

Mobile Internet in Stereo: an End-to-End Scenario

Henna Warma¹, Tapio Levä¹, Lars Eggert², Heikki Hämmäinen¹ and Jukka Manner¹

¹Aalto University, Department of Communications and Networking,
P.O. Box 13000, 00076 Aalto, Finland

{henna.warma, tapio.leva, heikki.hammainen, jukka.manner}@tkk.fi

²Nokia Research Center, P.O. Box 403, 00045 Nokia Group, Finland
lars.eggert@nokia.com

Abstract. Multipath TCP (MPTCP) is an extension to regular TCP that exploits the idea of resource pooling by transmitting the data of a single TCP connection simultaneously across multiple Internet paths between two end hosts. Operating system vendors are traditionally in the key position to facilitate the deployment of new functionality, such as MPTCP, to user devices, but their motivation to invest in such enhancements is not self-evident. A scenario in which one party has the capability to deploy software changes on both the mobile devices and on the content servers helps to understand the potential first-mover advantages associated with the deployment of a new IETF standard protocol. In this study, we have built a quantitative techno-economic model to estimate the implementation costs for a content provider selling application downloads who is also able to implement MPTCP on both ends and compare them against the revenue MPTCP generates. The results suggest that even a relatively small increase in the number of downloads could make the business case profitable within five years.

Keywords: Multipath TCP, deployment, content provider, techno-economic modeling

1 Introduction

One intriguing change to the Internet protocol suite that has gained significant standardization traction recently is Multipath TCP (MPTCP). It is an extension to standard TCP that supports the simultaneous use of multiple network paths between two end hosts involved in a TCP connection. Thus, MPTCP is a manifestation of resource pooling principle that improves utilization of Internet resources by allowing separate resources to behave as if they were a single aggregate resource [1]. MPTCP uses a coordinated congestion control scheme [2] to balance transmission rates across the paths it transmits over in order to guarantee fairness to standard TCP while increasing connection throughput and resilience. The standardization of MPTCP is currently ongoing in the Internet Engineering Task Force (IETF) [3].

The deployment of any change to an Internet protocol is a complex process; the deployment of a new protocol even more so. Although a new protocol may bring

substantial verified benefits and may have a good technical design, these factors alone are often not sufficient to guarantee successful deployment of the protocol. The deployment of such protocols depends on the incentives for and the behavior of other stakeholders in the broader Internet market. When a new protocol is available for end users, they may never have to come to a conscious decision to start using it but the new protocol functionality can be enabled for them, *e.g.*, through software updates.

This paper investigates MPTCP deployment and adoption using a specific scenario. Of particular interest is whether MPTCP offers a profitable business case for a client-server type of stakeholder controlling the operating systems on both mobile devices and on the content servers. Multiple examples of this scenario exist. Nokia, Apple, Microsoft and Google all sell portable devices to consumers and also provide service portals for users to access and download content, such as music, movies and applications.

A techno-economic model, *i.e.*, a net present value (NPV) analysis, is used to quantify the business case by comparing the costs of implementing MPTCP against the revenue of the additional business that MPTCP may generate. The model considers a mobile application store and uses generic values that are based on real-life references to the Nokia Ovi Store and the Apple App Store. The results indicate that even a small increase in downloads due to MPTCP turns the investment profitable. Exploiting scale advantages is possible, because CAPEX dominates OPEX and in CAPEX, software implementation costs are significantly higher than hardware costs.

2 Background

This section introduces the reader to the existing work related to MPTCP deployment and techno-economic modeling.

2.1 MPTCP Deployment

Because MPTCP is a strict end-to-end protocol, an implementation needs only affect the end points. Ref. [4] presents the overall three-step deployment process using different technology and Internet standards adoption theories. First, operating system (OS) vendors need to make the protocol implementation available. Second, the protocol needs to be deployed on mobile devices and data centers. End users can either directly install MPTCP or an installation may happen indirectly, *e.g.*, through automatic software updates and new device purchases. Third, at least one of the two ends needs to multihome, *i.e.*, to connect to the Internet through multiple paths, which may require contracting with multiple ISPs.

