Towards the Future Internet
G. Tselentis et al. (Eds.)
IOS Press, 2010

doi:10.3233/978-1-60750-539-6-21

Business Aspects of Multipath TCP Adoption

Tapio LEVÄ[a,1], Henna WARMA[a], Alan FORD[b], Alexandros KOSTOPOULOS[c], Bernd HEINRICH[d], Ralf WIDERA[d], and Philip EARDLEY[e]
[a]*Aalto University School of Science and Technology, Finland*
[b]*Roke Manor Research, United Kingdom*
[c]*Athens University of Economics and Business, Department of Informatics, Greece*
[d]*T-Systems International GmbH, Germany*
[e]*British Telecommunications plc, United Kingdom*

Abstract. Multipath TCP (MPTCP) is a resource pooling mechanism which splits data across multiple subflows (paths) and ensures reliable data delivery. Although Multipath TCP is a relatively small technical change to the TCP protocol providing improved resilience, throughput and session continuity, it will have considerable impact on the value networks and business models of Internet access provisioning. In this paper, we evaluate the viability of different MPTCP deployment scenarios and present what new ISP business models might be enabled by the new flexibility that MPTCP brings. This allows the research community to focus on the most promising deployment scenarios.

Keywords. Future Internet, multipath TCP, resource pooling, business models, deployment scenarios, adoption incentives, virtual multipath operator

Introduction

The resource pooling principle [1] advocates making improved use of the Internet's resources by allowing separate resources (such as links or processing capability) to act as if they were a single large resource. One particular manifestation of this principle is the development of multipath transport protocols (specifically Multipath TCP [2]), whereby multiple paths between two endpoints can be pooled to appear to the application as a single transport connection through dynamic scheduling of traffic across the available paths. MPTCP provides higher resilience to link or node failure. Furthermore, it increases bandwidth and resource utilization due to coordinated congestion control for the multipath data transfers [3].

From a technical point of view, MPTCP needs only a relatively small change to the TCP/IP stack at the end hosts which is currently being standardized in the IETF [4]. Although, at first sight, MPTCP may seem simply to be about technology, it will also significantly change the value networks between stakeholders in the Internet connectivity market.

---

[1] Corresponding Author, email: tapio.leva@tkk.fi

Acknowledgement: This work has been supported by the European Commission within the INFSOICT-216372 TRILOGY Project (www.trilogy-project.org).

The main contribution of this paper is to present different scenarios for MPTCP deployment and evaluate their viability based on 1) end-user costs and benefits (Section 3), and 2) the business opportunities for ISPs (Section 4). In particular, we present some new ISP business models that will be enabled by the different deployment scenarios. We also propose a new market player, a virtual multipath operator.

1. Stakeholder analysis

The first thing to note about Multipath TCP is that its deployment but not necessarily its success is dependent upon endpoints only. There are no technical adaptations required at intermediate providers in order to support its deployment. That does not mean, however, that other stakeholders do not have an interest in MPTCP. Therefore we identify the stakeholders and their motivations regarding the deployment and use of MPTCP, and show how these motivations will impact each other.

Figure 1 visualises identified stakeholders along with their level of interest in MPTCP deployment. This two-axis separation illustrates at what point during the development and deployment of MPTCP the stakeholder becomes interested, and whether they are actively involved in its development and deployment, or whether they are just passive observers.

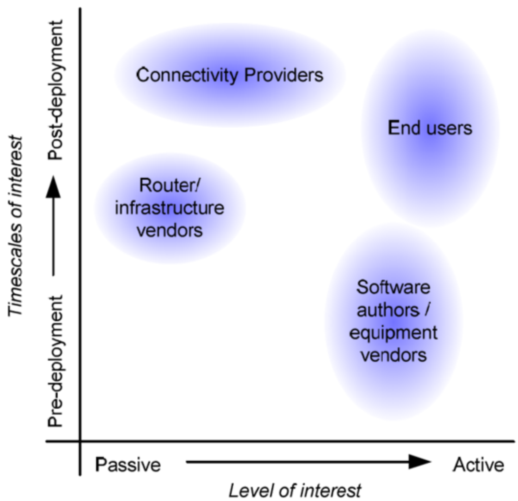

Figure 1. Stakeholders' interests in MPTCP deployment

End-users

The term "end-user" is widely used to cover not only individuals like domestic and mobile users but also large sites, such as corporate or academic sites, and content/service providers. These stakeholders have an active interest at "run time" (deployment and post-deployment), but also have an interest before deployment since they are in a position to influence other stakeholders, particularly software authors, to provide the MPTCP solution. For the further analysis in Sections 2 and 3 end-users are classified as 1) light, 2) heavy or 3) content provider since each of them will have a different notion to deploy MPTCP. The classification is based on users' intensity of Internet usage and their resilience requirements.

