Measurement and Analysis of End-to-end Dependability Characteristics in the Global Network

Eugene S. Myakotnykh, Bjarne E. Helvik
Centre for Quantifiable Quality of Service in Communication Systems (Q2S) [1]
Norwegian University of Science and Technology, Trondheim, Norway

Otto J. Wittner, Olav Kvittem, Jon Kåre Hellan, Trond Skjesol, Arne Øslebø
UNINETT [2]

Abstract—**This paper presents an analysis of measurements on a global route between Norway and China. Measurement statistics was collected every 10 milliseconds during a three-month period in October-December 2009. We provide a detailed characterization of observed loss and delay patterns especially focusing on downtime periods on the sub-second time scale (from 50 milliseconds to 1 second long). Generally, these downtimes are more difficult to measure, but they may also have a noticeable negative effect on quality of real-time services and multimedia applications. The paper investigates causes for all these events and finds that the plausible causes most of observed communication problems in the given study are either link or router related (level 2 or 3). Packet loss events caused by congestion could also be distinguished, but they happened very rarely. Using the collected data set, we also study several other questions related to routing and network availability modeling.**

Keywords- dependability, failure analysis, network measurements, modeling, Quality of Services, routing

I. INTRODUCTION

A growing number of real-time communications services and applications like voice-over-IP, videoconferencing, on-line gaming and distributed music collaboration are sensitive to irregularities in the transport service. Transmission outages in the order of seconds or minutes have for long been considered as significant; downtimes in the order of 100 ms may also have a noticeable negative effect on quality of interactive applications. To improve the quality of real-time services, it is important to understand in detail reasons for packet loss, delay and jitter caused by the network.

On the operational arena, the measurement of quality has often been done on an elementary level like study of the dependability of individual network elements, and also on the convergence of interdomain and intradomain protocols. However, even if there are redundant physical paths within and between networks, the link and network levels would incur some time to detect errors, for route recalculation and to distribute changes. Such fault handling will influence traffic forwarding, and will also be experienced by the end systems as increased packet loss, delay and jitter. To what extent such disruptions can be seen, what downtimes they do cause and how significant they are compared to downtimes caused by other reasons like congestion, is of particular interest.

In this study, active measurements between two systems located in Trondheim, Norway and Beijing, China were performed and downtime statistics was continuously collected during a three-month period in October-December 2009. Measurement experiments lasting during a long period of time (weeks or months) and having such a high "granularity" (packets between the end-points are sent every 10 milliseconds) are quite rare. We can detect and measure not only long outages between the two end-points (in order of seconds or minutes), but also investigate impairments with sub-second duration. Some results of this study like network availability or packet delay patterns must be with care generalized (they depend on a chosen network path), but investigated reasons for downtimes and the approach to distinguish and to classify these events are likely to be the same for other global paths.

The paper is organized as follows: the next section provides necessary background information and briefly reviews some previous work related to network measurements. Section 3 describes the measurement methodology used in the project. Section 4 present numerical results, Sections 5 focuses on deeper analysis and classification of observed delay and loss patterns. Conclusions are drawn and future work is discussed in Section 6.

II. BACKGROUND AND PREVIOUS WORK

Only downtime events exceeding 50 ms are analyzed in this paper. It is assumed that shorter loss events can be effectively compensated on the receiver side. And 50 ms is commonly regarded as the largest permissible time for interruptions from fault handling [1]. Outages in order of 50-100 ms may already cause noticeable degradation of service for many interactive applications like VoIP, videoconferencing, and gaming.

[1] "Centre for Quantifiable Quality of Service in Communication Systems, Centre of Excellence" appointed by The Research Council of Norway, funded by the Research Council, NTNU and UNINETT. http://www.q2s.ntnu.no

[2] UNINETT is a group of companies which supplies network services for universities, university colleges and research institutions in Norway and handles other national ICT tasks. The Group is owned by the Norwegian Ministry of Education and Research. http://www.uninett.no

In the measurement experiment described in this paper, packets are sent between Norway and China in both directions. It is well known that Internet routes are not symmetric: the forward and reverse paths between two nodes may not be the same due to policy-based interdomain routing and other traffic engineering mechanisms [2]. For this reason, packets are sent and statistics is collected on the both end-points.