Although the involvement of ISPs is not mandatory for either the implementation or the deployment of MPTCP, Ref. [5] describes how ISPs may support the adoption of MPTCP in different deployment scenarios. The model of this paper does not discuss how end users have contracted with ISPs to become multihomed, but the role of ISPs in the middle of MPTCP communication is discussed in Section 6.

In addition to the direct implementation by end hosts, MPTCP can also be deployed via proxies [6][7]. MPTCP proxies can be deployed upstream of end hosts that only implement standard TCP; they terminate inbound MPTCP connections and relay the communication over standard TCP in a way that is transparent to both ends. The result is that an MPTCP host believes it is communicating with another MPTCP host, while the other peer believes the communication uses standard TCP. Proxies can also be deployed in front of both end hosts [8], in which case a communication session across the network can be “upgraded” to MPTCP without modifying the end hosts at all.

2.2 Techno-economic modeling

The techno-economic modeling exploits future forecasting, technology design and investment theories in evaluating the economic feasibility of a new technology. Usually, a net present value (NPV) analysis is used, because it takes the time value of money into account and therefore lends itself well to estimating future cash flows. Different authors have developed spreadsheet-based tools specifically intended for the telecommunications industry. Katsianis *et al.* [9] use techno-economic modeling for estimating the evolution of mobile network operators, whereas Smura *et al.* [10] use it to model the financial position of mobile virtual network operators (MVNO). Because techno-economic modeling includes forecasting and usually many of the assumptions are uncertain, it should always be followed by a sensitivity analysis.

3 Implementation Scenario

Because MPTCP is an end-to-end protocol, both clients and servers require MPTCP support before the protocol can be used. This section identifies the implementation requirements for both ends and discusses how MPTCP can be implemented by a content provider.

Two high-level approaches exist for content providers to implement MPTCP. The first option is to implement MPTCP through a multihomed proxy that is installed at the border of the data center. MPTCP is then used between the mobile end user and the proxy, which forwards the communication to the actual content servers over a regular TCP connection. The second option is to update the load balancers at the edge of the data center to support MPTCP, deploy MPTCP on the servers in the data center, and also multihome those servers. The load balancers will need to be updated, because otherwise they may forward the individual subflows of one MPTCP connection to separate servers in the data center, therefore rendering MPTCP ineffective.

This paper assumes that MPTCP is deployed on the data center side using proxies. Fig. 1 illustrates the technical architecture of this scenario. The dashed line shows the multipath connection. A gradual rollout of MPTCP proxies is possible but the full benefits of MPTCP are realized only after all content transfer use MPTCP. An advantage of the proxy implementation is that the internal service infrastructure of the

data center can remain unchanged and that the load balancers, which forward user requests to different content servers, do not require any changes.

Although Fig. 2 depicts the MPTCP proxy as having multiple network connections, multihoming is not a mandatory requirement at the content server side, *if* the mobile user is multihomed (which this paper assumes). However, data center multihoming increases the number of possible connection subflows and many data centers are multihomed for fault tolerance already.




Fig. 1. Technical architecture of the scenario

A mobile end user can obtain MPTCP support either by installing MPTCP to her end host and acquiring multiple access connections to become multihomed, or through an MPTCP proxy provided by her ISP [5]. The scenario where an end user is multihomed is more promising, because the involvement of ISPs is not needed and MPTCP functionality can be implemented solely on the mobile device. Therefore, this paper does not consider proxying on the client side.

The basic functionality of MPTCP is implemented as an extension to TCP in the transport layer of the operating system kernel. In order for MPTCP to remain compatible with existing applications, an MPTCP implementation needs to operate effectively when accessed through the standard TCP API.

Although most existing hardware platforms already support multiple network access technologies, some operating systems do not expose the availability of multiple accesses above the IP layer. This prevents higher-layer protocols and applications from using multiple paths. Such operating systems require additional modifications to enable simultaneous use of multiple access connections.