Connectivity providers (ISPs)
Connectivity providers can be categorised as either access providers (tier 2) or transit providers (tier 1). Although both can be passive to the end-user driven deployment of MPTCP, it is likely that they will have significant post-deployment interest because of the changes in traffic or business models due to MPTCP use. Tier 1 ISPs are likely to be passive beneficiaries of resource pooling, in that MPTCP should shift load away from congested bottlenecks if alternative paths are available. Tier 2 ISPs, however, have much higher interest since MPTCP's prerequisite for multihoming increases the demand for connectivity and enables new business models (see Section 4).

Software (OS) authors and equipment vendors
These are the stakeholders who actually make the decision to implement and supply MPTCP support. By its very nature this interest lies in pre-deployment. There must be a motivation, however, for the software authors to provide MPTCP support, and this pressure will come primarily from end-users who wish to make use of multipath functionality. The benefits of adding MPTCP will initially be as a product differentiator, and later followed as a necessity to maintain the same functionality as competitors. Open source software will likely provide a significant driver here. Additionally, vendors who are also content-providers (i.e. they have control over hosts at both client and server side), e.g., Nokia with its Ovi services or Apple with its iTunes Store and App Store, may see business opportunities in MPTCP. However, this interesting business case is left for a future research topic.

The deployment scenarios illustrated in Section 3 (except one version of Scenario 1) presuppose that MPTCP support is available in end-user hosts, and thus the role of the software and equipment vendors is not discussed further in this paper.

Router and infrastructure vendors
Although no changes are required at routers to support MPTCP, router and infrastructure vendors may have a vested (passive) interest in its deployment. For example, if ISPs are concerned about multipath TCP having an impact on their traffic engineering or business models, and vendors who supply boxes to support traffic engineering see MPTCP as a threat, then they may act to prevent MPTCP deployment.

2. Costs and benefits of multipath TCP

Before diving deeper to the pros and cons of different deployment scenarios in Section 3, it is necessary to understand what are the general costs and benefits of MPTCP deployment for end-users and ISPs.

2.1. End-user's costs and benefits

In order to get MPTCP benefits, there are potentially two separate costs for end-users. Firstly, there is a cost involved with the deployment of MPTCP software, potentially monetary but primarily due to time and effort involved. Deploying MPTCP may take the form of installing a specific extension (i.e. the end-user is aware and makes an explicit choice), or it may come bundled as part of an operating system (in which case the end-user may not be aware, or may need to enable/configure it). Early adopters will have the highest costs in deployment due to relative complexity. As MPTCP becomes available and enabled by default in operating systems, this cost will tend towards zero.

Secondly, there is the cost of additional physical connectivity. In order for hosts to get resource pooling benefit across a MPTCP connection, at least one endpoint must be multihomed to create the requisite multiple paths. This can create an additional cost, e.g. through buying an additional DSL or WLAN connection. However, many large providers and mobile end-users are already multihomed, and with a little network reconfiguration they could quite easily make use of existing multi-addressing.

The benefits of MPTCP are seen differently by different end-user groups:

- *Higher resilience:* This is relevant to all end-users; however the larger sites will typically see a larger benefit since there are more people that could be inconvenienced in the case of failure. Furthermore, companies' need for business continuity increases their demand for resilience.
- *Session continuity:* In addition to resilience, in the event of deliberate changes in connectivity, MPTCP will allow the session to continue. This is predominantly of interest to (light or heavy) mobile users.
- *Higher throughput:* Most relevant to heavy users (domestic/sites, i.e. humans), however indirectly it is also of benefit to providers as it gives a perception of a better service. For content providers, higher throughput also potentially means more sales.

It is important to note that these benefits do not directly create revenue for end-users. They are rather enhancements that, through greater availability or better perceived service, bring in more business for suppliers, less hassle for domestic users, and more productive time for companies.