It is a challenging task to distinguish reasons for downtime events, which may happen on different levels, in different places (inside of an Autonomous System (AS) or in interdomain), have different causes and, as a result, various durations. Interior Gateway Protocols (IGPs) such as OSPF and IS-IS are commonly used to select a path to route traffic intradomain (within an AS). Border Gateway Protocol (BGP) is used to route traffic between neighboring ASes. Unfortunately, the links used in IP networks are not perfectly stable and link failures or router-related problems are common events. Markopoulou [3] extensively studied causes and durations of intradomain failures analyzing measurements data (IS-IS routing updates) from the Sprint backbone IP-network. Detected failures were classified based on observed patterns and parameterized using well-known distributions. Studies of Francois [4] and Shaikh [5] indicate that the convergence time for intradomain routing protocols (time required by all routers in an AS to go back to steady state operation after a change in the network state) is within less than one second. Other techniques may be used to achieve a faster intradomain link restoration. For these techniques, the target is usually to restore a failure within 50 ms. Brief overview of these approaches and references to other papers and standards can be found in [6].

In addition to intradomain link failures, interdomain communications may also experience outages. The BGP protocol recovers communications in this case and this process may take tens of seconds or even minutes [7, 8]. A number of measurement studies, for example, [9, 10], shows that even if there are redundant physical paths within and between networks, the link level and routing protocols would incur some time to detect errors and to distribute changes.

The objective of this paper is to analyze data collected during the three-month measurement experiment, to see how frequently certain events happen and possibly to explain the nature of the downtimes on a path across multiple domains (congestion, BGP updates, link failures, etc.). The data are also used to evaluate existing approaches of failures modeling.

III. Measurement Methodology

The set-up constitutes of measurement systems located at UNINETT facility in Trondheim, Norway, and at CERNET (China Education and Research Network) in Beijing. UNINETT and CERNET are interconnected through the Global Research Network crossing several Research Networks like NORDUnet (40 Gbps backbone between Norway, Sweden, Denmark and several other countries), Geant2 (10 Gbps in Europe), and TEIN3 (Trans-Eurasia Information Network, 2.5 Gbps channel exists between Europe and China). Fig. 1 shows roughly the most frequent route between the end-points. In case of traffic rerouting, which may be caused by different reasons, alternative paths can be used. An approximate distance between the end-point is around 12 hops including hops in the UNINETT and CERNET access

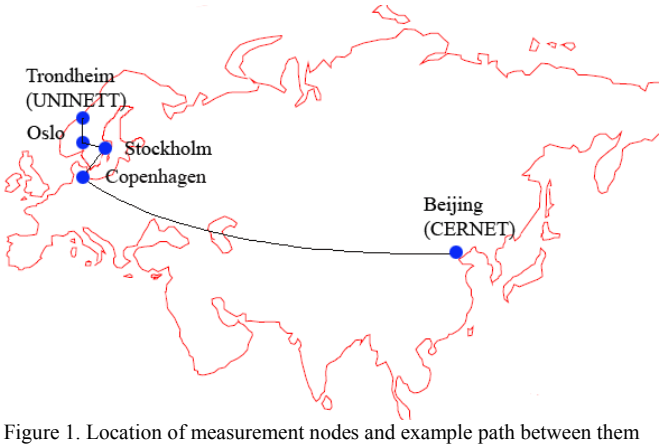

Figure 1. Location of measurement nodes and example path between them

networks. More details about other possible paths between the measurement systems will be presented in Section 5.The data are currently collected with active measurements. The main means of observation is to send an evenly distributed stream of packets between the nodes and measure delay, jitter and intervals with lost packets. The interval between probe packets was chosen to be 10 ms. We will thus be able to measure with acceptable accuracy the failure handling time goal of 50 ms for SDH [1, 4]. The end systems are synchronized using NTP protocol and close to Tier 1 clock sources. The programs used for the measurement streams were rude/crude [11]. Each packet is timestamped both on sending and receiving and also has a sequence number. Hence jitter, loss and reordering can be obtained.

