Journal of Universal Computer Science, vol. 18, no. 14 (2012), 2071-2092
submitted: 30/10/11, accepted: 27/7/12, appeared: 28/7/12

Modeling the Value of End-to-End Multipath Protocols

Henna Suomi
(Aalto University, Espoo, Finland
henna.suomi@aalto.fi)

Kalevi Kilkki
(Aalto University, Espoo, Finland
kalevi.kilkki@aalto.fi)

Heikki Hämmäinen
(Aalto University, Espoo, Finland
heikki.hammainen@aalto.fi)

Abstract: Recently, adding multipath capability in the Internet protocol suite has attracted increasing interest. By letting end hosts discover several paths to communicate, end-to-end multipath protocols aim to improve utilization rate of Internet resources. Although many proposals for end-to-end multipath communication exist, they have not reached significant deployment. Since the multipath protocols are mainly designed for open multi-stakeholder environments, understanding their economic impact is important. This paper introduces a model for assessing the value of the end-to-end multipath protocols from the end user perspective. Without a net benefit of the end user, the end-to-end multipath communication only results in the reallocation of costs and benefits in the Internet connectivity market. The model indicates that wireless devices having access to multiple independent access operators via similar or dissimilar access technologies are crucial in achieving end user value out of multipath communication. Initially, the end user value seems higher when the radio interfaces to access operators are active one at a time but later on, along with higher-energy batteries and lower-energy protocols, full benefit of multipath communication can be achieved. The value of multipath protocols depends on the effective path diversity and available capacity on the Internet.

Keywords: Multipath communication, path diversity, net benefit, end user
Categories: H.1.0, D.4.4

1 Introduction

The global Internet traffic is growing rapidly and putting pressure on scalability not only in access links but also in the Internet core. Stakeholders in the Internet connectivity market search for solutions to respond to the increasing demand. One key component of improving Internet scalability is load balancing, which promotes the efficient usage of resources.

Multipath capability is an example of a load balancing solution and refers to technologies which enable hosts or entire sites to use multiple access links or entire paths simultaneously. This results in statistically better resource utilization. If several paths exist, multipath capability potentially increases throughput and resilience of connections while singlepath communication leaves some resources underutilized.

Multipath technologies would balance the traffic load and thus improve the utilization rate [Wischik, 08].

Multipath capability can be implemented in different parts of networks and with several technologies. Internet service providers (ISP) and large content providers are using traffic engineering to increase the utilization rate within their networks and data centers (see, *e.g.*, [Awduche, 99][Greenberg, 09]). They can also engineer the traffic between the peering networks by using inter-domain routing protocols. Similarly, academic and corporate sites are constantly developing their networks to improve the usage of resources. Typically, they use network address translators (NAT) or similar middleboxes (see, *e.g.*, [Guo, 04]), for running scheduling algorithms to balance end users' traffic when relayed to the Internet. One potential alternative is to leave the control of the traffic load balancing to the end hosts which means implementing the multipath capability in the transport or application layer of the Internet protocol suite.

Load balancing mechanisms differ in terms of strength, speed and scope of response to congestion or link failures. The load balancing schemes implemented by a single administrative authority, such as an ISP, can be well optimized within the domain but they cannot control the congestion or address link failures in other parts of the end-to-end path. The advantage of end-to-end load balancing is that the end hosts are able to monitor end-to-end paths and address failures instantly without losing the ongoing session. This is called *Internet-wide end-to-end load balancing*, which is the focus of this study.

While intra-domain load balancing solutions are relatively common, the inter-domain end-to-end multipath solutions have not reached commercialization. Compared to the amount of technical efforts, the number of studies conducted on economic impacts of multipath communication remains low. The authors in [Tang, 08] used game theory for investigating load balancing between heterogeneous radio access network providers. They took the perspective of a single network provider and did not consider the overall market. The objective of this paper is to take a more holistic approach and shed light on the social value of end-to-end multipath protocols. This study introduces a functional modeling method for evaluating the net benefit of end-to-end multipath technologies as a function of measurable parameters.

The authors in [Joseph, 07] were the first to apply functional modeling to analyze the value of Internet protocols. They studied the value of IPv6 protocol to see how various factors such as standalone benefits and converters affect the utility of a new communication protocol. Functional modeling was also used in [Iannone, 10] to analyze how the deployment of locator/ID split (LISP) protocol would decrease the operational costs of the operators. The idea of functional modeling is to find the most relevant costs and benefits of a protocol, which allows estimating the total utility numerically. Both studies mentioned above estimate the utility of the protocol as a function of protocol adopters.

This paper takes a user-centric approach. We explicitly measure the incremental costs and benefits of an end-to-end multipath protocol during a user session compared to traditional singlepath communication. The costs and benefits perceived by the end user depend on the user download behavior, multihoming configurations and the state of network resources on the Internet. The proposed method culminates in an improved end user experience since the positive net benefits of end users will also benefit other stakeholders in the Internet connectivity market.

The rest of the paper is structured as follows. Chapter 2 introduces the central terminology used in this paper while Chapter 3 discusses path diversity and available bandwidth which are essential in achieving the intended end-to-end benefits of multipath protocols on the Internet. Chapter 4 introduces the developed model and Chapter 5 presents its application in chosen use cases. Chapter 6 discusses the model applicability and Chapter 7 finally concludes the study.

2 Multipathing and multihoming

Multipath communication, or multipathing, is mostly associated with multihoming but, in fact, both capabilities can exist independently. These terms are often used without further definitions and in order to avoid confusion, we start by defining them. By multipathing we refer to a technical feature which is used by an entity to split the data flow into smaller chunks, and to send them through separate access links or entire paths simultaneously. Therefore, multipathing aims to increase the communication performance by increasing throughput. The entity may be a single end user or a network, such as a corporate site.

Many studies have aimed to integrate multipath capability in the Internet protocol suite. Adding the capability in the transport layer was first proposed in [Huitema, 95], which allowed the set of addresses used by a TCP connection to change over time. Since then, multiple proposals have been made with slightly different implementations (see for instance, [Zhang, 04] and [Iyengar, 06]). These approaches allow the aggregation of the available bandwidth on different paths. The most recent proposal is called Multipath TCP (MPTCP) [Ford, 11], which is currently being standardized in the Internet Engineering Task Force (IETF). It is an extension to and backwards compatible with the regular TCP. MPTCP uses coupled congestion control, proposed in [Key, 06], which dynamically shifts traffic from congested to underutilized paths. MPTCP not only aggregates the available bandwidth but it also adapts to the congestion state in the network.

IETF has also made efforts to develop multipathing in other layers. A protocol called Multipath Real-time Transmission Protocol (MPRTP) [Singh, 11] in the application layer is an example of protocols which enables the end user to communicate over multiple paths over the Internet. In addition, proprietary solutions for multipath communication exist. For example, Real Time Media Flow Protocol (RTMFP) as a part of Adobe’s Flash Plugin, is an approach for providing multipath capability in peer-to-peer (P2P) communication [Kaufman, 09].