2.2. ISPs' costs and benefits

Due to higher visibility of throughput and network performance enabled by MPTCP, ISPs may become more comparable in real time. This increased competition and reduction in provider lock-in may be seen as a "cost" by ISPs, although innovative ISPs should see this as an opportunity to attract more customers through better or more MPTCP-tailored services.

The main benefit for ISPs is indirect and based on increased demand for connectivity, and new business opportunities. However, MPTCP may also allow ISPs to run closer to capacity, since spikes on load can be spread across links more effectively. Additionally, support calls are a major cost for ISPs, and if MPTCP can mask occasional, unexpected outages then it is likely that ISPs will have fewer such calls to deal with.

3. Potential MPTCP deployment scenarios

In this section, we consider different technical architectures which illustrate how MPTCP can be deployed in the access part of the network. We concentrate on the end-user perspective and at the end of the section we evaluate the attractiveness of each deployment scenario to end-users.

With the deployment of MPTCP, an end-user can be connected to either one or many ISPs through one or many access technologies. To simplify the following analysis, we restrict the number of physical connections to ISPs and the number of ISPs that an end-user is connected to two here, and we state that two can be generalized to n.

3.1. Scenario 1: End-user with single physical access to one ISP

Figure 2 illustrates the technical architecture of a scenario where the end-user has only one physical connection, such as a DSL access link, to a single ISP. This kind of deployment scenario also requires the involvement of the ISP1 in order to have MPTCP deployed because it is the ISP1 that shares the traffic into separate routes.

It is up to the ISP1 how it splits the end-users' traffic into multiple paths. First opportunity for the ISP is to provide the multipath feature through a proxy, splitting the end-user traffic into two independent paths to different transit providers. In this case, the end-user's host does not have to be MPTCP capable but there has to be a way how the different MPTCP flows are separated from each other in the network. If the end-user's host is MPTCP capable, the ISP can give the end-user multiple IP-addresses and provide many sessions over the single access link. The ISP ensures that traffic flows originating from the different end-user source addresses get routed to different transit providers to enable disjoint multipath connections through the Internet.

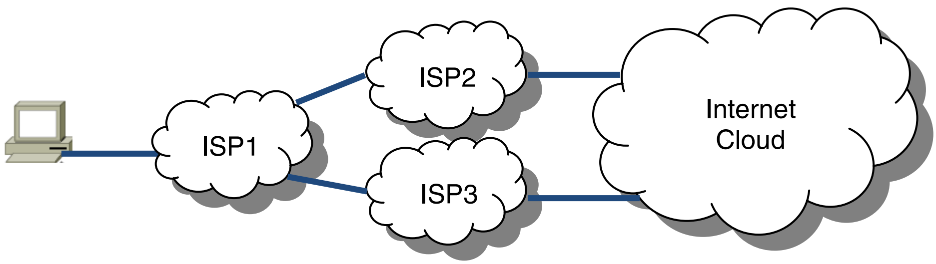

Figure 2. Technical architecture of deployment scenario 1.

This implementation provides the benefits of MPTCP, such as improved throughput. However, the throughput increment depends highly on the performance of the ISP. On the other hand, improvements in resilience only exist further into the network, beyond ISP1, because the traffic is operated between the end-user and its ISP through only one link. If the single access link fails there is no other path where the traffic could be re-routed. Also, in the proxy case the end-user does not have control over the load balancing - ISP1 decides on routing of MPTCP traffic flows.

3.2. Scenario 2: End-user with dual physical access to one ISP

Another possible scenario is the end-user having at least two separate physical connections towards a single ISP. The physical links can be either similar, such as two DSL connections, or dissimilar, such as a DSL connection and a 3G connection. Figure 3 illustrates the technical architecture of the deployment scenario with two different access technologies. In this case, the end-user host has to have its protocol stack updated to support MPTCP. This type of scenario can be deployed without the involvement of the ISP because the MPTCP functionality is implemented at the end-users' host.

In this kind of scenario the end-user has the control of the links and the host is able to balance the load between the access links. Also, the end-user is given transparency of the multipath links provided by MPTCP and has the power to assess the performance of the links. But if the disparity between bandwidths of the access links is large, the throughput advantage of acquiring two access links is limited. With two physical connections the user gets better resilience, since the access link is normally the least reliable part of the end-to-end path.