4 Modeling Provider Costs and Benefits

Evaluating the profitability of an investment compares the costs against the revenue and other benefits. This section discusses the costs and benefits for the content provider deploying MPTCP. The basic assumptions of the techno-economic model are also presented.

4.1 Costs

Implementing the MPTCP functionality in operating systems causes the largest expenses to content providers. Content servers typically run different or at least differently optimized operating systems than mobile devices. MPTCP needs to be implemented for each operating system in use, which is costly, but distributing the software to devices incurs only low costs.

A software development project consists of three main steps: feature specification, implementation and testing. The implementation phase is the most time consuming phase and requires more workforce than the other two phases. In the model, the implementation phase takes three months and the other phases two months. Software specification and testing require three people working full time, whereas implementation requires eight people. If the total monthly cost of an employee is 10 000 €, completing these three phases costs 440 000 €. In addition, installation of the software as well as testing equipment incurs some costs for the provider, so an assumption for implementing MPTCP for one operating system is 500 000 €. The model presumes that on the proxy side, the MPTCP software is implemented once and on the end user side, for three operating systems. After the software has been implemented once, costs of developing other versions are lower, because expertise has been acquired during the first implementation.

In addition to the software development costs, deployment of MPTCP on proxies incurs purchase, installation and maintenance costs. The level of these costs depends on the complexity of the content server infrastructure and the number of required proxies. Additional network connections may be required to satisfy the multihoming requirement, but this paper presumes that server farms are already sufficiently multihomed, so that this cost component can be neglected.

The number of required proxies depends on the capacity of each proxy. Modern servers appear to be capable of handling at least 20 000 simultaneous TCP sessions (*e.g.*, [11]). Because an MPTCP connection consists of multiple TCP subflows, the capacity of an MPTCP server is assumed to be 10 000 simultaneous MPTCP sessions. Interviewed experts in the field suggested that a proxy with this capacity costs approximately 5 000 €. The installation costs are included in the price. In order to avoid a single point of failure in the content service, the required number of proxies is doubled in the model.

The operational expenditures of a server include maintenance work such as fixing fault situations and updating the software on occasion. The energy consumption cannot be neglected and it presumably dominates OPEX compared to the maintenance costs. Ref. [12] suggests that the annual energy costs of a small server are approximately 1 500 €. The server type used in the model has larger energy consumption. The annual maintenance costs are initially 30 % and the energy costs are 50 % of the purchase price of a server. The model assumes that maintenance costs increase in relation to the growth rate of MPTCP users in the content provider's network, but energy costs remain stable during the study period, because the servers do not support power control. Table 1 summarizes the cost assumptions of the model.

Table 1. Cost assumptions of the model

Item	Assumption
CAPEX	
Implementing MPTCP (datacenter side)	500 000 €
Implementing MPTCP (end user side)	1 200 000 €
Proxy purchasing	5 000 €
OPEX	
Initial maintenance costs of a proxy	1 500 €
Energy costs of a proxy (per year)	2 500 €

4.2 Benefits

Analyzing the benefits for the content provider is not as straightforward as analyzing the costs. This is because end users may not be willing to pay directly for MPTCP functionality on their mobile device. Therefore, a content provider should not expect any revenue to come from simply providing such an update to the end users. However, content users experience increased throughput and resilience when using MPTCP for accessing content. These benefits translate into shorter download times and more robust connectivity, especially when mobile. Thus, the fundamental assumption of the model is that deploying MPTCP will increase demand, *i.e.*, the number of content accesses.

On the other hand, using two radio connections simultaneously may also increase battery consumption and thus decrease the number of downloads. However, in light of [13], the effect of battery consumption on the number of downloads should not be significantly negative. Until MPTCP has been implemented and tested in real deployments, the benefits of MPTCP cannot be verified.

Although MPTCP will likely increase the quality of experience (QoE), the question is whether this will lead to additional content downloads. Ref. [14] indicates that an increase in bandwidth has a positive effect on time spent online. This does not mean that end users would spend all this additional time downloading content, but the assumption is that a fraction of time is used for that purpose.