End hosts can discover multiple paths on the Internet by advertising all allocated IP addresses. Basically, this means that the minimum requirement to find separate paths on the Internet is that one of the end points has at least two IP addresses so that address pairs can be formed. This means that also singlehomed hosts are capable of communicating over multiple paths, assuming that another peer has multiple addresses. Also, port numbers can be exploited when discovering between the end points. For instance, equal-cost multipath routing (ECMP) routes packets differently based on a hash of port numbers. By varying the port numbers, a host could discover multiple paths towards a destination [Raciu, 11]. However, ECMP routing is deployed only for intra-domain purposes, not for inter-domain.

In the literature, multihoming may refer to the capability of connecting to multiple access operators or using several access technologies. The access operator and technical multihoming should always be kept separate for clarity. A multihomed end host or site is configured to use multiple IP addresses which can be associated with a single or several physical interfaces. An example where an end user is multihomed over a single interface is having a virtual tunnel the employer's virtual private network (VPN) [Gleeson, 00] and another IP address for regular web browsing over a wireless LAN (WLAN) access point. IETF has recently established a working group called multiple interfaces (MIF), which aims to alleviate the address configuration problems experienced by multihomed terminals [Blanchet, 10]. In addition, IETF has proposed a protocol called Host Identity Protocol (HIP) [Nikander, 10] to facilitate seamless handovers of a mobile host between different access network technologies. Similarly to the end host case, the physical access link of a single site may be shared between access operators. For sites multiple technologies already exist for implementing multihoming in practice. The traditional solutions include NATs, or boarder gateway protocol (BGP) which is used to propagate the address changes to the rest of the Internet.

The fundamental difference between multipathing and multihoming is that multipathing aims to increase the fine-grained runtime performance, *i.e.*, session throughput. Multihoming as such does not necessarily increase the session-level performance but it can be used to increase reliability or to extend network coverage. Depending on the context where multipathing or multhoming are deployed, they may use different technological configurations and fulfill different stakeholder requirements. Therefore, explicit definitions of the terms should always be provided.

In this paper, we only concentrate on multipath communication between two end hosts on the Internet. A host is multihomed when it is configured to use multiple access operators (and potentially multiple access technologies) and is accordingly assigned an IP address by each operator. We regard the communication between the hosts as multipath if, at least, one hop of the alternative end-to-end paths deviates.

3 Path diversity and available bandwidth

The fundamental idea of end-to-end multipath communication is to exploit several paths, with varying degree of available bandwidth, between the end hosts. On the Internet, a host can try discover paths by associating different IP addresses but it typically cannot control whether the subflows are actually traversing separate paths. If path diversity does not exist in the network, multipath protocols will result in a data split into smaller chunks which are merely transferred through the same path without additional benefits.

Inside a domain, path diversity is easier to implement since the topology of the network is fully controlled by a single administrative authority. However, a single entity has hardly any control over the inter-domain topology and the number of available paths depends on the evolution of stakeholder contracts as well as routing policies on the Internet. Depending on the multihoming configuration of communicating peers and the interconnectivity of Internet domains, the degree of end-to-end diversity varies. Therefore, we divide the concept of end-to-end diversity into *degree of multihoming* and *Internet path diversity*.

3.1 Degree of multihoming

The first component of the end-to-end diversity is the degree of multihoming. This component can be affected by the user. Wireless devices supporting several access technologies are becoming more and more prevalent which increses the potential of using several access operators. Most commonly, smart phones support two wireless interfaces such as 3G and WLAN. Current laptops also support wireless connections in addition to Ethernet connectivity. Further, the usage of mobile network connectivity through external modems has become prevalent. Multipath communication enables these different access technologies to be used simultaneously but the decision on the number of interfaces activated remains to be made by the user.

If a user device is multihomed, at least the first hops along the path are disjoint. Whether the packets will continue traversing the separate paths depends on the path diversity in the Internet core. The underlying access technologies and the amount of congestion in different parts of the end-to-end path dictate where the diversity is most needed. Traditionally, the bottleneck has been in the radio interface. However, the development of access technologies has consequently increased the capacity requirements inside and between autonomous systems (AS).

The decision regarding the number of interfaces to use in multipath communication depends on the access network performance. If the user prefers to communicate via WLAN connection, shifting traffic to capacity-constrained 3G interface for load balancing purposes may not be reasonable. Moving traffic from 3G interface to WLAN, however, is more attractive from the user perspective. In constrained network conditions, the usage of multiple 3G radios may also become an attractive alternative if several access operators are available. The emergence of 4G technologies, such as Long Term Evolution (LTE), will increase the multihoming potential of mobile devices.

3.2 Internet path diversity

The estimation of the degree of path diversity is more difficult than that of multihoming. Firstly, diversity depends on the physical connectivity between network devices such as routers. Secondly, the routing policies between these interconnected devices affect the Internet path diversity. Since the Internet is a complex constellation of networks, path diversity varies according to the level of abstraction and its evaluation is relatively challenging.

Basically, the only way of estimating the path diversity between two hosts on the Internet is by measuring. The first effort to estimate the path diversity on the Internet was made by [Texteira, 03]. They took a dualistic approach where they studied the path diversity inside a single AS, *e.g.*, an ISP network and across multiple ASs. They concluded that there is a high potential for path diversity inside ISP networks but the benefits depend on ISPs' capability to engineer the traffic. Therefore, the degree of path diversity may also vary from an ISP to another.

A more recent and profound study is presented in [Han, 06]. The authors conducted an extensive set of measurements to quantify the path diversity in multihomed networks and examine the impact of path diversity in overlay networks. They used the data set to analyze the overlapping routers in the path-, edge- and AS-levels. They wanted to separate edge and AS diversity, since a packet traverse through a same AS does not necessarily mean overlapping routers.

Although [Han, 06] reports a significant effort to study Internet path diversity, the usage of the results in further modeling is challenging. The measurement approach used in [Han, 06] is applicable, but this study requires the average path diversity to be expressed in a more general format. The identification of the path lengths between end points allows the normalization of the average Internet path diversity between zero and one.

Since the focus of this study is inter-domain multipath protocols, this paper claims that a decent level of abstraction for the average path diversity on the Internet is the AS-level. This yields a relatively conservative estimate for the diversity but we argue that this level of abstraction offers the sufficiently measurable approximation of the diversity of the Internet topology. We propose the following formula for estimating the relative inter-AS diversity on the Internet:

$$d = \frac{1}{m}\sum_{i=1}^{m}\frac{h_{diverse}}{h_{total}} \qquad (1)$$

The variable h_{total} refers to the number of inter-AS links that packets traverse during one measurement round. If the packets sent through ISP A_1 to a destination traverse for example three inter-AS links and the packets sent through ISP A_2 traverse four links before reaching the same destination, h_{total} equals seven. Similarly, the disjoint links on each path can be calculated. By dividing the number of disjoint hops by the number of total hops, one obtains the diversity for a specific combination of an origin and a destination. By conducting several measurements (*m*) with various origins and destinations around the globe, the average Internet path diversity can be estimated.

