

Using Concurrent Multipath Transmission for Transport Virtualization: Analyzing Path Selection

Thomas Zinner*, Kurt Tutschku[†], Akihiro Nakao[‡], Phuoc Tran-Gia*

*University of Wuerzburg, Institute of Computer Science, Wuerzburg, Germany
eMail: [zinner|trangia]@informatik.uni-wuerzburg.de

[†]University of Vienna, Professur "Future Communication",
Universitätsstrasse 10/T11, 1090 Wien, Austria. eMail: kurt.tutschku@univie.ac.at

[‡]University of Tokyo, Graduate School of Interdisciplinary Information Studies,
7-3-1, Hongo, Bunkyo-ku, Tokyo, 113-0033, Japan. eMail: nakao@iii.u-tokyo.ac.jp

Abstract—The concept of Transport Virtualization (TV) enhances the capabilities of future networks. TV enables transport mechanisms with arbitrary resource usage independent of the underlying transport system. The simplest form of TV can be achieved by collecting multiple transport resources (even from different virtual networks or providers) and selecting the best resources for exclusive or concurrent use. However, the selection and application of concurrent paths is complex and its impact on the transmission is non-intuitive. Path length diversity of different concurrent paths inevitably introduces out-of-order packet delivery. We present and discuss a mathematical model for the analysis of the fundamental behavior and influence factors for packet re-ordering in concurrent multipath transmissions. Our model facilitates the understanding of path selection algorithms for multipath transport virtualization.

I. INTRODUCTION

Nowadays, data transport is achieved mainly by the Internet Protocol (IP). Its routing feature achieves a scalable connection of local subnets and end systems. IP assumes that the interconnecting nodes (i.e. the routers) and the actual transport resources (i.e. the links) are stable and change only in case of failures. The Future Internet will consist of interconnected subnets and interconnecting nodes too. However, the availability of these resources is expected to be provided for a certain lifetime and their capabilities are highly variable, cf. [1]. Networks that are based on the model of temporal leasing of variable resources are referred to as *federated networks*, and technologies that enable the safe sharing of such network resources are denoted as *Network Virtualization (NV)* mechanisms.

A major task when creating federated networks for data transport is the selection of resources. The selection has to consider the temporal availability of the resources. Therefore, a measurement-based scheme for selecting these resources is needed. This mechanism should adapt the federated transport network to the variable demand for resources as well as to the currently available resources. Of course, the selection scheme has to scale sufficiently so that it can be applied in large networks.

The concept of Transport Virtualization (TV) enhances the capabilities of future networks. TV can be considered

as an alternative mode of Network Virtualization [2]. While NV typically facilitates the sharing of resources, TV creates virtual resources (e.g. virtual links) based on the aggregation of resources [2]. The simplest form of TV is achieved by collecting multiple transport resources (even from different virtual networks or providers) and selecting the best resources for exclusive or concurrent use. The type of the resources used for aggregation does not matter; they can be either of physical or of virtual nature.

The use of *concurrent multipath (CMP)* transmissions will bring exceptional advantages to networks, such as higher throughput and increased resilience. However, CMP will introduce additional complexity which has to be understood. First, CMP transmission will inevitably introduce out-of-order packets due to different stochastic packet delay characteristics on the paths. The re-ordering can be compensated by buffering at the destination, possibly leading to increased end-to-end delay but still being transparent to the transport protocol. Second, the different stochastic delay processes on the paths can amplify each other in their negative effects on out-of-order packets. Third, the strength and occurrences of such combination effects are highly non-intuitive.

We study these effects by an *analytical model*, which is able to calculate the *re-sequencing buffer occupancy probability distribution*. This model will be verified by event-based simulations. Originally, the model was introduced for the investigation of a many-to-one transmission scenario of parallel downloads for streaming applications, [3], [4]. In the context of TV, the model is used for a one-to-one transmission over multiple paths. By using the analytical model, "what-if" scenarios and path selection strategies can be easily evaluated.

The paper is structured as follows. First, Section II outlines the idea of TV. Section III discusses work related to the suggested TV. In Section IV we present the analytical model for the analysis of the re-sequencing buffer occupancy probability. Section V provides comprehensive discussions towards a path selection strategy based on the investigation of the re-sequencing buffer occupancy probability distribution. Section VI concludes with a brief summary.