However, more user activity does not necessarily translate into more profits. End users may be charged according to different pricing models that also affect the usage and consequently provider revenue. Three different pricing models exist in the content delivery business. Flat rate pricing encourages a user to access content as much as he desires but the increasing usage increases only costs – not revenue. In transaction-based pricing, a content provider gets additional revenue from each additional download, but users consider more carefully what content they access. In the third case, a content provider offers free services to customers, and gets revenue from somewhere else, typically through advertisements. Although advertisers are charged according to the service usage, the revenue cannot be easily attributed to downloads.

The model considers only application downloads, because they are charged per transaction and their revenue can therefore be easily verified. Another argument for choosing applications is that users most probably download them straight to mobile devices. Music, on the other hand, is presumably often first downloaded to a regular computer and from there transferred to a mobile device.

Browsing Apple's [15] and Nokia's [16] offerings reveals that only a small fraction of applications are being charged for. Therefore, a parameter describing the amount of chargeable downloads is included in the model. Ref. [17] suggests 25 % of downloaded applications from Apple's App Store require payment. This percentage seems quite substantial and MPTCP will unlikely increase the number of chargeable downloads by the same ratio. A user's buying decision is affected by many factors and the impact of QoE is seen to be a matter of secondary importance. Therefore, the proportion of chargeable downloads of all increased downloads is 1 % in the model.

Ref. [18] summarizes the average application prices in different application stores and reveals that the average price of all chargeable applications is 3.62 € for Apple and 3.47 € for the Nokia. An average price of 3.50 € has been used in the model. The entire price of a sold item does not go to the content provider. Ref. [17] suggests that the application developers get a 70 % cut of revenue while the provider gets 30 %. Table 2 summarizes the assumptions concerning the revenue.

Table 2. Assumptions concerning the revenue

Parameter	Value
Percentage of downloads chargeable	1 %
Average price of a chargeable application	3.50 €
Provider's profit of an application price	30 %

Deploying MPTCP can benefit a content provider also in ways that are not directly related to an increased usage of content. The increased throughput and resilience of MPTCP connections may increase device sales, especially if MPTCP becomes more widely adopted. Additionally, if MPTCP can optimize the battery consumption required for Internet connectivity, the device manufacturer can use the conserved energy to improve the quality of user experience, *e.g.*, through extended battery life or new attractive features. However, these benefits are not considered in the model.

5 Techno-Economic Model

This section proposes a simplified, quantitative, techno-economic model that integrates the different incremental costs and benefits to estimate the profitability of MPTCP for a content provider. MPTCP is only used for content downloads; all other communication between the end user and the content provider, such as signing in to the service, uses regular TCP. The basic assumptions of the model have been presented in Section 4, but this section explains how the traffic volumes in the content

provider's network have been calculated. The model is supposed to be general and therefore does not represent the business case of any specific content provider.

The model assumes a study period of five years and an interest rate of 10 %, both of which are often used in techno-economic modeling. Both the network dimensioning and revenue calculations are based on the estimated number of downloads per day from the content service. Daily download figures of Nokia's Ovi Store (1.5 million) [19] and Apple's App Store (6.6 million) [20] from the first half of 2010 have been used as references. The initial number of downloads per day is an average of these two (4 million).

The growth of application stores has been accelerating heavily. For example, [21] shows that the number of downloads from Apple's App Store has been increasing by approximately 500 % per year. An explanation of such a heavy growth rate is that sales of iPhone devices have also been growing considerably. Once the device sales have been saturated the growth rate of application downloads will be much smaller. The model in question presumes an annual linear growth of 25 %.

