

Towards Multipath TCP Adoption: Challenges and Opportunities

Alexandros Kostopoulos, Henna Warma, Tapio Levä, Bernd Heinrich, Alan Ford, and Lars Eggert

Abstract—Successful deployment of new network protocols on the Future Internet is not a trivial task. Deployable protocol design is necessary but not sufficient condition for protocol’s success, unless it takes all stakeholders involved in the deployment process into account. This paper investigates the challenges of deploying a new transport protocol on the Internet, using Multipath TCP – a TCP variant that transmits along several network paths at the same time – as an example and proposes a framework for its adoption process based on diffusion theory. The paper distinguishes the roles of adopters and other stakeholders in the deployment process, and presents scenarios that enhance Multipath TCP deployment and adoption. One key finding is that the role of end users is not of significant importance for Multipath TCP deployment, because they are not necessarily in a position to make a conscious adoption decision.

Index Terms—Future Internet, Multipath TCP, protocol design, deployment, adoption

I. INTRODUCTION

THE Internet has become one of the world’s most remarkable engineering and social phenomena. Despite its success, it is beginning to reach some fundamental capability limits [1]. Future Internet research aims to develop new architectures and protocols that address these known and emerging technical deficiencies in a way that is cognizant of the competing technical, economic and social demands.

The increasing numbers of users, providers and services stress the scalability of current Internet [2]. At the same time,

Manuscript received April 2, 2010. This research was supported by Trilogy (<http://www.trilogy-project.org>), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Programme. The views expressed here are those of the authors only. The European Commission is not liable for any use that may be made of the information in this document.

A. Kostopoulos is with Athens University of Economics and Business, Department of Informatics, Athens, Greece (corresponding author, e-mail: alexkosto@aueb.gr).

H. Warma is with Aalto University School of Science and Technology, Helsinki, Finland (e-mail: henna.warma@tkk.fi).

T. Levä is with Aalto University School of Science and Technology, Helsinki, Finland (e-mail: tapio.leva@tkk.fi).

B. Heinrich is with T-Systems International GmbH, Germany (email: B.Heinrich@t-systems.com).

A. Ford is with Roke Manor Research, United Kingdom (email: alan.ford@roke.co.uk).

L. Eggert is with the Nokia Research Center, Helsinki, Finland (email: lars.eggert@nokia.com).

the user’s performance and resilience requirements are increasing. Hence, focus of ongoing research is to design solutions that deliver effective and efficient control of resource sharing.

At present, all standard Internet transport protocols select only a single path between a source and a destination when transmitting, which limits the achievable throughput. Even SCTP [3], which has standardized mechanisms to recognize and use multiple paths, only shifts a connection from one path to another upon an outage and does not use multiple paths simultaneously (a non-standard, experimental extension [4] does introduce concurrent multipath transfers). The Multipath TCP protocol (MPTCP) [5], which is currently being standardized by the Internet Engineering Task Force (IETF), uses multiple paths at the same time to transmit the data belonging to a single TCP connection. In the case of congestion along a path, or even a complete path failure, MPTCP will make greater use of less congested paths in order to fairly balance network load [6]. This increases reliability, flexibility and throughput.

The successful deployment and adoption of a new transport protocol such as MPTCP depend on several factors. A minimum requirement for a new protocol to be deployable is that it meets a real need and solves an identified problem better than previous and competing approaches. For example, MPTCP could outperform other transport solutions when downloading large files, such as videos or applications. Additionally, the design of the new protocol directly affects its deployability. A classic counter-example is IPv6, whose design negatively affected deployment incentives, due to lack of backwards compatibility with IPv4. Another challenge in protocol design is that firewalls and other middleboxes can reject packets which are not using TCP or UDP. The middlebox problem has affected the deployment of other transport layer protocols, like SCTP and DCCP [7], and inevitably MPTCP has to overcome this challenge as well.

The transport protocol deployment differs from the diffusion of end user centered innovations, like consumer products. The Internet is a complex system with diverse end-systems, not all of whose aspects are under the direct control of the respective end users or service providers. For MPTCP deployment in particular, operating system vendors play a major role, because users cannot directly select network stacks for their end systems. Consequently, although MPTCP requires only end system changes for deployment, other

stakeholders play a major role. End systems may in fact become MPTCP capable without direct user incentives, simply because the user may upgrade his operating system due to other motivations. This is in contrast with existing adoption models discussed in Section II, which assume that customers (known as adopters) make conscious decisions to adopt an innovation. One lesson from this is that the dynamics of the deployment process including the required deployment steps and involved stakeholders is of key importance.

The main goal of this paper is to increase understanding and give a broader picture of the issues surrounding MPTCP deployment, especially related to new challenges for the involved stakeholders. Section II classifies the basic factors that affect adoption of a new technology and provides a brief overview of related work. Section III proposes a framework for analyzing MPTCP deployment and Section IV presents the key factors that make MPTCP deployable. Section V presents the deployment process and the role of the involved stakeholders. Finally, Section VI suggests possible scenarios that facilitate the required steps to support MPTCP adoption.

