Fluctuations and Lasting Trends of QoS on Intercontinental Links

Tomasz Bilski

Poznan University of Technology,
Pl. Sklodowskiej-Curie 5, 60-965 Poznan, Poland
tomasz.bilski@put.poznan.pl

Abstract. The paper presents an analysis of short- and long-term changes in the QoS of intercontinental connections. First we will show that despite fast and numerous advances in physical layer, link layer, router capacity and new telecommunication cables deployment, QoS measures are hardly progressing in long-term (years) perspective. Transatlantic (North America – Europe) connections will be thoroughly analyzed. Next we will show that due to submarine cable breakages temporary network performance is unpredictable. It may be much poorer than average and sometimes drops below the acceptable level – case study is provided. Even if the links are fully operational, due to the rerouting the QoS may deteriorate in the case of cable fault in another part of the World. The research is based mainly on data taken from IEPM (Internet End-to-end Performance Measurement) database.

Keywords: QoS measurement, large scale networks, performance, reliability.

1 Introduction

It is well known that the term “quality of service” is used in many meanings ranging from the user’s qualitative perception of the service to a set of quantitative connection parameters (RTT, jitter, throughput, packet loss rate) necessary to achieve particular service quality. In the paper we will mostly use the second meaning of the term. This meaning is consistent with IETF approach presented in RFC 2386 [4].

It is relatively easy to provide high quality of service in short-distance, local connections. Small number of network devices, usually homogenous and managed by single service provider facilitates optimization with such techniques as MPLS, RSVP, header compression, TCP and web acceleration, redundancy and so on.

On the other hand accomplishing high QoS in intercontinental connections is particularly tough problem. First of all packet delay is much longer, with approximately 4.5 μs per each kilometer of fibre and more network devices introducing delays. Long-distance communication channels are usually shared by many users and institutions. They are heterogeneous: they carry different types of traffic: data, phone calls, ATM (Automated Teller Machine) transactions, they consist of diverse network devices. All that means:

- transient nature of IP behavior,
- more complex management,

N. Bartolini et al. (Eds.): QShine/AAA-IDEA 2009, LNICST 22, pp. 251–264, 2009.
© Institute for Computer Science, Social-Informatics and Telecommunications Engineering 2009

- more bottlenecks
- more points of failure,
- more difficult and time consuming diagnosis and fixing after a failure.

Repairing intercontinental submarine cable after breakage (especially in the case of earthquake) may take weeks and even months. Achieving high reliability with redundancy is very expensive. In addition network resources are administered by various network providers, with diverse and occasionally conflicting objectives.

There are many services for Internet performance measurement[1]. Some of them are integrated with databases with present and past measurements. The example is IEPM PingER, a service monitoring performance of Internet links, developed at SLAC (Stanford Linear Accelerator Center) and operating since 1995 (with data stored since 1998). Monitoring is based on more than 300 distributed hosts. Hosts send periodically pings for each tested connection. The measurement results are written to database. PingER database is used in the next parts of the paper.

2 Long-Term QoS Changes in Transatlantic Connections

2.1 Transatlantic Cables Deployment

Long distance connections are based on telecommunication cables and satellite transceivers. Satellite transmission QoS is relatively low, with significant delay and small bandwidth (e.g. one way signal propagation via geostationary satellite at 36 000 km altitude takes 260 ms [6]) and shorter design life (10-15 years compared to about 25 years for cable). So long distance communication is based mainly on cables. It may be assumed that 95% of all intercontinental links are based on cables.

Transoceanic cables are used since 1858. In 1988 first fibre-optic cable TAT-8 had been laid in North Atlantic with capacity of 280 Mbit/s. Since that year many

Table 1. Transatlantic (North America-Europe) cables operated in 1998 [www.iscpc.org]

Year of operation start	Name	Bandwidth [Gbit/s]
1988	TAT-8	0.28
1989	PTAT-1	0.42
1992	TAT-9	0.56
1992	TAT-10	0.56
1993	TAT-11	0.56
1994	CANTAT 3	2.5
1994	Columbus-II	1.68
1996	TAT-12/13	15
1998	Gemini	115
1998	AC1	40

---

[1] For example: http://www.internettrafficreport.com/ or http://visualroute.visualware.com/

Table 2. New transatlantic (North America-Europe) cables deployment in 1999-2008 [www.iscpc.org]

Year of operation start	Name	Bandwidth [Tbit/s]
1999	Columbus III	0.04
2000	AC-2	0.64
2001	Hibernia Atlantic	0.16 (lit) 1.92 (designed)
2001	TAT-14	1.87 (lit) 3.2 (designed)
2001	VSNL Transatlantic	5.12
2001	FLAG Atlantic 1	4.8
2003	Apollo	3.2 (designed)

transoceanic cables were deployed in North Atlantic with growing capacity per cable.