This formula enables the quantification of path diversity also in specific contexts. Instead of defining the average Internet path diversity on a global level, diversity can be defined with regard to a single destination or a specific origin-destination pair. In this case, one of the parameters is kept constant while the other is varied. In addition, the path diversity of a specific geographical location can be measured by using local origin and destination ASs. Figure 1 shows two examples of quantifying path diversity on the Internet.

Since the current inter-domain routing protocol (BGP) limits the routers to choose only one route to each destination prefix, the Internet path diversity is typically presumed to be low. This can be seen from the measurements on [Han, 06] which claimed that multihoming or overlay networks do not guarantee path diversity during communication sessions.

However, the Internet path diversity can be affected by intentional actions of different stakeholders. For example, multipath inter-domain routing proposed by [Xu, 06] would increase flexibility of current BGP routing by controlling the packet forwarding to different paths according to predefined policies. However, the retention of the routing information of several paths towards a destination increases complexity and costs to the operators. Therefore, deploying multipath routing by default does not seem compelling. [Elena, 10] conducted a study which measured the AS-level path deviation, *i.e.*, the deployment of multipath extension in BGP. They noticed that only a fraction of monitored destinations experience path deviations, and AS-level load balancing is not widespread.

ISPs can also negatively affect path diversity. If they see the multipath communication potentially decreasing their revenue or control of the traffic flows in the network, they might have an incentive to start blocking the multipath traffic by using, *e.g.*, deep packet inspection (DPI). The potential multipath traffic blocking by individual ISPs will affect negatively to the overall degree of Internet path diversity.

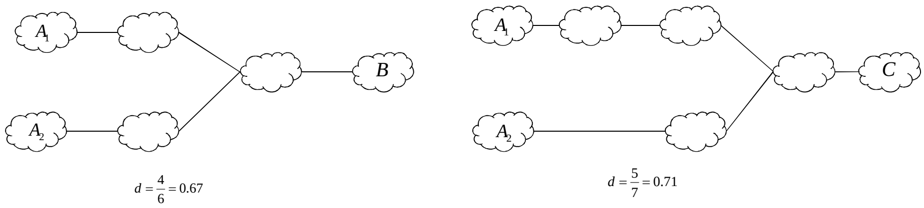

Figure 1: Examples of inter-AS path diversity (end hosts not included)

In addition, end users can potentially increase the usage of alternate paths by source routing. In source routing, the end user (or possibly AS) can decide which path the packet should traverse, see, for example [Yang, 03] or [Argyraki, 04]. Both IPv4 [Postel, 81] and IPv6 [Deering, 98] introduce strict and loose source routing. Strict source routing specifies the whole route that the packets will traverse, while the loose source routing sets only one node which the path should cover on its way to the destination. Although source routing can increase diversity, it introduces a security concern since attackers might use source routing, *e.g.*, to bypass firewalls. Therefore, this protocol option seems not to get widely used in the current Internet.

3.3 Available bandwidth on the Internet

The existence of the Internet path diversity does not alone guarantee the benefits of end-to-end multipath communication. Because the end-to-end multipath communication aims to balance the load on the entire Internet, load levels of alternative links and paths should deviate. If the load levels of Internet links are equal, no further benefits can be achieved by using multipath communication. On the other hand, high difference in available bandwidths introduces a high potential for load balancing. As Internet path diversity, also Internet path congestion can be estimated by measuring.

[Wischik, 09] claimed that the link congestion could be considered as an integral part of path diversity. The resource poolability index measures how easily the traffic can be shifted away from a specific resource when it experiences a data surge. To calculate the resource poolability index the capacities for each parallel link as well as data volumes passing through each link should be measured. However, defining the resource poolability for entire paths becomes complex. In addition, we see that separating the path diversity and available bandwidth is beneficial when evaluating the value of end-to-end protocols.

Because our interest is in modeling the value of multipath protocols from the end user perspective, we propose the usage of user-centric measurements to estimate the available capacity in each path. The available capacity in end-to-end paths, *i.e.*, the maximum TCP throughput which does not affect the rate of existing flows, has been thoroughly investigated and several tools have been proposed for measuring it [Lu, 05]. Some of the tools do not only expose the average available capacity, but also reveal the range in which the available capacity varies over time. Typically, these tools utilize the information on end-to-end delay (either one way or two way) to evaluate the load or available bandwidth in end-to-end paths. We propose that methods such as the one proposed in [Jain, 03] should be used to estimate the available bandwidth of different paths on the Internet.

4 Value model

We build the model by using MPTCP as an example protocol [Ford,11] since it aims to agile Internet-wide load balancing which would overwhelm slower congestion management mechanisms on the Internet. The adoption of MPTCP has been previously studied in [Warma, 11] and [Kostopoulos, 10]. These studies assume MPTCP to increase throughput and resilience of Internet connections but they do not present the derivation and the degree of MPTCP benefits. By introducing a novel way of modeling the benefits of MPTCP, the intention is to shed light on the value creation potential of multipath protocols in the Internet connectivity market.

The benefits of MPTCP can be seen from two perspectives. Firstly, MPTCP yields performance benefits to the users when they deploy MPTCP. Assuming that the traffic on the Internet is unbalanced the users will be able to exploit the underutilized capacity in the network which they will experience as increased throughput. Secondly, MPTCP increases the bargaining power of end users since they can easily switch between the paths. This will enforce the operators to increase the quality of their network as well as lower the connectivity prices to retain customers.

As stated in [Warma, 11], a novel protocol needs to have relative advantage compared to the existing technologies in order to trigger the adoption process. Therefore, the model introduced in this chapter is based on the net benefit offered by the protocol to a human user. The reason for taking the human perspective is that the improved performance provided by MPTCP should lead to better quality of experience (QoE) of the end users [Kilkki, 08]. The total value of MPTCP is the aggregate net benefit of all MPTCP capable end users on the Internet.

Since MPTCP aims to utilize the network resources more efficiently, the protocol will balance the traffic flows among the network providers by moving traffic from congested to underutilized networks. If the benefits to the end users remain non-existent, the protocol only results in a reallocation of the costs and benefits, and the total value for the society will not increase. If MPTCP is capable of increasing the throughput and resilience of connections so that the users will save time, the social welfare also increases.

4.1 Assumptions

The costs and benefits of MPTCP are measured during a user session which consist of several downloads. An example of a user session could be, *e.g.*, a time frame during a web surfing or downloading music pieces or applications from the content provider's online store. Assuming that both end points support the protocol, MPTCP can be used for client-server communication as well as in peer-to-peer (P2P) traffic. The model is applicable for both communication types.

We argue that if MPTCP is capable of bringing value to the end users, their value will also turn into the benefit of content providers. In [Warma, 10] the proposal was that the improved QoE in content providers service would result in some additional application downloads. The limitation of the model is that users were assumed to increase the application sales of a specific content provider. However, the consumption increase of any type of communication service due to improved online QoE is more likely. Therefore, we do not consider the server side benefits explicitly.

