Helsinki University of Technology Department of Communications and Networking

Teknillinen korkeakoulu Tietoliikenne- ja tietoverkkotekniikan laitos

Espoo 2009 Report 1/2009

TECHNO-ECONOMIC ANALYSIS OF INTERNET ARCHITECTURE EVOLUTION

Heikki Hämmäinen, Pekka Nikander, Jörg Ott, Mikko Heikkinen

Helsinki University of Technology
Faculty of Electronics, Communications and Automation
Department of Communications and Networking

Teknillinen korkeakoulu
Elektroniikan, tietoliikenteen ja automaation tiedekunta
Tietoliikenne- ja tietoverkkotekniikan laitos

Distribution:
Helsinki University of Technology
Department of Communications and Networking
P.O. Box 3000
FIN-02015 HUT
Tel. +358-9-451 5300
Fax +358-9-451 2474

ISBN 978-951-22-9992-8 (pdf)
ISSN 1797-4798 (pdf)

Espoo 2009

Foreword

Heikki Hämmäinen, Pekka Nikander, Jörg Ott, Mikko Heikkinen

This report is a collection of best student papers from a postgraduate level seminar, “Techno-Economic Analysis of Internet Architecture Evolution”, held at the Helsinki University of Technology. The idea for the topic came up from discussions driven by Dr. Pekka Nikander with Professors Heikki Hämmäinen and Jörg Ott. The key observation was that the fundamental price ratio trends in the context of Internet are poorly studied and published. Still the slow changes in fundamental price ratios are working indisputably, and gradually influencing the power positions of competing architectures and players. Therefore it seems that demand exists for this kind of studies while their supply is low.

One key enabler of the success of Internet has been the rapid technology performance advances and the resulting steady decrease of cost of system components. However, the relative cost of system components is changing over time -- in other words, the cost of some elements decreases more rapidly than others. For example, the relative cost between processing, transmission, and storage capacities is changing over time. Similar price ratio changes are also observable on higher level when different architectures are compared (for example wireless versus wired access price ratio) or when different usage scenarios are analyzed (for example client-server model versus peer-to-peer model of data exchange).

The ratio changes have fundamental impacts on the evolution of the technical and business architecture of the Internet. As price ratios change, some technologies, applications and systems substitute others possibly resulting in disruptions in value chains, business models, and industry structure. Therefore, analysis of the trends in price ratios can provide valuable insights to the possible future evolution scenarios of the Internet architecture.

Many interesting issues resulting from price ratio changes are discussed in the following papers. One such topic is the emergence of a memory wall (i.e. processing rates growing faster than memory access times) which is having and will have profound impact on computer architectures. The fact that memory storage prices are decreasing more rapidly than transmission prices is leading to interesting substitutes for information distribution such as Internet caching and utilizing content distribution networks or even preloading content when buying new external hard drives. Another observation is that a key enabler for packet switching, and eventually the entire Internet, was the fact that at one point of time the price of computing became lower than the price of transmission.

Another interesting observation is that the different elements discussed in the papers are often interrelated. For example the concept of caches can be seen both on the level of a single computer, and also on the level of the Internet with content distribution networks which can be thought of as geographical caches. Furthermore, a network of computers, or even the entire Internet, could be seen as an enormous system enabling distributed and cloud computing, where some of the same analogies as for a single computer apply, especially in terms of fundamental concepts like information processing power, transportation capacity and storage capacity.

Traditional techno-economics would assume cost-revenue calculations for chosen services, technologies, and players over predefined time periods. The papers of this seminar remain far from that level of financial detail. They rather focus on a certain price ratio change and try to draft the related systemic forces. Some papers succeed to describe the forces as system dynamic diagrams which turned out to be a useful method for documenting and communicating the related systemic complexities.

This report contains seven papers with topics ranging from basic hardware technologies to different aspects of communication systems and their applications. We have divided the papers into three categories:

Fundamental Concepts of Price Ratios

Storage vs. Processing by Somaya Arianfar
Arianfar investigates the different cost and technology trends in storage and processing, and analyzes the feasibility of early choices and rules of thumb in the current situation. She employs different metrics to enable a fair comparison and discusses related new business cases.

