

Dynamics of Communication Protocol Diffusion: the Case of Multipath TCP

Henna Warma

Aalto University, Department of Communications and Networking

P.O. Box 13000, 00076 AALTO, Finland

+358 50 3412284

henna.warma@aalto.fi

Tapio Levä

Aalto University, Department of Communications and Networking

P.O. Box 13000, 00076 AALTO, Finland

Howard Tripp

Roke Manor Research Ltd

Salisbury Lane, Romsey, SO51 0ZN, United Kingdom

Alan Ford

Roke Manor Research Ltd

Salisbury Lane, Romsey, SO51 0ZN, United Kingdom

Alexandros Kostopoulos

Athens University of Economics and Business, Department of Informatics

76 Patission Street, Athens 10434, Greece

Accepted manuscript. Please cite this article as:

Henna Warma, Tapio Levä, Howard Tripp, Alan Ford, Alexandros Kostopoulos (2011). Dynamics of Communication Protocol Diffusion: the Case of Multipath TCP. Netnomics, vol. 12, no. 2, pp. 133-159. DOI: 10.1007/s11066-011-9063-8..

ABSTRACT: During the last decade the Internet has faced an architectural stagnation due to lack of wide scale adoption of new communication protocols. A significant reason for non-adoption is that the conflicting interests of networked stakeholders involved in the diffusion process are not understood or taken into account during the protocol development. This paper increases understanding of the dynamics of communication protocol diffusion and provides feedback to protocol development by studying the case of Multipath TCP (MPTCP). Firstly, we introduce a protocol development process which builds on the existing diffusion of innovation theories. Secondly, a quantitative analysis using system dynamics is provided to evaluate the criticality of the factors affecting the MPTCP diffusion. The diffusion of communication protocols is found to follow three adoption models differentiated by the basis of adoption decision. The key finding is that unintentional adoption, alongside device acquisitions or operating system updates, adds a new dimension to the diffusion of innovations theory

and may have a significant impact on protocol diffusion. The cross-side network effects between different adopter groups play also an important role and may lead to either market pull or technology push type of diffusion depending on which stakeholder starts to adopt first. Although MPTCP is used as an example protocol, the findings can be also applied to the diffusion studies of other communication protocols.

Keywords: Protocol diffusion, adoption, multipath TCP, system dynamics

1. Introduction

The proliferation of Internet usage puts a lot of pressure on scalability, especially when new access technologies increase traffic volumes in the Internet backbone. The answer to the problem of ever-increasing demand should not only be to increase bandwidth capacity, but to also find more efficient ways of utilizing network resources. Although solutions for improving Internet scalability exist, they are not widely adopted. Indeed, the core protocols of the Internet have not changed much for more than a decade [1]. The diffusion of new communication protocols within network devices seems to be very challenging because it requires actions of multiple stakeholders with varying interests; consider Internet service providers (ISPs) and content providers (CPs).

Diffusion of innovations and its role in the innovation development process are conceptualized by Rogers [2]. Communication protocols have, however, some special characteristics, which has incentivized us to re-evaluate Rogers' generalized theories. A communication protocol is mostly seen as a product feature, (where the product is, for example, a network device or an operating system of a personal computer) which is useless without the product itself and which can diffuse to end users differently depending on the product that the protocol is integrated into. By understanding the dynamics of protocol diffusion, strategies can be developed to accelerate the process: relevant stakeholders can be incentivized and feedback to protocol design can be provided, for instance.

This paper uses the case of Multipath TCP (MPTCP) to study not only the diffusion of MPTCP itself, but to also gain understanding of protocol diffusion in general. MPTCP is a backwards compatible extension to regular TCP which aims for better utilization of resources in the network by exploiting multiple paths between hosts [3]. This is an example of resource pooling, enabling traffic to be shifted from congested links to under-utilized ones. By distributing network traffic throughout the network, increases in overall throughput and resilience can be achieved, without the need to increase capacity [4]. Like many other communication protocols, the benefits of MPTCP are highly dependent on the scale of its

diffusion. Diffusion of MPTCP within proprietary networks (*e.g.*, data centers) could result in operational cost savings [5] but does not improve the overall Internet architecture. Only as the number of capable nodes in the public Internet increases, the benefits of a flexible and scalable Internet architecture will emerge.

MPTCP has been studied from many perspectives including, for example, analysis of protocol design, implementation requirements and identification of potential business scenarios of different stakeholders [6, 7, 8]. None of the studies, however, has conceptualized the interdependencies of stakeholder actions, nor evaluated their criticality quantitatively. This study concentrates on analyzing how a protocol can diffuse in an ecosystem consisting of various stakeholders and shows the interactions between the selected market stakeholders in the protocol diffusion process. *System dynamics (SD)* [9] is utilized in finding the most important factors in the protocol diffusion process and in evaluating their criticality. The focus of the study is to analyze how *cross-side network effects* and stakeholder actions affect the value of the protocol rather than explicitly comparing the costs against the benefits which the protocol incurs to a specific market player.

System dynamics is a methodology for considering the underlying interactions within complex systems, and hence is potentially applicable to diffusion studies of Internet-related innovations. For example, a study carried out by Kelic [10] develops a SD model of the adoption of fiber-to-the-home to aid policymakers in decision-making. Thun et al. [11] apply system dynamics to understand the diffusion of information goods showing network effects, such as electronic mail, which has some analogy to the diffusion of end-to-end networking technologies. However, there has not been extensive research on applying system dynamics to the diffusion of communication protocols.

The paper is structured as follows: Section 2 discusses the related work and introduces a protocol development process. Section 3 identifies and describes the different adoption models for communication protocols. Section 4 introduces the diffusion model of MPTCP using system dynamics, and Section 5 presents the quantitative results of the model. Section 6 discusses the simulation results and model limitations, whereas Section 7 finally concludes the study with generalized findings.

2. Protocol development process

Communication protocols are technological innovations which follow Rogers' general innovation development process. The process consists of all the decisions, activities, and their impacts that occur in the six stages during which a recognized need is transformed through research, development and commercialization into an innovation that diffuses to potential adopters and causes consequences to an individual or a social system [2]. The focus is on a single innovation (*i.e.*, a protocol) and how it advances from an idea to actual usage in an ecosystem of different kinds of organizations. In addition to the Rogers' product-centric perspective, the field of innovation management studies innovation development from a perspective of, and as a process within a single organization [12]. Standardized protocols are typically developed together by multiple companies, and therefore we do not restrict our view on a single organization and how they manage innovation. Nevertheless, our results may have interesting implications to innovation management which are discussed in Section 6.

Even though our main interest is regarding the protocol diffusion, the events and decisions occurring before this point, as well as the envisioned consequences, have to be understood since they have a considerable impact on the diffusion process. Therefore, this section describes briefly the other relevant steps of the protocol development before we focus on diffusion in Section 3.

When the specification of a communication protocol begins, the need or problem has already been identified and the basic research has been done. Therefore, the protocol development process starts from the **development** step. During this step, the protocol is designed to solve at least one problem based on the research results. Standardization guaranteeing interoperability and thus communication between different implementations is typically an integral part of the development step. In the **commercialization** step the protocol is implemented in the commercial products, typically operating systems (OS) or applications depending on the layer in which the protocol is operating. Finally, the protocol **diffuses** to the device base consisting of either client and server devices in the case of an end-to-end protocol or routers and middleboxes in the case of core protocols. These three steps form the protocol development process presented in Figure 1 and described below in more detail. Although consequences are not seen as a part of the process, monitoring and analyzing both expected and realized consequences is crucial when a protocol is being developed.

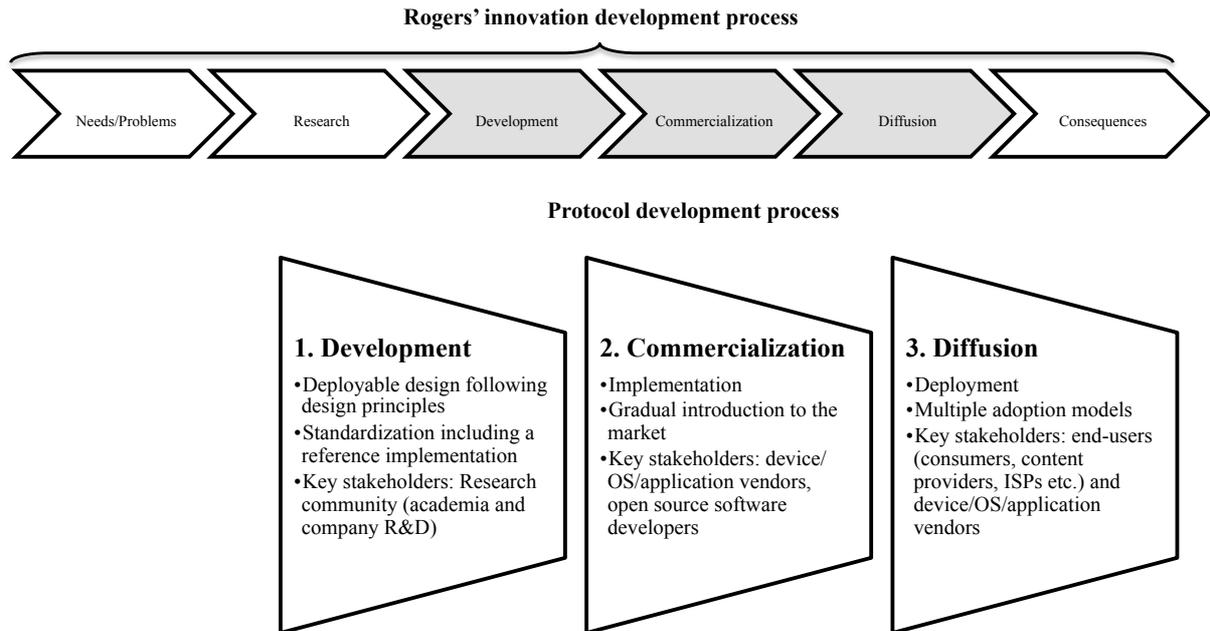


Figure 1. Protocol development process refined from Rogers' innovation development process.