II. RELATED WORK

This section presents an overview of the theories for studying the adoption of new technologies and identifies how these methods apply for analyzing MPTCP deployment.

The Internet Architecture Board (IAB) has identified the most important factors that enhance or limit the success of a protocol based on several case studies [8]. Although a protocol design will not necessarily be able to incorporate all the proposed success factors, experience indicates that following some of them will improve the probability of success. The most important factors for the initial success are *filling a real need* and being *incrementally deployable*.

Additional design principles and guidelines have been proposed to help researchers and engineers in designing successful protocols. Ford *et al.* [9] present a set of new design principles that help to design flexible and tussle-aware solutions. The “design for tussle” principle proposed by Clark *et al.* [10] suggests that protocol design should accommodate to an environment where multiple stakeholders with varying interests interact. Thus, a new protocol that follows a set of “good” design principles may see benefits, even if that alone is not sufficient to ensure a short-term success. Ahlgren *et al.* [11] propose a complementary methodology, motivated by the view that evolution and interworking flexibility are determined not so much by the principles applied during initial design, but by the choice of fundamental components or “design invariants” in terms of which the design is expressed.

Apart from the abovementioned factors, classical diffusion theory has increased our understanding of how innovations (e.g., a new protocol) spread within populations. Rogers [12] diffusion of innovations (DOI) theory breaks the adoption process down into five stages. In the *awareness* stage, the individual is exposed to the innovation, but lacks complete information about it. In the *interest or information* stage the

individual becomes interested in the new idea and seeks additional information about it. The next stage is *evaluation*, where the individual mentally applies the innovation to his present and anticipated future situation, and then decides whether or not to try it. In the *trial* stage the individual employs the innovation. Finally, in the *adoption* stage the individual decides to continue the full use of the innovation.

Rogers also presents five characteristics of an innovation. The *relative advantage* is the degree to which the new technology is better than a preceding one. *Compatibility* is the consistency with existing values, past experiences and needs. *Complexity* is the difficulty of understanding and use. A new technology is more likely to be adopted if it is compatible with existing practices of adopters, and is relatively easy to understand and use. *Trialability* is the degree to which it can be experimented with on a limited basis. Finally, *observability* is the visibility of its results. Even though this paper does not argue that adoption of MPTCP will follow Rogers DOI model, the listed characteristics are still helpful for our analysis.

The diffusion phenomenon has also been studied from a community point of view, focused on the economic value an innovation brings to potential adopters. This economic value to an adopter depends on the size of the existing network of adopters and the potential network of adopters. Katz and Shapiro [13] analyze the adoption of a new technology for cases, where network externalities are significant. Adoption becomes more likely when the number of current adopters in the network increases.

The theories presented above have been initially used for studying the adoption of consumer products. However, the adoption of new Internet protocols is more complex than that of consumer products, and therefore requires more elaborated modeling. Several attempts have been made at studying the adoption of new Internet protocols.

In particular, Hovav *et al.* present a model of Internet standards adoption [14] that identifies additional concepts that influence adoption of a new technology. Development of a related technology infrastructure, economies of scale and amount of information available could also help a new protocol to spread. Moreover, the presence of sponsorship could decrease the risk of adoption.

In [15], an economic model based on user utility is used to study the adoption of new network architectures. The model incorporates various factors, such as user and network benefits, and switching costs, and discusses the impact of converters on the adoption of new network architectures. Key findings include that new network architectures need to withstand a period of decreasing total system utility till a critical mass of users is reached.

III. FRAMEWORK FOR MPTCP ADOPTION PROCESS

This section proposes a new framework for a successful adoption process of MPTCP. It is important to make a distinction between the concepts of *deployment* and *adoption*. By deployment, we are referring to MPTCP being deployed in

the required network equipment. Adoption is dependent upon deployment, with the additional step that end users are actually sending traffic using the protocol.

Multiple vendors control the operating systems of different end systems used in the Internet. Those OS vendors have the power to decide whether they want to implement MPTCP in their OS or not. When an OS vendor decides not to implement MPTCP, there is almost nothing that an end user or service provider can do that will let him use MPTCP, short of switching to a different operating system. This consideration is what complicates adoption models for MPTCP and other Internet protocols.

Our framework consists of three different steps. As in Rogers' model, we assume that the new solution already exists. In our case, MPTCP is a solution for increased demand of bandwidth and resilience. Fig. 1 illustrates the steps for MPTCP adoption.