Transmission vs. Processing by András Zahemszky
Zahemszky studies the trade-off between transmission and processing in communication networks. He argues that the change in the price ratio between the cost of transmission lines and the cost of routers enabled packet switching, but memory access speed may become a bottleneck in future IP routers with 40 Gbps or faster interfaces.

Storage vs. Transmission by Matti Peltola
Peltola reviews the influence of storage and transmission price ratio development on the architecture of the Internet. He shows that the prices of both memory elements and transmission devices have decreased continuously and discusses the implications of caching in the future Internet architecture.

Price Ratios in Communication Systems

Transit and Access Prices – Evolution and Interdependencies by Tapio Levä
Levä studies how transit prices affect access prices by conducting an analysis how they have developed during the past eight years. He develops a model on access cost and price formation. He also investigates the effects of traffic growth on transit cost.

Price Ratio Evolution of Cellular Data and WiFi by Thomas Casey and Juuso Karikoski
Casey and Karikoski model the price ratio evolution of cellular data and WiFi both in the past and future. They use system dynamics to form an overview of the evolution and the critical issues affecting it.

CAPEX vs. OPEX from the Perspective of a Mobile Operator by Heidi-Maria Rissanen
Rissanen describes the different cost types forming the total cost of ownership of a mobile operator. She discusses the mechanisms used to decrease both capital and operational expenditures, including heterogeneous networks, femtocells and site sharing.

Price Ratio Evolution Implications in Usage Scenarios

System Dynamics of Client-Server and Peer-to-Peer Content Distribution by Mikko Heikkinen and Mikko Särelä
Heikkinen and Särelä study price ratio evolution in the context of client-server and peer-to-peer content distribution paradigms by developing a system dynamics model. The model preliminary confirms that client-server distribution is beneficial to ISPs and peer-to-peer distribution to content providers.

Storage vs. Processing

Somaya Arianfar

sarianfa@cc.hut.fi

Abstract

Our current computer systems and in turn communication systems, are designed economically feasible by the cost and technology ratios of early system design days. Technology and cost trends has changed a lot during the years. The goal of this paper is to look at the different cost and technology trends in storage and processing and analysis the feasibility of early choices and rules of thumb in current situation. We also keep in mind the effect of changes in the amount of data and data requests in digital world. We look at different metrics important for a fair comparison; having an eye on historical trend of the technology and cost, we try to figure out new business cases that can emerge the technology or can be emerged by technology.

1. Introduction

The amount of data in the world is increasing so quickly and it is now measured by Exabytes (10^{18} bytes)[5,6]. Huge part of these data are stored digitally. Handling this much digital data requires computing platforms that at least can deal with terabyte-level (1,024 gigabytes) workloads [7]. The growth of the data is a trigger for new technological innovations and therefore there is huge business market on that.

Computing industry is emerged to look at different ways of handling this huge amount of data. For instance Intel is practicing the Era of Tera [7], when people need teraflops (a trillion floating point operations per second) of computing power, terabits (a trillion bits) of communications bandwidth, and terabytes (1,024 gigabytes) of data storage to handle the information all around them.

The industry is seeing a convergence of digital technologies around consumer electronics, communications and computing platforms. Memories, disks and storage are gaining more capacity and getting cheaper all the time. Faster processors are introducing to the market everyday. Considering the possibility of having lots of data one may think having just bigger storage systems may help. However, processing power may not be able to handle bigger storages. History demonstrates that computation capabilities do not always advance linearly; rather there have been several "leaps" in computing capabilities that have had profound impacts.

In this paper we discuss the technology improvement trends. Then we discuss the price trend of different parts of the storage and processing. At the end we discuss how both of these trends can affect the market, design of the new products and systems as well as administration decisions.

2. Background

We live in a constantly changing world. Computer systems technology is not an exception. Price and Performance of different components of the system changes all the time. If everything changes in the same way, then actually no change is happening. If changes in price and performance of some components are much faster/slower than others, then that is a real thing, which will affect the design and administration choices and most importantly the future market.

Some factors are not changing much, such as cost of the people or speed of the light; however some other factors such as capacity are changing a lot. Therefore, ratios are changing and as a matter of fact one can see that today, latency lags bandwidth[14] and bandwidth lags capacity. This ratio changes affect the way one, design or administrate a system. To make the case more obvious we use this example, where pipelining (prefetch) become more and more popular as result of latency lagging bandwidth, because prefetching can somehow hide the latency.