Fig. 1: Selection of one path out of several possible paths

II. TRANSPORT VIRTUALIZATION

The idea of *Transport Virtualization* constitutes an abstraction concept for data transport resources. In TV, an abstract data transport resource is combined from one or more physical or overlay data transport resources. Such a resource can be, e.g., a leased line, a wave length path, an overlay link, or an IP forwarding capability to a certain destination. The resources can be used preclusive or concurrently.

A. Path Selection and Resource Pooling

First let us consider ordinary path selection. In this case, a single path out of a set of candidate paths is selected. This is case is depicted in Figure 1. The path characteristic has a direct impact on reliability, security or on QoS parameters between the start point and the end point of the transmission. Therefore, a selection function has to exist which enables the appropriate choice of the used path. For instance, the different resources (i.e. paths) could be classified by the path delay.

Moreover, the concurrent usage of different paths might have appealing advantages. The reliability of the end-to-end connection is increased and even the overall throughput might be higher if concurrent paths are used for the data transmission. This is illustrated in Figure 2. In this example two paths are pooled and appear as one virtual link. However, there has to be a mechanism distributing the data among the used paths and which is selecting the appropriate set of concurrent paths. This feature will be discussed in Section IV. The set of used paths can consist of paths offered by different Internet or overlay providers. So the data transport is not bounded to one of these providers any more. It is even possible to choose the cheapest or fastest resources and pool them together to a high capacity pipe. The pooling of resources out of different networks is shown in Figure 3.

TV achieves the pooling by a mechanism that is placed on nodes that have access to multiple paths or networks, e.g. end hosts which are multi-homed or routers running exterior gateway protocols. At the origin of the transmission, the mechanism decides which resources, i.e. which paths, are




Fig. 2: Selection of a couple of paths out of several possible paths




Fig. 3: Selection of a couple of paths out of several possible paths from different overlays

selected, whether they are used preclusive or concurrently, and how the data is distributed among these paths. If necessary, the split data is recombined at the destination to a consistent flow. Due to the ability to choose between different transmission options, TV achieves an independence from a specific physical transport resource.

Even if one of the used resources fails, TV using CMP may be able to continue partly the transmission service. Other advantages of CMP transmissions are higher capacity due to parallelism, avoiding head-of-the-line blocking, spatial separation of control and data connections, differentiations between traffic types, and leveraging multiple processors in routers [5]. On the downside, CMP transmissions require additional overhead and memory for leveraging packet re-ordering.

B. Availability of Concurrent Paths

A major concern for the applicability of TV is the availability of different transport resources. Recent investigations of tier 1 topologies [6] reveal that multiple paths between many destinations readily exist [7]. Moreover, measurements in PlanetLab [8] revealed that up to 25% of certain Internet routes may violate the *triangle inequality*. That means a shorter delay can be experienced when another path through an intermediate node is used.

C. Selection of Pooled Resources

An important question for TV is the selection of pooled resources, i.e. the choice of a potential set of paths. Typically, a *good* path in the set has a short transmission delay [9]. Thus, the mean packet delay on a path is an initial candidate as selection criteria. However, the interaction between concurrent paths is expected to be complex. Though, the selection has to be based on a more detail path description. This will be evaluated in Section V. For that we will use the *packet delay distribution* on a path as a main characterization for the path behavior.

III. RELATED WORK

Multipath transmission mechanisms have been suggested for IP networks a while. The Stream Control Transmission

Protocol (SCTP) [10] enables multi-homing but doesn't facilitate CMP transfer and its advantages. As seen in [11] typical problems with packet reordering occur. While [11] proposed to enhance SCTP such that it can react to packet reordering this paper presents a mechanism to suppress reordering as far as possible. As a result transport protocols like SCTP or TCP do not have to be modified. The *pTCP* [12] and the *Multi-Path TCP* [13] proposals exploit concurrent transmissions, however, they focus on flow control and the coordination among flows. Our work complements these studies since we are aiming at the selection of paths.

Path selection is often investigated on network layer in context of multipath routing, e.g. [14]. We amend these works by considering the transmission mechanism on transport layer.

A major work, which addresses the network and transport layer concurrently, is the DaVinci architecture [15]. In DaVinci, the paths are selected such that the creation of bottlenecks is avoided. The investigation in our paper goes beyond the results presented in the DaVinci architecture. The DaVinci proposal considers mainly the mean path delay. However, we will see in Section V that this assumption might be not sufficient. Our paper discusses the fundamental performance issues in selecting paths according to their detailed statistical characteristics.