It may be assumed that total capacity of North Atlantic cables operated in 1998 was equal roughly to 150 Gbit/s (table 1). Approximate, total bandwidth of several North Atlantic cables deployed from 1999 to 2008 is 15 Tbit/s (table 2). Data are based on ICPC (International Cable Protection Committee) resources.

So we have approximately 100-fold increase in total submarine connections bandwidth in 10 years time. This is consistent with the rate in which DWDM fibre links are improved.

2.2 Lasting Trends in Performance

Introduction. IEPM database[2] preserves Internet performance data for many connections (pairs of monitoring sites) throughout the World [3]. The number of connections monitored at the given date varies in time. In the period 1998-2008 among 47 and 154 connections between Europe and USA were monitored. The connections are tested with packet sizes of 100 and 1000 bytes. In the next part average values of QoS parameters of the connections with 100 bytes per packet are provided.

Average, yearly (1998-2008) values of QoS parameters for Europe to USA connections are presented in table 3. The communication channels are asymmetric, so the parameters for reverse direction (USA to Europe connections) are different (table 4). Generally USA to Europe parameters are worse than Europe to USA. For example, more of the time, Europe to USA throughput is greater than USA to Europe (figure 1). The difference is in the range of about 140-600 kbit/s or 20-25%.

All analyzed QoS parameters have improved in a given period of time. Nevertheless the improvement pace is not smooth and rather slow compared to other computer and network performance indicators (e.g. bandwidth of Ethernet or Moore's Law).

---

[2] http://www-wanmon.slac.stanford.edu/cgi-wrap/pingtable.pl

Table 3. Long-term changes of QoS on transatlantic (Europe to USA) connections

Year	RTT [ms]	Jitter [ms]	Throughput [kbit/s]	Packet loss [%]
1998	219	0	530	5.0
1999	213*	160	531	4.3
2000	171	143	840	1.8
2001	153	3.7	1385	0.8
2002	151	149	1525	1.0
2003	150	227	1788	0.4
2004	146	82	1899	0.6
2005	155	348	2472	1.6
2006	155	84	2271	1.9
2007	154	1.9	1958	5.5
2008	154	2.1	2985	0.7

* Value calculated after excluding 2 anomalous numbers of RTT from the IEPM PingER database. Without excluding the numbers RTT for 1999 would be at the level of 9120 ms.

Table 4. Long-term changes of QoS on transatlantic (USA to Europe) connections

Year	RTT [ms]	Jitter [ms]	Throughput [kbit/s]	Packet loss [%]
1998	237	0.0	391	3.7
1999	233	853	429	5.1
2000	206	154	609	2.8
2001	252	69	1026	1.2
2002	227	46	1294	0.9
2003	183	3.8	1712	0.8
2004	169	3.5	2622	0.6
2005	158	66	2523	0.4
2006	171	1.6	2301	0.1
2007	172	1.4	1716	0.2
2008	177	1.2	2371	0.5

> *"(...) it would be easy to telegraph from Ireland to Newfoundland at a speed of at least from eight to ten words per minute." Samuel F.B. Morse, 1856*

Throughput. Throughput is a fraction of bandwidth. In the given period of time we've observed roughly 100-fold increase in total available bandwidth between North America and Europe (see 2.1). On the other hand, average throughput enhance is much smaller and irregular. In 1998 and 1999 throughput remained at the level of about 400-500 kbit/s. The following years brought relatively low increase to 2.4-3.0 Mbit/s (figure 1). So we've observed about 6-fold increase. This 2.4-3.0 Mbit/s throughput is enough for mail, web access and several VoIP channels but is much below threshold necessary for video applications with PAL-MPEG2 coding.