From the economical aspect, changes in techno/economical ratio introduces new business cases which are interesting for us. The ratio changes sometimes require a serious business case for investment in new technologies and policies. For example, huge amount of digital data requires more and more investment on capacity or other way around, huge amount of capacity requires creating use cases for the market. As another example, exponential growth in capacity and the lag of bandwidth and latency has introduced the new flash memory business to the market, or caching and backup policies has changed.

We look at the more detailed technology and cost changes and their emerging business cases further on in the rest of this paper.

3. Metrics

Standard storage metric, includes both the capacity and performance, though the pricing mostly follows only the capacity metric. Traditionally secondary storage devices with capability of providing higher capacity are cheaper but with the trade off of performance degradation. The traditional performance metric for storage systems are access time (Latency) and transfer rate (throughput/bandwidth). Table I shows an approximation of the current capacity and performance attributes of 3 different widely used storage devices.

Table 1: Different storage attributes and their approximate current values

Type	Capacity	Base price	Latency	Throughput
RAM	4 GB	100$/GB	1-100ns	1-10GB/s
Disk	700GB	500$/TB	5-15ms	50MB/s
Tape	400GB	300$/TB	30s	50MB/s

Considering the changing world, performance criteria are introducing new storage metrics to the market. These new storage metrics include three categories [8]. Kaps: How many kilobyte objects served per second. It is actually a file server or transaction processing metric. Maps: How many megabyte objects served per sec which is a multi-media metric. SCAN: How long to scan all the data which is a data mining and utility metric. And finally obviously Kaps/$, Maps/$, TBscan/$ are important factors for final decisions. Later on we discuss a little bit about Kaps and its introduction to the storage world, which in turn will give a clue about Mpas and SCAN as well.

Though in this paper we look at the processing mostly from the storage performance metrics point of view, but we also discuss about the processors in the Era of Tera. Important factors for the processors are latency, MIPS (bandwidth) and MIPS/$.

Of course, the overall performance and cost of a computing system is related to both storage and processing metrics and their relationships via I/O subsystem.

4. Technology Trends

Processor speed and storage capacity has moved forward rapidly in recent years. Though, comparable advances have not occurred necessarily in other parts of the systems such as I/O subsystems. Historically most of the metrics in storage and processing follow the Moore’s Law or at least are compared with that. Original Moore’s law predicts a 18-month periodic doubling in the number of transistors per chip, due to scaling and in part to larger chips; recently, the doubling period for memories has been 22-24 months [15]. By original law, 18-month periodic doubling means for example memories get four times larger each three years, or about 100x per decade. Moore’s law originally applied only to random access memory (RAM). It has been generalized to apply to microprocessors and to disk storage capacity. It is also good to note that, nothing forever can follow exponential growth of Moore's law, before some fundamental barriers limit the actual growth rate; but at least in most of the cases the historical trend is comparable by Moore's law and it is logical to consider when the Moore's law is not applicable anymore there would be some advances in other criteria which will make the proceeding possible.

Looking at disk capacity, indeed, it has been improving by leaps and bounds; it has improved 100 fold over the last decade. Fig. 1 [9] shows the areal density[1] improvement for hard disk drives since 1956. The trends has been changing all the time, today’s CGR (compound growth rate) is essentially 100 percent or doubling every year. Possibly the growth continues at this rate, maybe a little slower, anyhow, it is growing faster than Moore’s law which is originally almost 60% improvement per year.

Processor speed also has grown quite fast, though it is going to slow down a lttle bit. As can be seen in Fig. 2 [10], in 2006, performance is a factor of three below the traditional doubling every 18 months that has occurred between 1986 and 2002. Based on the [3], doubling of uniprocessor performance may now take 5 years, which nicely shows the fundamental barriers affecting the growth rate by Moore's law. That's the place where in case of processor, market starts to be prepared for multi core processors more and more[2]; of course it has its own problems which affects the overall processor speed but we are not going to discuss it here.