Like any other new innovation, MPTCP is adopted gradually. Ref. [22] shows that the diffusion of new features on mobile handsets follows an "s-shaped" curve. The authors in [22] have chosen Gompertz model [23] to analyze the diffusion patterns. The same model is adapted here to correspond to assumed progress of MPTCP diffusion. The suggested model has two parameters that define the shape of the curve. The displacement parameter b has been chosen to be 0.8 and the growth rate parameter c is set to 5. After five years, the penetration of MPTCP increases up to 90 %. The diffusion curve models relatively fast penetration, which is reasonable because as an end-to-end provider the OS vendor can significantly increase the penetration rate of MPTCP.

The required number of proxy servers can be calculated when the maximum number of simultaneous MPTCP sessions per second and the capacity of a proxy are known. The maximum number of simultaneous sessions has been deduced from the average number of downloads per second using a dimensioning rule that the average traffic load is 40 % of the maximum load. Based on different statistics of traffic amounts on different Internet links, 40 % appears to be a valid estimate. The traffic is assumed to spread equally over all geographical areas. The average number of simultaneous sessions is calculated using Equation 1.

$$Avg\ number\ of\ sessions/s = \frac{Number\ of\ downloads\ per\ day * Avg\ duration\ of\ a\ session}{Seconds\ per\ day} \quad (1)$$

The average duration of a session can be further deduced from the average size of an application and the average throughput of a session. Ref. [15] is used to estimate the average size of an application (10 MB) and measurements with a mobile phone supporting 3G and WLAN to define the average throughput of a session (4000 Kb/s).

5.1 Results and Sensitivity Analysis

The model reveals that in order for the NPV of the investment to be zero, MPTCP should increase downloads by approximately 3.5 % per MPTCP user. If a content user

currently downloads one chargeable application per month, he spends an average of 42 € per year on downloading applications. Consequently, MPTCP should make the user spend less than an additional 2 € on applications per year.

The number of daily application downloads reaches 12.2 million during the study period. This traffic volume can be managed with two proxies and thus both CAPEX and OPEX of proxies remain small. In addition, the model shows that software implementation costs dominate CAPEX instead of hardware costs.

To have a perception how MPTCP affects a provider's revenue, the increasing demand of applications is illustrated in Fig. 3. The figure shows how the provider's NPV varies as a function of MPTCP impact on downloads. The solid line shows the base scenario with the values presented above.




Fig. 3. NPV of the investment as a function of MPTCP impact on downloads

Because several assumptions of the model come with significant uncertainty, a sensitivity analysis is performed. The analysis reveals that the most critical parameter in the model is the percentage of the chargeable items among additional downloads. An intriguing fact is that this is also the most uncertain parameter in the model. To illustrate the effect of this parameter, Fig. 3 also shows how NPV is affected when the percentage of chargeable items varies.

Five years may be a relatively long period for an investment like MPTCP, because it is only one software update among many others. With a payback time of 3 years, the basic scenario approximately results in a 12 % increase in downloads due to MPTCP. If we assume the same behavior from the user as earlier, each MPTCP user should spend around 5 € more on applications per year.

The sensitivity analysis suggests that the provider's revenue may be much bigger than shown in Fig. 3. The reason is that the increased demand of applications affects the revenue more strongly than the costs of the provider. In addition, the other sources of revenue have not been considered in the model, such as increased sales of music pieces or mobile devices, which can further increase the provider's revenue.

6 Discussion

Although the presented techno-economic model suggests that MPTCP is a profitable investment for content providers, it relies heavily on the assumption that increased bandwidth increases content downloads. Even if this would be the case with free content, the assumption can be questioned with paid content. The selection of content is naturally the most important factor but also the price, ease of paying and general usability of the service affect the purchase decision probably much more than download speed.

The paper also assumes that end users are always interested in maximizing throughput. This is reasonable with application downloads but, for example, for video streaming, the bandwidth only needs to exceed the video peak rate, and faster download rates are not very useful to any stakeholder. Thus, depending on the service, users may actually be more interested in, for example, minimizing Internet connectivity costs or energy consumption. This paper assumes flat rate pricing in all access connections, but if, for example, a 3G connection is usage- or block-priced, end users will favor flat rate priced WLAN connections over 3G, if they are available.