Standard workstations are connected close to the backbone path. In the future, passive measurement equipment (workstations with special "wire-tap" interface cards) will be used to verify the accuracy of the NTP-based measurements.

Because the measurement points are located close (a few hops) from the backbone, the effect of access networks is considered to be insignificant. The route goes through the Global Research Network, which has high capacity and does not (at least, should not) experience congesting traffic loads. So we do not expect to detect a large number of downtime events, i.e. series of packet losses, due to high traffic load. Thus we can analyze causes for every detected outage individually. Observations and conclusions presented in the rest of the paper could help us to understand the nature and the variety of network downtime events.

IV. Network Downtime Characteristics

The statistics of downtime periods exceeding 50 ms, collected in October-December 2009 in both directions is summarized in Table 1.

The results are stable and similar for both directions and for each month except December (UNINETT-CERNET). In December 2009, five periods with bursty packet losses and with durations from 8 seconds to 1 minute were detected on the path from Norway to China, i.e. only in one direction. A more detailed analysis of these singular events will be provided in the next Section. If these long bursty loss patterns are not included, the number of downtimes during the month is 32, which is shown in the Table in parentheses and which is also similar to the data collected in October and November.

TABLE I. OVERVIEW OF OBSERVATIONS.

OCTOBER - DECEMBER 2009, NUMBERS IN PARANTHESES DO NOT INCLUDE FIVE DETECTED PERIODS WITH BURSTY PACKET LOSS

Outage period	Number of events					
	UNINETT-CERNET			CERNET- UNINETT		
	Oct.	Nov.	Dec.	Oct.	Nov.	Dec.
50 ms – 100 ms	13	19	73 (9)	4	6	12
100 ms – 1 s	9	6	40 (7)	7	9	5
1 s – 10 s	9	4	49 (9)	8	2	10
More than 10 s	1	2	7 (7)	1	2	4
Total # of events exceeding 50 ms	32	31	169 (32)	20	19	31
Total duration of all the events (s)	189.7	255.6	539.4 (472)	182.3	197.9	209.0
Path availability (%)	99.993	99.990	99.979	99.993	99.992	99.992

The total number of the detected events is not very large; about one event per day on average. The total duration of all outage periods exceeding 50 ms is around 200 seconds per month (except December), which gives the service availability about 0.9999. The data are not exactly symmetric: the path from Norway to China experienced more downtime events with duration in the range of 50-100 ms. The two longest downtimes were around 165 seconds and occurred simultaneously on the both routes in October and November. The minimum end-to-end delay between the end-point in Norway and China along the given paths is around 100 ms.

A convenient way of presenting network availability as a function of acceptable downtimes was proposed by Norros [12]. Instead of just a single number like 0.9999 defining network availability, it was proposed to use downtime-frequency curves that characterize the frequency of each down-period length separately. Using this approach, network operators in their SLAs can specify network availability not only as a single number, but also as a function of downtime durations. The SLA curves should lie above the curves defining the actual system availability. Fig. 2 shows the curves for the both paths during the three-month period. T – downtime duration in seconds (log-scale); Unavailability = 1 – Availability = P(W>T) – network unavailability for downtime events exceeding a certain threshold T, log-scale also. The Figure shows that downtimes in the order of tens of seconds most affect the availability. But, the relatively large number of loss periods in the sub-second range will also affect the perceived quality of real-time services.

In the literature, for instance, in [3], a cumulative density function (CDF) of the distribution of downtime events is approximated with CDF of the Weibull distribution [14] $F(x) = 1–exp(–((x–0.05)/\alpha)^{\beta}),\ x \geq 0.05$ (only losses exceeding 50 ms are analyzed in the paper). We have collected a large volume of measurement data and can verify this model. The five periods with bursty loss are not included to this analysis. We feel that these events have a totally different nature: they may be caused by operation, but not by random loss processes in the network. Matlab *cftool* fitting toolbox was used for the