Increasing processor speed also affects the overall system speed, at least during last decade it was so. The last decades doubling of performance was because of the independence of the system performance from the limitations of the external attributes, such as I/O speed. But now the performance is somehow in a threshold that can be reached without being limited by external attributes. So any improvement in the performance now is quite dependent on external factors such as the memory wall and power wall limitations [3], mentioned briefly in background section.

Latency and bandwidth improvement trend which affects the processor and overall systems performance are shown in Fig. 3 [14]. In a comparison between bandwidth and latency, it shows that at the time latency improved 10x, bandwidth improved about 100x-1000x. Looking at the same historical trend, one should realize

1 A traditional measurement for disk drives, determines capacity, internal (media) data rate, and ultimately price per unit of capacity

2 Thanks to Risto Mononen for pointing this out

the ratio changes are not something be overcome but should be adapted for taking advantage.

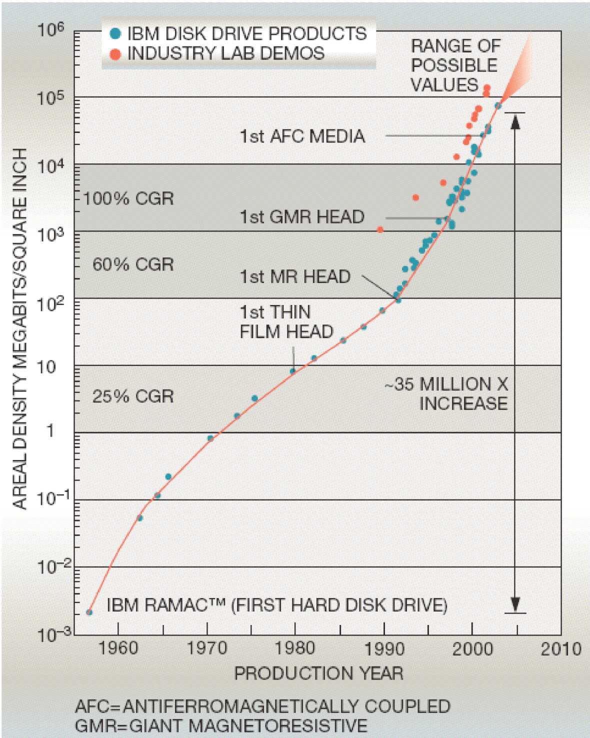

Fig. 1: Hard disk drive areal density trend

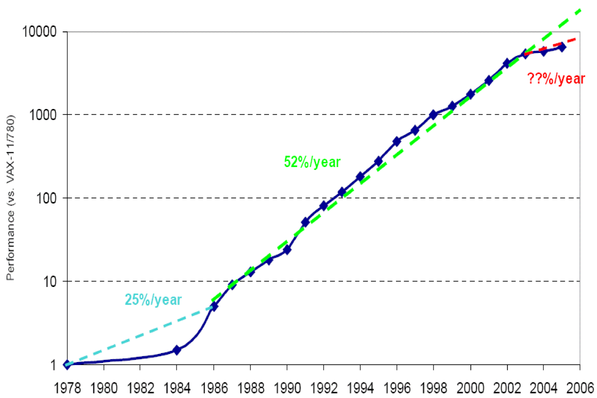

Fig. 2: Log-log plot of processor performance improvement trend

Comparing Fig. 1 and Fig. 3, plus considering other data [8], [14] also shows for example for disk and memory at same time the capacity has improved much faster than both of these performance metrics (bandwidth and latency). What is interesting for us from technology point of view is the effect of this different trends on each other and whole system. Consider this example: In 1990, disks offered 50 Kaps (kilobyte accesses per second) to 1GB of data, and 5-minute disk scan times. In 2000, disks could only offer 120 Kaps to 80 GB of data with a 45-minute scan times. This is 1 Kaps per 20MB in 1990 vs. 1 Kaps per 500MB in 2000. [8] The trend continues in 2006, where disks offered 150 Kaps (kilobyte accesses per second) to 750GB of data, and 5 hours disk scan times. This is 1 Kaps per 10GB in 2006. Overall it means data in modern disks is getting colder and colder, and backup/restore takes longer time.