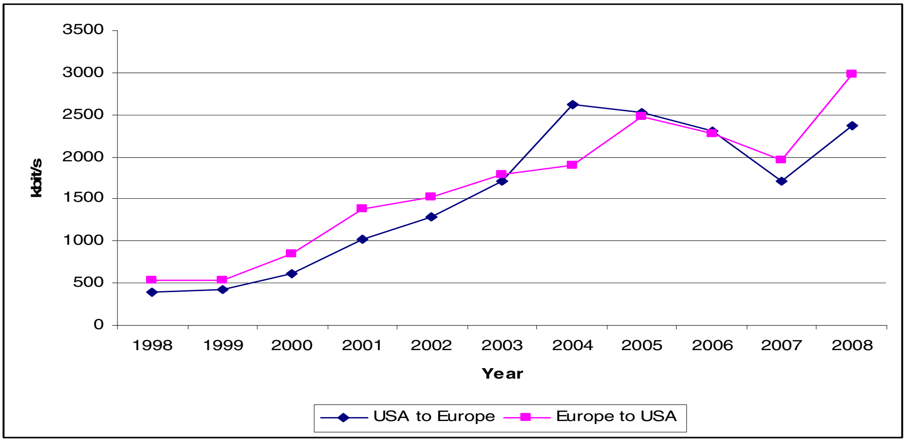

Fig. 1. Long-term changes of average throughput on transatlantic connections

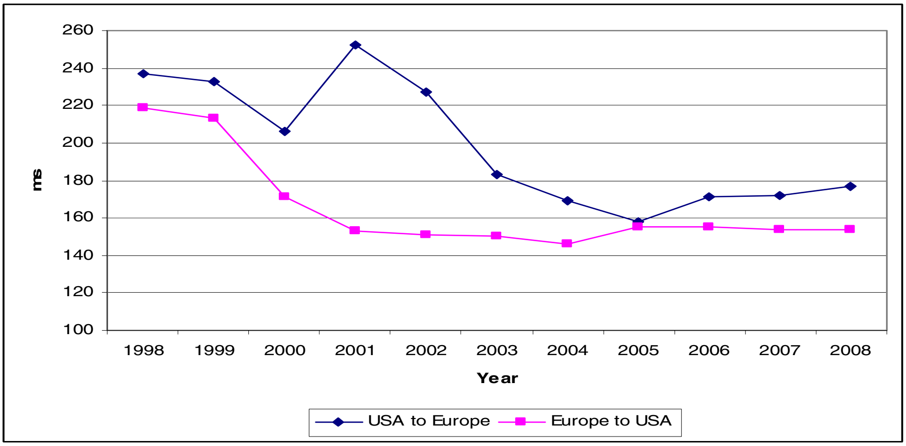

Fig. 2. Long-term changes of average RTT on transatlantic connections

RTT and Jitter. RTT is total propagation time needed to transfer IP datagram from one host to another and back. It includes delays in a fibre and all intermediate network devices (routers, switches, repeaters, transmit and receive terminals). Typical length of North Atlantic cable is between 6000 and 7500 km. The speed of infrared waves in a fibre[3] is close to 220 000 km/s, so the signal travels one way through the cable in 27-34 ms. In the worst case transmit terminal in submarine optical fibre system adds 13 ms delay and receive terminal 10 ms delay [6].

In 1998 average RTT remained at the level of roughly 220-240 ms. It decreased to the level of 150-180 ms in the first 4 years of 2000's and stayed at this level to 2008 with minor fluctuations (figure 2). So, one-way datagram delay is 75-90 ms. This

---

[3] In fact the speed is dependent on wavelength but the variations are negligible in this case.

value is well below ITU-T 200 ms threshold for absolute (mouth to ear) delay in VoIP applications in which users are very satisfied [6]. On the other hand the requirements for control and feedback in medical operations are hardly satisfied (it is assumed that one-way delay should be less than 80 ms in this case [5]).

Assuming that average one way delay between North America and Europe (in 2008) is in the range 75-90 ms, the cable propagation time (27-34 ms) is about 1/3 of the total delay and 2/3 is introduced by network devices.

The greatest fluctuations of average values are observable in jitter. It changes from below 2 ms up to several hundreds ms.

Packet Loss Rate. Average packet loss rate slightly improves in time. In 1998 it was at the level of 4-5%. It decreased to the level below 1% and stays at this level with some temporary deteriorations (e.g. elevated 5.5% rate in 2007 for Europe to USA connections). Packet loss rate below 1% is recognized as acceptable for interactive applications such as VoIP[4].

