore, we can “measure” how much affected the connections between the probing sources are by IX failures.

To assess and compare the effects of a disturbance we simulate different IX failure. Each scenario lasts over 10 runs. The IXs learn the network’s topology for 100 time steps. Afterwards packet generation starts. The time steps are counted from the start of packet generation onwards. At time step 119 the disturbance occurs. After another 100 time steps the observation period starts and lasts until the 500th time step. The system needs about 100 time steps to stabilize. Thus, we wait 100 time steps between packet generation start and disturbance and between disturbance and observation. On average each backbone has a utilization of 14% without any disturbance.

## 5 Results and discussion

Our analysis follows a top-down approach. We start at the system level analysing the packet loss and the number of lost probing packets. Then an important phenomenon called “remote effect” is presented. Afterwards we look at the probing connections.

### 5.1 Packet loss on system level

In the base-line scenario with no disturbance we observe some packet loss due to jammed backbones. This is an “artificial” packet loss, because we feed the packets in a simulated network, which is smaller than a real one. The figures are used as a reference scenario, because our focus is on IX disturbances and their consequences for the packet loss compared to the base-line scenario.

Compared to the base-line scenario a disturbance at IXs in Miami, San Francisco & Los Angeles, San Francisco and Tokyo does not significantly increase the transport packet loss if traffic is doubled (significance: >= 0.063).[12] For disturbances in IXs in Frankfurt, London, Paris and Los Angeles the transport packet loss is significantly higher than in the base-line scenario (significance: <=0.002).

---

[11] The probing sources are located in IX Auckland, New Zealand, Beijing, China, Belfast, Northern Ireland, Buenos Aires, Argentina, Calgary, Canada, Kansas City, US, Kuala Lumpur, Malaysia, La Coruna, Spain, Rome, Italy, St. Petersburg, Russia, Warsaw, Poland.

[12] The used significance level is 5% for all tests.

Each IX disturbance changes the number of undeliverable packets significantly at traffic of 100% (significance: <=0.002). For traffic of 200% the number of undeliverable packets are no longer significant for IX failures in Frankfurt and Miami (significance: >=0.075). The network is under heavy stress at a packet production of 200% causing constant advertising and withdrawal of congested routes ("flipping"). Furthermore, there is no more reserve capacity at the IX. The artificial packet loss mentioned becomes worse, because some backbones are too small. These effects together wear away the consequences of the disturbance: Such an incidence in Tokyo reduces the number of undeliverable packets significantly because the disturbance hinders Asian packets from entering the US-backbone network lowering the utilization of the network. Furthermore, Tokyo is a strong source of packet production in Asia. Their failure also lowers the network's utilization increasing the chances of other Asia packets being delivered successfully. A failure in IX London has the most serious effects. The packet loss is even higher compared to a scenario in which both IX Los Angeles and San Francisco are disturbed together. London links most of the British backbones to the European backbone network. Some links to the European network bypass London, but in terms of capacity London is the most important gate to the rest of the world: Thus, London is a bottleneck for large parts of the British backbone network. Furthermore, many transatlantic backbones terminate in IX London. Neither Paris nor Frankfurt has such a critical role.

In general, increasing traffic reduces the increment of total packet loss (Figure 2). Thus, the effect of a disturbance is *lower* under heavy traffic conditions, because the network is under stress even without an IX failure, leading to route flipping, causing packet loss - the "artificial" packet loss becomes worse.