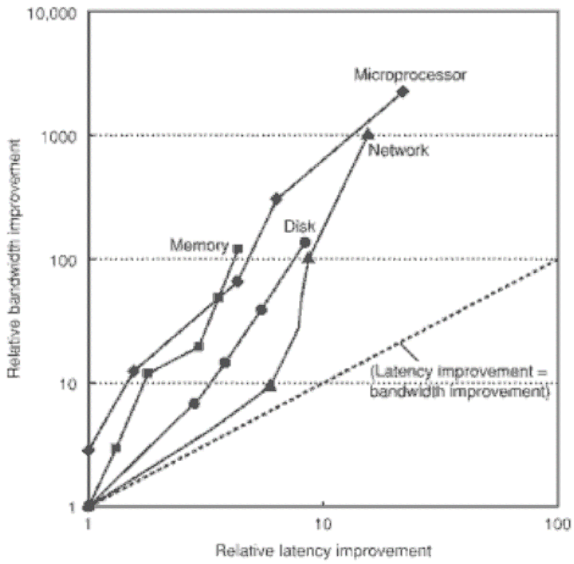

Fig. 3: Log-log plot of Bandwidth vs. Latency improvement trend

5. Cost Trends

The price trends are expected to go hand in hand with technology changes. As shown in Fig. 4 [9], a dramatic drop in cost-per-megabyte in storage prices is evident. Following the trend in hard disk price drop and looking at Fig. 1, one can clearly see the dependency between technology and price trends. Since areal density growth within the last 10 to 15 years were 60 percent to 100 percent per year, it would be expected that price declines would somehow follow the same trend. Because of the sensitivity of the market to different attributes, price reduction though follows less dramatic changes in its curve, on average 37 percent to 50 percent, respectively to technology trend per year [9]. The trends depicted in Fig. 4 show this rate of the price declines corresponding to the areal density CGR increases shown in Fig. 1. It is expected that any future trend of areal density increase would also affect the rate of price decline for storage systems. One had to pay almost 200$ to get one MB of hard disk in 1980 but around 1998 with same amount of money it was possible to get one GB of hard disk and now it is possible to buy almost 200 GB of hard disk with same amount of money. $/Kaps etc has improved 100x.

One nice thing to observe in Fig. 4 is the fact that though different storage prices are dropping there is no cutting edge in changing ratios. Another observation from the figure is the astonishing fact of having cheaper digital storage than traditional paper and film storage space. It can introduce currently emerging market of using digital storage instead of paper, for instance e-invoicing companies can be emerged instead of paper invoicing.

The same price drop rate for semiconductor based storages, also can be seen in Fig. 4. The cost/MB of semiconductor memories also declines with time, about 100x per decade. Since disk and RAM have a 1:100 price ratio, this price decline suggests that whatever you pay for hard disk now, in a decade would be economically possible to pay for same amount of RAM.

Following the same price declining trend, Fig. 5 [2] shows the price drop rate for microprocessors.

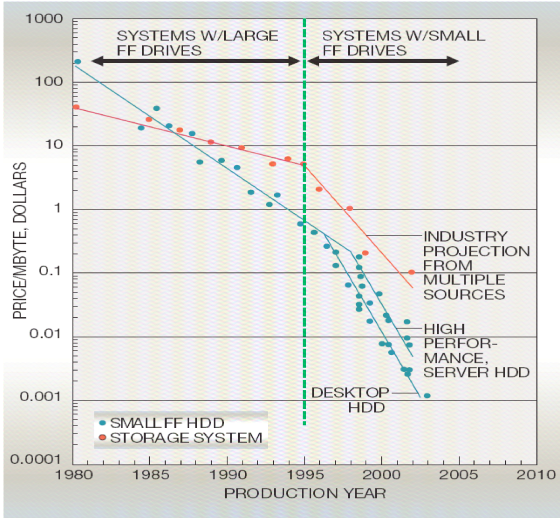

Fig. 4: Storage price trend based on \$/MB

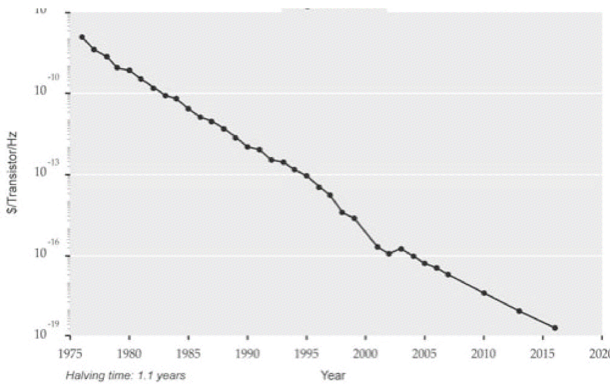

Fig. 5: Microprocessor cost per transistor cycle

6. Discussion

The main challenge for storage systems is to keep their capabilities in a cost-effective way to service the storage and access requests. For instance, in many cases, the reduction in execution time is achieved at the expense of an increase in cost per solution, but sometimes it is possible with changing the pre assumptions.