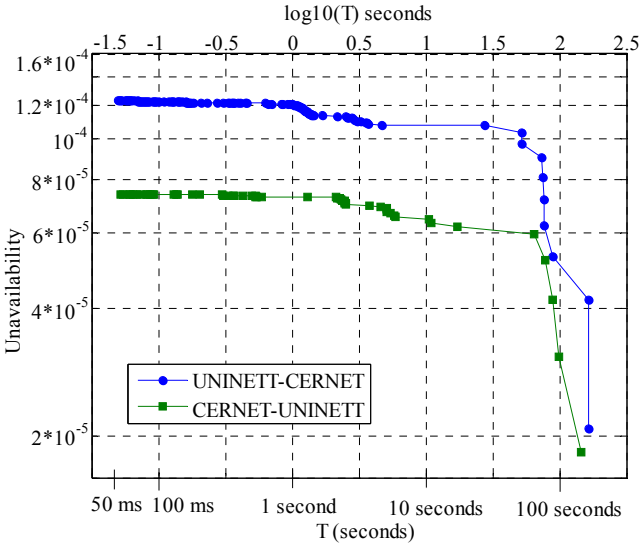

Figure 2. Downtime-frequency curves built based on the measurement results

rest of the data to find the coefficients. The results are presented in Table 2.

Further analysis showed that the model fits relatively well for downtimes periods not higher then, approximately, 5 second (Fig. 3). For longer outage periods, the fitting is noticeably less accurate. Fig. 4 shows the same fitting curve as in Fig. 3, but for the range of downtimes exceeding 5 seconds. This can be explained by the different nature of failure events: most of intradomain downtimes are relatively short (in most cases, up to several hundred milliseconds; see Section 5), long outages may be caused by operational and maintenance activities and interdomain rerouting. They can last tens of seconds or even minutes.

TABLE II. WEIBULL DISTRIBUTION FITTING

Path	α	β	R-square	RMSE
UNINETT-CERNET	0.52	0.30	0.980	0.037
CERNET- UNINETT	1.08	0.38	0.973	0.049

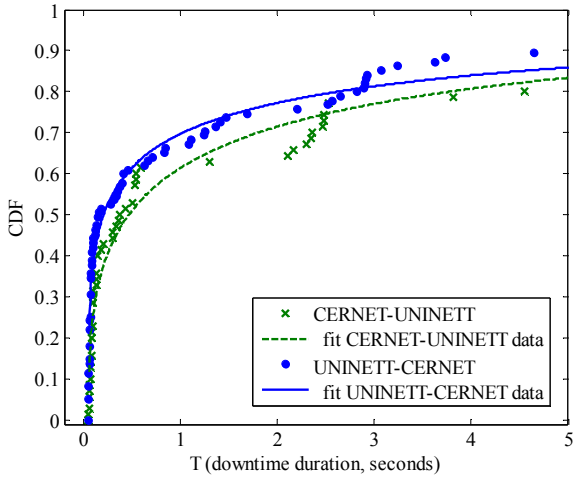

Figure 3. Fitting measurement data with Weibull CDF function (for downtimes in the range 0.05 – 5 second)

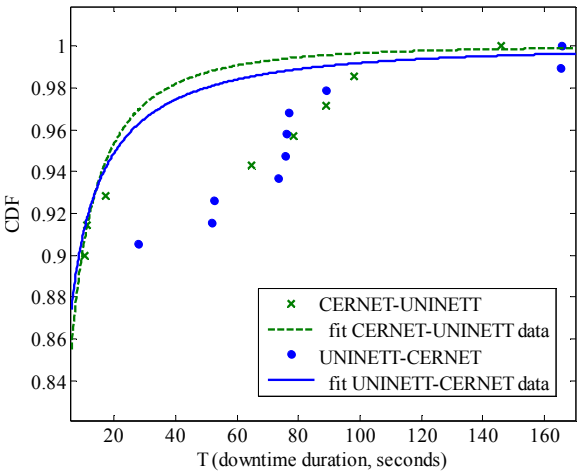

Figure 4. Fitting measurement data with Weibull CDF function (for downtimes exceeding 5 seconds)