3 Short-Term Fluctuations in Normal Operation

Section 4 presents large changes to average daily QoS in the case of submarine cable fault. Before we present and analyze the data we should see how QoS parameters fluctuate (table 5) during normal operation time (without cable faults). The same

Table 5. Short-term fluctuations of QoS on exemplary Europe to India (n2.cern.ch to n1.cnieds.bangalore.in) connection during normal operation (16-28 January 2008)

Date	RTT [ms]	Jitter [ms]	Throughput [kbit/s]	Packet loss [%]
Dec 2007 average	184	7.9	1244	0.3
Jan 16	185	7.7	1027	0.4
Jan 17	180	6.8	564	1.3
Jan 18	193	9.2	1910	0.0
Jan 19	194	10.3	1901	0.0
Jan 20	179	7.1	2059	0.0
Jan 21	194	9.4	1022	0.3
Jan 22	205	9.2	1801	0.0
Jan 23	213	16.0	893	0.4
Jan 24	201	10.3	1333	0.2
Jan 25	197	10.3	461	1.7
Jan 26*	-	-	-	-
Jan 27	187	6.3	811	0.6
Jan 28	193	8.0	1394	0.2
Jan 29	190	10.4	524	1.4

* Data for January 26 are not available in the PingER database.

---

[4] Assuming other measures (delay, jitter) are at the satisfactory levels. Additionally, voice quality is related to the character (random or bursty) of the losses.

exemplary Europe to India[5] connection is evaluated in both cases (between n2.cern.ch and n1.cnieds.bangalore.in).

Fluctuations of QoS in the given period of regular operation time are visible. Minor fluctuations are observable in RTT and packet loss rate. RTT changes in the range of 179-213 ms. Packet loss rate do not exceeds 1.7%. Jitter stays below 16 ms (this is 2x more than average December 2007 jitter). Highest increases and decreases are noticeable in throughput. On January 20 it reaches more than 2 Mbit/s (166% of December 2007 average), but on January 25 it drops below 0.5 Mbit/s (37% of December 2007 average).

On the sidelines we may analyze weekdays fluctuations. Some minor fluctuations are visible in the average QoS of particular days of week. Throughout a given period of time (e.g. year) Saturdays and Sundays have highest average QoS and Mondays have lowest (it may be seen also in table 5: in the period 16-28 January 2008, highest throughput and lowest RTT are on Sunday, January 20). Nevertheless in a given week of the year it is not possible to predict a day with the highest QoS level.

4 Short-Term Fluctuations in the Case of Cable Fault

4.1 Submarine Cable Faults

Network reliability and performance is disrupted by intentional and unintentional factors: malicious software, spam, hackers and disasters (e.g. earthquakes). First 3 factors are widespread. Damages imposed by the factors are usually restricted to single services and are short-lived. On the other hand disasters are uncommon but their impact is long-lived. They disrupt Internet services, telephone calls and ATM transactions.

Intercontinental cable faults[6] are relatively infrequent (in 2003 annual fault rate was at the level of one fault per 10000 km of cable [7]). In the case of faults channel redundancy plays an important role. In the Atlantic, cable breaks happen repeatedly (more than 50 cable repairs are yearly in the Atlantic) but due to the high level of redundancy (about 20 cables connect nowadays Europe with North America) they are almost invisible to end users. On the other hand disasters on Mediterranean inflict big impact on Internet services. Submarine cables are prone to being affected by earthquakes, storms, fishing and anchors. 70% of faults are attributed to human activity with fishing as the major cause [7]. Usually, earthquake extent of the damage is much greater than anchor cut, the cable may suffer several breaks. The severed ends could be buried by deep-sea landslides or washed kilometers from their previous positions. It may take many days to just find the cable and months to full service restoration.

In November 2003 TAT-14 (USA-Europe) cable fault occurred. Many Internet service providers in UK experienced some problems. As TAT-14 is a dual, bi-directional

---

[5] India is an IT outsourcing centre, with Bangalore as India's Silicon Valley, so the connections to India's networks are particularly important for US and European businesses.

[6] Satellites are also vulnerable, they may move away from orbit or may collide (e.g. on February 11 2009 Iridium Satellite collided with Russian Cosmos 2251 satellite, an incident resulted in limited disruptions of Iridium service).

ring of cable, a single serious fault should not be enough to break it, as traffic would still be able to flow between the countries on the ring. Unfortunately, a part of the cable near the USA coast had already suffered a technical fault few days earlier, which meant there was no built-in redundancy.