Finally, three efforts of IETF working groups (WGs) should be mentioned. The Application-Layer Traffic Optimization (ALTO) WG [16] aims at providing P2P CDNs with information to perform better-than-random initial peer selection. This selection protocol might be extended to path selection. The Low Extra Delay Background Transport (LED BAT) WG [17] investigates multi-path TCP connections in order to better saturate bottlenecks. At last, Transport Area WG discusses currently the combination of BGP and Multi-path TCP [18].

IV. MECHANISMS AND PERFORMANCE MODELS

TV can be facilitated using *one-hop source routing overlays* [19]. Such an overlay is implemented by the *SORA architecture* [9] where a *striping* mechanism is used.

Besides the striping of the data to different paths, the TV mechanism is selecting the pool of concurrent paths. The candidate paths for pooling might be provided either by a *path oracle* (e.g., [20] for operator supported oracles) or by the start point of the multi-path connection, if this start point measures the paths by itself.

Next, we will outline the performance model for the suggested striping mechanism. We will start with the mechanism and then detail the analytical model of the re-sequencing buffer occupancy. We assume a *quasi-stationary* case for the path selection for the model. This means the analytical model computes steady state probabilities under the assumption that a certain set of paths is currently used. However, the set can be changed by the selection mechanism when a used path degrades or becomes unavailable during operation. The adaption of the used set of path is of particular interest in federated networks since here the resources are only available for a certain lifetime and are variable in their performance.

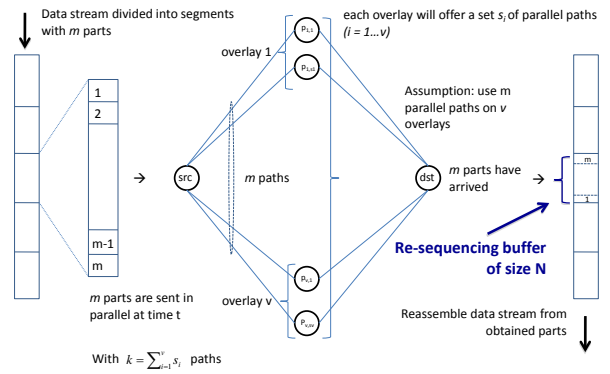


Fig. 4: Striping Mechanism

Since we are interested in a simple model which facilitates path selection we assume a connectionless transmission between the end points. So the model can focus on the analysis of different path delay distributions without the influence of interfering control mechanisms like, e.g., the TCP control loop, on the data transmission.

A. Striping Mechanism

Figure 4 shows a detailed model of the striping mechanism. The data stream is divided at the source (*src*) into segments which are split into m smaller parts. The parts are transmitted in parallel on m different overlay paths. The destination (*dst*) reassembles these parts. Due to different stochastically varying delays on the paths, the parts can arrive at the destination out-of-order. In the model we assume that the re-ordering of parts can happen only between different paths while the order on a path is maintained.

Part or packet re-ordering may have a severe impact on the application performance. In order to level this behavior, the destination maintains a finite *re-sequencing buffer*. A high buffer occupancy stands for a high number of parts waiting for other parts and implies unnecessary delay. Moreover, if the buffer overflows, part loss will occur which leads either to data loss or complex retransmission. In case of retransmissions, the delay is increased even further, possibly leading to more part loss and a corruption of the complete data transport. Therefore, an important objective is to *minimize the re-sequencing buffer occupancy*.

B. Analytical Model for Re-sequencing Buffer Occupancy

In literature initial analytical and approximative methods exist for estimating the re-sequencing buffer occupancy in case of multipath downloads or transmissions. We adopt the analytical model [3] to our system and perform an analytical performance evaluation upon it. The enhanced model enables an appropriate parameter choice in case of a multitude of possibilities.