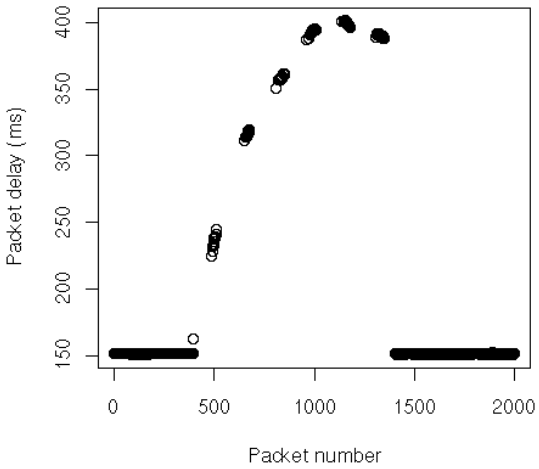

Figure 5. Example of a busty packet loss event

V. Downtime Periods Classification

In this Section, we provide a more detailed analysis of all the detected downtime periods. Examining delay patterns before and after the downtimes, we try to explain plausible causes for these events.

As it was mentioned in the previous Section, five periods were detected in December 2009 on the UNINETT-CERNET route with durations from 8 seconds to 1 minute and with a very significant loss in the range 35-95% during the intervals. They result in 137 out of the 169 loss events. Further analysis demonstrated that delay patterns in all these cases are very similar and example of received packet delays during one of the periods is shown in Fig. 5.

During the considered 8-second period, 15 packet loss events occurred with duration between 70 ms and 1.5 seconds. This delay-loss pattern is likely to be explained by congestion in the network and the Random Early Detection (RED) [13] queue management algorithm. RED is widely implemented in major commercial routers, especially in edge routers. The steady increase in packet delay between the end-points indicates congestion in a certain place of a path between Norway and China. The sequences of dropped packets are probably caused by the queuing mechanism. Also, in four cases out of the five, the fixed part of delay has slightly changed after the bursty loss and a new path was used to send traffic. This may indicate the use of certain load balancing mechanisms. Although the hypothesis about congestion and the queue management algorithm sounds reasonable, this question will be further investigated in a lab environment.

To identify causes of other measured downtime events, it is also helpful to know a behavior of packet delays before and after each downtime interval. 20 packets (200 ms) before and after each event were used; their average delays and standard deviations of the delay (indicator of jitter) were calculated. Then, change in the network delay was computed as the absolute value of difference between the delays. If this change is statistically different from zero, this indicates us that a new route was chosen in intra- or interdomain to deal with an interruption. Table 3 contains results of this analysis during the whole period of measurements. It shows not only differences in delay for both directions, but also how these changes correspond to the duration of outage periods. The used notation is x / y; x – statistics for the UNINETT-CERNET route, y – statistics for the CERNET-UNINETT route. The column "others" contains statistics of special interesting cases demonstrated, for example, in Figures 6-11 and which will be discussed separately. In these cases, the difference in delay was not statistically different from zero with 95% confidence. The five bursty periods discussed above, are not included to the statistics.

Table 3 shows that in approximately 20% of cases (21 out of 95 and 12 out of 70) no change in packet delay was detected after a downtime interval. Example of such scenario is shown in Fig. 6. Note that, for simplicity, Figures 6-10 show

TABLE III. Absolute change in Delay before and After Downtime Periods

October - December 2009, UNINETT → CERNET / CERNET → UNINETT

Change in delay / Outage period	0 ms	> 0 ms – 10 ms	10 ms – 100 ms	> 100 ms	others	Total
50 ms – 100 ms	10 / 3	24 / 19	0 / 1	0 / 0	7 / 0	41 / 23
100 ms – 1 s	9 / 5	10 / 11	2 / 1	1 / 0	0 / 3	22 / 20
1 s – 10 s	0 / 2	3 / 5	8 / 2	11 / 7	0 / 3	22 / 19
More than 10 s	3 / 2	3 / 1	0 / 2	4 / 2	0 / 1	10 / 8
Total	22 / 12	40 / 36	10 / 6	16 / 9	7 / 7	95 / 70

downtime intervals equal to 10 packets (100 ms). In most cases when the delay before and after a down-period did not change, the loss duration was in the range up to 1 second (up to 150 ms in more than 50% of cases).