The access pricing model also affects the ISPs' stake in MPTCP. Using multiple access paths simultaneously can reduce the load per single path, which can either be a desired or an undesired effect. If both access connections are purchased from the same ISP, the ISP will benefit if the traffic moves from a more congested access link to a less congested one. The pricing model dictates the impact of load: with flat rate pricing, the operator's profit increases when traffic decreases, and with usage-based pricing, more traffic increases revenues. Because ISPs are in the middle of an end-to-end MPTCP communication, they have also the possibility of interrupting or blocking MPTCP traffic, if it seems harmful for their business.

Finally, it should be noted that the transport layer is not the only option for implementing benefits of MPTCP. If different access connections have highly asymmetrical throughput, the usefulness of using multiple connections at the same time can be questioned. A solution that would provide prioritizing of access connections and would enable smooth handover can offer similar benefits without the need to change the networking stack. Furthermore, new radio access technologies will increase the available bandwidth but the bottleneck can also be located elsewhere than in the access network. The availability of multiple paths affects the usefulness of MPTCP, because multiple access paths are ineffective if all share the bottleneck link.

MPTCP-like functionality could also be implemented directly by applications. A multi-server HTTP mechanism [24] increases the efficiency of network resource utilization and Adobe's proprietary Real-Time Media Flow Protocol (RTMFP) [25] provides a real-time, peer-to-peer bulk transport protocol over UDP. RTMFP is provided as part of Adobe's Flash plugin for web browsers, which currently makes it unsuitable for general-purpose applications, but significantly eases deployment and adoption. MPTCP is a more comprehensive solution that requires operating system support but is available to all applications on the host. This tradeoff between generality and deployability poses an interesting question in which layer the benefits of resource pooling principle should actually be trialed and implemented.

7 Conclusion

This paper has analyzed the profitability of an MPTCP deployment scenario where a single provider controls both the mobile devices and content servers, and implements and adopts MPTCP. The model suggests that even a relatively small increase in download numbers due to MPTCP could make the business case profitable for such a provider. Although the results of the model are very uncertain, this is a positive signal for MPTCP. All the example providers of the scenario are corporations with wide customer bases and getting these players “on board” to implement MPTCP is hence highly important.

Although the model is simple, it gives a good understanding of different cost and revenue components incurred by MPTCP and serves as a good base for further research. The model should be extended to also cover other types of content and pricing schemes. We chose to currently limit the scenario to only a single provider, because the linkage between costs and revenue can be easily investigated. Taking into account multiple content providers as well as other deployment scenarios are worth of further study.

A fundamental question regarding the current and potential future models is whether MPTCP really can improve the quality of experience for Internet services. Although the benefits can be verified, it is unclear whether better service quality will lead to an increase in service use. Actual measurements as well as studies on user behavior would be needed to answer these questions.

Acknowledgments. This paper was partly funded by *Trilogy*, a research project supported by the European Commission under its Seventh Framework Program (INFSOICT-216372). The research is also linked to the European COST IS 0605 / Econ@Tel project.

References

1. Wischik, D., Handley, M., Bagnulo Braun, M.: The Resource Pooling Principle. ACM SIGCOMM CCR 38(5), pp. 47–52 (2008)
2. Wischik, D., Raiciu, C., Handley, M.: Balancing Resource Pooling and Equipose in Multipath Transport. Submitted to ACM SIGCOMM (2010)
3. Ford, A., Raiciu, C., Barre, S., Iyengar, J.: Architectural Guidelines for Multipath TCP development. IETF Internet-Draft (work in progress), draft-ietf-mptcp-architecture-00 (2010)
4. Kostopoulos, A., Warma H., Levä, T., Heinrich, B., Ford, A., Eggert, L.: Towards Multipath TCP Adoption: Challenges and Opportunities. 6th Euro-NF Conference on Next Generation Internet, Paris, France (2010)
5. Levä, T., Warma, H., Ford, A., Kostopoulos, A., Heinrich, B., Widera, R., Eardley, P.: Business Aspects of Multipath TCP Adoption. Future Internet Assembly Book (2010)
6. Balakrishnan, H., Seshan, S., Amir, E., Katz, R. H.: Improving TCP/IP Performance over Wireless Networks. 1st Annual International Conference on Mobile Computing and Networking, Berkeley, USA (1995)