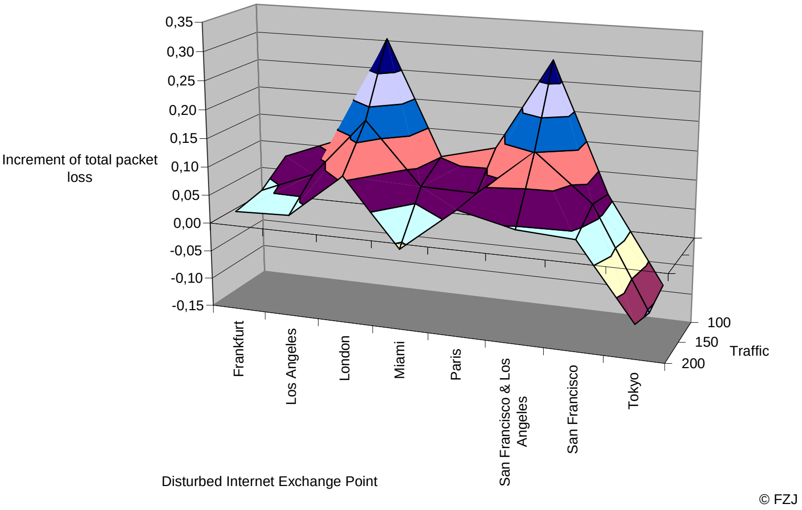

Figure 2: Increment of total packet loss under different disturbances and traffic conditions.

Several studies of the survivability of the Internet focus on topology [e.g. GOM03; FFF99; VPV02]. They operate on the assumption that the number of adjacent backbones of an IX indicates how critical a failure of the IX would be. If this assumption holds true, then there will be a relationship between the number of adjacent backbones and the total packet loss. Figure 3 shows the total packet loss divided by the number of linked backbones for each disturbed IX. This quotient is not constant revealing that there is no connection between total packet loss and the number of adjacent backbones.

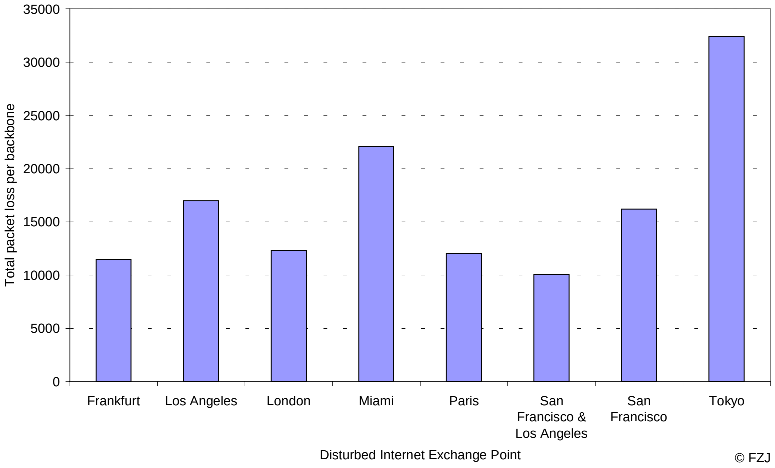

Figure 3: Comparing the number of undeliverable packets in 10 runs divided by the number of linked backbones for traffic of 100%.

## 5.2 Detecting the “remote effect” (“Fernwirkung”)

Packet loss is an important criterion for classifying the impact of a disturbance on the entire Internet. But it does not show how many IXs are affected, i.e. which IX lost the packets. A connection which does not use any affected IXs does not suffer from the failure. In other words, the ratio of IXs with significantly increased packet loss to the total number of IXs is an indicator of the impact of the disturbance on connections.

In the remaining part of the section we concentrate on the traffic of 100%, because one can observe the effects of disturbances more clearly than by using heavier traffic conditions. To assess the effects on an IX the mean number of undeliverable packets at this IX in the base-line scenario is compared to the mean number of undeliverable packets at the same IX in a scenario with a disturbance. If the two means differ significantly over all runs for each scenario, then we term it an effect, which has two dimensions: If the mean number of undeliverable packets is reduced the IX is *positively* affected. If the mean number of undeliverable packets increases the IX is *hampered* (negatively affected).

Disturbances in Paris or in Los Angeles & San Francisco increase the number of undeliverable packets significantly at more than 102 IXs, about 50% of all IXs . On the other hand, a disturbance in London produces the highest number of undeliverable packets, but hampers only 43 IXs. But because the number of adjacent IXs is smaller than the number of hampered IXs, the impact of a disturbance is not limited to the neighbourhood of the disturbed IX. A larger part of the network could be affected.