Development

Designing the protocol forms the core of protocol development since the prerequisite for diffusion is that protocol itself is sufficiently deployable. This means that the design is compatible with the protocols and devices already used in the Internet, and that it matches the business interests of all the stakeholders who are necessary for the protocol diffusion. Therefore, understanding expected consequences that formulate stakeholder attitudes towards a given protocol is important already during the development. Using proven design principles and architectural guidelines of the Internet [13, 14], like the classic end-to-end principle [[15]] and refinements such as “design for tussle” [16, 17], helps in designing protocols that are feasible for each relevant stakeholder. Although following such principles does not ensure the success of a certain protocol, they are important for the long-term evolution of the Internet [18].

For communication to take place, protocols have to be agreed upon. Therefore, the development of communication protocols is typically standardization-driven¹. MPTCP, as most of the other protocols used in the Internet, is standardized in the IETF [3]. IETF standardization is based on “rough consensus and running code” meaning that the functioning

¹Also multiple proprietary communication protocols exist, especially in the application layer, e.g., protocols used in Skype or BitTorrent. Diffusion of proprietary protocols requires typically that a single stakeholder is able to control the diffusion process. This paper focuses on the standardized protocols, even though findings are at least partly valid also to proprietary protocols.

of a protocol needs to be proved by a reference implementation before the final version of the standard is published [19].

Commercialization

To reach the target market, the protocol implementation needs to be made available for end user devices. This requires implementation of the protocol directly into devices, operating systems or applications, depending on the layer in which the protocol is operating. MPTCP, for example, is a transport layer protocol which requires implementation in the network stack of the operating system. Therefore, OS vendors can be seen as a bottleneck in the development process of MPTCP.

Typically device, OS and application vendors introduce a new feature (such as a new protocol) gradually, so that its market potential and functionality can be tested with decreased risk. For example, mobile phone manufacturers, like Nokia, bring new features first to their high-end devices, whereas application vendors, like Google, put them first in the beta or development versions of their software. Usually, the commercialization takes some time, and may never happen entirely if the end user experience is not satisfactory. The commercialization stage does not need always be led by a commercial entity – free and open source software developers can play an equally significant role in this phase through their choice to implement new features using community contributions.

Diffusion

When the protocol implementations are available, they spread to the end user device base in the process called diffusion. In [20] the authors divide research on diffusion into two main streams. *Classical diffusion of innovation (DoI)* literature, conceptualized especially by Rogers [2], has focused on identifying the attributes of innovations and their influence on the decision to adopt. This relates closely to the protocol design principles and suggested protocol adoption models in Section 3. The *economic perspective* on innovation adoption suggests that the diffusion depends on the economic value an innovation brings to potential adopters. This value depends on both the size of the existing network of adopters and the potential network of adopters, *i.e.*, network externalities [21]. In general, adoption becomes more likely when the number of existing adopters in the network increases. The economic perspective is highly relevant for communication protocol diffusion and therefore we focus on that in Section 4.

Diffusion terminology, especially the terms *diffusion*, *adoption* and *deployment*, are sometimes mixed without further definitions, which may confuse the reader. To avoid this confusion, these key terms and their usage in the context of communication protocols are defined here.

Roger's [2] definition for diffusion is used as a basis for our definition: *Diffusion is the process by which an innovation is communicated through certain channels over time among the members of a social system.* So by that definition diffusion considers the patterns of how an innovation spreads, focusing on the impact of communication, *i.e.*, creating and sharing information about an innovation, among people. Adoption is similar to diffusion except that it deals with the psychological processes an individual goes through, rather than an aggregate market process. Adoption follows from five step innovation-decision process during which individuals (or other decision-making units) seek and process information to decrease uncertainty about innovation to make a decision to adopt or reject it [2]. Based on this, intentional decision-making is implicitly assumed to precede adoption. However, this does not always apply when the diffusion of communication protocols is concerned, as is explained in the following section. Therefore, in this paper, the protocol has **diffused** to an individual, *i.e.*, an individual has **adopted** the protocol, when she has started using it.

Deployment, for one, is a term which is used especially by engineers to describe all of the activities that make an innovation available for use by the end users. This does not mean, however, that the innovation would actually be used by individuals. Therefore, by **deployment** we are referring to protocol being put in the required networking equipment. Adoption is dependent upon deployment, with the additional step that end users are actually sending traffic using the protocol. This kind of distinction may not be relevant in other fields of diffusion research, where the deployment (or even acquisition preceding deployment²) of the innovation is used as an adoption event even though the innovation may end up never being used. When communication protocols are concerned, using deployment as an adoption event is not sufficient because the benefits and the network effects that affect diffusion rate are realized only during actual usage of the protocol.

² Fichman and Kemerer [22] have studied the significance of this concept known as the *assimilation gap*, meaning the gap that exists between the cumulative adoption curves associated with the alternatively conceived adoption events, *e.g.*, acquisition, deployment, or usage.

3. Adoption models of communication protocols

Innovation literature has focused on the diffusion process during which an individual (or organization) makes an intentional adoption decision. This applies also to earlier utility-based diffusion studies of communication protocols [20, 23, 24] which neglect the reality that the adoption is not necessarily based on rational decision-making; adopting individuals may also start using software features (*i.e.*, adopt a protocol) unintentionally. And even if the adoption involved intentional decision-making, the adoption decision may concern directly the protocol itself, or indirectly a product or service which benefits from using the protocol. Therefore, a protocol can often be considered as a product feature that diffuses along with a device or an operating system. For example, Kivi et al. found that the diffusion of mobile handset features is supply-driven meaning that the manufacturers make the decision which features to bundle in a specific device, which leads to a nested model of feature diffusion among the handset population [25, 26]. Due to these special characteristics we can recognize three different adoption models which are presented below and summarized in Table 1 using MPTCP as an example.

3.1 Direct adoption

Direct adoption matches to the general process of adoption decision-making. The adopter becomes aware of the protocol, evaluates its benefits against its costs and then makes a decision to adopt or reject it. The necessary actions are then taken to acquire the protocol, typically either purchasing a new device or updating the software of the existing device. The distinctive characteristic from the indirect adoption is that the adopter can link the perceived benefits, for example, increased throughput and resilience in the case of MPTCP, to the protocol itself.

3.2 Indirect adoption

Indirect adoption is similar to direct adoption but now the actual adopted innovation is a product or a service which contains the protocol as an integral part. This indirection means that the adopter may be completely unaware of the protocol's existence, even though she observes and makes her decision based on the perceived benefits that are caused by the protocol. It can be argued that this is a more natural model of communication protocol diffusion than direct adoption because adopters (especially consumers) are seldom informed about or even interested in communication protocols. They are just means to enable attractive applications and services. MPTCP could diffuse, for example, alongside a new version of the

popular P2P file sharing application uTorrent offering shorter download times, or a “FastYouTube” service providing better quality of experience for watching video clips, if the content providers would deploy the protocol and distribute it to their customer base.

3.3 Unintentional adoption

In unintentional adoption the adopters may be fully unaware of the existence of the protocol, or the decision preceding the adoption may not even be based on the protocol at all. For example, adopters replace their devices every now and then, and the selection of the device may be based on completely different features and attributes than the communication protocol or its expected benefits. Additionally, many consumers outsource the updating of their OS to the automatic update procedure controlled by the OS vendor, due to which they get protocol support automatically if the OS vendor decides to include it to the update patch. Therefore, OS and device vendors play a key role in the unintentional adoption and thus also in the protocol diffusion. For unintentional adoption to happen, the protocol needs to be enabled by default, since the adopter’s intentional decision to turn the protocol on in her OS, would be the direct adoption model.

Table 1. Adoption models of communication protocols.

	Unintentional adoption	Indirect adoption (intentional)	Direct adoption (intentional)
Adoption route	Device acquisitions and OS updates (either full update or OS patch) including the protocol	Acquisition of a service application including the protocol	Acquisition of the protocol support by updating the device or its software
Driver of the adoption	Updating the software due to other reasons (the protocol does not matter at all)	Adopter perceives benefits of using a service but may not be aware of the source causing the benefits	Adopter perceives benefits of using the protocol and links these benefits to the protocol
Key decision and decision-maker	<i>Device / OS vendor</i> enables the protocol by default in its products	<i>Application service provider</i> takes the protocol into use in (and distributes it with) its application service software	<i>Adopter itself</i> acquires the protocol support by updating her device
Adopter’s awareness of the innovation	No awareness (or user is indifferent to the protocol and its benefits)	Partial awareness (user knows the benefits of the protocol)	Full awareness
Example from the MPTCP case	MPTCP comes with the OS update	Google introduces “FastYouTube” encouraging users to install MPTCP patch, or MPTCP would be added to the latest version of uTorrent	Adopter installs the MPTCP patch from the OS support site

4. MPTCP diffusion model

Diffusion of a communication protocol is essentially a complex process and modeling the whole system helps to understand the interaction between the market stakeholders, which further serves a long term goal to assist with the protocol diffusion. System dynamics is an applicable method for modeling such problems. This section first explains the principles of system dynamics modeling, articulates the problem of the MPTCP diffusion model and then

explains its qualitative and quantitative assumptions. The values presented in this section are for the baseline scenario, and they are varied during the sensitivity analysis of Section 5.