The considered model assumes a continuous data stream for the multipath transmissions over m concurrent paths. The delay on the paths is described by discrete delay distributions

with a resolution of one time unit. We further consider paths with equal capacity, and that the transmission rate on each path is equal to one packet per time unit. The delay for different paths are independent and may follow different delay distributions. A detailed explanation of the mathematical model can be found in [3], [4]. Here we only present the central ideas and adopt them to the given mechanism. The used model ensures that no packet re-ordering on a single path can occur. That means that packets send over one path can not overtake each other. To facilitate the explanation, the following notations are used:

The packets transmitted at time 0 over path 1, 2, ..., m are packets 1, 2, ..., m respectively. After transmitting the first m packets, the packets are appointed to the sources in a round-robin manner. Thus, at time t packets $1 + mt, 2 + mt, \dots, m + mt$ are transmitted. We further use the term *minimum valued packet (mvp)*, as introduced in [3], denoting the lowest indexed packet at time t that has not arrived at the destination by time t . For instance, if packets 1 through 5 arrived, but packet 6 did not, packet 6 is the mvp. Thus, the resequencing buffer occupancy at time t is exactly the number of packets indexed higher than the mvp that have arrived by time t and no packet stored in the re-sequencing buffer was transmitted via the path of the mvp, since packets transmitted on every path arrive in transmission order. We denote the index of the path of the mvp by s_n and $\delta_{X,t}$ as the time passed since the last packet received via path X was transmitted at time t . For brevity we refer to $\delta_{X,t}$ as δ_X in the following. With this notation the re-sequencing buffer occupancy can be computed as:


$$P(B = k) = \sum_{i=1}^m \sum_{x=0}^{\infty} P(B = k, s_n = i, \delta_i = x). \quad (1)$$

The right hand side of the equation denotes the buffer occupancy probability for each path transmitting the mvp and each possible value for the time passed since the last packet transmitted over this path was received. As discussed in [4], this yields to :

$$\begin{aligned} & P(B = k, s_n = i, \delta_i = x) \\ &= P\left(\sum_{j=1, j \neq i}^m \delta_j = (m-1)x + 1 - i - k, \right. \\ &\quad \left. \delta_j < x \forall j < i, \delta_j \leq x \forall j > i\right) \\ &= \sum_{S_{i,x,k}} P_i(x) \prod_{j=1}^{i-1} P_j(l_j) \prod_{j=i+1}^m P_j(l_j) \end{aligned} \quad (2)$$

where $S_{i,x,k}$ defines the delay configuration on path before the arrival of the mvp:

$$\begin{aligned} S_{i,x,k} = \{ & l_1, \dots, l_{i-1}, l_{i+1}, \dots, l_m : l_1 < x, \dots, l_{i-1} < x, \\ & l_{i+1} \leq x, \dots, l_m \leq x, \sum_{j=1, j \neq i}^m l_j = (m-1)x + 1 - i - k \}. \end{aligned}$$



(a) Used path delay distributions (b) Influence of the delay on the re-sequencing buffer occupancy

Fig. 5: Comparison of analysis and simulations

With this formula we can compute the re-sequencing buffer occupancy in case of a transmission over m paths with equal transmission rate.

V. TOWARDS PATH SELECTION FOR END POINTS

Next, we discuss towards the design of path selection strategies of end points, i.e. of end host and routers which are capable to choose the set of paths used for the CMP transmission. A major aim of such a selection strategy is to minimize the negative effects of CMP transmission, i.e. the re-ordering of parts due to over-taking among different paths. This behavior results from different *path delay distributions*.

As outlined before the *re-sequencing buffer occupancy* characterizes the negative effects and is chosen as the main performance criteria for the presented CMP mechanism.

Two main scenarios are considered for the discussions of the performance of the CMP mechanism: scenario a) has two concurrent paths and scenario b) three concurrent paths. Scenario a) is sufficient to show the impact of the type of different packet delay distributions on the buffer occupancy. However, the complexity of path selection can be addressed in a "three concurrent paths" scenario. This scenario allows for the discussion of more complex "what-if" settings required for path selection.

Before discussing the detailed performance evaluation we validate the correctness of the analytical model by comparison with a separately implemented, *event-based simulation* of the presented mechanism.

A. Validation of the Analytical Model

The simulation model used for the validation was outlined in [7]. The correct simulation of the packet delay on a single path is not trivial when requiring that packets do not overtake each other. Therefore, a special technique for deriving the appropriate packet delay distribution, as proposed in Section IV.D of [7], was applied.