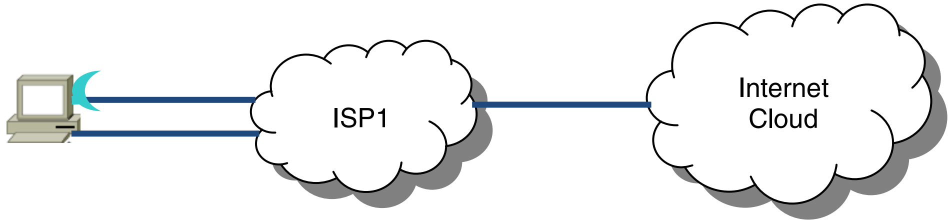

Figure 3. Technical architecture of deployment scenario 2.

Because end-users need to contract with their ISPs for using multiple physical connections, the increased costs of Internet connectivity may be a drawback of this scenario. Additionally, an end-user still relies on a single ISP which might become a bottleneck.

3.3. Scenario 3: End-user with dual physical access to different ISPs

Finally, we introduce a deployment scenario where the end-user has multiple similar or dissimilar access links to, at least, two different ISPs. In Figure 4, we present the technical architecture of the scenario where the end-user uses different access technologies to connect different ISPs.

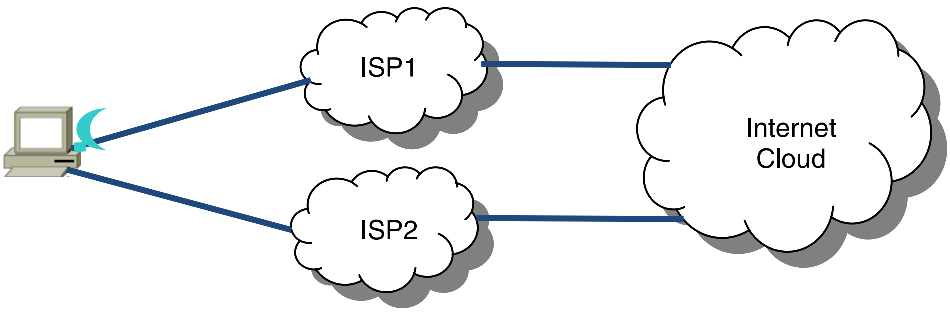

Figure 4. Technical architecture of deployment scenario 3.

In this scenario, the end-user needs to contract with at least two different ISPs for using their access links. On one hand, this gives more bargaining power to the end-user but on the other hand it entails the burden of handling many contracts and paying multiple bills. The end-user is not dependent on only one ISP and the Internet connection is not lost even though one of the operators fails.

If the end-user uses dissimilar types of technologies to connect to the ISPs, the links have different characteristics in reliability and throughput. Thus the end-user loses some of its ability to compare the ISPs. This scenario also increases the end-users' monetary costs of Internet connectivity. Additionally, for mobile users, the use of multiple radio interfaces at the same time will increase the power consumption and shorten the battery life.

3.4. Viability of deployment scenarios

All the deployment scenarios we have presented in this section have their costs and benefits located at the end-user. In this section, we evaluate the viability of different scenarios by presenting the attractiveness for the different types of end-users and the most important advantages and drawbacks of each scenario. Table 1 summarizes our evaluation.

A typical light user is a domestic consumer who is not interested in deploying MPTCP if it brings additional costs such as effort or monetary costs. Heavy users or content providers, on the other hand, may have incentives to invest time and money for getting the benefits entailed by MPTCP. Thus heavy users will want to enjoy the best possible performance and full resilience of MPTCP, while light end-users prefer lower monetary costs and a simple deployment.

In some cases, the attractiveness of MPTCP to an end-user depends on whether the scenario with similar or dissimilar physical connections is in question. We find that scenarios where the end-users are able to use multiple access technologies are more attractive compared to a deployment with multiple similar connections. Current laptops and mobile devices already support multiple access technologies and MPTCP allows more efficient usage of the available connectivity resources. 3G enables ubiquitous connectivity while WLAN, e.g. in the office, offers higher bandwidth. Also seamless handover, when, e.g., a mobile user returns home and starts using WLAN for higher bandwidth, is a remarkable benefit.

If MPTCP capability is not rolled out to end-user hosts, the ISP has to provide this functionality through an additional proxy as we stated in Section 3.1. This deployment scenario is attractive only to light end-users and it is questionable if ISPs want to take the effort addressing this user group which is reluctant to pay more for the increased network throughput and reliability. All other scenarios require the implementation of MPTCP at the end-users host. Hence, this issue should be driven with software providers.