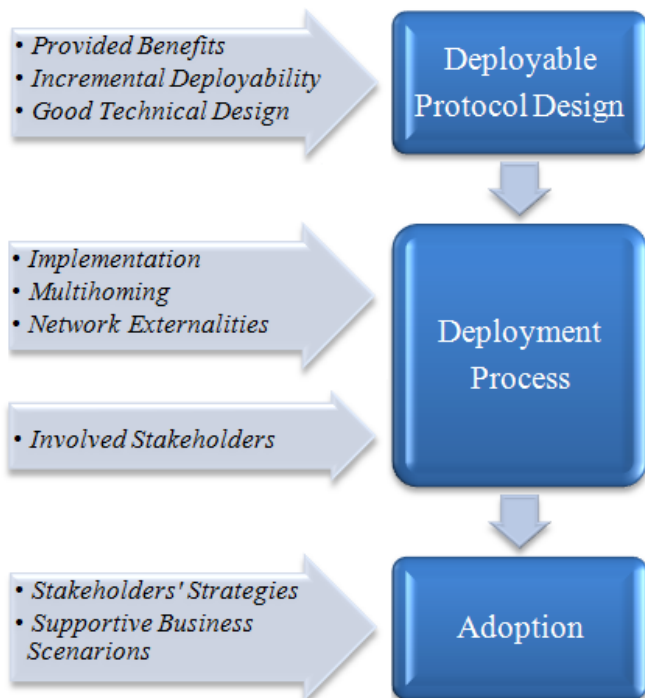


Fig. 1. Framework for MPTCP Adoption Process

The first step towards adoption is that the protocol itself is sufficiently deployable. The minimum requirement is that MPTCP provides some advantage compared to the previous solution (e.g., TCP), but this alone may not be sufficient condition for adoption. The Internet is a complex system and the deployment of a protocol usually requires the involvement of several stakeholders. This is why the design of MPTCP should be feasible for each relevant stakeholder.

If the protocol itself is deployable, the actual deployment process can be considered. This step identifies the actions required from different stakeholders to allow the protocol being deployed. Three requirements need to be satisfied before MPTCP can be fully used. First, MPTCP has to be implemented by the operating system vendor and installed to end systems. Second, at least one of the end systems must have

multi-homing capability, i.e., both the operating system needs to allow the usage of multiple Internet connections at the same time, and the end user has to obtain simultaneous Internet access across different network interfaces. This can, for example, require several service contracts with multiple ISPs. Third, *other* end systems in the Internet need to become MPTCP-capable, so that MPTCP can be actually used end-to-end.

Deployment of the protocol does not necessarily mean that it will be actually adopted. Thus, a further investigation is needed, concerning the key stakeholders and potential business models [16] that could boost MPTCP adoption.

IV. DEPLOYABLE PROTOCOL DESIGN

This section focuses on protocol design itself. In particular, it investigates whether MPTCP is an incrementally deployable protocol that provides new advantages (compared to older or competitive technical proposals) and whether it is designed in a manner that follows “good” design principles, based on the design principles outlined in Section II.

A. Provided Benefits

One necessary condition for successful protocol deployment is that a real need is met and that a new protocol solves an identified problem better than previous or competing approaches. MPTCP provides several such benefits to users.

An MPTCP connection uses several paths for a single connection at the same time. This results in several benefits. First, in case of severe congestion or a failure along one path, MPTCP can make greater use of less congested alternate paths. Thus, MPTCP continues to provide a useful – albeit potentially somewhat reduced – service, whereas traditional TCP often fails to adequately support the user’s transport needs in these cases. Second, because MPTCP pools the available capacity along all paths for a single connection, it can support faster transfers than traditional TCP.

Furthermore, MPTCP uses coupled congestion control [6] for controlling the sending rates it uses along different paths. Key and Massoulié [17]-[19] investigate the benefits of such coordinated congestion control schemes and show that when a user opens multiple independent TCP connections with uncoordinated congestion control, the total throughput is not maximized. With coordinated congestion control, however, the total throughput is maximized. In order to investigate in more depth the benefits of coordinated congestion control provided by MPTCP, a simplified approach of the optimization framework presented in [19] is given in the appendix.

It should be noted that coordinated congestion control is also useful for other reasons; namely, for fairness between TCP and MPTCP users. Without coordinated congestion control, an MPTCP connection across a network that causes several paths to have a joint bottleneck can use an unfair share of the bottleneck capacity. Several proposals address this issue. The coordinated controllers proposed in [20] and [21] use different approaches to limit the resource use of an MPTCP connection on a joint bottleneck to approximate that

of a single traditional TCP connection. However, a similar problem appears in peer-to-peer applications, where peer-to-peer users open multiple independent TCP connections, hence being more aggressive than MPTCP users.

For mobile battery-powered devices, the benefits of MPTCP lie elsewhere. Sending and receiving data across multiple radio interfaces increases the energy consumption of network communication, even when the aggregate data rate is not higher than that of a normal, unipath TCP connection. This mode of operation is hence useful only in cases where capacity pooling is required to satisfy bandwidth demand. However, an extremely interesting feature of MPTCP for mobile use is the ability to switch an established connection between different paths. This allows a mobile end system to aggressively switch an MPTCP connection to the most energy-efficient path based on its current data rate.

B. Incremental Deployability

According to [22] the deployment of a new technology is encouraged when related technologies already exist. On the other hand, a well-established infrastructure with a large installed base burdens the deployment of a new architecture, due to inertia and sunk costs. In the case of MPTCP, this is not necessarily a concern. MPTCP is not a completely standalone new protocol; it is a backward compatible extension of standard TCP. MPTCP offers an unmodified *sockets* API to applications, which means that applications do not need to be modified or even recompiled to run over MPTCP.