We consider for the validation a CMP mechanism with two concurrent paths and verify the analytical model in two cases: a) the packet delay on both paths follow a Gaussian distribution (with mean path delay $\mu = E[D = d] = 25$ and standard deviation $\sigma = 10$) and b) the packet delay on both paths is uniform distributed (in the range of $[0, \dots, 50]$). The delay distributions are shown in Figure 5a). Figure 5b) shows the re-sequencing buffer occupancy probability $P(B = k)$ for the investigated scenarios computed by the analytical model




Fig. 6: Buffer occupancy probability $P(B = k)$ for distributions with same coefficient of variation and different skewness

and measured in the event-based simulation. The analytically computed distributions are depicted by the solid lines and the probability values obtained from simulations are given by a statistical confidence intervals for a level of 99%.


Figure 5b) shows that there is a very close match of the buffer probabilities obtained by simulation and by the analytical model. Consequently, we consider the analytical model as trustworthy and use it in the reminder of this study.

B. Impact of Type of Delay Distribution

Next, we focus our attention to the impact of the type of delay distribution on the CMP performance. We investigate a scenario with two paths for making the effects clearly identifiable.

Probability distributions are typically described by type, *mean* and *coefficient of variation* c_v . The asymmetry of a probability distribution is described by the *skewness* of the distribution $\nu = \mu_3/\sigma^3$, where μ_3 is the third moment about the mean and σ is the standard deviation. A negative skewness denotes that most of the probability mass is located at values higher than the mean value. In case of a positive skewness the mass of the distribution is concentrated on values smaller than the mean. Thus, the right tail is longer indicating that delay values higher than the mean are located in a greater distance from it. For the investigations, we consider the delay distributions for both paths to be of either *negative-binomial* (*nbin*) type or of *poisson* type. The poisson distribution was selected since it may reveal a symmetrical behavior and inherits a strong relationship to a Normal distribution, which is often used for the approximation of path delay in the Internet [21]. The negative-binomial distribution was chosen because of its asymmetry, i.e. a high skewness, while being able to set the variance. We use for both distributions a coefficient of variation of $c_v = 1$ and mean values of $\mu = [50|100|150]$.

The Figure 6 depicts the buffer occupancy probability $P(B = k)$ for two concurrent paths. Of course, the mean of the path delay has remarkable impact. However, although the means and coefficient of variations are fixed to the same values for both distributions, Figure 6 shows that the type of distribution, presented also by different skewness values, leads to very different buffer fillings. The buffer occupancy is much lower and the performance better when poisson-like paths are used. In consequence, distributions with a high asymmetry should be avoided when selecting paths.



(a) Used negative-binomial distributions with equal standard deviations and different mean values. (b) Impact of different path combinations with equal standard deviations on the buffer occupancy for three paths

Fig. 7: Complexity of path selection

C. The Complexity of Path Selection

The complexity of path selection will be investigated by a CMP setting applying three concurrent paths. The number of three paths permits already the investigation of "what-if" scenarios where two paths are already selected and a third path will complement the other paths.

The used distributions for this setting are Negative-binomial distributions with mean $\mu = E[D] = [25|50]$, standard deviation $\sigma = 20$ and skewness $\nu = [1.55|0.75]$. It has to be noted that the distribution with lower mean has a higher skewness. The distributions are depicted in Figure 7a).

Naive Combination: First, we now consider four cases of paths combinations: *a)* all three path delay distributions have the same low mean delay $\mu = 25$, *b)* all three delay distributions have the same high mean $\mu = 50$, *c)* two paths have low mean delay $\mu = 25$ and one a high mean $\mu = 50$, and *d)* one path has a low mean delay $\mu = 25$ and two paths have high mean of $\mu = 50$. Figure 7b) shows the CDF of the re-sequencing buffer occupancy for the four cases. The figure reveals that case *a)* and *b)* have almost similar buffer occupancy distributions. This indicates that the pure delay has no impact on the buffer occupancy. The slight difference is the result of the different skewness of the distributions, cf. [7].

For the cases *c)* and *d)* it is remarkable that case *d)* with two high path delays performs better in terms of buffer occupancy than case *c)*. Intuitively more paths with lower delay should result in a better performance, but this is obviously not true as shown in case *d)*.

Let us consider case *c)* in greater detail. Therefore, we assume a single packet from the high delay path which is much overdue. Until the arrival of this overdue packet, the low delay paths can easily increase the occupancy of the re-sequencing buffer. Thus, the buffer is filled quickly by the low delay paths. This example shows that the high delay path becomes more dominant over low delays paths in terms




Fig. 8: Impact of the path characteristics of an arbitrary 3rd path on the mean re-sequencing buffer occupancy

of buffer occupancy. The selection of the paths should level the variation of the range of mean delays. Moreover, recent measurement studies in Voice-over-IP systems confirm that a constant high or low delay may have no significant impact on the Quality of Experience of the application [22]. Thus it might be better in CMP to choose a path with a higher mean delay and lower delay variation in order to relieve the re-sequencing buffer and avoid packet loss.