Table 1. Evaluation of deployment scenarios based on their attractiveness, advantages and drawbacks.

Deployment scenario	**Attractiveness**			**Advantages**	**Drawbacks**
	Light end-user	**Heavy end-user**	**Content provider/Site**		
1. ISP acts as a multipath operator	Medium	Low	Low	Requires the least effort from the end-user	Limited benefits due to single access line, e.g., limited resilience
2. Disjoint connectivity to a single ISP	Low	High	Medium	Double access under one contract	End-users are still dependent on one ISP
3. Disjoint connectivity to different ISPs	Low	High	High	Power to race ISPs	Burden of having multiple contracts and bills

To conclude our evaluation we can say that for a light end-user the deployment of MPTCP is not that attractive. For heavy end-users and content providers the scenarios which offer additional incentives or where multihoming capability exists already for other reasons are most interesting. For example, mobile end-user variants of Scenarios

2 and 3 with high capacity fixed or nomadic access supported with low capacity mobile access have potential because mobile devices already support multiple access technologies and different access connections complement, not substitute, each other. Similarly, the fixed multihoming architecture of Scenario 3 is already today a reality for content providers and large sites, so harvesting additional benefits through MPTCP adoption would be reasonable for them.

To assess whether the ISPs are eager to encourage adoption of this technology by the end-users, we have to evaluate the possible ISP business models enabled by each deployment scenario, which is done in the following section.

4. ISP business models

In this section we move the focus to ISPs and their business opportunities related to the introduction of MPTCP. We present a specific business model enabled by each deployment scenario. The attractiveness of a business model and the corresponding deployment scenario from an ISP's perspective is evaluated by using SWOT analysis.

4.1. ISP offers MPTCP as a value added service (enabled by Scenario 1)

Upgrading a single access link with a multipath feature will bring a unique selling proposition to the ISP. With the proxy solution, ISPs can give end-users an early start to gain MPTCP's benefits before the roll-out of MPTCP in computer operating systems. This advantage will be chargeable to end-users, together with the MPTCP inherent benefits of connection reliability, compensating for the ISPs efforts to provide the proxy. There is also chargeable added value to the end-user in the other case, where ISP provides many sessions (i.e. multiple IP addresses) over a single access link.

The customer target groups of these scenarios are end-user's in need of multipath functionality, but not willing to pay for a second access link, such as consumers and SMEs (Small/Medium Enterprises) with a single DSL connection. However, also mobile scenarios of consumers or business users are possible. The ISP can offer this multipath-upgrade as a premium feature to its current customer base to retain customers with increased demands.

Table 2. ISP SWOT analysis of "MPTCP as a value added service" business model

Strengths • Ownership of customer access line • Existing customer relationship will be maintained • Improved resilience and thus service quality because of multiple peering ISPs	**Weaknesses** • May require additional peering contracts if ISP1 was previously connected only to one Tier 1 ISP • Cost to establish a solution facilitating disjoint MPTCP paths • Potential dependency on software vendors to develop MPTCP-proxy functionality • Throughput dependencies on the performance of the upstream ISPs
Opportunities • Revenue by access provision • Unique selling proposition for MPTCP provisioning • Warranty of reliable throughput can be charged to customers	**Threats** • Possible visibility of bad network or proxy performance may lead to complaints and loss of customers • Customers may not be willing to pay for the service to sufficiently compensate for the ISP's additional costs

4.2. Access connection bundling (enabled by Scenario 2)

Although Scenario 2 does not require ISP involvement, the ISP can attract customers through lucrative access connection bundles. Target customers could be, for example, professional/heavy end-users or large sites with two DSL connections, or end-users with a fixed (or WLAN) and mobile access from the same provider. The ISP can market the bundle with MPTCP features like flexibility, seamless transition from fixed or nomadic to mobile access, and resilience against the failure of one of the access connections. Further, the ISP can offer full resilience by using the routing techniques described in Section 3.1 to achieve disjoint paths beyond the access connections.

What makes access bundling especially interesting is that by offering multiple connections itself the ISP can avoid the increased competition and comparability of Scenario 3 where the end-user has multiple connections from different ISPs. Furthermore, bundling allows sustaining or even increasing the provider lock-in. When mobile end-users are concerned, traffic can be moved away from the congested mobile access links to fixed access links when available.