In May 2003 6.8 magnitude earthquake with epicenter near Algiers damaged five submarine cables. All Algerian voice, mobile and Internet traffic was disrupted. The last repair completed only 6 weeks after the earthquake.

In June 2005 SMW 3 cable had been cut off Karachi. Pakistan lost all terrestrial Internet connectivity. The outage of services lasted 12 days.

On December 17, 2006, CANTAT-3 cable connecting Iceland with Europe and Canada was damaged. Most notable effects of the event was a temporary shut-down of data-communications by Iceland's universities and hospitals which rely exclusively on CANTAT-3's services. It took more than 7 months, until July 29, 2007 before service was fully restored.

On December 26 2006 a 7.1-magnitude Hengchun earthquake south of Taiwan knocked seven submarine cables out of service, impairing communications from North America to China, Taiwan, Japan and Korea as well as inside North and Southeast Asia. Cables accounting for 90% of telecommunications capacity of the region had been broken. It took 49 days to repair all the cables.

On January 30 2008 a series of accidents with Mediterranean cables started (p. 4.2).

On December 19 2008 five submarine cables (including SMW 3, SMW 4, FLAG, Seabone) had been cut near Sicily due to 5.3 magnitude quake in the central Mediterranean. The cut disrupted Internet and telephone services in the region and in parts of the Middle East and South Asia [2].

4.2 Case Study

Introduction. January 2008 Mediterranean accident is an interesting research subject from many points of view. It consisted of several events. It occured in the area of high traffic and relatively small connection redundancy. It showed that cable faults have high and distributed impact on Internet performance. It demonstrated several problems with current Internet infrastructure. The accident is well documented. News networks provided many reports. IEPM preserved data (RTT, throughput, jitter) on net performance during the accident. RIPE [9] (Réseaux IP Europeens) and Renesys [8] researched IP route changes.

Accident Timetable. The accident was a series of events. Cables from Europe to Middle East and Asia were affected: January 30, 2008 about 4:30[7] SMW 4 cable (deployed in 2005, total capacity 1.28 Tbit/s) near Alexandria was damaged, January 30 about 8:00 FLAG cable (deployed in 1997, total capacity 10 Gbit/s) near Alexandria was damaged. The two cables carry about 70% of the traffic between Europe and the Middle East. February 1 about 6:00 FALCON cable near Dubai was damaged, February 8 SMW 4 cable was repaired, February 9 FLAG cable was repaired, February 10 FALCON cable was repaired [9]. It is assumed that the cables were damaged by anchors.

---

[7] All times are UTC.

Table 6. Fluctuations of QoS on exemplary Europe to India (n2.cern.ch to n1.cnieds.bangalore.in) connection during cable fault

Date	RTT [ms]	Jitter [ms]	Throughput [kbit/s]	Packet loss [%]
Dec 2007 average	184	7.9	1244	0.3
Jan 28	193	8.0	1394	0.2
Jan 29	190	10.4	524	1.4
Jan 30*	-	-	-	-
Jan 31	614	6.8	40	23.0
Feb 1	563	66.5	38	30.9
Feb 2	502	13.0	54	18.4
Feb 3	479	6.8	131	3.5
Feb 4	526	14.7	106	4.4
Feb 5	658	8.5	561	0.0
Feb 6	619	22.1	434	0.2
Feb 7	486	6.7	194	1.5
Feb 8	369	5.9	564	0.3
Feb 9	184	1.4	786	0.6
Feb 10	193	8.7	874	0.5
Feb 11	195	10.6	1897	0.0

* Data for January 30 are not available in the PingER database.

IP Rerouting. Three general types of effects to IP routing are possible in the case of submarine cable cut: immediate loss of connection between some networks after the failure, much smaller number of available AS (Autonomous System) paths and likely rerouting with a use of backup paths. These backup paths are longer and offer poorer performance. A good indication of the impact is the number of IP address network prefixes, that are announced in BGP messages. If prefix to a given network is not announced to routing peers then the network is not reachable. In some countries (Egypt, Kuwait, Sudan), immediately after the failure, more than 30% of the prefixes were removed from BGP announcements. The total number of AS paths decreased and many networks disappeared from the routing tables. The number of changes in AS paths rapidly increased in the region. Not more than 10% of paths change daily in normal operation routing mode. This measure grew, in the region, to above 60% on January 30. This higher than normal and fluctuating (between 10–30%) paths change percentage persisted up to February 18. Additionally, average AS path length (the number of different ASes between two distant networks) increased slightly, around the time period of the cuts from 5.5 to about 6.