To the network, each flow of an MPTCP connection looks like a single standard TCP connection that uses some new TCP option headers. An MPTCP connection starts as a normal TCP connection with an additional option indicating that the sending host is MPTCP-capable. If one of the endpoints does not support MPTCP, the connection remains a standard TCP connection. The connection is upgraded to a MPTCP connection only when both end systems are MPTCP-capable. As shown in Fig. 2, during the lifetime of the connection, new flows can be added to it or removed from it as needed.

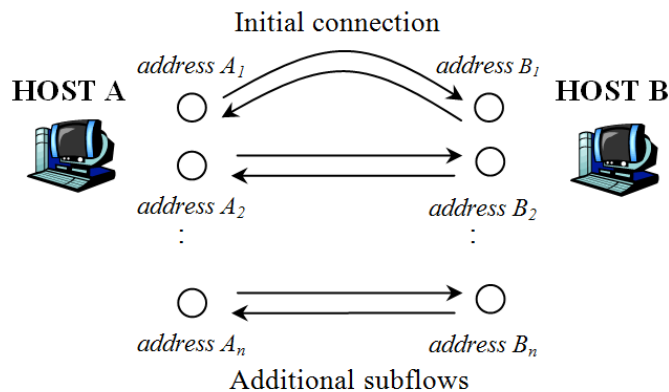


Fig. 2. A Multipath TCP connection establishment

Because the network traffic generated by an MPTCP connection looks like a bundle of regular TCP connections carrying new options, it should also operate correctly through

most of the existing middleboxes¹ (i.e., NATs, proxies, firewalls) and work well with logging and other operational procedures.

These arguments support the claim that the MPTCP design is cognizant of TCP, which increases the chances of its deployment [23].

C. Good Technical Design

Designing a protocol that follows “good” principles outlined in Section II could enhance deployment and interoperability. Even if a good technical design is not the most important success factor, tussle-aware protocol designs have better chances at deployment in the long-term [24].

Ford *et al.* [9] present a new set of design principles for Future Internet architectures with a particular focus on enabling socio-economic tussles between stakeholders. This section investigates how these principles can be applied when designing Internet protocols and, specifically MPTCP.

Resource Pooling

The “resource pooling” principle [25] suggests that when resources in a network can be pooled, effectiveness and efficiency of the network will be improved. The concept of pooling describes a system that makes a set of resources appear as a single resource of aggregate capacity. Multipath TCP is a resource pooling mechanism that sends data along multiple paths and uses a coordinated congestion control algorithm that is designed to facilitate resource pooling [6]. This approach allows the traffic load to be relocated to or spread over several paths.

Information Exposure

This principle suggests that sufficient information about resource usage should be exposed to support an effective and efficient allocation of that resource.

MPTCP monitors the congestion signals on each individual subflow, in order to respond appropriately to resource usage and congestion by shifting load between the subflows. Thus, it facilitates building systems that have a higher degree of information sharing than those built on standard TCP. Additionally, the *Trilogy* architecture [26], which MPTCP is one component of, has an explicit information exposure mechanism similar to Re-Feedback [27].

Separation of Policy from Mechanism

This principle recommends allowing local choices of a network entity according to its priorities (policy), which is separate from the standardized implementation (mechanism).

The separation of policy and mechanism is integral to MPTCP, as it is possible to specify the protocol without having to specify how the end systems decide which paths to use for a given connection and traffic volume. This allows MPTCP to be deployed in various situations depending on the needs of an end host and the path characteristics of the paths available to it. This policy is entirely separate from the

¹ Note that some middleboxes may strip new TCP option headers.

mechanism (the protocol standardized across all parties).

Although transmission policies are typically determined globally by the operating system on an end host, possibly based around a simple user preference (e.g., maximize throughput or minimize monetary cost expressed by Explicit Congestion Notification (ECN) marks [27]), an application could express its own transmission preferences for traffic to varying degrees of granularity.

Fuzzy Ends

This principle suggests that the end points should be allowed to explicitly delegate some functions to the network. Although MPTCP is designed for use by end hosts in an end-to-end way, the proposed architecture is sufficiently extensible to allow the development of MPTCP proxies [16]. Such proxies could be placed within the network in order to provide multipath benefits without the need of endpoints to be multi-homed themselves.

V. DEPLOYMENT PROCESS

Deployment of MPTCP involves multiple stakeholders who need to take actions during the deployment process. This paper takes a pragmatic view and identifies the required steps for an end user to be able to use MPTCP, as well as the role of different stakeholders in taking these steps.

The fundamental requirements for MPTCP deployment are:

1. An MPTCP implementation for operating systems is available,
2. An MPTCP-capable OS is installed on an end system (i.e., device is MPTCP-capable), and
3. The end user is multi-homed (i.e., it connects to the Internet via multiple paths at the same time).