A user session consists of multiple TCP or MPTCP flows. The number of flows in a session depends on the applications which are used during the session. For example, webpages consist of multiple objects which are embedded to the main webpage, see, *e.g.* [Svoboda, 08]. The objects may locate in different servers which require opening multiple TCP flows to retrieve the contents of the desired webpage downloaded. The current MPTCP standard is fully transparent to the application layer which means that MPTCP does not support opening new subflows to fetch web objects located in different servers [Ford, 11]. The implementation of an extended API would allow applications layer protocols to better exploit MPTCP features but also reduce the compatibility with existing applications. In the model, we assume that the API is transparent to the applications and a piece of content is retrieved from one server. By the piece of content we refer to a full webpage or application which can be utilized by end users.

End-to-end communication protocols are extremely prone to network effect [Katz, 86] since the communication is not possible without capable peers. As stated in [Kostopoulos, 11] the network effect, which increases the value of the protocol as a function of protocol users, has a great impact on the MPTCP diffusion. However, when the benefits of a single end user are explicitly concerned, the value of network effects is only perceived through the improved QoE not the potential of communicating with other MPTCP capable peers. Therefore, the costs and benefits of MPTCP are not explicitly dependent on other MPTCP users in the network, which was the assumption in [Joseph, 07] and [Iannone, 10].

The proposed model considers only run time costs and benefits. The capital expenditure required for installing software patch of MPTCP and acquiring multihoming configuration are excluded from the model.

4.2 Benefits

The benefits of MPTCP derive from using several subflows simultaneously instead of one TCP flow. In this paper, the modeling has been restricted on two subflows although the existing specifications of MPTCP support several subflows [Ford, 11]. The reason for the assumption is that more than two subflows are unlikely to bring any significant benefit, if that is not the case with two subflows. Therefore, the value creation of dual path communication is essential to elaborate.

The main strength of MPTCP communication compared to the regular TCP is the capability of decreasing the probability for low throughputs during a user session. If the congestion window in the first subflow reduces but the second subflow maintains higher throughput the overall QoE remains satisfactory. The aggregate throughput can be calculated by using expected value for each throughput level.

The effective throughput on each subflow is affected by the congestion control scheme which adapts to the load level on different paths. The current congestion control scheme in the MPTCP standard aims to move the traffic from more congested to less congested paths, *i.e.*, proportionally to the available bandwidth on each subflow [Ford, 11]. To estimate the aggregate throughput experienced by the end user, approximations of the throughputs on subflows are needed. The maximum throughput on each flow depends on the available bandwidth of the tight link [Jain, 03] by which we refer to a link with the minimum available bandwidth on the path. Typically, the tight link is the radio interface since the fixed links in the core network support higher capacities.

We propose a two-level approximation model to estimate the aggregate MPTCP throughput. Depending on the path congestion, the user gets higher throughput with probability p_1 on the first subflow and with probability p_2 on the second subflow. The rest of the time the user experiences lower throughput on each subflow which is a fraction c_1 or c_2 of the higher throughput. Parameters p_i and c_i can be estimated by carrying out measurements similar to [Jain, 03]. By monitoring the congestion on different paths the average load level and the degree of fluctuations can be analyzed. If the measurements do not expose any spare capacity in the parallel path, MPTCP should be useless and parameters p_2 and c_2 decrease to zero and no transmission is possible in the parallel path. Figure 2 shows the approximation of subflow throughputs.

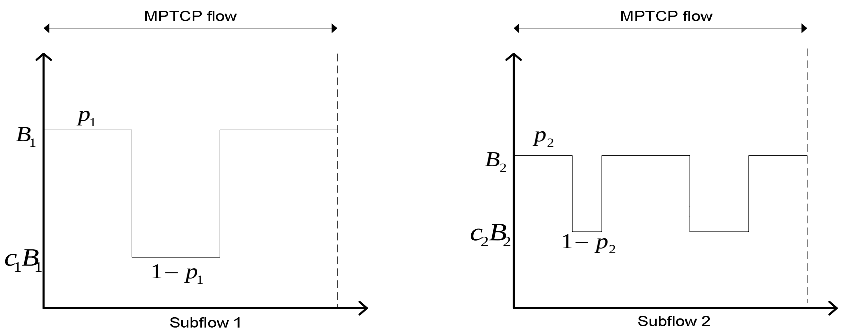

Figure 2: Throughput approximation of subflows

Let us consider that the user has a certain download profile when downloading content from the Internet. He downloads *n* webpages of size f_i and reads each webpage time t_i between each download. Depending on the throughput, the download delay l_i varies. The user can start reading the document only after the full piece of content has been downloaded. The duration of the session is reduced if MPTCP increases the aggregate throughput compared to the regular TCP because the time consumed on reading the web pages is assumed to remain unchanged. Figure 3 illustrates an example of a user session.

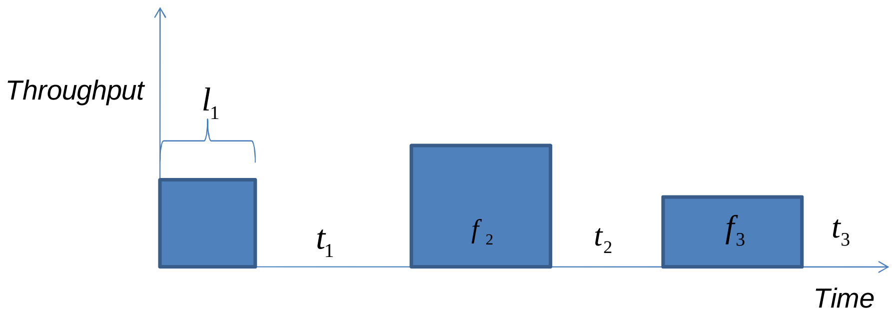

Figure 3: Example of a user download profile in a session

Once the estimation for users' download behavior has been achieved, the expected duration of the session can be calculated as follows:

$$E = \sum_{i=1}^{n} l_i + t_i \qquad (2)$$

To be able to estimate the degree of MPTCP benefits we need to set a reference benefit level. The natural choice for the reference benefit level is the user gain of using single-path TCP communication. The approximation for expected transmission delay with TCP is

$$l_i = p_1 \frac{f_i}{B_1} + (1 - p_1)\frac{f_i}{c_1 B_1} \qquad (3)$$

Similarly, the expected session duration of a user downloading content over MPTCP can be calculated. The expected transmission delay with MPTCP can be estimated by using Equation 4. The equation holds only if the throughputs of the subflows are independent. This assumption does not necessarily apply in practice since the data packets on different subflows may traverse same links and thus affect each other's throughputs. However, the assumption is made for simplicity reasons.

$$l_i = p_1 p_2 \frac{f_i}{B_1 + B_2} + p_1(1 - p_2)\frac{f_i}{B_1 + c_2 B_2} + (1 - p_1) p_2 \frac{f_i}{c_1 B_1 + B_2} + (1 - p_1)(1 - p_2)\frac{f_i}{c_1 B_1 + c_2 B_2} \qquad (4)$$

The throughput on the second subflow depends on the average end-to-end path diversity which is discussed in Chapter 3. If the average inter-AS path diversity on the Internet is low, a host can increase the end-to-end diversity by multihoming which guarantees that the first hops are disjoint. We propose that the throughput on the second subflow can be estimated as follows:

$$B_2 = dB_2^* \qquad (5)$$

where d is the average Internet path diversity normalized between zero and one. The equation says that if the path diversity in the network is 0.5, the links which deviate can expose 50% of the maximum bandwidth on the second path. The maximum value B_2^* depends on the *tight link* on the second path which can be found out by measuring as explained in Chapter 3. Although the relation between the path diversity is somewhat simple, it indicates the importance of Internet path diversity in multipath communication.