This was caused by rerouting of some connections. Europe-Asia connections were rerouted through the SMW 3 (deployed in 1999, total capacity 20 Gbit/s) cable or fibres taking the way around the globe (Europe-USA-Asia). Due to the limited bandwidth and traffic increase on these new routes it was difficult to quickly converge routing tables on alternate topologies [9]. This is an effect of distance-vector algorithm for BGP routing optimization. It will be shown (table 7) that traffic rerouting had significant impact not only on rerouted connections but also on some of the USA to Asia connections. IP routes modifications resulted directly in significant changes of RTT, throughput and packet loss rate. Parameters returned fully to their values from before accident only several days after the time in which all repairs were completed.

Fluctuations of performance parameters on exemplary Europe to India connection (between n2.cern.ch and n1.cnieds.bangalore.in) are presented in the table 6 and discussed in the next subsections. It is evident that the connection was directly hit by the cable cut.

Since the incident had indirect impact on QoS in other parts of the World, network parameters on exemplary USA to India connection (between n8.doe.gov and n1.cdacmumbai.in) are presented in the table 7 and discussed in the next subsections.

Table 7. Fluctuations of QoS on exemplary USA to India (n8.doe.gov to n1.cdacmumbai.in) connection during cable fault

Date	RTT [ms]	Jitter [ms]	Throughput [kbit/s]	Packet loss [%]
Dec 2007 average	257	7.5	223	4.2
Jan 28	255	4.1	203	5.1
Jan 29	254	4.4	224	4.2
Jan 30*	456	12.5	39	43.3
Jan 31	515	18.6	34	45.2
Feb 1	401	18.7	53	30.5
Feb 2	421	14.9	58	22.7
Feb 3	382	15	128	5.7
Feb 4	369	13.8	98	10.4
Feb 5	357	9.2	145	5.1
Feb 6	348	6.2	114	8.7
Feb 7	356	11.5	79	17.2
Feb 8	362	24.1	110	8.6
Feb 9	387	32.9	168	3.2
Feb 10	320	15.8	203	3.2
Feb 11	257	2.7	220	4.3

Throughput. Average monthly throughput (assuming TCP connections with 100 bytes per packet) for Europe to India connection was in December 2007 at the level of 1244 kbit/s. It rapidly decreased on January 31 to 40 kbit/s. That means 30-fold worsening. It should be observed that this is unacceptable throughput level for VoIP services. Throughput remained below 1000 kbit/s up to February 10 (figure 3).

Similarly, average monthly throughput for USA to India connection was in December 2007 at the level of 223 kbit/s. It rapidly decreased on January 30 to 39 kbit/s. Throughput remained well below 200 kbit/s up to February 9 with temporary improvement (to 145 kbit/s) on February 5 (figure 3).