These three requirements are not enough to enable the use of MPTCP if desired communication peer is not MPTCP-capable (multihoming not a necessary condition; partial benefit can be derived from partially disjoint paths). Consequently, the fourth requirement concerning the impact of network externalities to the deployment decision is:

4. Other MPTCP-capable end systems of interest exist (i.e., systems the user has an interest connecting to).

The key stakeholders in the deployment process are:

- Operating system vendors that implement MPTCP in operating systems for use on end systems.
- End users (i.e., individual users, but also service providers, content distribution networks, etc.) that own end systems.
- Internet service providers (ISPs) that provide connectivity for multi-homing.

The role and the motives of each key stakeholder are discussed in the following sections.

A. Availability of OS Implementation

MPTCP requires changes only to the TCP/IP stack of end systems, which in practice means that an OS update that adds support for MPTCP needs to be available. The availability of

such an update fully depends on OS vendors which are the key stakeholder in this stage. These vendors can, in addition to altruistic reasons, have four specific motives for implementing MPTCP in their operating systems:

1. **Pressure from end users:** This is only relevant if end users are aware of MPTCP, or if they demand solution to a problem that MPTCP alleviates. At the time of OS implementation, large corporations and content providers are the most probable end user groups to have sufficient influence to drive this implementation reason.
2. **Pressure from application developers:** Application developers may influence on OS vendors if they see that their products would be enhanced by MPTCP support. At the time of OS implementation, it is likely that only the developers of highly important applications have sufficient influence to drive this implementation reason.
3. **Own business interest:** If OS vendors also take the end user or application developer role, they may get direct business benefits from implementing MPTCP. For example, Microsoft could be interested in using MPTCP in their Windows update service, or Nokia could deploy MPTCP in their platform for Internet services called Ovi.
4. **Competitive environment:** If MPTCP is implemented in other operating systems, an OS vendor has a higher incentive to implement it as well. The “leader role” of open source operating systems can be significant in incentivizing commercial OS vendors to implement MPTCP.

Most end users will use their devices and operating systems “out-of-the-box”. Therefore, concerning actual usage of MPTCP, availability of OS implementation will not be enough if MPTCP is not enabled by default in the shipping configuration. However, getting OS vendors to do this should, not be a problem, since MPTCP is backward compatible with traditional TCP.

B. Installation of MPTCP-Capable OS to End Systems

An available operating system update with an MPTCP implementation needs to be installed to end systems. Consequently, the end users are the key stakeholders at this stage, because they have the ultimate control over their devices. End users can be divided into those who make a conscious decision to deploy MPTCP and those who get the MPTCP unbeknownst to them.

Conscious end users, e.g., content providers interested in increasing the perceived quality of their services, or private users with large traffic volumes, will be willing to install an OS update with MPTCP support for the sole reason of said support. For those, the five stages of Rogers’ diffusion process is applicable concerning the active adoption of an MPTCP.

However, large part of end users, especially most consumers, is not interested or even aware of particular Internet protocols and thus will not make an active decision to install MPTCP. Thus, the typical adoption models that assume conscious end user decisions are not applicable to this adoption case, and the role of OS vendors increases. For these

unaware end users a new OS feature, such as MPTCP, can be offered either when they purchase a new device or a new version of the OS, or through automatic operating system updates. In both cases, end users are not making a conscious MPTCP deployment decision, but are simply updating their devices and operating systems for other reasons. MPTCP deployment may be slowed down significantly due to these end users.

Finally, if MPTCP is not rolled out in operating system updates, ISPs may foster MPTCP use by providing an MPTCP proxy service that intercepts standard TCP traffic generated by end systems and translates it to MPTCP, as presented in [16]. The availability of this deployment results from the MPTCP design following the fuzzy ends design principle. However, this approach does not provide full resilience on the access link, and it depends on MPTCP support, or at least a similar proxy setup, on the service provider side.

C. Multi-homing

Fulfilling the multi-homing prerequisite is primarily in the hands of end users, because they need to make conscious decision to acquire additional access connections to the Internet. End users are hence key stakeholders at this stage. Because acquiring additional Internet access connections normally will involve subscriptions, ISPs also play a major role. Although an end user's interest for multi-homing can directly come from the desire to run MPTCP, most probably other incentives for multi-homing play a role, e.g., the need for ubiquitous access for a mobile user, or the desire to have backup connections for content providers.

Many end users may already have multi-homing capability available. For example, large enterprise or academic campus networks are often multi-homed, and users at such sites will probably have an interest in MPTCP for its increased throughput and resilience with no required hardware updates. In addition, consumer devices such as mobile phones already support both WLAN and 3G technologies. 3G is a wide-area access technology that can be used everywhere, so the users of such devices only need one additional WLAN connection to benefit from MPTCP. Such users are often at locations where such WLAN connectivity is available to them, such as their homes or workplaces, or wherever "free WLAN" is offered.