It is natural for the routing protocols to wait for a certain period of time (called as Dead timer for OSPF [15] and as Hold time for IS-IS [16]) before switching to a new forwarding table computed for the network without the failed link. This period of time depends on a router configuration, but it is in the order of several hundreds milliseconds (a default value for IS-IS Hold timer is 1 second [16]). If the failed connection is restored during the period time, the same "old" route is used. That is why, the pattern is observed when a communication disruption has occurred, but the traffic route did not change. This loss pattern may also happen simply because of transmission error in the underlying transport network/layer. In this scenario, it is also observed that in certain cases the detected failures exceed one second. This may have happened due to events involving interdomain links.

In the remaining 80% of cases, delays after measured loss periods have changed (either increased or decreased) compared to packet delays before the events. In most of cases, the difference in delay was statistically different from zero, i.e. it is likely that a new route was chosen for the traffic. Examples of such scenarios are shown in Figures 7-8.

Step changes in delay profiles are very important events. They are often followed by changes in routing configuration caused by link failures, power shutdowns and routing exchange protocols misbehaving [17]. These cases are not caused by network congestion because no increase in delay is seen before the loss events. Fig. 7 demonstrates the case when a difference in packet delay is significant. Differences up to 400 ms were observed in our measurements.

It is interesting to see that even small changes in packet delay order of 1 ms, which most likely happened in intradomain (inside an Autonomous System), are also detected (Fig. 8). It is reasonable to assume that short changes in packet delay (up to several ms), which are detected most often, occur in intradomain (inside of an AS). Delay difference in order of tens or hundreds of milliseconds may be caused by both interdomain and intradomain route changes.

Table 3 also shows that in certain number of cases (9 on the UNINETT-CERNET route and 7 in the opposite direction; listed as “others”), a packet delay variation was significant and observed pattern was different from those in Figures 6-8. These can be divided into three groups: a) spikes as in Fig. 9 (7 out of 16); b) gradual failures as in Fig. 10 (4 out of 16), and c) congestion similar to Fig. 11 (5 out of 16 during the three-month measurement period). These events are not frequent, but also interesting to observe. The continuous decrease in packet delay in Fig. 9 is most likely caused by route queuing. The packets in the spike arrived “back-to-back”, with a very small interarrival time between them. Fig. 10 indicates that at certain cases, failure happens not instantly, but during a relatively