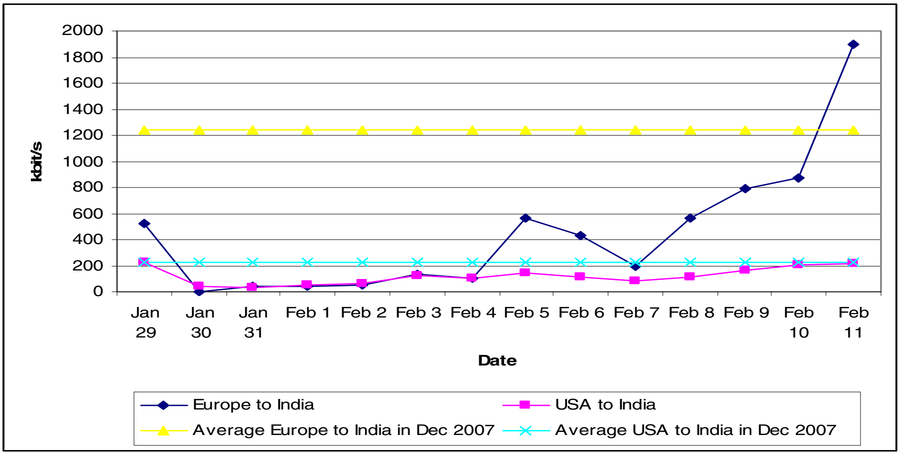

Fig. 3. Fluctuations of throughput during cable fault

RTT and Jitter. Average monthly RTT in December 2007, for exemplary Europe to India connection, was 184 ms. Connections between Europe and India had been significantly affected: RTT increased 3x on January 31 to above 600 ms. Changes were inflicted by traffic rerouting through USA networks which started to carry additional traffic. In consequence connections between USA and India had been affected too: average RTT (USA to India) increased 2x on January 31 from 257 ms (average in December 2007) to 515 ms. RTT increased due to longer routes, more routers and longer queues in routers, which have to carry additional traffic. Average RTTs started to improve after January 31 and returned with fluctuations (especially on February 4-5 on Europe to India connection) to their values from before the accident after 12 days (figure 4).

Average monthly jitter for Europe to India connection was (in December 2007) at the level of 7.9 ms. It slightly increased during first 2 days to 10.4 ms and sharply increased on February 1 to above 66 ms. Next days it started to improve with oscillations. Average monthly jitter for USA to India connections was in December 2007 at the level of 7.5 ms. It increased during first 2 days to 19 ms. The highest jitter 24–33 ms appeared on February 8 and 9, at the time of cables repairment.

Packet Loss Rate. Average monthly packet loss for Europe to India connection was in December 2007 at the level of 0.3%. It increased after cable fault up to 30.9%. This level of packet loss means total unavailability of interactive services including e.g. control and feedback for remote medical operations [5]. Average monthly packet loss for USA to India connections was in December 2007 at the level of 4.2%. It increased on January 31 to 45%. Packet loss rate temporarily recovered in next few days. On February 3 we observe decrease for Europe to India and USA to India connections. In both cases the level not greater than about 6% has been achieved. In the next few days it fluctuated between 0 and about 20% and returned to its values from before the accident on February 9.

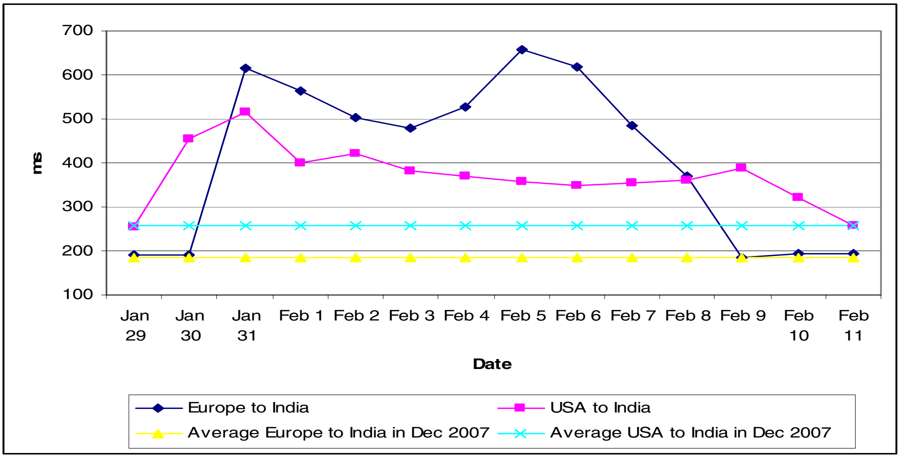

Fig. 4. Fluctuations of RTT during cable fault

Cable Cut Impact on QoS. Internet survived the accident. Nevertheless the disaster's impact on QoS was: significant, long-lived and widespread (cable cut near Alexandria degenerated USA to Asia connections which do not use the cable). The performance parameters during the accident decreased to unacceptable (for interactive applications) levels. Internet performance have been changing in unpredictable way (it is seen for example on throughput and RTT data). The effects of submarine cable cut accident are notably different from the effects of hacker attack on Internet server. The analyzed accident was not an exception. Similar accidents in the future should be expected. Generally we are not able to predict time and place of the accidents (e.g. earthquakes are hardly predictable) [1].

5 Conclusions

Intercontinental connections performance is upgrading very slowly and irregularly. Compound growth of throughput for exemplary transatlantic connections in 10 years is just 6-fold on average. The increase rate is lagging behind other computer and network performance indicators. At the same 10-years time total North Atlantic bandwidth increased approximately 100 times. This is related to fiber transmission capacity and DWDM link speed, which grow by a factor of about 200 in the decade. Similarly, according to Moore's Law computer power increase in the same period is 100-fold. The router capacity, which takes advantage of the Moore's Law increases at approximately the same rate. Hard disk data areal density grows also at the similar rate. Ethernet bandwidth is growing exponentially from 10 Mbit/s in 1989 to 10 Gbit/s in 2002 and 100 Gbit/s in 2010 (2010 is feasible year of the new standard ratification) – this means 10000-fold increase in 20 years. Discrepancy between different factors influencing IT performance is visible.