Table 3. ISP SWOT analysis of "Access connection bundling" business model

Strengths • Ownership of customer access lines • Existing customer relationship will be maintained • Provider lock-in can be increased	**Weaknesses** • Full benefits of MPTCP require offering routing functionalities
Opportunities • Revenue by access provision • Warranty of reliably throughput can be charged • Part of the load on congested mobile access links can be moved to fixed links when available	**Threats** • Possible visibility of bad performance of individual paths may lead to complaints and loss of customers • Customers unlikely to be willing to pay double price for dual access

4.3. Virtual multipath operator (enabled by Scenario 3)

We also see the potential of a new type of service provider, Virtual Multipath Operator (VMPO), entering the ISP market. VMPO will provide multipath access to end-users by bundling and reselling access products of other ISPs. The VMPO is not in the possession of its own IP access or backbone network. However, an ISP may also decide to act as a VMPO through a separate business unit.

Even though a VMPO can be involved in all deployment scenarios, it suits most naturally Scenario 3 where it can remove end-users' burden of multiple contracts while offering all the benefits of full multihoming. The simplified value network is presented below in Figure 5.

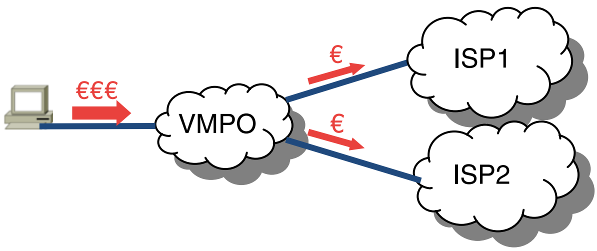

Figure 5. Value network of VMPO business model in the deployment scenario 3.

It is likely that a VMPO could negotiate significantly reduced price deals with the ISPs by purchasing access provisioning in bulk, as well as eliminating the ISP's costly interactions with the end-user. This scenario opens an additional, though indirect, sales channel for the ISP to the MPTCP user market, without having to put additional technical changes to its infrastructure. The main disadvantages for the ISP is that it has no direct customer relationship and it will be under price pressure from the VMPO. The power of the VMPO will depend on the development of the market. Will ISPs address the multipath market themselves by other presented business models? Or will VMPOs be the first innovators and have a unique selling proposition for a long time?

Table 4. ISP SWOT analysis of "VMPO" business model

Strengths • Ownership of at least one customer access line • ISP is in a good position for becoming a VMPO player itself	**Weaknesses** • Loss of direct customer relationship • Limited control over throughput allocation between two independent links • Inability to provide multipath without involvement of a second access ISP
Opportunities • Revenue by wholesale access provision in addition to direct sale	**Threats** • Because of visibility of network performance the VMPO can compare ISPs and terminate the contract with a bad performing ISP • VMPO may put pressure on wholesale access pricing, and could control a large customer base

5. Conclusion

Although MPTCP is a relatively small technical change to the TCP protocol, it will have considerable impact on the business models of Internet access provisioning. The direction of the impact will depend on the way end-users choose to set up their access links. Different deployment scenarios suit different types of end-users and the possible market pull depends on how much each group value the benefits of MPTCP. Consequently, introducing MPTCP is most feasible in the scenarios which offer additional incentives to end-users or where multihoming capability exists already for other reasons.

Even though the ISPs are not required to have any technical involvement to deploy MPTCP, they can support the deployment in all the scenarios, for example, by offering MPTCP as a value added service or introducing multihoming through attractive access connection bundles. Additionally, a virtual multipath operator, a new type of market player, can offer the maximized MPTCP benefits to end-users without the burden of multiple contracts. In the end, the success of these business models depends on the end-users' willingness to pay for the benefits they offer.

References

[1] D. Wischik, M. Handley, and M. Bagnulo Braun, The Resource Pooling Principle, *ACM SIGCOMM CCR*, volume 58, issue 5, (October 2008), 47–52.

[2] A. Ford and C. Raiciu, TCP Extensions for Multipath Operation with Multiple Addresses, *IETF Internet-Draft (work in progress)*, draft-ford-mptcp-multiaddressed, (July 2009).

[3] C. Raiciu, M. Handley, and D. Wischik, Coupled Multipath-Aware Congestion Control, *IETF Internet-Draft (work in progress)*, draft-raiciu-mptcp-congestion-00 (October 2009).

[4] Multipath TCP BOF, http://trac.tools.ietf.org/area/tsv/trac/wiki/MpTcpBofDescription, (June 2009).