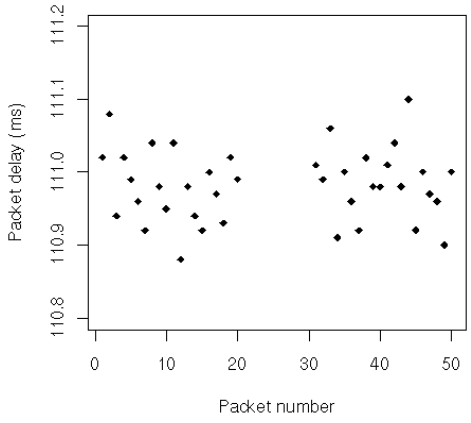

Figure 6. No change in delay after a downtime

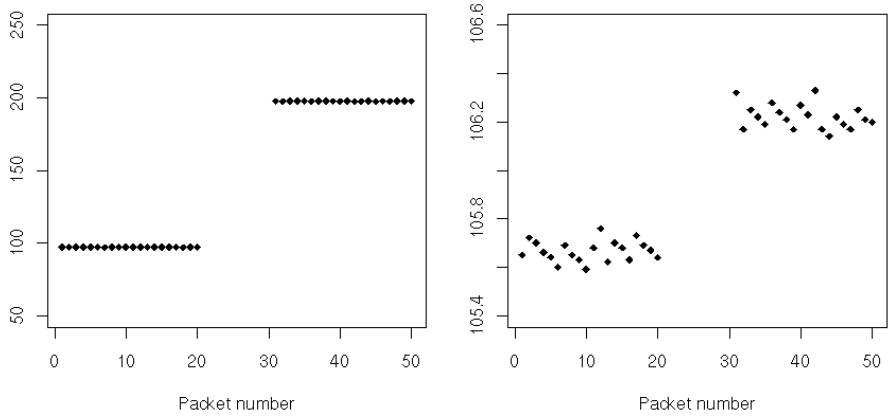

Figures 7, 8. Change in the fixed part of packet delays

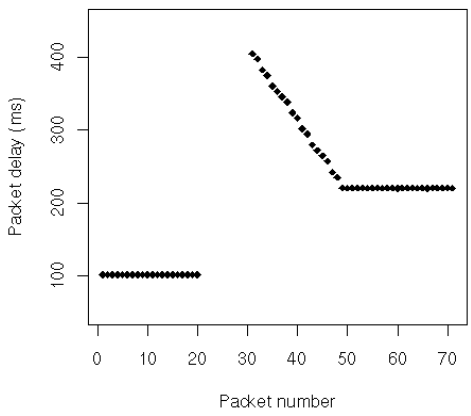

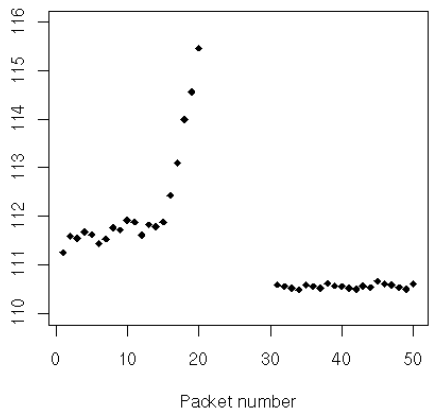

Figures 9, 10. Example of a packet delay pattern

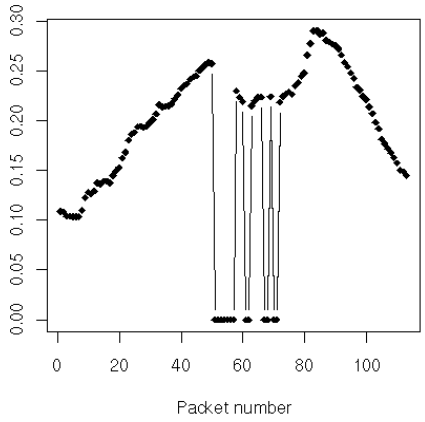

Figure 11. Example of loss event caused by congestion

short (50-100 ms) interval. Fig. 11 shows an example of downtime happened likely due to congestion: continuous increase in packet delay is seen and it is followed by a series of packet drops (the dropped packets are shown as packets with zero delay). Then, a continuous decrease in delay is seen.

Durations between downtimes are calculated for both directions and presented in Fig. 12. Most of the time periods between failures are in the order of 10^5 seconds (about 1 day). The process is more bursty than a Poisson process.

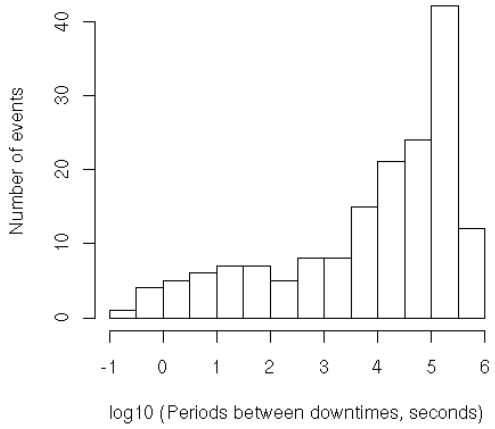

Figure 12. Histogram of periods between downtimes for both routes together

Table 4 contains autocorrelation test results. The 95% confidence limits for the correlograms are ± 0.20 and ± 0.23 respectively. It demonstrates that detected failure events do not correlate and successive failures are likely to be independent.

TABLE IV. AUTOCORRELATION FOR PERIODS BETWEEN OUTAGES

Lag	UNINETT-CERNET	CERNET- UNINETT
1	-0.004	-0.052
2	-0.062	-0.128
3	0.229	0.081
4	-0.161	0.137
5	-0.032	-0.120

VI. CONCLUSION AND FUTURE WORK

This paper presents analysis of measurement results collected during a three-month period on the global research network between Norway and China. The statistics was collected every 10 ms and this allowed detecting events with a relatively small duration. Detailed analysis of the received data is provided and parameterized. Although it is a challenging task to define a cause of a certain outage based on end-to-end interdomain measurements, the duration of a failure and change of delay before and after a downtime may provide some hints about the possible cause. The analysis shows that the plausible causes most of observed communication problems in the given study are either link or router related (level 2 or 3).