System dynamics [27] combines system thinking and mathematical simulations. A benefit of systems modeling is that a complex system is considered as a whole rather than studying its parts separately. This enables the identification of feedback loops and interactions between different parts of the system to estimate the behavior of the system and its components. Causal relationships between different model parameters are positive or negative which form reinforcing (R) or balancing (B) feedback loops, respectively. The positive reinforcement loop indicates a feedback which amplifies the behavior of the system. This means that if a quantity³ grows, the rate at which the quantity grows also increases. The negative reinforcement/balancing loop represents the opposite feedback which counteracts or decreases the rate of further growth. Although both feedback loops act instantaneously, they may have different strengths at different times, which gives rise to a dynamic system. The power of the approach is that properties of the various system entities can be varied, to understand their effect on the overall system, which would otherwise be too complex to determine.

The problem behind the MPTCP diffusion model can be articulated as follows:

Which factors have the most significant impact on the diffusion rate of MPTCP among consumers and content providers?

We tackle this question by simulating a model which estimates the number of consumers and content providers (CP) over time. Although system dynamics simulations are done over time, the intention is not to forecast the exact timescale for the diffusion of MPTCP. The focus is on identifying the most relevant factors affecting the diffusion and evaluating the relative diffusion rates of the protocol.

Due to its end-to-end nature, MPTCP needs to be deployed both in client and server devices in order to become widely used in the Internet. Without getting content providers to adopt, MPTCP will be unlikely to diffuse widely. Therefore, we have chosen content providers and consumers as adopters in the model. By a content provider we refer to a stakeholder with business interests which makes any type of content available on its server(s) to be accessed

³ For instance, the number network nodes that are MPTCP capable.

by consumers. A consumer uses a device (mobile or fixed) to access content in the Internet. MPTCP can also be used for P2P traffic. During recent years, however, P2P traffic proportion has been declining due to extending usage of client-server type of content (especially video) distribution over HTTP, including one-click file sharing, progressive downloading (YouTube) and Adobe Flash [27]. The model could also be applied to P2P traffic but it is not seen as a significant driver for MPTCP diffusion.

The minimum requirement for MPTCP communication is that both end points have the software support and at least one of the end points is multihomed, *i.e.*, has multiple network interfaces and a dedicated IP address for each interface to be used simultaneously. Many content providers already have multihoming capability, or would be willing to acquire or reconfigure multiple access lines if they decide to adopt MPTCP. Technically advanced content providers may look to optimize their choice of providers in order to maximize the benefit from MPTCP by using disparate routes. Many consumers, however, may never become aware of MPTCP, and it will not be a driver for consumers acquiring multihoming support. On the other hand, many mobile handsets already have multiple network interfaces (3G and Wireless LAN), and this may provide an adoption route for multihoming. In this study, a consumer adopter implies a user who has the software support for MPTCP (and may optionally have multihoming support), whereas a content provider adopter has both software and multihoming support. The deployment of multihoming as well as the protocol itself may incur many complexities to different stakeholders, such as contracting with multiple access operators. Since we restrict the modeling to motivations stemming primarily from the market, we do not consider the influence of these complexities for a single stakeholder in the model.

The availability of the MPTCP software implementation is different for content providers and consumers. Content providers are more likely to make an intentional choice to enable MPTCP, possibly by implementing the software themselves, enabling OS features that are not on by default, or pressuring their operating system vendors to provide the required software patch. For consumer devices, the software availability increases gradually depending on OS vendors who have implemented it.

The proposed causal diagram of MPTCP diffusion process is shown in Figure 2. The model shows the external or primary reasons why different stakeholders may begin to use MPTCP. In case of content providers, for example, they first need to become interested in a certain

technology before they start to do any investment calculations or decisions. The actual decision-making process where stakeholders consider deployment cost or complexity is not considered because they are internal factors of each stakeholder and represent a second layer of detail which is not appropriate here. The fundamental assumption is that MPTCP inherently has some benefits which are increasing as a function of other adopters in the system. Therefore, we have omitted the detailed cost and benefit comparison in the model.

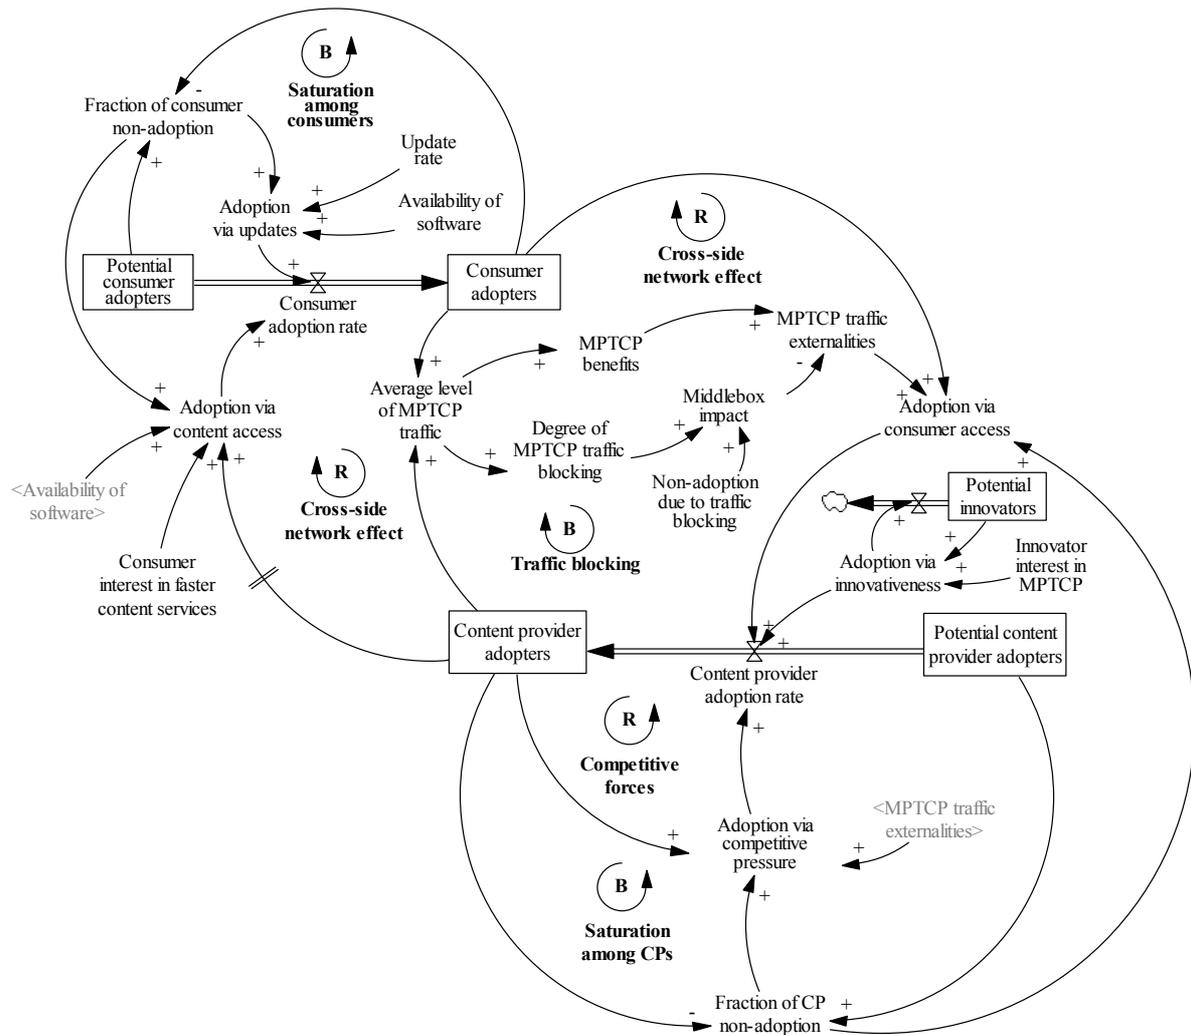


Figure 2. Diffusion of MPTCP to consumer devices and content provider networks.

4.1 Content provider adoption

We assume that the content providers are always fully aware of the new technology and its expected benefits, which means that only *direct adoption* is applicable. Content providers can base their decision on different motives, which are:

1. Adoption via innovativeness;
2. Adoption via consumer access;
3. Adoption via competitive pressure.

“Adoption via innovativeness”: Content providers, who expect to benefit from MPTCP and, regardless of non-existent customer base, are willing to adopt the protocol are called innovators. The adoption motivations for innovators stem from the expected customer base, willingness to benefit from MPTCP before competitors and possibly drive the adoption. In the model, innovators are assumed to be 2.5% of the total number of potential adopters as proposed by Rogers [2]. The rate at which potential innovators become adopters depends on their interest in MPTCP and the rate decreases in proportion to the number of potential innovators.

“Adoption via consumer access”: Content providers who wish to offer better service quality for their existing customers become interested when the installed base of MPTCP capable consumers is large enough, *i.e.*, there is a market that can be accessed. They may also be driven by the interest in attracting new customers with their improved service. We propose that the mapping between consumer adopters and content provider interest follows an s-shaped curve. As the diffusion of innovation typically follows an s-curve [2], the potential adopter interest is assumed to behave correspondingly. The rate of CP interest (the gradient) increases slowly at the beginning, is assumed to be highest when the number of MPTCP capable consumers is 50%, and slows down thereafter (see Figure 3). In the model, all content providers may become interested in, and can therefore become adopters of, MPTCP, although in reality this may not be the case.