ISPs cannot prevent end user from multi-homing, but they can improve support for it in multiple ways, as identified in [16] and summarized in Section VI. Their motives for this support are mainly monetary, because multi-homing offers them a possibility to sell more access connections. However, MPTCP can also help ISPs to balance the load in their networks, although this may necessitate ISPs adjusting their traffic engineering strategies in light of MPTCP's adaptation to path failures at the transport layer (instead of relying on re-routing at lower layers). Especially interesting is the possibility that the traffic could move away from congested mobile access links to fixed links through WLAN hotspots in a transparent way. Nevertheless, ISPs would probably prefer a solution that would give them more control over this off-loading, because

off-loaded traffic does not generate revenue for them.

D. Other End-points and Network Externalities

Network externalities are the positive or negative effects on a user using a product or service when others are using the same or compatible products or services. In case of MPTCP positive network externalities are especially important because if one of the endpoints is not MPTCP capable, resource pooling cannot be exploited. Thus, it is obvious that if a large number of users adopt MPTCP, the probability of a successful MPTCP connection establishment is increased (i.e., the well-known "network effect").

Since network connections can be established either between two clients or between a client and a server, we notice that there are two different types of network externalities: network externalities between *two clients*, and network externalities between *a client and a server*. The level of MPTCP network externalities increase with different speeds in these two different cases. MPTCP availability on the client-side depends on the actions of each individual user of a certain client device. Servers, on the other hand, are clustered in the network and one provider has the control over the updates of all servers in the domain. If Google, for example, would make the decision to deploy MPTCP in its servers, it would probably update a significant fraction of its servers at the same time. This tends to increase the network externalities in jumps.

However, a specific end user in reality only cares, that those peers he mostly connects to, are MPTCP-capable. For example, if he often accesses a specific service, it is important to him that the particular service is MPTCP-capable. Whether other end systems or services are MPTCP-capable typically matters very little to a specific end user.

Moreover, MPTCP internalizes the negative network externalities of all Internet users. Network resources are efficiently allocated, hence congestion is reduced. In the cases where a MPTCP connection shares a bottleneck with a single-path TCP connection, the congestion control algorithm will ensure that MPTCP acts fairly on other users and does not take more bandwidth than a legacy, single-path TCP.

VI. POTENTIAL SCENARIOS SUPPORTING ADOPTION

This section investigates potential scenarios that could accelerate the MPTCP deployment and adoption.

A. Both Ends in one Hand

In this scenario multi-homed devices and content or application servers are under the control of one stakeholder, which could enhance the deployment and adoption of MPTCP. For example, companies that provide a mobile device for their employees to use company applications remotely over WLAN or 3G could significantly benefit from MPTCP. Another scenario would be an end user accessing content using WLAN and 3G from a provider which controls both end user devices and content servers, such as Nokia or Apple both delivering devices and services/content (e.g., Nokia Ovi and Apple App Store). This type of scenario could significantly accelerate the

deployment and adoption of MPTCP.

Consumers will probably be MPTCP-unaware, but may become opportunistic adopters of MPTCP when it is implemented by device manufacturers. The deployment in the client devices (OS vendor's enabling MPTCP by default) is the key driver that leads to the adoption on the client-side if the end user is multi-homed already.

B. Lobbying

An important factor for the deployment of MPTCP will be lobbying towards OS vendors who have to implement the new protocol in their network stacks. Key players like Microsoft who by some count holds a market share of around 85% of personal computer operating systems worldwide [28], need to be convinced of the merits of MPTCP. Other OS vendors will probably follow each other after the first ones decide to implement MPTCP. Especially organizations that represent end users with a vital interest for MPTCP deployment need to take on the lobbying initiative.

C. Killer Applications

The utilization of MPTCP in widely distributed applications can also be considered as a scenario for accelerating adoption. The measurements of Labovitz [29] reported an increase in global IPv6 traffic when *BitTorrent* application *uTorrent* took IPv6 into use in their 1.8 release. If the MPTCP API would be available and widely distributed applications would decide to implement enhanced multipath support, this could have a similar effect².

D. End user decision

End users that transfer lots of data and operators of large content sites will have a direct interest in the increased resilience and throughput provided by MPTCP. Once the protocol has been made available by OS vendors, they may take a conscious decision for adopting MPTCP. The decision will depend on the involved cost for OS upgrade installation and potentially additional physical access lines for multi-homing if not already in place for fault tolerance or load balancing. The adoption may also depend on the availability of MPTCP enabled clients or peers.

E. ISP Support

A considerable barrier to MPTCP adoption by consumer end users will be the requirement of multi-homing which may cause additional costs for additional connections. To overcome this problem, [16] proposes that ISPs offer access bundles (e.g., for DSL plus 3G access), at a price that is cheaper than offering the two individually. The incentive for the ISP would be customer retention, lock-in, and potentially improved traffic engineering control. In addition to the usual MPTCP benefits, the end user would enjoy the ability to seamlessly roam between fixed and mobile access networks.

Such solution is proposed in [16], where a Virtual Multipath

Operators (VMPO) could offer such bundles by buying access lines, potentially of different kinds, from other ISPs. The increase in competition due to such VMPOs could be a driving factor for ISPs to offer their own price-reduced bundles.