7. Henderson, T.R., Katz, R.H.: Transport protocols for Internet-compatible satellite networks. *IEEE Journal on Selected Areas in Communications*, 17(2), pp. 326–344 (1999)
8. Border, J., Kojo, M., Griner, J., Montenegro, G., Shelby, Z.: Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations. IETF RFC 3135 (2001)
9. Katsianis, D., Welling, I., Ylönen, M., Varoutas, D., Spicopoulos, T., Elnegaard, N. K., Olsen, B. T. & Burdry L.: The financial perspective of the mobile networks in Europe. *IEEE Personal communications*, 8(6), pp. 58-64. (2001)
10. Smura, T., Kiiski, A. & Hämmäinen, H.: Virtual Operators in the mobile industry: a techno-economic analysis. *NETNOMICS*, 8(1-2), pp. 25-48. (2008)
11. Dodownload.com: Faststream IQ Proxy server. Available at <http://www.dodownload.com/servers/other+servers/faststream+iq+reverse+proxy.html>, accessed 2010-04-27
12. Electronics Cooling: In the data center, power and cooling costs more than the it equipment it supports. Available at <http://www.electronics-cooling.com/2007/02/in-the-data-center-power-and-cooling-costs-more-than-the-it-equipment-it-supports/>, accessed 2010-06-10
13. Nurminen J.K.: Parallel Connections and their Effect on the Battery Consumption of a mobile phone. 7th IEEE Consumer Communications & Networking Conference, Las Vegas, Nevada (2010)
14. Anderson, B.: The Social Impact of Broadband Household Internet Access. *Information, Communication & Society*, 11:1, pp. 5–24 (2008)
15. Apple App Store. Available at <http://www.apple.com/iphone/apps-for-iphone/>, accessed 2010-04-14
16. Nokia Ovi Store. Available at <http://store.ovi.com/>, accessed 2010-04-14
17. GigaOm: Apple App Store Economy. Available at <http://gigaom.com/2010/01/12/the-apple-app-store-economy/>, accessed 2010-04-02
18. Distimo report: Apple App Store, BlackBerry App World, Google Android Market Nokia Ovi Store & Windows Marketplace for Mobile. Available at <http://www.distimo.com/uploads/reports/Distimo%20Report%20-%20December%202009.pdf>, accessed 2010-06-11
19. Symbian Freak: Ovi Store Statistics: 22 Application downloads per second. Available at http://www.symbian-freak.com/news/010/03/ovi_store_stats.htm, accessed 2010-04-12
20. MediaMemo: Apple's Apps Flying Off the Virtual Shelves: 6.6 Million downloads per day. Available at <http://mediamemo.allthingsd.com/20090928/apples-apps-flying-off-the-virtual-shelves-6-6-million-downloads-per-day/>, accessed 2010-04-12
21. Macrumors.com: Apple announces 3 billion App Store Downloads. Available at <http://www.macrumors.com/2010/01/05/apple-announces-3-billion-app-store-downloads/>, accessed 2010-04-26
22. Kivi, A., Smura, T. & Töyli, J.: Diffusion of Mobile Handset Features in Finland. 8th International Conference on Mobile Business, Dalian, China. (2009)
23. Gompertz, B.: 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 Transaction Series I*, 115, pp. 513–583 (1825)
24. Ford, A. & Handley, M.: HTTP Extensions for Simultaneous Download from Multiple Mirrors. IETF Internet-Draft (work in progress), draft-ford-http-multi-server-00 (2009)
25. Kaufman, M.: RTMFP Overview. IETF-77, Anaheim, CA, USA. Available at <http://www.ietf.org/proceedings/10mar/slides/tsvarea-1.pdf>, accessed 2010-04-12