Based on the equations presented above, the gross benefit of MPTCP can be defined as the difference of the expected session durations with and without MPTCP.

$$\Delta E = E_{TCP} - E_{MPTCP} \qquad (6)$$

The reason for considering time as the most important component in the model is that people undoubtedly value time. A more disputable question is the extent to which they value it. [Pohjola, 07] proposed that the value of time in communication services can be seen as an opportunity cost, *i.e.*, the lost net benefit of alternative activities. While downloading the piece of content, users typically lose the potential to do more valuable activities. The studies presented in [Niida, 10] show that users perceive waiting time in communication services uncomfortable. However, the context dictates the extent to which users tolerate waiting.

In addition to saving in time, MPTCP may yield other benefits for the user. Depending on the access connection tariffs, the user may experience monetary savings. If the user is connected via 3G which is often usage-priced, the user might be willing to shift the load more on the flat rate-priced (or free) WLAN interface. Balancing the traffic from a flat rate-priced connection to the usage-priced interface is unlikely to happen unless the user experiences a significant gain in the throughput. The savings in the connectivity costs can be estimated and taken into account in the model when the price to transfer a megabyte of data is known. Since the proposed approach considers the benefits as a function of transferred data, the connectivity cost savings can be evaluated when the shifted data volume is known.

As proposed in Chapter 4.1, we argue that the end users do not explicitly value the number of other MPTCP capable peers in the network. They rather perceive the network effect through improved service quality. The more MPTCP capable peers exist in the network the more they are able to use the protocol for downloads. If not all the web services, that the end user accesses during the session, support MPTCP the expected duration of the session should be calculated as presented in Equation 7.

$$E_{MPTCP} = \sum_{i=1}^{n_{TCP}} E_{TCP} + \sum_{i=1}^{n_{MPTCP}} E_{MPTCP} \qquad (7)$$

4.3 Costs

The most obvious cost component of MPTCP is the increased energy consumption of the device. Increased energy consumption can be turned into a monetary cost since the devices needs to be charged more often. Firstly, the establishment of several subflows increases energy consumption. The more subflows are established the more processing power is needed. If the transferred file during the MPTCP flow is short, the control overhead might outweigh the benefits of MPTCP. [Raciu, 11] measured the MPTCP performance and they proposed that MPTCP outperforms TCP in a data centre environment when the file size exceeds approximately 220 kilobytes.

Secondly, multihoming configuration impacts the energy consumption of the device. The study in [Wang, 10] characterized the energy consumption of different radio technologies found in mainstream smart phones. The study shows that when a radio interface is fired up, the energy consumption is excessive. The number of bits which is transferred in the radio interface raises the energy consumption much less. This means that energy-wise multipath transport only makes sense if the transmission rates in the both subflows are high. The authors in [Miettinen, 10] analyzed the processing to communication ratio with mobile cloud services. According to their measurements, WLAN is very energy efficient with low data rates while 3G becomes more energy efficient when the throughput is high.

The measurements in the above-mentioned studies were conducted by measuring a single interface at a time. The overall energy consumption of parallel radio usage during the session is a key issue from the user perspective. What should be measured is whether the overall energy consumption in a session increases or decreases since the data is transmitted and received faster and allowing a quicker turn off of the radios. The simultaneous usage of two radio interfaces requires further study since it may create interference problems.

Algorithms, such as the one developed by [Pluntke, 11], can be used to optimize the energy consumption in devices. This kind of algorithms take into account the overall energy consumption of radio interfaces including the stand-by energy consumption. However, optimizing the energy consumption in multipath communication, results in a decrease in throughput. This paper assumes that the end user wants to maximize its throughput instead of minimizing the energy consumption.

In addition to the energy consumption, struggling with the MPTCP configurations might introduce an additional cost to the end user. If the MPTCP implementation is not working fully transparently and requires, for example, user decisions on which interfaces to use during the session, the time saving of MPTCP may be overwhelmed by the time consumption of configuration tasks. Finally, MPTCP may yield monetary costs if the user starts using usage-priced interface instead of free or flat-rate priced interface.

4.4 Total benefit

Having identified the benefit and cost componenets of the end user, the net benefit and the total value of the multipath protocol can be formulated. The net benefit of user j can be defined as follows:

$$NB_j = \Delta E_j v_j - (e_{p,j} + e_{m,j})u - t_j v_j \pm f_j p \quad (8)$$

where ΔE_j is the time saving calculated by Equation 6 and v_j is the value of time for the user *j*. Increased energy consumption due to MPTCP processing and the usage of two radios are marked as e_p and e_m, respectively. The increased energy consumption can be turned into a monetary value by multiplying e_p and e_m by an estimate of energy unit price *u*. Parameter t_j is the time period elapsed for configuring MPTCP and its lost value can be estimated by using v_j. The increased or decreased connectivity charges are taken into account in Equation 8 by using the current access prices of operators *(p)*, and the data volumes which the user transfers through a certain interface (f_j). If the number of MPTCP capable users on the Internet is *m* the total value of MPTCP is proposed to be the sum of their net benefits as stated in Equation 9.

$$TB = \sum_{j=1}^{m} NB_j \qquad (9)$$

5 Model application

Having discussed the costs and benefits of end-to-end multipath communication, this chapter presents practical applications of the model. As proposed in the previous chapter, the evaluation of the net benefit is divided into two parts: gross benefit and inconvenience or costs incurred by the multipath transport. Two use cases are presented as examples and compared to the case of singlepath TCP.

The modeling starts from estimating the user download behavior. Handset-based measurements and data analysis (see, *e.g.*, [Verkasalo, 09]) are possible ways to increase understanding of users' online behavior. By monitoring the behavior of mobile users, statistics of a certain mobile service, such as browsing, can be extracted from the data. For example, the number of HTML requests and the number of user sessions per day can be used to figure out the data volume transferred in a session. A more straightforward way is to extract the transferred data volumes per day as done by [Falaki, 10]. The main conclusion of [Falaki, 10] was that the handset usage highly fluctuates from the user to another and an average mobile user does not exist. Therefore, the value of multipath communication should be analyzed separately for different end user groups depending on their online behavior.

Let us assume that the user is communicating over 3G HSDPA interface. The maximum download speed of the user's 3G interface is 2000 kbit/s [Wang, 10] but he only gets maximum of 1200 kbit/s when communicating over singlepath TCP. The number of downloaded webpages is three. According to [Svoboda, 08], 6 kilobytes was the mean of a requested webpage in UMTS networks in 2008. However, the sizes of mobile webpages are presumably increased and browsing through several analyzing tools reveals that today the sizes of mobile optimized webpages are measured to be tens of kilobytes, see, for example [Website Optimization, 2012].