In general, most of the previously mentioned trends for storage and performance are positive ones: with each new generation of fabrication technology, processor capability, memory bandwidth, and memory latency and their cost all improve. However, the most important consequences of these technology trends are the differences between them. When one metric changes at a different rate than another, it requires rethinking the assumptions behind the system design and administration and it may introduce new market opportunities.

From another point of view, it is also important to note that the willingness of the market also affect the development of different metrics and ratio changes.

We discuss this issues more in detail in the rest of this section. We first start to look at the different kind of storages and their capabilities for current and emerging market, then we look at the performance issues that come with the ratio changes and business cases that they emerge.

6.1. Market emergence: storage hierarchy

One should consider all factors when choosing the storage type. On the other hand, storage producers also need to have a look at the market requirement when they want to invest on new kinds of storage. It is good to keep in mind that hardware is driven by mass market economics.

The application partially determines the success of a memory technology. Over the past decade, there has been a major shift in the use of memories, away from enterprise and desktop systems towards more diverse mobile applications. However, the memory hierarchy model has not been updated to reflect these changes.

Normally storage devices are organized in hierarchies with the fastest performance at the top, and the lowest cost per bit at the bottom. The model also represents the data flow in a system: from faster storage to slower, cheaper storage. Traditionally this hierarchy starts from DRAM at the top, HDD at next step and tape at the bottom. This view also represents the design of enterprise systems, and to a lesser extent, desktop systems. Emerging market of for instance, mobile devices and mobile applications can not be directly presented by this hierarchy, though considering huge market of these devices, usage of new kind of storage devices also should be considered in the hierarchy. It is well known that the 2.5 billion cell phones in use worldwide significantly outnumber the 500 million wired PC's on the Internet. As smart phones and PDA's with high-speed cellular and WiFi service proliferate, the number of Internet transactions from mobile devices may be expected to surpass those from wired network PC's over the next 5-10 years. Since mobile devices have size and power consumption limitations which makes them different, they require revisiting the storage hierarchy for their special needs.

Given the high growth rate of the mobile segment, it is also likely that mobile will become the driver for many memory technologies. For instance the huge market opened with the need of less battery draining storage systems such as NAND flash. NAND flashes introduce nonvolatile kind of memory with low power consumption. NAND flash has faster read performance than a HDD, but significantly slower small block random write performance. Low write operation does not limit the choice of using flash in mobile hand held devices and digital cameras. This low write operation is sufficient for the applications of NAND flash, such as in digital cameras and audio devices, where the active duty cycle is low and the work load is mostly sequential. The pressure of this new market, has caused faster development for flash memories than any other kind of memories in recent years. The price decline is much faster than DRAM and it may even reach to the HDD price level in a few years.

Currently, from both base price and (read) performance point of view NAND flash memories locate between DRAM and HDD for enterprise systems and desktop servers. Though the low write operation of NAND flash and its low maximum capacity (currently 64GB) limits its use cases in these systems.

However, as electricity prices increase and power management has become a critical piece of today’s systems, using flash memories seems to be more and more economical, if their capacity and performance problems can be adapted.

6.1.1. New market: nano cells

The programmable metallization cell, or PMC, is a new form of non-volatile computer memory. PMC is one of a number of technologies that are being developed to replace the flash memory, providing a combination of longer lifetimes, lower power, and better memory density. It store a terabyte of information–more than most hard drives hold today. It is claimed that no other technology can deliver the orders-of-magnitude improvement in power, performance and cost that this kind of memory can.