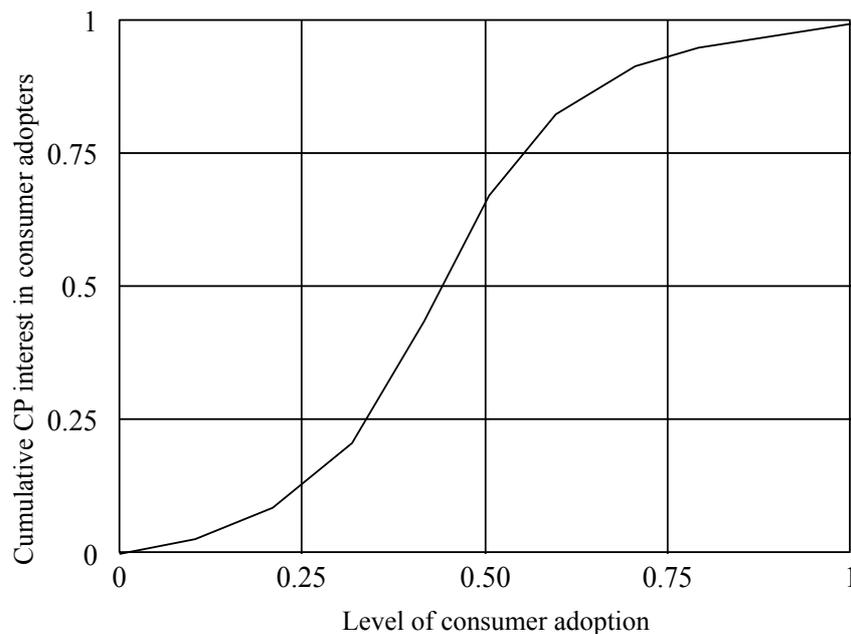


Figure 3. Cumulative interest of content providers as a function of MPTCP capable consumers.

“Adoption via competitive pressure”: Some content providers in the market tend to follow the market leaders. Their motive for upgrading to a new technology is because other providers have already done it. Especially if the benefits seen from MPTCP are significant, the followers need to adopt the technology in order to stay competitive. The competitive adoption is affected by the number of MPTCP capable content providers since the pressure to adopt increases in proportion to the number of adopters.

The second and third adoption channels are scaled by the fraction of content providers who have not yet adopted the technology which assures that none of the potential adopters is considered to adopt twice.

4.2 Consumer adoption

The model assumes that consumers are not aware of MPTCP at the time they adopt it. They may be interested in services which use the protocol, or the protocol may become installed in their devices without the consumer knowing it. Therefore, the channels for consumer adoption are *indirect* or *unintentional* adoption. The model identifies these channels as:

1. Adoption via content access;
2. Adoption via updates.

“Adoption via content access” is adoption whereby a consumer becomes interested in a certain content service which uses MPTCP. Before starting to use the service, such as our hypothetical example of “FastYouTube”, a consumer first needs to install the protocol in order to experience the benefits. The content provider’s site would instruct the consumer where to download and how to install the MPTCP OS patch before being able to start using the service. The model assumes a three-month delay (illustrated as a cross-line in the arrow in Figure 2) in service provision which illustrates the time that it takes from content providers to advertise its service to consumers. The number of available services is directly proportional to number of MPTCP capable content providers.

A component that could be explicitly taken into account in this adoption channel is the word-of-mouth within consumers. The consumer adopters leveraging information of better functioning applications to each other might be one stimulating factor in protocol adoption. However, we see that the number of available content services is the primary enabler for the indirect adoption and, therefore, the effect of word-of-mouth is only considered implicitly in

the parameter called “Consumers interest in faster content services”. In the early phase of the protocol diffusion the interest of consumers derives from advertising of improved applications while later on the consumer interest is probably mainly affected by word-of-mouth.

“Adoption via updates” consists of operating system updates. OS update can be either full (e.g., from Windows Vista to Windows 7) or partial (i.e. OS update patches). Full OS updates may happen by installing a new version of the software or by purchasing a device which contains a new OS. If the version of the updated OS supports MPTCP the consumer automatically becomes an MPTCP adopter, and therefore the OS update rate needs to be scaled by a factor which represents the degree of the OS availability. In the model, the software availability follows an s-shaped curve. The availability of the software implementation at initial time step is 0.1% and follows a Gompertz model [29]. The Gompertz model is used since it has been proven to be applicable for the diffusion of handset features, see, for example [25]. The upper asymptote of the model is 90% because MPTCP is unlikely to be implemented for old operating systems like Windows 95. The availability of MPTCP software reaches the saturation level in five years, which is similar to IPv6 which had OS support for the majority of installed systems within five years after the publication of the standard [30]. Figure 4 shows the increasing availability of the MPTCP software for consumer devices relative to time (in months).

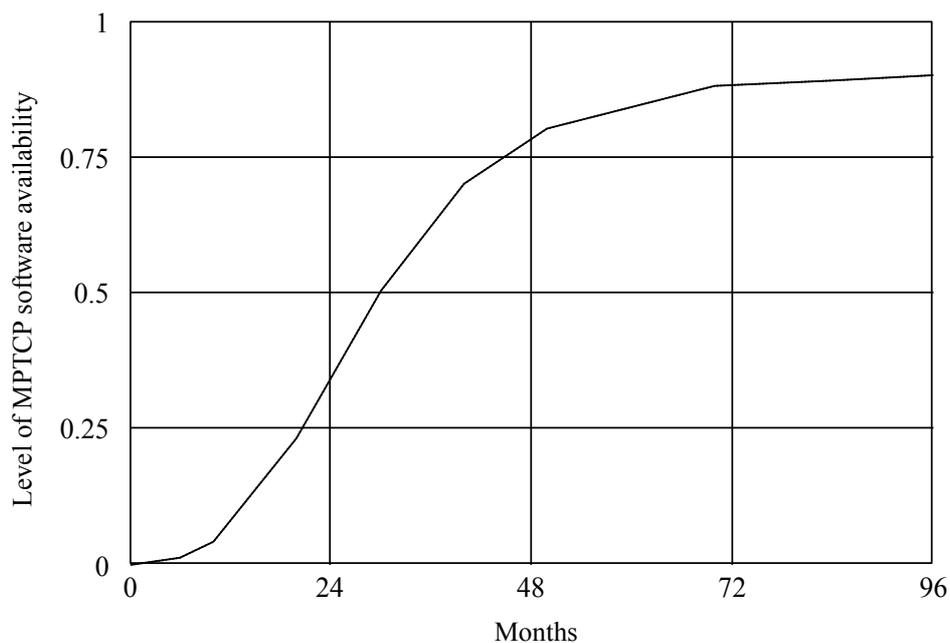


Figure 4. Availability of MPTCP software as a function of time (Gompertz parameters $a=0.9$, $b=1$, $c=7$).

The increasing availability of the software implementation affects also the adoption via content services. If consumers are interested in a certain content service but the OS implementation is not available for their device, the OS cannot be updated and users cannot start using the service with MPTCP. As in the case of content providers, each type of adoption is scaled by the fraction of consumer who have not yet adopted. The word-of-mouth within consumers does not apply in this adoption model since the users are adopting the protocol fully unintentionally.

4.3 MPTCP traffic externalities

As the number of MPTCP adopters increases and consequently MPTCP traffic starts to emerge in the Internet, both positive and negative externalities may unfold. On the positive side (the intention), the more MPTCP traffic exists, the more the congestion in the links is predicted to decrease. Therefore, the benefits to end users (such as improved throughput) increase, which further fuels content providers' interest in MPTCP. In our model, the benefits of MPTCP are assumed to follow a convex function. This means that when the level of MPTCP traffic is small the benefits of the MPTCP also remain small since only a few MPTCP users are unlikely to relieve the congestion in the Internet significantly. However, once the traffic level starts increasing and the traffic in the network becomes more balanced, the benefits of MPTCP increase more rapidly.

On the negative side, there is a possibility that MPTCP traffic might get blocked for two reasons. Firstly, misconfigured or poorly implemented middleboxes or routers may unintentionally block the MPTCP traffic due to unrecognizable header options. The authors in [[31]] measured the interaction between transport protocols and middleboxes and concluded that due to unconventional header options in TCP SYN packet in approximately 15 % of the cases the option is ignored. Since the measurements in [[31]] are relatively old, we assume that the unintentional blocking has increased to some extent. Secondly, MPTCP may have unexpected consequences in ISPs' internal traffic engineering and may even be seen to be harmful for their business⁴. Being in the middle of MPTCP communication, ISPs are potentially able to intentionally block the MPTCP traffic, although it may be questionable from a regulatory point of view. The model assumes that the overall traffic blocking increases in the early phase of the diffusion, since the intentional blocking emerges. As MPTCP becomes widely adopted, however, ISPs will adapt as negative publicity and regulatory

⁴ The role of ISPs in MPTCP communication is discussed in detail in [7] and [8].

interventions become more probable. Additionally, as with OS updates, over time the natural process of equipment upgrade should reduce the number of misconfigured or badly implemented middleboxes. In our model, at first MPTCP traffic is heavily blocked, but after a tipping point (where the average MPTCP traffic level is 30%) traffic blocking starts decreasing. Figure 5 shows the level of MPTCP traffic blocking as a function of average MPTCP traffic level.