This paper described a part of a large measurement project. Future work will be continued in several directions. First, experiments in a lab environment will be performed to verify certain assumptions states in this paper. Second, the active measurements between the end-point will be continued. In addition to the academic route, similar measurement experiments will be performed in between Norway and other remote locations and will also use backbone channels of commercial ISPs. Third, the passive measurements will be performed on the end-points with special "wire-tap" interface cards. This will help to verify accuracy of the active measurements. And finally, deeper analysis of the collected traceroute statistics will be provided to better understand causes of failures in inter- and intradomain.

VII. REFERENCES

[1] ITV-T Rec. G.841, "Types and Characteristics of SDH Network Protection Architectures," 1996.

[2] Z. M. Mao, L. Qiu , J. Wang , Y. Zhang, "On AS-level path inference", Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems, June 06-10, 2005, Banff, Alberta, Canada.

[3] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. Chuah, and C. Diot, "Characterization of failures in an operational IP backbone network", IEEE/ACM Transactions on Networking, Vol. 16, No. 4, pp. 749-762, Aug. 2008.

[4] P. Francois, C. Filsfils, J. Evans, and O. Bonaventure, "Achieving sub-second IGP convergence in large IP networks," ACM SIGCOMM Comput. Commun. Rev., vol. 35, no. 3, pp. 35–44, 2005.

[5] A. Shaikh and A. Greenberg, "OSPF monitoring: Architecture, design and deployment experience," in Proc. USENIX 1st Symp. Networked Systems Design and Implementation (NSDI '04), San Francisco, CA, Mar. 2004, pp. 57–70.

[6] O. Bonaventure, C. Filsfils, P. Francois, "Achieving Sub-50 Milliseconds Recovery Upon BGP Peering Link Failures", IEEE/ACM Trans. on Networking, Vol. 15, No. 5, Oct. 2007.

[7] F. Wang, L. Gao, J. Wang, and J. Qiu, "On understanding of transient interdomain routing failures," in Proc. IEEE Int. Conf. Network Protocols (ICNP'05), Boston, MA, 2005, pp. 30–39.

[8] A. Sahoo, K. Kant, and P. Mohapatra, "Characterization of BGP recovery under Large-scale Failures," in Proc. ICC 2006, Istanbul, Turkey, Jun. 11–15, 2006.

[9] F. Wang, Z. M. Mao, J. Wang, L. Gao, and R. Bush, "A measurement study on the impact of routing events on end- to-end internet path performance," SIGCOMM Computer Communication Review, vol.36, no. 4, pp. 375–386, 2006.

[10] H. Pucha, Y. Zhang, Z. M. Mao, and Y. C. Hu, "Understanding network delay changes caused by routing events," in Proceedings of the 2007 ACM SIGMETRICS international conference on Measurement and modeling of computer systems. New York, NY, USA: ACM, 2007, pp. 73–84.

[11] "Rude/crude real-time udp data emitter/collector," http://rude.sourceforge.net/

[12] J. Kilpi, I. Norros, and U. Pulkkinen, "Downtime-frequency curves for availability characterization", in The 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN2007), Edinburgh, U.K., Jun. 2007.

[13] S. Floyd, and V. Jacobson, "Random early detection gateways for congestion avoidance", IEEE/ACM Transactions on Networking, 1(4), 397–413, 1993.

[14] Weibull Distribution, http://www.weibull.com/LifeDataWeb/weibull_probability_density_function.htm/

[15] R. Perlman, "A comparison between two routing protocols: OSPF and IS-IS," IEEE Network, 5(5), pp. 18--24, Sep. 1991.

[16] H. Gredler , W. Goralski, "The Complete Is-is Routing Protocol", SpringerVerlag, 2004.

[17] V. Paxson, "End-to-end routing behaviour in the Internet", IEEE/ACM Transactions on Networking, 1997, 5, (5), 601–615