One can presume that the hampered IXs are the adjacent IXs of the disturbed IXs. But this is not the case. Figure 4 shows the distribution of affected and unaffected IXs in relation to their distance to the disturbed IX. The percentage of positively affected IXs decreases with the distance. The neighbours of the failed IX are usually relieved by the disturbance because the traffic is rerouted around them. The majority of hampered IXs are three HOPs away from the disturbance and the majority of unaffected IXs four HOPs. The positive and negative effects of a disturbance concentrate on a region around the disturbed IX. But they are not limited to this region. We call this phenomenon "**remote effect**" ("Fernwirkung"). A disturbance affects all parts of the network, because redirected traffic increases the networks utilization far away from the failed IX, thus causing more packet loss. The probability of being hampered for an IX decreases with the distance to the failed IX (Figure 5). For a distance of four HOPs the probability of being hampered nearly drops below 20%. But the probability remains above zero until nine HOPs. This finding reinforces the observed remote effect.

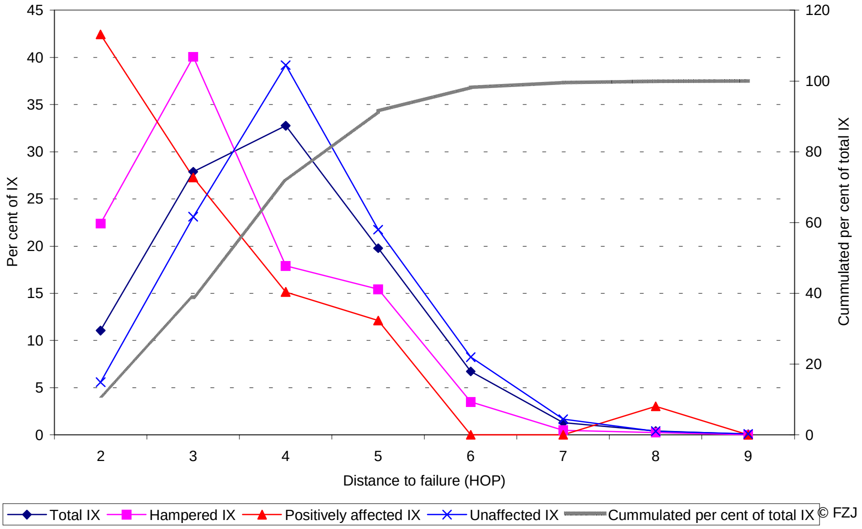

Figure 4: Distribution of affected and unaffected IX versus distance to disturbed IX.

Disturbances always affect the entire network, but with different strengths. Thus, high aggregate figures do not reveal which specific connections are hampered. Therefore we analyse the probing connections in more detail.

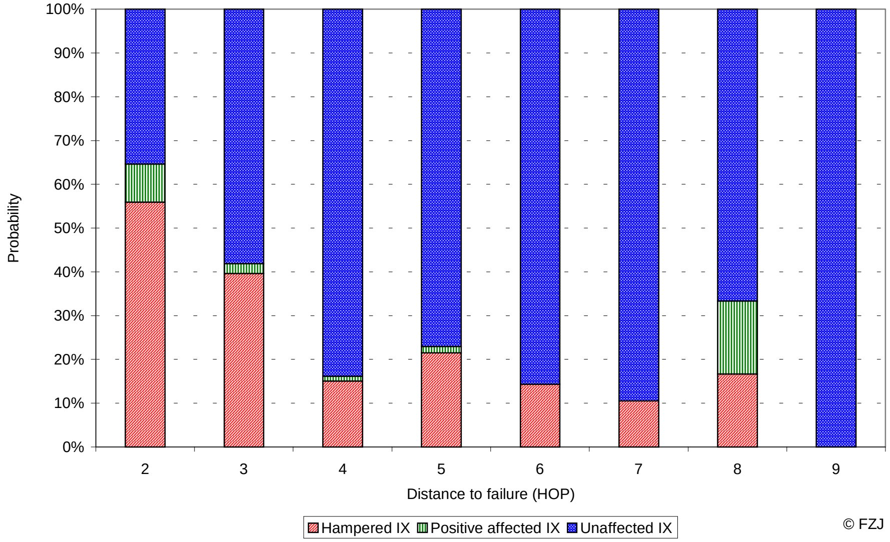

Figure 5: Probability of an IX being hampered in relation to the distance to the failed IX.