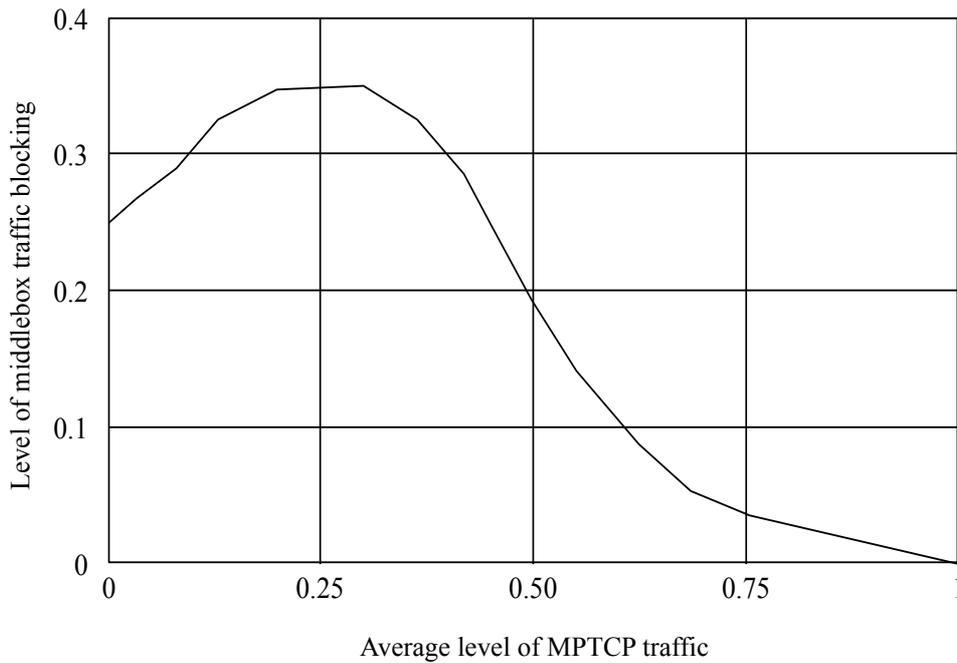


Figure 5. Level of MPTCP traffic blocking as a function of average level of MPTCP traffic.

The traffic externalities are only affecting “Adoption via consumer access” and “Adoption via competitive pressure” since innovators are not aware of realized traffic externalities when they make the adoption decision.

4.4 Model quantification

The quantification of the model is done implicitly. All parameters are abstracted to normalized unitless quantities that vary between zero and one. The sum of adopters and potential adopters equals 1 at each time step, and the initial value of adopters is zero. The stock that measures the number of adopters is calculated as an integral of adoption rates as shown in Equation 1.

MPTCP Adopters:

$$MA(t) = MA(t_0) + \int_{t_0}^t MA(t)dt \quad (1)$$

In the SD modeling, the most common ways to model the effect of different variables on a particular output are additive and multiplicative formulations. In this model the consumer and content provider adoption rates are calculated by using additive formulation while multiplicative formulations are used for calculating the adoption volumes.

Content provider interest in consumer adopters and the level of traffic blocking have been implemented as “Lookup” tables in the model. The “Lookup” tables of content provider interest and the degree of MPTCP traffic blocking get input and output values shown in Figures 3 and 5, respectively. “Degree of MPTCP benefits” follows a convex function so that 10% of the maximum MPTCP benefits are achieved when stakeholders start adopting MPTCP. When MPTCP traffic level increases the benefits increase quadratically as proposed in Section 4.3. The maximum benefits from MPTCP are achieved when the average traffic level is 100%. The total impact of MPTCP traffic externalities on content provider adoption is calculated as a joint probability of the degree of MPTCP benefits and the middlebox negative impact on adoption. These two factors are scaled to be in the same order of magnitude which means that content providers put equal emphasis on MPTCP benefits and the middlebox traffic blocking. The availability of MPTCP software has been set to follow the s-shaped function shown in Figure 4. We summarize all the functions used in the model in Table 2.

The system in the model is affected by six parameters which are assumed to stay constant over the simulation time. The parameters are consumer OS update rate, consumer interest in faster content services, innovator interest in MPTCP, and non-adoption due to traffic blocking fraction of innovators of CP adopters and delay in consumer service provision.

Table 2. Analytical expressions of the SD model.

Parameter	Formulation
CONTENT PROVIDER ADOPTION	
Content provider adoption rate	Adoption via innovativeness + Adoption via competitive pressure + Adoption via consumer access
Adoption via innovativeness	Potential Innovators * Innovator interest in MPTCP
Adoption via competitive pressure	Content provider adopters * MPTCP traffic externalities * Fraction of CP non_adoption
Adoption via consumer access	CONSUMER LOOK UP(Consumer adopters) * MPTCP traffic externalities * Fraction of CP non_adoption
Fraction of CP non-adoption	$\frac{\text{Potential content provider adopters}}{(\text{Content provider adopters} + \text{Potential content provider adopters})}$
CONSUMER LOOK UP	See Figure 3
CONSUMER ADOPTION	
Consumer adoption rate	Adoption via content access +Adoption via updates
Adoption via content access	Fraction of consumer non_adoption * Consumer interest in faster content service * Content provider adopters * Availability of software * STEP(1, Delay in service provision)
Adoption via updates	Fraction of consumer non_adoption * Availability of software * Update rate
Fraction of consumer non-adoption	$\frac{\text{Potential consumer adopters}}{(\text{Consumer adopters} + \text{Potential consumer adopters})}$
Availability of software	See Figure 4
MPTCP TRAFFIC EXTERNALITIES	
Average level of MPTCP traffic	If $\text{Consumer adopters} \wedge \text{Content provider adopters} > 0$ $\frac{\text{Consumer adopters} + \text{Content provider adopters}}{2}$ Else 0
MPTCP benefits	$(\text{Average level of MPTCP traffic})^2 + 0.1$
Degree of MPTCP traffic blocking	TRAFFIC LOOK UP(Average level of MPTCP traffic)
Middlebox impact	Non_adoption due to traffic blocking * Degree of MPTCP traffic blocking
MPTCP traffic externalities	MPTCP benefits * (1 – Middlebox impact)
TRAFFIC LOOK UP	See Figure 5

A previous study on mobile handset evolution [26] reveals that a mean lifetime of a mobile handset is approximately 2.6 years which has been achieved by fitting data from mobile handset sales and population in a Weibull distribution (e.g., [32]). The study assumes that the discard probability increases as a function of device lifetime and therefore the discard rate

varies over time but in our model, however, the update rate remains constant for simplification. Assuming the same mean lifetime for mobile handset and setting the Weibull shape parameter to one (discard probability constant over time) we get the 3% as an approximate monthly discard rate for mobile phones. Our study covers also laptops and home computers which usually have slightly longer lifetime and will decrease the overall discard rate. On the other hand, our study does not only cover the device replacements but also other means of OS updates which again will increase the OS update rate. Therefore, initially, OS update rate is set to 3% which means that 3% of the potential consumer adopters update their OS so roughly once in 2.5 years which should be a relatively modest assumption.

The consumer interest in faster content services represents how many consumers per month become interested in new MPTCP-enabled services. Some new services may attract multiple consumers while others may attract none at certain time step. To be easily comparable with the OS update rate, 3% of the consumers are assumed to gain interest in MPTCP services per month. “Innovator interest in MPTCP” and “Non-adoption due to traffic blocking” are set to 10% initially. These figures can be interpreted so that 10% of the remaining innovators adopt MPTCP per time step while the “Non-adoption due to traffic blocking” represents the degree that content providers are disturbed of MPTCP traffic blocking. The parameter values used in the model baseline are summarized in Table 3.

Table 3. Quantification of the SD model parameters in the baseline scenario.

Parameter	Baseline value
Fraction of CP innovators	2.5 %
Innovator interest in MPTCP	10 %
OS update rate	3 %
Consumer interest in faster content services	3 %
Non-adoption due to traffic blocking	10 %
Delay in service provision	3 months

5. Quantitative analysis

Several tools exist for system dynamic simulations, see, for example [33, 34, 35]. One of the most commonly used simulation software is called Vensim [33]. The Vensim PLE version is used as a modeling tool in this study. To obtain a baseline scenario for the quantitative analysis, the parameters introduced in Section 4 are set as initial values. In this section, a sensitivity analysis is provided to evaluate the criticality of model parameters. The impact of

different model parameters is invariant to the relative adoption rates regardless of their order of magnitude and instead of forecasting the exact adoption times of MPTCP, these relative adoption rates are of key interest.

5.1 Content provider adoption

First, the sensitivity analysis focuses on the parameters affecting primarily content provider adoption. To compare the impact of each adoption motive of content providers the rate of each adoption channel can be set to zero at a time. The analysis shows that the lack of consumer access significantly delays the content provider adoption. In the other cases, the content provider adoption reaches over 80% level in roughly eight years but the absence of consumer access results in content provider adoption of only 10% in this timescale.

To see how consumer adoption is affected by the content provider adoption the diffusion curves of both stakeholders can be compared in the different scenarios. An intriguing observation is that the number of consumer adopters is not significantly affected by the number of content providers, which implies that the impact of network effects on consumers is relatively weak. Figure 6 shows how the numbers of content provider and consumer adopters evolve when different CP adoption channels are effective.