Influence of Mean Delay and Delay Variation: Now we investigate the impact of the mean delay and of the delay variation. In this setting we assume that two paths are already selected. The selected paths are of negative-binomial type, one with a low mean delay $\mu = 25$ and one with a higher mean delay of $\mu = 50$. We consider three cases where these paths have a standard deviation of $\sigma = [10|15|20]$. In case *a*), the delay on the third path, which is completing the other two paths, is distributed according to a Negative-binomial distribution with $\sigma = 20$ and a mean in the range of $[25, \dots, 75]$, in case *b*) the delay distribution is also Negative-binomial but with $\sigma = 15$, and in case *c*) we use a Normal or a Negative-binomial distribution for all paths with low delay variance and $\sigma = 10$. This case should compare the influence of the skewness, which is much lower for the Gaussian distribution.

Figure 8 shows that the mean delay of the third path has a significant impact on buffer occupancy. If the mean path delay of the third path increases beyond the one with the higher mean delay then the mean buffer occupancy increases strongly. However, if the mean path delay of the third path is between the other two means then even a minimum exists. This effect is taking more shape when σ is low (case *ba* and *c*). We also conclude that a high delay variance degrades the CMP performance. Case *c* shows that the skewness has only little impact on the mean re-sequencing buffer occupancy.

All-in-all, the above presented investigations show that a path selection strategy should consider the second moment of path delay. For a deeper analysis of the re-sequencing buffer occupancy beyond a mean value analysis, the skewness might have a significant impacts.

VI. CONCLUSION

Our contribution is to present an analytical model for *concurrent multipath (CMP)* transmissions for *transport virtualization (TV)*. CMP transport has many appealing advantages such as higher throughput and increased resilience. However, its application increases also the complexity of the system. In particular, the use of concurrent paths introduces inevitable out-of-order packet delivery. This effect has to be leveled by a re-sequencing buffer at the destination and careful selection of paths. We discussed the impact of different path delay distributions on the *re-sequencing buffer occupancy probability distribution*. Our model facilitates the understanding of path selection algorithms for multiple-path TV. The obtained results show that a path selection using the mean delay as sole criteria is not sufficient. Further criteria like higher moments and even the type of the delay distribution have to be considered. Since

the analytical model is rather simple it might be used as a first step toward the on-line estimation of the quality of a certain set of paths. Such an on-line estimation is of particular interest in federated networks.

Further, future work should address the coordination of TV and CMP mechanisms among different overlays. In addition, the findings presented in this paper have to be validated by real measurements. Another extension should be an analytical approximation of the real end-to-end delay which is not yet included in the presented model.

ACKNOWLEDGMENTS

The authors would like to thank Dirk Staehle and Tobias Hoffeld for the fruitful discussions and their support during the course of this work. The sponsorship of this research by the European FP7 Network of Excellence "Euro-NF" through the Specific Joint Research Project "Multi-Next" is thankfully acknowledged.