One should wait and see the mass production of this memory and the improvement trend that it is going to bring. Though just introducing this kind of memory can show instability of flash memory future market. [1]

6.2. Technology emergence: storage performance

The interesting point to discuss here is, what are we going to do with this huge amount of storage which is both technically and economically available, are we going to change the way we use them, are they going to change the storage hierarchy in enterprise and desktop system, how we are going to fill them and backup them, do they open new market opportunities, what about increasing performance of micro processor and their dropping cost, are technicalities such as lagging latency from bandwidth and bandwidth from capacity going to be problematic in long run and etc.

6.2.1 HDD

Disks are getting infinitely large, the amount of the data also is getting doubled every year. This increasing amount of data is because of different reasons, one is the ability to economically store petabytes of data on disks in end points; another reason is the Internet and computational Grid that makes all these archives accessible to anyone anywhere, allowing the replication, creation, and recreation of more data.

But managing these data somehow can go out of control when the other properties of the storage systems, such as bandwidth and latency don’t grow with same trend. The MB/$ ratio in disks is rapidly going up, while the number of disk heads per MB is going down. [4] Disk capacity has been improving by leaps and bounds; it has improved 100 fold over the last decade. Disks spin three times faster now, but they are also 5 times smaller than they were 15 years ago, so the data rate has improved only 30 fold. [8]

So storing the data in disks is becoming cheaper and cheaper but accessing them is becoming more expensive. The ratio between disk capacity and disk accesses per second is increasing more than 10x per decade. Also, the capacity/bandwidth ratio is increasing by 10x per decade. These changes have two implications: (1) disk accesses become more precious; and (2) disk data must become cooler (have fewer accesses per byte stored).

One knows what is important is accessing and using the data, not just putting them somewhere, content on its own does not have money, what has money is pipelining and accessing the content, connectivity is more important than content. So from economical aspects one should try to gain from putting more resources and intelligence into computing and networking the content and data [12].

The question that comes to mind here is that, is it so that disk price should go down just based on capacity or bandwidth and latency also needs to affect the price. In this way managing the data can be handled in same proportion with data. Where Storage capacities increase 100x/decade and storage costs drop 100x per decade without considering the fact that storage throughput increases 10x/decade. Isn't it better to have new pricing attitude such as Kaps,Maps,SCAN/\$.

Another suggestion could be using faster devices than HDD. Since dropping rate of the storage devices are generally high, it should be economically feasible to use faster devices than HDD to handle the data. For instance, the cost/MB of RAM declines with time: about 100x per decade. Since disk and RAM have a 1:100 price ratio, this price decline suggests that during computation time what is economical to put on disk today will be economical to put in RAM in about 10 years.

However RAM is a volatile memory and to find a nonvolatile replacement for HDD market should invest on other kinds of nonvolatile storage which are faster than HDD but still cheap enough to be economically feasible. It is also to note that from Fig.6 [10] it is obvious that though both RAM Currently flash memory is the emerging technology in the market which tries to fulfill the nonvolatile replacement problems of HDD and fill the gap between the RAM and HDD speed. It is still more expensive than HDD and it can not provide same capacity so it is not a good replacement for HDD. However there are some business cases to use flash and HDD technology as a hybrid storage especially for laptops, which can help to fill the latency gap between RAM and HDD.

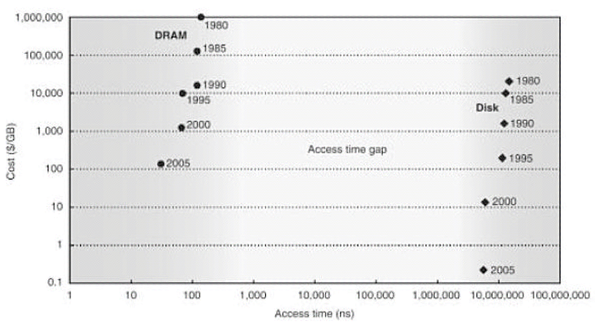

Fig.6 Latency gap between RAM and HDD

6.2.2 HDD vs. Tape

The problem with large disks is we just can’t fill them and we can’t backup/restore them easily. It’s Hard to Archive a Petabyte and it takes a LONG time to restore it. Traditionally tapes are used for backuping the data on hard disks. Tape capacities are expected to improve faster than tape speed, and access time is expected to stay about the same, making the access problem even more problematic: several days to scan a tape archive.