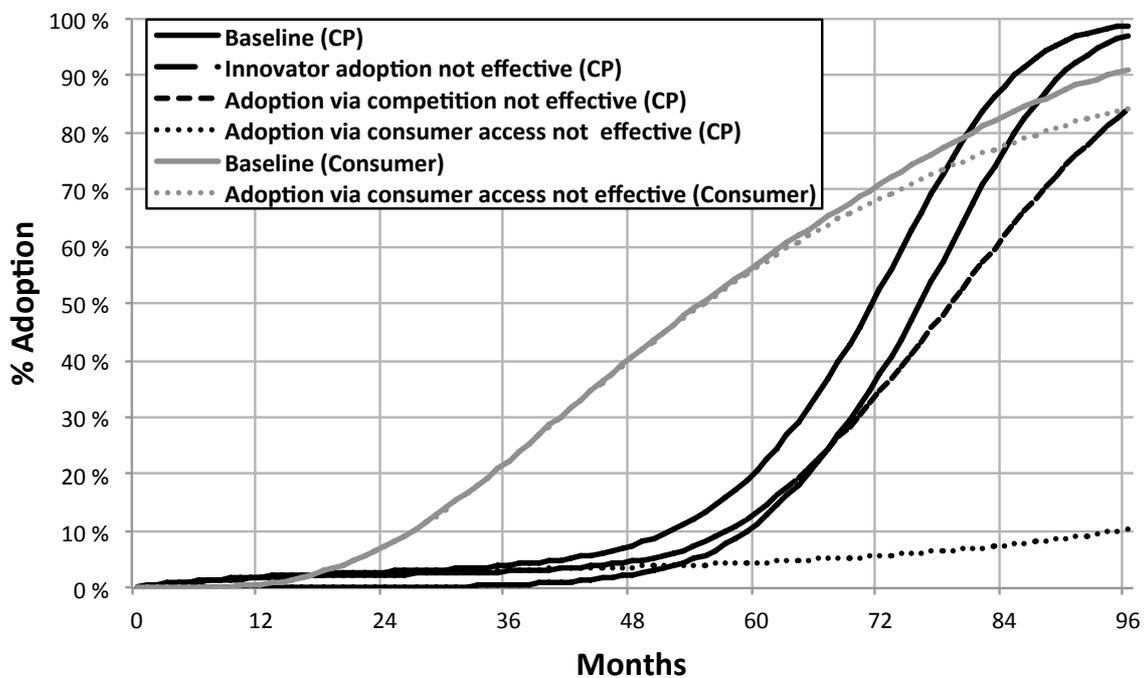


Figure 6. The impact on the number of consumer and content provider adopters over time by excluding the effect of one content provider adoption channel at a time.

The figure also shows that the innovator adoption is the only effective channel until roughly 36 months. This clearly illustrates the uncertainties which affect content provider adoption at the beginning of the protocol diffusion. The possible threat of middlebox blocking and uncertain benefits received from MPTCP may keep the adoption rate of content providers low for a relatively long time.

The sensitivity analysis also reveals that neither the adoption rate of content providers nor consumers is significantly affected by the cumulative interest curve of content providers. The baseline scenario assumes that the content provider interest in MPTCP capable consumers follows the curve shown in Figure 4. Shifting the gradient to reach its maximum value when 70% of consumers has adopted, decreases the content provider saturation level roughly 3% and the effect on the consumer penetration level is even less.

5.2 Consumer adoption

In order to understand the diffusion of MPTCP to consumer devices and how it affects the content provider adoption a sensitivity analysis can be carried out on two parameters that directly affect the consumer side. The first is the “Update rate” which models adoption through automatic operating system updates or other “side-effects” of consumer actions. In the baseline scenario this is set at 3% of devices per month, and Figure 7 shows the results of varying this parameter over an order of magnitude. The graph shows that the system is quite sensitive to this parameter, as it takes roughly 20 months in the best case or 48 months in the worst to reach a consumer adoption level of only 15%. Therefore, further work should be undertaken to characterize this parameter accurately, and improve the reasoning and modeling behind it.

The effect of OS availability was also tested in the sensitivity analysis. It revealed that a high level of OS support (say 80%) on day one significantly decreases the protocol diffusion time among consumers and therefore also among content providers. This consolidates the assumption that OS vendors can significantly delay the protocol diffusion process by not implementing the software. Figure 7 also shows how the maximum adoption level of consumers may increase above the level of OS availability. There will be devices on the market for which the MPTCP support is never made available (*e.g.*, Windows 95 computers). If we assume that at some point of time all the other operating systems have MPTCP support, the availability of MPTCP still increases as the market share of Windows 95 decreases

through renewal of potential adopters' device base. Therefore, in theory, all consumers may become adopters even though the MPTCP software would not be available for all operating systems.

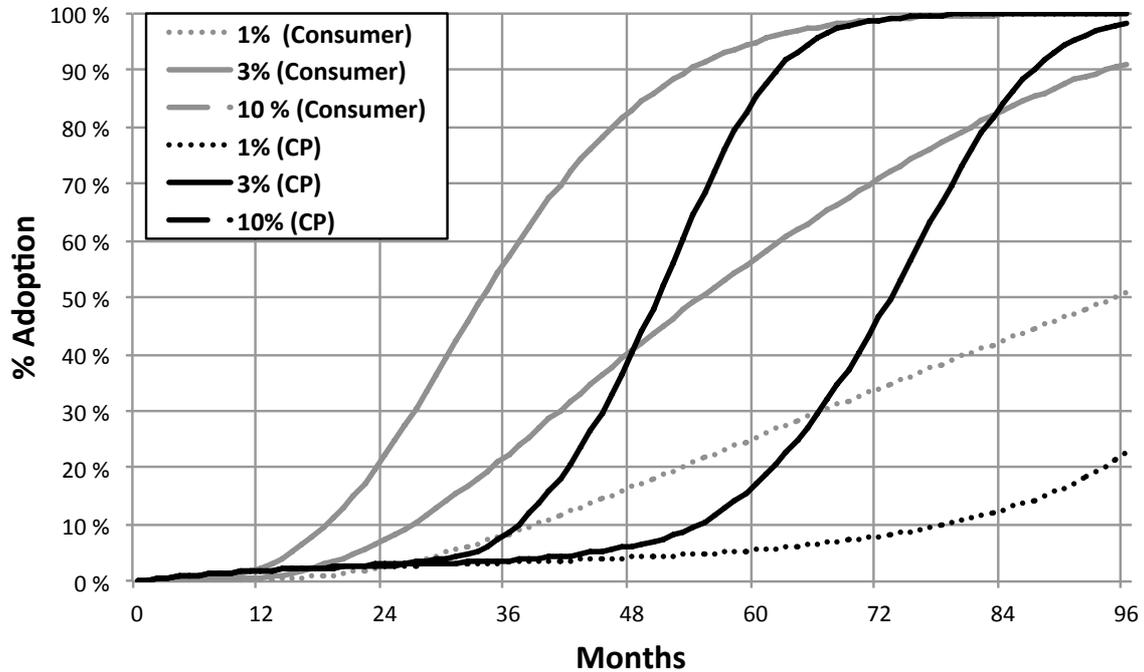


Figure 7. The impact on the number of consumer and content provider adopters over time by varying the "Update rate" parameter.

The second parameter is the "Consumer interest in faster content services", which is used to model indirect adoption, where consumers install a patch themselves⁵, to gain access to a particular service. This parameter is represented as a percentage of consumers who would indirectly adopt an MPTCP service, set at 3% in the baseline. Figure 8 shows the effect of varying this parameter over an order of magnitude. In contrast to the update rate, the system is reasonably tolerant to changes in this value. So, for example, even an order of magnitude change in the intentional adoption parameter results in only a small (less than six-month) shift to reach the relatively high 60% consumer (or content provider) adoption level.

The update rate and the consumer interest parameters can be also varied simultaneously. As expected the results reveal that if both parameters are low, the diffusion of MPTCP takes a double hit. Reaching the 40% consumer adoption level would take many decades, which to all intents and purposes represents failure. Although these results may be intuitive, they do represent a validation of the proposed system dynamics model.

⁵ Perhaps, provided directly by a content provider.

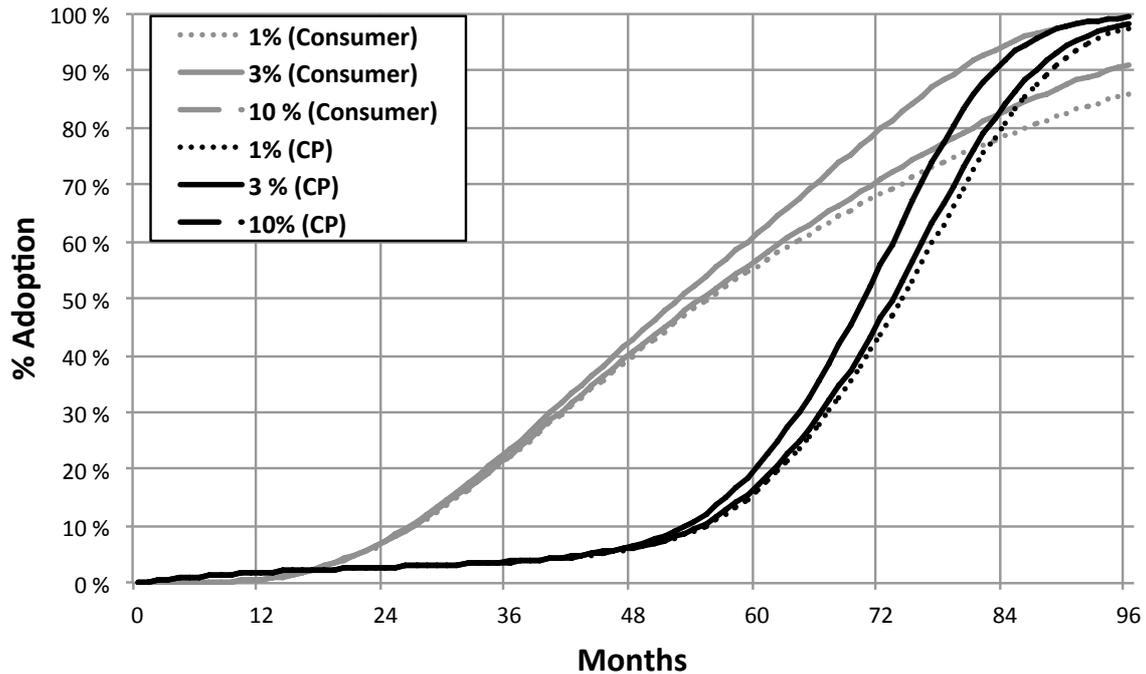


Figure 8. The impact on the number of consumer and content provider adopters over time by varying the “Consumer interest in faster content services” parameter.