Typically, the size of a normal web page not optimized for mobile usage, can be several hundreds of kilobytes. We set 600 kilobytes (4800 kbit) as the size of a piece of content.

The user reads the document 30 seconds before downloading the next. The user regards the reading as desirable and the waiting time as uncomfortable activity. The paths used for transferring the data have similar path characteristics. MPTCP users experience better throughput 60 % of the time while 40 % of the time they experience a throughput which is half of the maximum. To estimate the total value of MPTCP, 100 000 MPTCP capable users on the Internet are assumed. Table 1 summarizes the assumptions.

Model parameter	Symbol	Value
Content downloads per session	n	3
Size of the piece of content	f	4800 kbit
Probability of higher throughput	p_1, p_2	0.6
Decline in throughput	c_1, c_2	0.5
Reading time	t	30 s
Number of MPTCP capable users	m	100 000

Table 1: Model assumptions

Applying the methodology proposed in Chapter 4 we obtain the reference level for the expected session length.

$$E_{TCP} = n\left[p_1 \frac{f}{B_1} + (1-p_1)\frac{f}{c_1 B_1} + t\right] = 3\left[0.6\frac{4800kbit}{1200kbit/s} + 0.4\frac{4800kbit}{0.5 \cdot 1200kbit/s} + 30s\right] = 106.8s$$

5.1 Multipath over 3G

When using MPTCP over the same physical interface the server which the user communicates with should be multihomed. Otherwise, it is unlikely for subflows to find any alternative paths. Since the maximum download speed of the 3G interface is assumed to be 2000kbit/s, MPTCP can increase the throughput 800 kbit/s at most. The average path diversity on the Internet is assumed to be low (d=0.3). Equations 2 and 4 yield 103.5 seconds for the expected MPTCP session lenght.

The results show that MPTCP decreases the uncomfortable time of the session 3.3 seconds. If the Internet path diversity has its maximum value and occasionally the end user is able to achieve the maximum possible reception rate on the second subflow, the expected duration of the session is reduced to 99.5 seconds which is approximately 7.3 seconds less than with the regular TCP. If no time is lost because of configuring the protocol, this is the time that the user saves for using MPTCP.

By using Equations 8 and 9 the total value of the protocol can be evaluated but first the estimates for value of time, increased energy consumption and energy price are needed. Let us say that the value of time is 0.5 €/min which is the value used by [Pohjola, 07]. The energy consumption due to increased processing might increase for example 0.02 W during the content reception compared to the TCP session. The price of the energy equals 15 cents/kWh which was the average price of energy for household consumers in Finland in the first half of 2011 [Eurostat, 11]. If no other costs or benefits incur the user net benefit during a session is 2.722 cents with lower diversity and 6.06 cents with higher diversity. Assuming 100 000 MPTCP capable users with similar download behavior on the Internet, the total value of MPTCP reaches 2722 € and 6060 € respectively.

5.2 Multipath over 3G and WLAN

Let us assume that the user starts using the WLAN interface for the second subflow. The maximum throughput in the interface is 10 Mbit/s [Wang, 10]. Path characteristics remain the same as in the previous example. If the path diversity remains low (d = 0.3) the expected duration of the session is 94.6 s while the best possible diversity would reduce the session length to 91.8 s. Thus, in the best case, MPTCP is able to reduce the session lenght very close to the total reading time.

In addition to the energy consumption due to MPTCP processing, the increased energy consumption using two radios could be, for instance 1 W. Since the user starts communicating over two interfaces, the configuration increases the length of the session with 5 seconds, for example. If the energy price and value of time remain the same as in Chapter 5.1, the total benefit is 6021 € with low path diversity conditions and 8365 € with the highest path diversity. Without the time consumed in configuring the interfaces, the total benefit values are 10 186 € and 12 530 € respectively.

If the user is charged per usage in the 3G interface, the net benefit increases even more when he changes to fixed-priced or free WLAN interface. Usage-based pricing is common, for example, when users are roaming outside their home country. European Union has set an upper limit for the roaming charges which is 50 cents/Mbyte but we assume 10 cents/Mbyte in the calculations. This is the amount of money that user saves by sending data on the second subflow through WLAN interface. Taking into account these cost savings, the net benefit of a MPTCP user increases up to 19.74 cents and the total net benefit of all MPTCP users up to 19 740 € with low path diversity.

WLAN and 3G access interfaces have relatively asymmetric throughputs which raises a question whether the user should switch the whole traffic to the WLAN interface when he comes in the coverage area of a certain access point rather than communicating over two radio interfaces. Assume that the user communicates only over WLAN interface with similar path characteristics as before. He gets, for example, a maximum throughput of 6000 kbit/s with singlepath TCP which results in an expected session length of 93.36 seconds. This is very close to the session length of multipath TCP and it is achieved without the cost of increased battery consumption or the struggle of configuring physical interfaces. Calculating the total value of singlepath TCP over WLAN results in a total value of 11 190 €. This indicates that singlepath TCP can be outperform MPTCP in some conditions when only a slight increase in energy consumption due to processing is assumed.

6 Discussion

The model in this paper evaluates the economic potential of end-to-end multipath protocols on the Internet. It explicitly identifies the factors which affect the adoptability of MPTCP from the end user perspective in the early phase of the adoption process. To better communicate these factors to the readers, unnecessary complexity is avoided.

The model could be extended to cover more precise and even continuous probability distributions. During the model development, TCP throughput distributions measured in [Franceschinis, 05] and [Halepovic, 08] were approximated both with two- and eleven-level probability distributions with the same average value and standard deviation. The analysis showed only a marginal difference between the MPTCP gains. Only if the throughput has occational drops close to zero, the two-level distribution yields over-estimated results for MPTCP. To cover this case, a more precise probability distribution should be applied.

If the singlepath TCP or combined MPTCP throughput transiently drops to zero, our modeling as such cannot be applied (due to a zero value in the denominator). However, from authors perspective this seems unlikely but if it happens, the case of zero transmission should be considered separately in the value analysis.

The proposed model only considers interactive data exchange. The benefits of services using streaming technologies, such as real-time video, cannot be estimated with the proposed model. However, similar modeling method could be developed also for estimating the value of real time multipath communication which could be worth of further study.

A limitation of the model is the relation between the path diversity and the increased throughput. The model assumes that Internet path diversity increases the throughput on the second subflow linearly. However, the path redundancy in the shared bottlenecks might be more valuable than in the links which do not often experience congestion. If many disjoint paths exist in the first hops but many hosts share one link to a certain destination in the middle of the route, MPTCP may not be able to alleviate the congestion problem. The relations between the path diversity and increased throughput can be refined after the required measurements have been conducted, and the understanding of the Internet topology and congestion state has increased.

Chapter 5 applied the proposed methodology in basic use cases of MPTCP and compared its value to the regular TCP communication. The results indicate that with the current energy price MPTCP does not introduce a significant monetary cost to end users. Bigger pain is caused to the end user if the battery drains in the middle of an important session. However, the impact of the energy cost may change in the future when energy prices are increasing globally. Therefore, low energy consumption may become a critical feature not only for multipath but also for other networking protocols.