REFERENCES

- [1] Larry Peterson, Soner Sevinc, Jay Lepreau, Robert Ricci, John Wroclawski, Ted Gaber, and Stephen Schwab, "Slice-Based Facility Architecture," 2009, Information available at <http://svn.planet-lab.org/attachment/wiki/GeniWrapper/sfa.pdf>.
- [2] Thomas Zinner, Kurt Tutschku, Akihiro Nakao, and Phuoc Tran-Gia, "Re-sequencing Buffer Occupancy of a Concurrent Multipath Transmission Mechanism for Transport System Virtualization," in *Proc. of the 16. KiVS 2009*, Kassel, Germany, Mar. 2009.
- [3] Yoav Nebat and Moshe Sidi., "Resequencing Considerations in Parallel Downloads," in *Proc. 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002.
- [4] Y. Nebat et al., "Parallel Downloads for Streaming Applications – a Resequencing Analysis," *Performance Evaluation*, vol. 63, no. 1, pp. 15–35, Jan. 2006.
- [5] R. Penno, S. Raghunath, R. Woundy, V. Gurbani, and J. Touch, "LEDBAT Practices and Recommendations draft-ietf-ledbat-practices-recommendations-00.txt," Mar. 2009.
- [6] Michael Liljenstam, Jason Liu, and David Nichol, "Development of an internet backbone topology for large-scale network simulation," in *Proceedings of the 2003 Winter Simulation Conference*, New Orleans, December 2003.
- [7] Thomas Zinner, Kurt Tutschku, Akihiro Nakao, and Phuoc Tran-Gia, "Performance Evaluation of Packet Re-ordering on Concurrent Multipath Transmissions for Transport Virtualization," to be published in *JCNDS Special Issue on: Network Virtualization - Concepts and Performance Aspects*, 2011.
- [8] Suman Banerjee, Tomothy Griffin, and Marcelo Pias, "The Interdomain Connectivity of PlanetLab Nodes," in *Proc. of the 5th Passive and Active Measurement Workshop (PAM2004)*, Antibes Juan-les-Pins, France., Apr. 2004.
- [9] John Russell Lane and Akihiro Nakao, "Best-Effort Network Layer Packet Reordering in Support of Multipath Overlay Packet Dispersion," in *Proc. of Globecom08*, New Orleans, LA., Nov./Dec. 2008.
- [10] L. Ong and J. Yoakum, "RFC 3286: An Introduction to the Stream Control Transmission Protocol (SCTP)," May 2002.
- [11] Janardhan R. Iyengar, Paul D. Amer, and Randall Stewart, "Concurrent Multipath Transfer using SCTP Multihoming over Independent End-to-End Paths," *IEEE/ACM Trans. Netw.*, vol. 14, no. 5, pp. 951–964, 2006.
- [12] Hung-Yun Hsieh and Raghupathy Sivakumar, "pTCP: An End-to-End Transport Layer Protocol for Striped Connections," in *Proceedings of the 10th IEEE International Conference on Network Protocols*, Paris, France, Nov. 2002.
- [13] Huaizhong Han, Srinivas Shakkottai, Student Member, C. V. Hollot, R. Srikant, and Don Towsley, "Multi-Path TCP: A Joint Congestion Control and Routing Scheme to Exploit Path Diversity in the Internet," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1260–1271, Dec. 2006.

- [14] Ivan Gojmerac, Peter Reichl, and Lasse Jansen, "Towards Low-Complexity Internet Traffic Engineering: The Adaptive Multi-Path Algorithm," *Computer Networks*, vol. 52, no. 15, pp. 2894–2907, Oct. 2008.
- [15] Jiayue He, Rui Zhang-Shen, Ying Li, Cheng-Yen Lee, Jennifer Rexford, and Mung Chiang, "DaVinci: Dynamically Adaptive Virtual Networks for a Customized Internet," in *Proc. 4th ACM International Conference on emerging Networking EXperiments and Technologies (ACM CoNEXT 2008)*, Madrid, Spain, Dec. 2008.
- [16] John Peterson, Vijay Gurbani, and Enrico Marocci, "Charter of the Application-Layer Traffic Optimization (ALTO) Working Group," Mar. 2009.
- [17] Stanislav Shalunov and Murari Sridhavan, "Charter of the Low Extra Delay Background Transport (LEDBAT) Working Group," Mar. 2009.
- [18] James Polk and Gorry Fairhurst, "Charter of the Transport Area Working Group (TSVWG)," Mar. 2009.
- [19] Krishna P. Gummadi and Harsh V. Madhyastha, "Improving the reliability of internet paths with one-hop source routing," in *Proc. OSDI'04*, San Francisco, Ca., USA, Dec. 2004.
- [20] Christian Scheideler Vinay Aggarwal, Anja Feldmann, "Can ISPs and P2P Systems Co-operate for Improved Performance?," *ACM SIGCOMM CCR*, vol. 37, no. 3, Jul. 2007.
- [21] Tamas Elteto and Sandor Molnar, "On the Distribution of Round-trip Delays in TCP/IP Networks," in *Proceedings 26th Conference on Local Computer Networks (LCN'99)*, Lowell, Oct. 1999.
- [22] Tobias Hoßfeld, David Hock, Phuoc Tran-Gia, Kurt Tutschku, and Markus Fiedler, "Testing the IQX Hypothesis for Exponential Interdependency between QoS and QoE of Voice Codecs iLBC and G.711," in *18th ITC Specialist Seminar on Quality of Experience*, Karlskrona, Sweden, May 2008.