A delay in content provider service provision has a marginal effect on the consumer adoption. Initially, the delay is set to 3 months. In the sensitivity analysis the delay is first set to 6 months and then to zero but in both cases only a slight change in the penetration level is observed.

5.3 Middlebox impact and MPTCP benefits

Finally, a sensitivity analysis can be carried out for the MPTCP traffic externalities that may affect the content providers’ willingness to invest in MPTCP. We only consider the effect of traffic externalities on content provider adoption since the results in Section 5.1 show that consumer adoption is not significantly affected by the content provider adoption. In the baseline scenario the middlebox traffic blocking follows the curve shown in Figure 5 and the level of content providers’ non-adoption due to traffic blocking was set to 10% as explained above. Varying the parameter value from zero to one results in roughly a six-month shift to reach 60% of penetration (see Figure 9). When the traffic blocking curve is varied so that the maximum middlebox blocking is increased up to different levels and the tipping point is then shifted, only an insignificant alteration is observed. These results imply that the middlebox impact on the diffusion seems to be relatively small.

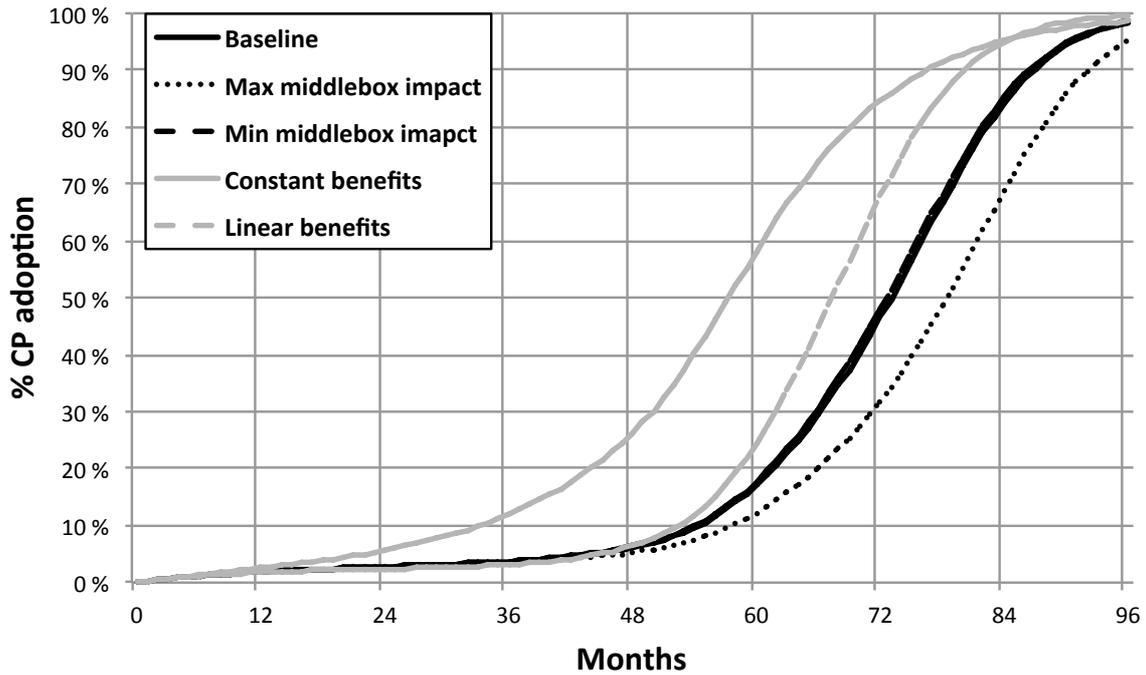


Figure 9. The impact of the MPTCP traffic externalities on the content provider adoption.

Since the evolution of MPTCP benefits over time is extremely difficult to estimate, different scenarios for MPTCP benefits can be tested. In our first scenario the benefits increase linearly (from zero to one), which represents a situation where each new adopter increases the MPTCP benefits, *i.e.*, has a positive effect on the traffic balancing of the Internet. In the second scenario the benefits remain constant (50% of maximum) over time. This means that the MPTCP benefits come mainly from multihoming whereupon the penetration of MPTCP and thus the Internet traffic becoming more balanced does not significantly further increase the benefits.

As shown in Figure 9 both linear and constant functions of MPTCP benefits seem to have a positive impact on the content provider adoption. The linear benefits result in roughly a six-month decrease when adoption level of 60% is considered, which however, is relatively small. Constant benefits seem to have even more positive effect on the content provider adoption. An implication of the results is that exposing the benefits of MPTCP especially in the early phase of the diffusion process is more important than having great benefits after achieving a certain critical mass.

The benefits of MPTCP come from exploiting the path diversity in the Internet. This is not under the control of a single authority but multihoming, for example, affects positively on

path diversity and may support the MPTCP diffusion. The reasoning behind the content providers' decision-making can be improved to show the impact of the MPTCP traffic externalities more accurately. Nevertheless, our model shows that if content providers are incentivized to adopt in many ways, the negative issues can be overcome so that the diffusion does not stagnate.

6. Discussion

This study is the first attempt to try to apply system dynamics to communication protocol diffusion and investigate whether it aids understanding of the dynamics of a protocol diffusion process. By building a high level model we have gained a perception of the most critical factors affecting the MPTCP adoption, which gives a good baseline for the future research.

In the modeling process we have focused on the reasons why different stakeholders would become interested in a novel protocol. The actual decision-making process, where the stakeholders evaluate the benefits against the costs of deploying and using the protocol, could be modeled separately with system dynamics or with other applicable modeling methods. When the reasoning behind each adoption channel has been studied in detail they can be integrated in the model and result in more accurate estimates about MPTCP diffusion.

The model, however, has limitations which should be considered in the future work. First, the content providers and consumers are assumed to be homogenous. No distinction has been made between small and large content providers in terms of customer base. If a large provider like Google was to adopt MPTCP, it would have much greater impact on the consumer adoption than a smaller, *e.g.*, national video streaming provider. Consumers are assumed to have similar interest in content services exploiting MPTCP although it may vary a lot depending on demographics and usage patterns.

The OS vendors who provide the required MPTCP software for consumer devices are all assumed to have a customer base of similar scale. In fact, a provider like Microsoft makes the software available for a significantly larger fraction of consumers than Linux, for example. In the real world, this leads to a situation where MPTCP software becomes available in chunks. One improvement could be to move to a more discrete model, to represent the fact that large

numbers of machines are updated simultaneously when a new operating system patch is pushed out.

Although MPTCP would become fully adopted by content providers and consumers, MPTCP would still not be universally used across the Internet. This model only considers traffic between consumers and content providers, and neglects other adopters such as companies and academic sites. These stakeholders might be interested in adopting MPTCP, for example, to provide better remote connections for their employees, and therefore the adoption by enterprises could accelerate others' adoption. The capabilities of MPTCP could also be implemented in the application layer which would make the deployment easier for a specific content provider but also restrict the protocol benefits to a specific application.

Additionally, the model does not consider the possibility of discarding the technology. Content providers make an investment when adopting MPTCP and prior to the adoption they carefully consider the benefits and cost of the business case, so they are unlikely to discard the protocol afterwards. Consumers are not making any monetary investments on MPTCP so they are also more susceptible to discard the protocol. Although the protocol would be enabled in the shipping configuration, consumers can possibly disable it if it impairs the user experience.

Despite the fact that the model is limited in many ways it offers a good perspective for MPTCP diffusion. The most interesting observation from the simulations is that due to unintentional adoption the diffusion rate of MPTCP among consumers is largely defined by the normal OS update rate which also has a significant impact on content provider adoption. Although OS vendors cannot be seen as the actual adopters of MPTCP (except in the case where OS vendor is also a content provider) they may have a prominent effect on MPTCP diffusion. OS vendors play an important gatekeeper role in the protocol development process and in-depth analysis of their decision-making is a relevant direction for further studies. In order to exploit the power of unintentional adoption the protocol needs to be enabled by default in the operating systems. If any extra configuration is required from the user, the update of the operating system cannot be seen as unintentional adoption.