QoS of intercontinental connections is hardly predictable. Day to day fluctuations are large, even in normal operation time. Irregular disasters with submarine cable faults make this predictability more complicated. In the case of such fault temporary level of performance may rapidly drop to unacceptable value (e.g. throughput 40 times lower than average). Some cable faults have high and distributed impact on Internet performance. Significant QoS deterioration may last days and sometimes weeks and months. Such disasters have impact on many telecommunication services: email transfers, web access, telephone calls and ATM transactions. There are differences between cable fault and intentional attack consequences. In a typical case damages imposed by hackers and malicious software are usually short-lived and restricted to single computer or service.

It is clear that many things should be done to improve QoS on long-distance connections. The changes ought to be implemented in networks as well as in hosts and applications.

Network (fibre) redundancy should be carefully planned at every infrastructure level. For example, for end user it is pointless to use two ISPs if both utilize the same international cable. Of course more cables are needed. Their location should be better planned. Existing cables should be upgraded so that they are operating at a percent of their potential capacities, leaving plenty of room not only for future traffic growth but also for rerouted (in the case of disaster) traffic. BGP protocol should be upgraded or replaced with completely new one EGP protocol. Routing table convergence time should become important optimization criterion. More efficient implementations of TCP (MulTCP, HighSpeed TCP) should be integrated with common operating systems.

Globally used services should be based on data mirroring, caching proxies, CDNs (Content Delivery Networks). Applications, which vary in their QoS requirements should be better profiled before determining appropriate classification and routing treatment. Both versions of IP are ready for such data classification. Routers should be aware and able to carry differentiated traffic. To save bandwidth real time applications should be based on more efficient codecs. RTP/UDP/IP header compression should be broadly utilized.

Of course it is easy to recommend actions, procedures, modifications and much harder to apply them. Suggested improvements are not easy to implement but many of the techniques, algorithms, protocols and their implementations are already available.

References

1. Bilski, T.: Disaster's Impact on Internet Performance - Case Study. In: Kwiecień, A., Gaj, P., Stera, P. (eds.) CN 2009. CCIS, vol. 39, pp. 210–217. Springer, Heidelberg (2009)
2. Cottrell, L.: Effects of Mediterranean Fibre Cuts December 2008, SLAC (2009), `https://confluence.slac.stanford.edu/display/IEPM/Effects+of+Mediterranean+Fibre+Cuts+December+2008`
3. Cottrell, L., Matthews, W., Logg, C.: Tutorial on Internet Monitoring and PingER at SLAC. In: SLAC 2007 (2007), `http://www.slac.stanford.edu/comp/net/wanmon/tutorial.html`
4. Crawley, E., Nair, R., Rajagopalan, B., Sandick, H.: RFC 2386 – A Framework for QoS-based Routing in the Internet. Network Working Group (1998)

5. Gutierez, D., Shah, A., Harris, D.A.: Performance of Remote Anatomy and Surgical Training Applications Under Varied Network Conditions. In: Barker, P., Rebelsky, S. (eds.) Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications 2002, pp. 662–667. AACE, Chesapeake (2002), `http://havnet.stanford.edu/pdfs/ed-media.pdf`
6. ITU-T Recommendation G.114. One-way transmission time, ITU-T 05/2003
7. Kordahi, M.E., Shapiro, S.: Worldwide Trends in Submarine Cable System Faults. In: SubOptic 2004 (2004), `http://www.suboptic.org/Viewdocument.aspx?id=381`
8. Popescu, A., Premore, B., Zmijewski, E.: Impact of the Middle East Cable Breaks. A Global BGP Perspective, Renesys Corp. (2008), `http://www.renesys.com/tech/presentations/pdf/nanog42-lightning.pdf`
9. Wilhelm, R., Buckridge, C. (eds.): Mediterranean Fibre Cable Cut – a RIPE NCC Analysis, RIPE (2008), `http://www.ripe.net/projects/reports/2008cable-cut/index.html`