### 5.3 Connection-based analysis

The loss of probing packets indicates that some of the probing connections are hampered by a IX failure. We focus on traffic of 100%. Again, even without any disturbance, some packet losses exist that hampers connections. Therefore, the base-line scenario is used as a reference case only.

21 connections are unaffected by all disturbances of IXs. 89 connections are significantly affected by at least one disturbance. If Tokyo is disturbed the probing source Kuala Lumpur is completely disconnected, because there is only one backbone to Tokyo. This leads to 20 failed probing connections. The failure of San Francisco and Los Angeles together disconnects the probing source Auckland. The impact of disturbing the IXs varies dramatically (Figure 6): A failure in IX Paris significantly affects 12 probing connections, but a failure at Tokyo 33. If the IXs of San Francisco and Los Angeles are disturbed together, 39 probing connections are significantly affected. A minority of affected probing connections show an improved performance revealing a “**Janus face”** of disturbances. This Janus face is due to the fact that disturbances could unburden the network and the connections, because packets are destroyed. As a consequence connections are “de-jammed” and packets can now go through. We call this effect the (relieving) “first Janus face”. The owners of the lost packets are confronted with the (grim) “second Janus face”, because they face a loss of their data.

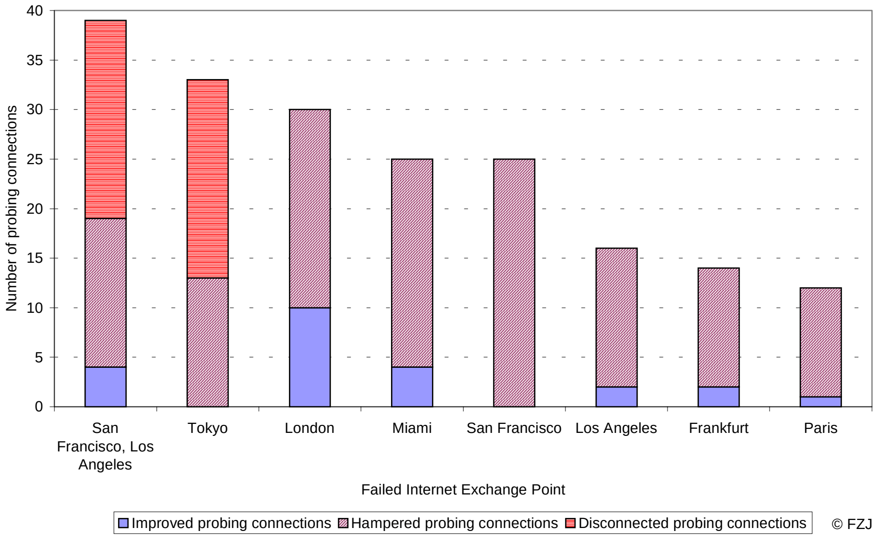

Figure 6: Number of affected probing connections with traffic of 100% sum of 10 runs.

The analysis of the probing connections reveals consequences due to the different features of the regional parts of the simulated network. An IX disturbance in Europe has a minor effect on the probing connections, whereas a disturbance in the USA has a strong effect for Asia and Latin America. Developed network regions (Europe and the USA), embedded in a continental and transatlantic mashed, dense backbone network with high backbone capacities, offer more alternative paths with suitable reserve capacities. Instead, Latin America and Asia have "star-type" backbone networks, mostly depending on a few IX in the USA. To increase their survivability more backbones organized in a meshed structure are needed.

# 6 Conclusions

The Internet's survivability depends on the relationship of topology, capacity and traffic. Simulation is the key technique to cope with these factors. Treating these factors separately neglects the nature of the Internet. Therefore, using the INESS simulation model we compare several disturbances of IX to reveal the consequences for the Internet.

We analyse disturbances in different IXs. The analysis of the probing connections reveals, that an IX disturbance in Europe has a much smaller effect on the probing connections than a disturbance in the USA has for Asia and Latin America. A meshed, dense and "capable" backbone network pays off in the case of a major IX disturbance.