Another way to accelerate the availability of a cost-effective MPTCP solution for end users is through ISPs that offer MPTCP-enabled access as a value-added service by providing a MPTCP proxy service to end users. This would not require a second access link, nor an MPTCP-enabled OS. Limited to one access connection, the solution will not realize the full potential benefits of MPTCP.

VII. CONCLUSION

This paper presented a framework to investigate MPTCP deployment and the related new challenges faced by the involved stakeholders. It concludes that the deployment and adoption of transport protocols differ from the diffusion of end user-centered innovations, such as consumer-products. The performance, reliability and flexibility improvements that MPTCP brings are beneficial, but they are unlikely to be the main drivers for adoption. Therefore, the role of end users in MPTCP deployment process is not of primary importance, because they are not necessarily in the position to make a conscious adoption decision.

The deployment of MPTCP will hence be mainly in the hands of software vendors, especially operating system vendors, which need to make the deployment decision of enabling MPTCP by default. The deployment of MPTCP-enabled OS will take different channels: roll out on new devices delivered with new operating systems, automatic software updates to the deployed base (often without awareness of the end user), and intentional installation by operators of large sites (e.g., content providers).

The paper presented the generic benefits of resource pooling being the motivation to implement MPTCP but it did not analyze what are the applications that would benefit from it. Thus, further research with applications is needed after a working prototype of the protocol is available. Also comparison to the MPTCP-like solutions in other layers is needed to understand if transport layer is the proper layer for implementing the properties of MPTCP.

APPENDIX

We present a simplified approach of the optimization framework presented in [19] to understand the benefits of coordinated congestion control provided by MPTCP.

[19] considers a triangle network topology (A-B-C), where there are two types of flows between any two pairs of nodes: along a direct path (i.e. A-B), and an indirect path (A-C-B). For simplicity, we assume that each node represents one user.

There are also capacity constraints for each link which are:

$$x_A + y_B + y_C \leq C$$

$$x_B + y_A + y_C \leq C$$

$$x_C + y_A + y_B \leq C$$

² Note, however, that in the case of *BitTorrent*, its peer-to-peer nature already results in some similar benefits to those delivered by MPTCP.

where

x_i : the throughput of user i in the direct link

y_i : the throughput of user i in each indirect link

C : the capacity of each link

The first scenario assumes that each user opens two independent TCP connections; hence the congestion control is uncoordinated. The optimization problem for the network is:

$$\max [U(x_A) + U(y_A) + U(x_B) + U(y_B) + U(x_C) + U(y_C)] \quad (1)$$

For simplicity, we do not take into account the round trip time, as in [17], and we assume that users' utility function is given by the formula $U(r) = -1/r$, where r is user's throughput in a specific link³. We also do not make any assumptions concerning the fairness of the allocation to each connection, as in [19].

Solving the previous maximization problem⁴, we will see that each user sends the most but not all of his traffic through the direct path and less traffic through the indirect path:

$$x = \frac{C}{1 + \sqrt{2}}, \text{ and } y = \frac{C}{2 + \sqrt{2}} \quad (2)$$

We observe that the total throughput is not maximized with uncoordinated congestion control, because less than the available capacity in each link is used:

$$x + y = \frac{3 + 2\sqrt{2}}{4 + 3\sqrt{2}} \cdot C < C \quad (3)$$

The second scenario assumes that each user opens one MPTCP connection. Thus, the congestion control is coordinated. The optimization problem for the network is:

$$\max [U(x_A + y_A) + U(x_B + y_B) + U(x_C + y_C)] \quad (4)$$

Solving the previous maximization problem we will see that each user sends all of his traffic through the direct path and no traffic through the indirect path:

$$x = C, \text{ and } y = 0, \text{ thus } x + y = C \quad (5)$$

We observe that the total throughput is maximized with coordinated congestion control, because all the available capacity in each link is used (C). Each user gets more than in the previous case. Thus, MPTCP provides more throughput than classic TCP does.

ACKNOWLEDGMENT

The authors would like to thank Costas Courcoubetis and Antonis Dimakis for their initial contribution and comments.

REFERENCES

- [1] M. Handley, "Why the Internet Only Just Works", in *BT Technology Journal* 24, 2006.
- [2] C. Kalogiros, M. Bagnulo, and A. Kostopoulos, "Understanding Incentives for Prefix Aggregation in BGP", *ACM CoNEXT ReArch '09*, Rome, Italy, December 2009.