Depending on the market, the model can be also applied in many alternative scenarios and contexts taking into account those aspects which are relevant for the use case in question. An example of an alternative scenario is the Indian mobile market which is known to be highly competitive [Sridhar, 12]. Users in India currently acquire multiple WAN connectivity subscriptions to find the cheapest and best quality voice connections in different locations. Therefore, India is one of the countries where people favour using dual-SIM phones [Sheshagiri, 07]. In the future, a dual-SIM capable device combined with MPTCP might allow using, for instance, multiple 3G interfaces simultaneously. The advantages of MPTCP are pronounced when the throughput on both radio links is limited. Enabling the concurrent usage of several limited access connections might result in an acceptable user experience.

As stated earlier, the proposed approach concentrates on modeling the value of the protocol in the early phase of the adoption process. In the longer term, when the number of MPTCP users increases, the network resources will be used more efficiently and the load level of the most congested links in the operator's networks will decrease. However, the available capacity released by MPTCP will eventually be filled with increased data volumes generated by the existing users and new subscribers. This means that the throughput gain achieved by MPTCP may remain short-term. However, MPTCP and other multipath protocols increase the bargaining power of the end users which increases the pressure for ISPs to improve their network quality and to decrease connectivity prices. Thus, the multipath protocols, if adopted, will eventually benefit end users. However, modeling the value of increased competition should take a yet another approach compared to this study.

7 Conclusion

By using MPTCP as an example protocol this study has modeled the net benefit of the end-to-end multipath communication in a user-centric manner. According to the model, multipath TCP seems initially more appealing to wireless end hosts when used over one access operator at a time compared to running several access operator connections in parallel. The concurrent usage of multiple radio connections increases the maximum throughput but the additional energy consumption or struggle with the protocol may outperform the benefits. As the battery technologies improve in the future, simultaneous usage of several radio interfaces might become a compelling alternative. The conclusion is that the Internet-wide load balancing may develop initially on a per session basis and later on proceed towards more real-time load balancing that requires specific multipath protocols such as MPTCP.

The overall value potential of Internet-wide end-to-end load balancing depends essentially on measurable parameters: path diversity and unused capacity. Our model allows quantification of the total net benefit provided that statistically sufficient field measurements can be performed. Such measurements are an obvious, although challenging next step in the research. We anticipate that measurements will show differences in value potential between parts of the Internet, for instance between geographical areas. Also, this type of user-centric and measurement-driven value modeling can be useful for protocols other than multipath, which opens another direction of future research.

From the market perspective, multipath communication is significant since it increases competition. More flexibility is given to end users to choose their access providers, which increases competition between operators and eventually may push access prices down. This raises an intriguing tussle between end users and operators about the control of (multipath) communication on the Internet.

Acknowledgements

This work was supported by TEKES as part of the Future Internet program of TIVIT (Finnish Strategic Centre for Science, Technology and Innovation in the field of ICT). The authors would like to thank Jaakko Korkeaniemi for the valuable comments.

References

[Argyraki, 04] Argyraki, K., Cheriton, D. R.: “Loose source routing as a mechanism for traffic policies”, ACM Workshop on Future Directions in Network Architecture, Portland, OR, USA, August 2004.

[Awduche, 99] Awduche, D.: “MPLS and Traffic Engineering in IP networks”, IEEE Communications Magazine, 37(12), 42-47, 1999.

[Blanchet, 10] Blanchet, M., Seite, P.: “Multiple Interfaces Problem Statement”, Internet Draft, Available at http://tools.ietf.org/html/draft-ietf-mif-problem-statement-07.html, accessed 8 June 2011.

[Deering, 98] Deering, S., Hinden, R.: Internet Protocol, Version 6 (IPv6) Specification, RFC 2460, December, 1998.

[Elena, 10] Elena, E., Rougier, J-L., Secci, S.: “Characterisation of AS-level Path Deviations and Multipath in the Internet Routing”, 6th Euro-NF Conference on Next Generation Internet (NGI 2010), Paris, France, June 2010.

[Eurostat, 11] Eureopean Commission, Eurostat statistics, available at http://epp.eurostat.ec.europa.eu/statistics_explained/index.php/Energy_price_statistics, accessed June 11th, 2012.

[Falaki, 10] Falaki, H, Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R., Estrin, D.: “Diversity in Smart phone usage”, 8th international conference on Mobile systems, applications, and services (MobiSyS 2010). San Fransisco, CA, USA, June 2010.

[Ford, 11] Ford, A., Raciu, C., Handley, M., Barre, S., Iyengar, J.: “Architectural Guidelines for Multipath TCP Development”, RFC 6182, March 2011.

[Franceschinis, 05] Franceschini, M., Mellia, M., Meo, M., Munafò, M.: “Measuring TCP over wifi: A real case”, 1st workshop on Wireless Network Measurements (WiNMee 2005). Lago di Gardi, Italy, April 2005.

[Gleeson, 00] Gleeson, B., Lin, A., Heinanen, J., Armitage, G., Malis, A.: “A Framework for IP Based Virtual Private Networks”, RFC 2764, February 2000.

[Greenberg, 09] Greenberg, A., Hamilton, J., Maltz, T. A., Patel, P.: “VL2: Scalable and Flexible Data Center Network”, ACM SIGCOMM Computer Communication Review, 39(1), 51-62, 2009.

[Guo, 04] Guo, F., Chen, J., Li, W., Chiueh, T.: “Experiences in Building A Multihoming Load Balancing System”, Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOMM 2004), March 2004.

[Halepovic, 08] Halepovic, E., Wu, Q., Williamson, C., Ghaderi, M.: “TCP over WiMAX: A Measurement Study”, IEEE International Symposium on Modeling, Analysis and Simulation of Computers and Telecommunication Systems (MASCOTS 2008). Baltimore, MD, USA, September 2008.

[Han, 06] Han, J., Watson, D., Jahanian, F.: “An Experimental Study of Internet Path Diversity“, IEEE Transactions on Dependable and Secure Computing, 3(4),273 – 288, 2006.

[Huitema, 95] Huitema, C.: „Multi-homed TCP“, Old Internet Draft, Available at http://tools.ietf.org/html/draft-huitema-multi-homed-01, accessed 8 June 2011.

[Iannone, 10] Iannone, L., Levä, T.: “Modeling the economics of Loc/ID Split for the FutureInternet”, In G. Tselentis et al. (Eds.) Towards the Future Internet, 21-30, IOS Press, 2010.

[Iyengar, 06] Iyengar, J.R. Amer, P.D., Steward, R.: “Concurrent Multipath Transfer Using SCTP Multihoming over Independent End-to-End Paths”, IEEE/ACM Transactions on Networking, 14(5), 951-964, 2006

[Jain, 03] Jain, M., Dovrolis, C.: “End-to-end Available Bandwidth: Measurement Methodology: Dynamics and Relation with TCP Throughput”, IEEE/ACM Transactions on Networking 11(4), 537- 549, 2003

[Joseph, 07] Joseph, D., Shetty, N.: “Modeling the adoption of new network architectures”, 3rd International Conference on emerging Networking EXperiments and Technologies (CoNEXT 2007), New York, NY, USA, December 2007.