The assumption that effects of disturbances are local around the failed IX is not correct. Instead, we identified a phenomenon called “remote effect”: A disturbance affects all parts of the network, because redirected traffic increases the network utilization far away from the failed IX, thereby causing more packet loss. The probability of an IX being hampered decreases with the distance from the failed IX, but remains above zero of nearly all distances. Neighbours have the highest probability of being hampered, but most hampered IX are three HOPs away. In consequence, a disturbance affects the entire network, but the strength of the effects varies. A second important impact of disturbances is the “Janus face”. Disturbances could unburden the network and the connections, because packets are destroyed. As a consequence connections are dejammed (the “first Janus face”). The owners of the lost packets are confronted with the “second Janus face”, because they face a loss of their data.

The effects of the disturbances differ in packet loss and in the number of IXs affected. Together they classify the impact of disturbances. Figure 7 classifies the simulated disturbances according to the packet loss (number of undeliverable packets) and the number of hampered IXs. According to this classification for example a disturbance in IX Tokyo has only a minor impact on the Internet. Disturbances in San Francisco and Los Angeles together are critical. Therefore, a major earthquake (“the big one”) would have a major impact on the Internet.

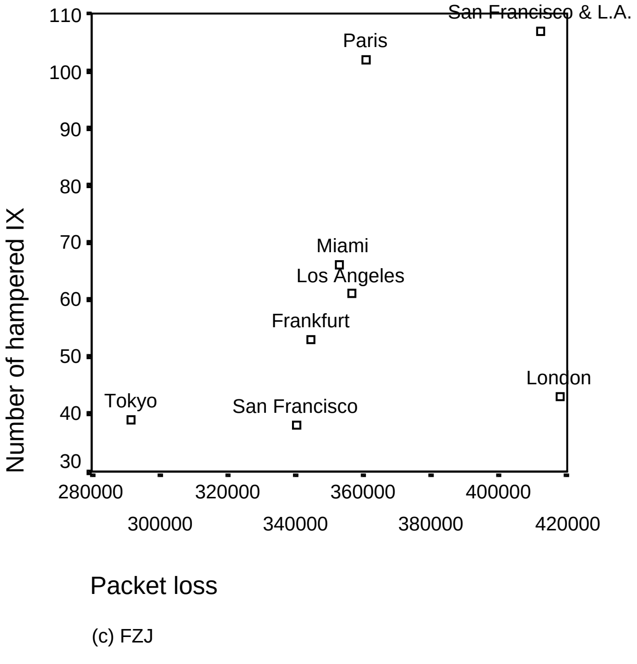

Figure 7: Classification of disturbance’ impact.

# References

[FFF99] Faloutsos, M.; Faloutsos, P.; Faloutsos, C.: On Power-Law Relationships of the Internet Topology. In Computer Communication Review, vol 29, no. 4, 1999.

[GOM03] Grubesic, T. H.; O’Kelly, M. E.; Murray, A. T.: A Geographic Perspective on Telecommunication Network Survivability. In Telematics and Informatics, vol. 20, 2003; pp. 51-69.

[LFV02] Lepperhoff, N.; Fischer, W.; Volst, A.: Simulation of Internet failures with INESS. In (Isaías, P. ed.): Proceedings of the IADIS International Conference WWW/Internet 2002, Lisbon, 2002. IADIS Press, 2002; pp. 276-282.

[Te01] Telegeography: International Bandwidth 2001. CD-ROM. Telegeography, Washington, 2001.

[Te02a] Telegeography: US Internet Geography 2003. CD-ROM. Telegeography, Washington, 2002.

[Te02b] Telegeography: Terrestrial Bandwidth 2003. CD-ROM. Telegeography, Washington, 2002.

[Te02c] Telegeography: Global Internet Geography 2003. CD-ROM. Telegeography, Washington, 2002.

[Te02d] Telegeography: Submarine Bandwidth 2002. CD-ROM. Telegeography, Washington, 2002.

[VLF02] Volst, A.; Lepperhoff, N.; Fischer, W.: Paketerzeugung und Datenverkehr im Internet. Forschungszentrum Jülich GmbH, Internal Report FZJ-STE IB-4/02. Forschungszentrum Jülich GmbH, Jülich, 2002.

[VPV02] Vázquez, A.; Pastor-Satorras, R.; Vespignani, A.: Internet topology at the router and autonomous system level. In arXiv cond-mat, 0206084 v1, 2002, URL: www.arxiv.org.