- [3] R. Stewart, "Stream Control Transmission Protocol", RFC 4960, September 2007.
- [4] J. Iyengar, P. Amer, and R. Stewart, "Concurrent Multipath Transfer using SCTP Multihoming over Independent End-to-End Paths", *IEEE/ACM Transactions Networking* 14(5), 2006.
- [5] A. Ford, and C. Raiciu, "TCP Extensions for Multipath Operation with Multiple Addresses", IETF Internet-Draft, draft-ford-mptcp-multiaddressed-03, 2010.
- [6] C. Raiciu, M. Handley, and D. Wischik, Coupled Multipath-Aware Congestion Control, IETF Internet-Draft, draft-raiciu-mptcp-congestion-01, 2010.
- [7] D. A. Hayes, J. But and G. Armitage, "Issues with Network Address Translation for SCTP", *ACM SIGCOMM CCR*, pp. 23–33, 2008.
- [8] D. Thaler, and B. Aboba, "What Makes for a Successful Protocol?", RFC 5218, 2008.
- [9] A. Ford, P. Eardley, B. van Schewick, "New Design Principles for the Internet", Proc. *International Workshop on the Network of the Future*, Dresden, June 2009, pp. 1–5.
- [10] D. D. Clark, J. Wroclawski, K. R. Sollins, R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", *IEEE/ACM Transactions on Networking* 13(3), 462–475, 2005.
- [11] B. Ahlgren, M. Brunner, L. Eggert, R. Hancock, and S. Schmid, "Invariants – A New Design Methodology for Network Architectures", Proc. *ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA 2004)*, Portland, OR, USA, pp. 65–70, 2004.
- [12] E. M. Rogers, "Diffusion of Innovations", Free Press, New York, 1983.
- [13] M. L. Katz, and C. Shapiro, "Technology Adoption in the Presence of Network Externalities", *Journal of Political Economics*, 94, pp. 822–841, 1986.
- [14] A. Hovav, R. Patnayakuni, and D. Schuff, "A model of Internet Standards Adoption: the Case of IPv6", *Information Systems Journal*, vol. 14, issue 3, pp. 265–294, June 2004.
- [15] D. Joseph, N. Shetty, J. Chuang, and I. Stoica, "Modeling the Adoption of New Network Architectures", International Conference on Emerging Networking Experiments and Technologies, December 2007.
- [16] T. Levä, H. Warma, A. Ford, A. Kostopoulos, B. Heinrich, R. Widera, and P. Eardley, "Business Aspects of Multipath TCP Adoption", in *Future Internet Assembly (FIA) Book 2010*, Valencia, Spain 2010.
- [17] P. Key, L. Massoulié, and D. Towsley, "Multipath Routing, Congestion Control and Load Balancing", Proc. *IEEE ICASSP*, Hawaii, USA, 2007.
- [18] P. Key, L. Massoulié, and D. Towsley, "Path Selection and Multipath Congestion Control", Proc. *IEEE INFOCOM* Anchorage, AL, USA, pp. 143–151, 2007.
- [19] L. Massoulié, and P. Key, "Schedulable Regions and Equilibrium Cost for Multipath Flow Control: the Benefits of Coordination", Proc. *IEEE CISS*, Princeton, NJ, USA, 2006.
- [20] M. Honda, E. Balandina, and L. Eggert, "Designing a Resource Pooling Transport Protocol", Proc. *IEEE Global Internet Symposium*, Rio de Janeiro, Brazil, 2009.
- [21] D. Wischik, M. Handley, and C. Raiciu, "Control of Multipath TCP and Optimization of Multipath Routing in the Internet, *NetCOOP 2009*.
- [22] Farrell, Joseph, and G. Saloner, "Installed Base and Compatibility: Innovation, Product Preannouncement, and Predation", *American Economic Review*, Vol. 76, pp. 940-955, 1986.
- [23] C. Shapiro, and H. Varian, "Information Rules: a Strategic Guide to the Network Economy", Harvard Business School Press.
- [24] C. Kalogiros, A. Kostopoulos, and A. Ford, "On Designing for Tussle: Future Internet in Retrospect". Proc. *EUNICE 2009*, LNCS 5733, pp. 98–107, Barcelona, Spain, September 2009.
- [25] D. Wischik, M. Handley, M. Bagnulo Braun, "The Resource Pooling Principle", *ACM SIGCOMM CCR* 38(5), pp. 47–52, 2008.
- [26] P. Eardley, L. Eggert, M. Ford, R. Karrer, P. Sarolahti, R. Widera, and R. Winter, "Towards a Unified Internet Control Architecture", Proc. *3rd International Conference on Future Internet Technologies (CFI08)*, Seoul, Korea, 2008.
- [27] B. Briscoe, A. Jacquet, C. Di Cairano-Gilfedder, A. Salvatori, A. Soppera, and M. Koyabe, "Policing Congestion Response in an Internetwork using Re-Feedback". Proc. *ACM SIGCOMM*, 2005.
- [28] Globalstats, <http://www.w3counter.com/globalstats.php>; February 2010.
- [29] C. Labowitz, "Internet Traffic and consolidation", a presentation at IETF 77, available: <http://www.ietf.org/proceedings/10mar/slides/plenaryt-4.pdf>, March 2010.

³ This utility function approximately models TCP Reno's rate control.

⁴ For simplicity we examine the symmetric case, where it is assumed that each node sends the same traffic: $x_A = x_B = x_C = x$, and $y_A = y_B = y_C = y$.