[Katz, 86] Katz, M. L., Shapiro, C.: ”Technology Adoption in the Presence of Network Externalities”, Journal of Political Economics, 94, 822-841, 1986.

[Kaufman, 09] Kaufman, M.: RTMFP Overview. IETF-77 meeting, Anaheim, CA, USA, available at http://www.ietf.org/proceedings/10mar/slides/tsvarea-1.pdf, accessed 6 June 2011.

[Kostopoulos, 10] Kostopoulos, A., Warma, H., Levä, T., Heinrich, B., Ford, A., Eggert, L.: „Multipath TCP Adoption: Challenges and Opportunities“, 6th Euro-NF Conference on Next Generation Internet (NGI 2010), Paris, France, June 2010.

[Key, 06] Key, P., Massoulie, L., Towsley, D.: ”Combining Multipath Routing and Congestion Control for Robustness”, 40th Conference on Information Sciences and Systems (CISS 2006), Princeton, NJ, USA, March 2006.

[Kilkki, 08] Kilkki, K.: “Quality of Experience in Communications Ecosystem”, Journal of Universal Computer Science, 14(5), 615-624, 2008.

[Lu, 05] Lu, D., Qiao, Y., Dinda, P.A., Bustamante, F.E.: “Characterizing and Predicting TCP throughput on the Wide Area Network”, 25^{th} IEEE International Conference on Distributed Computing Systems (ICSCS’05), Columbus, OH, USA, June 2005.

[Niida, 10] Niida, S., Uemura, S., Nakamura, H.: “User Tolerance for Waiting Time”, IEEE Vehicular Technology Magazine, 5(3), 61-67, 2010.

[Nikander, 10] Nikander, P., Gurtov, A., Hnederson, T.R.: “Host Identity Protocol (HIP): Connectivity, Mobility, Multi-Homing, Security, and Privacy over IPv4 and IPv6 Networks”, IEEE Communications Surveys & Tutorials, 12(2),186-204, 2010.

[Miettinen, 10] Miettinen, A. P., Nurminen J. K.: “Energy Efficiency of Mobile Clients in Cloud Computing”. 2^{nd} USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 2010), June 2010.

[Plunkte, 11] Pluntke, C., Eggert, L., Kiukkonen, N.: “Saving Mobile Device Energy with Multipath TCP”, 6^{th} ACM International Workshop on Mobility in the Evolving Internet Architecture (MobiArch 2011), Washington, D.C., USA, June 2011.

[Pohjola, 07] Pohjola, O-P., Kilkki, K.: “Value-based methodology to analyze communication services”, NETNOMICS, 8(1-2), 135 – 151, 2007.

[Raciu, 11] Raciu, C., Barre, S., Plunkte, C., Greenhalgh, A., Wischik, D, Handley, M.: “Improving Datacenter Performance and Robustness with Multipath TCP”, ACM SIGCOMM, Toronto, Canada, August 2011.

[Postel, 81] Postel, J.: Internet Protocol Specification, RFC 791, September 1981.

[Singh, 11] Sing, V., Kärkkäinen, T., Ott, J., Ahsan, S., Eggert, L.: “Multipath RTP (MPRTP)”, Internet Draft, October 2011, available at http://tools.ietf.org/html/draft-singh-avtcore-mprtp-03, accessed 11 June 2012.

[Seshagiri, 07] Seshagiri, S., Aman, S., Joshi, D.: “Connecting the “bottom of the pyramid”: an exploratory case study of Indian’s rural communication environment”, 16th International Conference on World Wide Web, Banff, Canada, May 2007.

[Sridhar, 12] Sridhar, V., Casey, T., Hämmäinen, H.: “Systems Dynamics Approach to Analyzing Spectrum Management Policies for Mobile Broadband Services in India”, International Journal of Business Data Communications and Networking, 8(1), 37-55, 2012.

[Svoboda, 08] Svoboda, P.: “Measurement and modeling of Internet Traffic over 2.5G and 3G Cellular Core Networks”, Doctoral Thesis, Vienna University of Technology, 2008.

[Tang, 08] Tang, L., Chen, Q., Zhang, B., Li, Y., Zheng, X.: “Load Balancing Model based on Stackelberg Game for Multi-homing in Heterogeneous Radio Access Networks”. International Conference on Advanced Infocomm Technology (ICAIT 2008), Da Meisha, Senhzen, China, July 2008.

[Texteira, 03] Texteira, R., Marzullo, K., Savage, S., Voelker, G.M.: “Characterizing and Measuring Path Diversity of Internet Topologies”, International Conference on Measurements and Modeling of Computer Systems (SIGMETRICS 2003) (extended abstract), San Diego, CA, USA, June 2003.

[Verkasalo, 09] Verkasalo, H.: “Handset-Based Analysis of Mobile Service Usage“, Doctoral Thesis, Department of Electronics, Communications and Automation. Helsinki University of Technology, 2009.

[Wang, 10] Wang, L., Manner, J.: “Energy Consumption Analysis of WLAN, 2G and 3G interfaces”, ACM International Conference on Green Computing and Communication (GreenCom2010), Hangzhou, China, December 2010.

[Warma, 10] Warma, H., Levä,T., Eggert, L., Hämmäinen, H., Manner, J.: „Mobile Internet in Stereo: an End-to-End Scenario“, 3rd Workshop on Economic Traffic Management (ETM 2010), Amsterdam, The Netherlands, September, 2010.

[Warma, 11] Warma, H., Levä, T., Tripp H., Ford, A., Kostopoulos, A.: „Dynamics of communication protocol diffusion: the case of MPTCP“, NETNOMICS Journal, 12(2), 133-159, 2011.

[Wischik, 08] Wischik, D., Handley, M., Braun, M. B.: “The Resource Pooling principle”, ACM SIGCOMM Computer Communication Review, 38(5), 47-52, 2008.

[Wischik, 09] Wischik, D., Handley, M., Raciu, C.: “Control of Multipath TCP and optimization of the multipath routing in the Internet”, 3rd Workshop of on Network Control and Optimization (NET-COOP 2009), Eindhoven, the Netherlands, November 2009.

[Webite Optimization, 2012] Webite Optimization LLC: Web Page Analyzer 0.98, available at http://www.websiteoptimization.com/services/analyze/, accessed May 7th , 2012.

[Xu, 06] Xu, W., Rexford J.: "MIRO: Multi-path Interdomain Routing", ACM SIGCOMM, Pisa, Italy, September 2006.

[Yang, 03] Yang, X.: "NIRA: A New Internet Routing Architecture", ACM SIGCOMM workshop on Future directions in network architecture (FDNA 2003), Karlsruhe, Germany, August 2003.

[Zhang, 04] Zhang, M., Lai, J., Krishnamurthy, A., Peterson, L., Wang, R.: "A Transport Layer Mechanism for Improving End-to-End Performance and Robustness Using Redundant Paths", USENIX Annual Technical Conference, Boston, MA, USA, June 2004.