Since MPTCP is an end-to-end type of protocol, cross-side network effects play a significant role in the protocol diffusion. Depending on which end user starts adopting first and which

market players are considered, the diffusion may lead to different diffusion patterns in the market. If MPTCP software becomes quickly available and the consumer adoption arouses the interest of content providers, the diffusion can be seen as market pull. If content providers become innovators and start providing the protocol along with their services, the diffusion may lead to technology push, assuming that the implementation for consumer devices is available. On the other hand, if the market between the OS vendors and adopters (both consumers and CPs) is considered, the unintentional adoption would resemble technology push. However, the unintentional adoption reduces the impact of cross-side network effects as seen in the simulations and the number of content provider adopters has little effect on the consumer adoption.

Although the unintentional adoption may have a significant role in the MPTCP diffusion process, the protocol itself needs to have verified benefits which intentional adopters' interest is based on. If MPTCP would end up in all consumer devices, but content providers did not see any benefit of adopting it, the diffusion would stagnate. Similarly, the potential obstacles for the diffusion, for example the threat of traffic blocking, should be identified and analyzed as early as possible in the protocol development process so that actions to remove them can be taken in the protocol design.

7. Conclusions

In this paper we have studied the dynamics of communication protocols and used MPTCP as an example protocol. The study shows that network effects combined with the unintentional adoption by stakeholders may have unexpected effects on MPTCP diffusion. The result is not limited to MPTCP, and the phenomenon should apply to other end user-centric protocols as well. Unintentional adoption adds a new dimension to the diffusion theories and can facilitate diffusion of communication protocols, which are often prone to network effects. It can also provide a route to reaching critical mass, and overcoming the “chicken and egg” problem that is typical in markets showing network effects.

The intentional protocol adoption, which may occur through direct technology adoption or indirect service adoption, is more sensitive to protocol benefits than the unintentional adoption. Without a real need and benefits for the stakeholders, the protocol is unlikely to become widely adopted. This has been seen with the IPv6 protocol, for example. If

stakeholders do not see any monetary or other incentive for adopting it, a new protocol diffuses very slowly.

Instead of concentrating on solely internal organization and procedures in innovation management, understanding the protocol development process as a whole is beneficial for companies. The analysis of the diffusion and consequences of the protocol usage should provide feedback to the preceding phases in the process, so that the protocol deployability can be improved and scenarios for successful commercialization of the protocol can be developed. Especially, content providers who also control operating systems (think Apple, Google or Microsoft, for example) may improve the quality of their strategic decisions by looking at the protocol development more holistically. These providers may, for example, increase their chances of becoming innovators by delivering a novel quality-improving protocol to their customer base. The decisions of content providers, after all, dictate how the protocol diffuses in the market.

The authors wish to thank Matthias Will and Philip Eardley for their constructive comments as well as the anonymous reviewers for their valuable insights and suggestions for improving the article.

The paper was partly funded by *Trilogy*, a research project supported by the European Commission under its Seventh Framework Program (INFSOICT-216372). Alexandros Kostopoulos is co-financed by the European Social Fund and National Resources (Greek Ministry of Education - HERAKLEITOS II Programme). The research is also linked to the European COST IS605/Econ@Tel project.

References

- [1] Handley, M. (2006). Why the Internet Only Just Works. *BT Technology Journal*, 24(3), 119-129.
- [2] Rogers, E. M. (2003). *Diffusion of Innovations*, 5th edition. Free Press, New York.
- [3] Ford, A., Raiciu, C., Handley, M., Barré, S. & Iyengar, J. (2011). Architectural Guidelines for Multipath TCP DEvelopment. RFC 6182. Available at: <http://tools.ietf.org/html/rfc6182>. Accessed 4 October 2011.
- [4] Wischik, D., Handley, M., & Bagnulo, M. B. (2008). The resource pooling principle. *ACM SIGCOMM CCR*, 38(5), 47-52.
- [5] Raiciu, C., Pluntke, C., Barré, S., Greenhalgh, A., Wischik, D., & Handley, M. (2010). Data Centre Networking with Multipath TCP. 9th ACM Workshop on Hot Topics in Networks (HotNets-IX), Monterey, USA (20-21 October 2010).
- [6] Levä, T., Warma, H., Ford, A., Kostopoulos, A., Heinrich, B., Widera, R., & Eardley, P. (2010). Business Aspects of Multipath TCP Adoption. In G. Tselentis, A. Galis, A. Gavras, S. Krco, V. Lotz, E. Simperl, B. Stiller & T. Zahariadis (Eds.) *Towards the Future Internet* (pp. 21-30). IOS Press, Amsterdam.
- [7] Kostopoulos, A., Warma, H., Levä, T., Heinrich, B., Ford, A., & Eggert, L. (2010). Towards Multipath TCP Adoption: Challenges and Opportunities. 6th Euro-NF Conference on Next Generation Internet (NGI 2010), Paris, France (2-4 June 2010).
- [8] Warma, H., Levä, T., Eggert, L., Hämmäinen, H., & Manner, J. (2010). Mobile Internet in Stereo: an End-to-End Scenario. 3rd Workshop on Economic Traffic Management (ETM 2010), Amsterdam, the Netherlands (6 September 2010).
- [9] Serman, J. (2000). *Business dynamics: Systems Thinking and Modeling for a Complex World*, McGraw-Hill/Irwin, 2000.

- [10] Kelic, A. (2005). Networking Technology Adoption: System Dynamics Modeling of Fiber-to-the-Home, PhD Thesis, Massachusetts Institute of Technology.
- [11] Thun, J.H., Grössler A., & Milling P. M (2000). The Diffusion of Goods Considering Network externalities: A System Dynamics-Based Approach. 18th International Conference of the System Dynamics Society, Bergen, Norway (8 August 2000).
- [12] Trott, P. (2008). Innovation management and new product development, Pearson Education, Harlow, England.
- [13] Carpenter, B. (1996). Architectural Principles of the Internet. RFC 1958. Available at: <http://www.ietf.org/rfc/rfc1958.txt>. Accessed 12 October 2011.
- [14] Bush, R., Meyer, D. (2002). Some Internet Architectural Guidelines and Philosophy. RFC 3439. Available at: <http://www.ietf.org/rfc/rfc3439.txt>. Accessed 12 October 2011.
- [15] Saltzer, J., Reed, D., Clark D. (1984). End-to-end arguments in system design. Transactions on ACM Computer Systems (TOCS), 2(4), 277-288.
- [16] Clark, D., Wroclawski, J., Sollins, K., & Braden, R. (2005). Tussle in Cyberspace: Defining Tomorrow's Internet. IEEE/ACM Transactions on Networking 13(3), 462-475.
- [17] Ford, A. Eardley, P., & van Schewick, B. (2009). New Design Principles for the Internet. International Workshop on the Network of the Future, Dresden, Germany.
- [18] Kalogiros, C., Kostopoulos, A., & Ford, A. (2009). On Designing for Tussle: Future Internet in Retrospect. 15th Eunice International Workshop (EUNICE 2009), Barcelona, Spain (7-9 September 2009).
- [19] Alvestrand, H. (2004). The Role of the Standards Process in Shaping the Internet. Proceedings of the IEEE, 92(9), 1371-1374.
- [20] Hovav, A., Patnayakuni, R., & Schuff, D. (2004). A Model of Internet Standards Adoption: the Case of IPv6. Information Systems Journal, 14(3), 265-294.
- [21] Katz, M. L. and Shapiro, C. (1986). Technology Adoption in the Presence of Network Externalities. Journal of Political Economics, 94, 822-841.
- [22] Fichman, R., & Kemerer, C. (1999). The illusory diffusion of innovation: an examination of assimilation gaps. Information Systems Research, 10(3), 255-275.
- [23] Joseph, D., Shetty, N., Chuang, J., & Stoica, I. (2007). Modeling the Adoption of New Network Architectures. International Conference on Emerging Networking Experiments and Technologies (CoNEXT'07), New York, USA.
- [24] Sen, S., Jin, Y., Guérin, R., & Hosanagar, K. (2010). Modeling the Dynamics of Network Technology Adoption and the Role of Converters. IEEE/ACM Transactions on Networking, 18(6), 1793-1805.
- [25] Kivi, A., Smura, T., & Töyli J. (2009). Diffusion of Mobile Handset Features in Finland. 8th International Conference on Mobile Business, Dalian, China (26-28 June 2009).
- [26] Kivi, A., Smura, T., & Töyli, J. (2012). Technology product evolution and the diffusion of new product features. Technological Forecasting and Social Change. In Press.
- [27] Forrester, J. W. (1961). Industrial Dynamics, Productivity Press.
- [28] Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., & Jahanian, F. (2010). Internet Inter-Domain Traffic, Sigcomm'10, New Delhi, India (30 August-3 September 2010).
- [29] Gompertz, B. (1825). On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. Philosophical Transactions Series I, 115, 513-583.
- [30] IPv6 Forum (2010). IPv6 Historic Timeline. Available at: <http://www.ipv6forum.org.au/timeline.php>. Accessed 30 November 2010.
- [31] Medina, A., Allman, M. & Floyd, S. (2004). Measuring Interactions Between Transport Protocols and Middleboxes, 4th ACM SIGCOMM Conference on Internet Measurement (IMC 2004), Taormina, Italy (25-27 October 2004).
- [32] Bayus B. L., Hong S., & Labe R. P. (1989). Developing and Using Forecasting Models of Consumer Durables, Journal of Product Innovation Management, 6(1) 5-19.
- [33] Ventana Systems Inc. (2010). Product webpage. Available at: <http://www.vensim.com>. Accessed 4 October 2011.
- [34] Powersim (2011). Product webpage. Available at: <http://www.powersim.com>. Accessed 5 October 2011.
- [35] Optisim (2010). Software webpage. Available at: <http://www.optisim.org/QLENG/>. Accessed 5 October, 2011.