A striking thing about these storage cost calculations is that disk prices are approaching near line tape prices which means the ratios are becoming close to 1:1. By using RAID (mirroring or parity), administrators sacrifice disk storage capacity to protect against disk media failures. Administrators are discovering that you may be able to backup a terabyte to tape, but it takes a very long time to restore a terabyte from tape.

Considering the huge dropping rate of HDD prices, what comes to mind is storage should be used to cache data to save disk bandwidth, network bandwidth, and people’s time. We should be able to use the "extra disk space" creatively and manage cold storage. For instance redundancy on disks can increase both availability and backup.

On the other hand, storing petabytes of data on disk remains prohibitively expensive compared to tape in terms of operating cost (power and cooling requirements), durability, and purchasing cost. Tapes do not consume power or generate heat when not in use, two increasingly important factors for storage systems especially at data centers.

6.2.3. DRAM and Processor

The DRAM industry has dramatically lowered the price per gigabyte over the decades, to \$100 per gigabyte today from \$10,000,000 per gigabyte in 1980 [10]. However, as mentioned earlier, the number of processor cycles to access main memory has grown dramatically as well, from a few processor cycles in 1980 to hundreds today. While the memory wall is a major performance limiting factor, considering Fig. 3, memory latency is a bigger problem than memory bandwidth for memory wall issues.

6.2.4. Processor

In this section we discuss the processors. They are important factors for the systems in the Era of Tera, when huge amount of data, kept in huge cheap storage spaces define new trend of requests for processing by the processors. Keeping in mind the trends shown previously, we show a number of guiding principles in system design considerations from [3], that explains how things are changing in techno/economical world of computing, specifically in processors.

1) Old rule: Power is free, but transistors are expensive.

New rule is the Power wall: Power is expensive, but transistors are "free". That is, we can put more transistors on a chip than we have the power to turn on.

2) Old rule: If you worry about power, the only concern is dynamic power.

New rule: For desktops and servers, static power due to leakage can be 40% of total power.

3) Old rule: Performance improvements yield both lower latency and higher bandwidth.

New rule: Across many technologies, bandwidth improves by at least the square of the improvement in latency [14]

4) Old rule: Multiply is slow, but load and store is fast.

New rule is the Memory wall [11]: Load and store is slow, but multiply is fast. Modern microprocessors can take 200 clocks to access Dynamic Random Access Memory (DRAM), but even floating-point multiplies may take only four clock cycles.

As a result of these new rules and the slow down of the improvement in uniprocessor, multi core design is becoming popular and gaining their new place in the market. Multi core architecture tries to exploits increased feature-size and density. It increases functional units per chip (spatial efficiency), limits energy consumption per operation and constrains growth in processor complexity [16] . Of course it has its own problems which can be find in more detail in Supercomputing 2006(SC06) Conference panel discussion [16].

An important fact about the processors (uni/multi core) is, they are infinitely fast but always waiting for memory. We can identify three major issues that will help drive better choices of the future for processor and system design: compute versus communicate, latency versus bandwidth, and power.

In the trade-off between communicate and recompute/cache, the decisions shows the tendency toward choosing the latter. The dropping cost of storage versus communication strengthen this decision.

First, Moore's Law continues (even if not for transistors on uniprocessor model but for number of multi cores), so we will soon be able to put thousands of simple processors on a single, economical chip . Tera-scale computing stresses the platform architecture with memory bandwidth being a likely bottleneck to processor performance that presents unique challenges to CPU packaging. So investments in CPU packaging should go toward solving this bottleneck to gain more profit.

Because of the latency issues, data transfer should be as low as possible. For instance just transfer the changes instead of whole page. Or if there is so much communication with data then move the processing next to the data. Minimizing remote accesses. In the case where data is accessed by computational tasks that are spread over different processing elements, we need to optimize its placement so that communication is minimized. The mapping of computational tasks to processing elements must be performed in such a way that the elements are idle (waiting for data or synchronization) as little as possible.

Power management has become a critical piece of today's systems, new generation of hardware has increased its power demand. One need increasing amounts of custom or power-aware considerations for power-hungry parts of the system; and more sophisticated cooling management. [13] For instance, although smaller transistors require less power than larger ones, the number of transistors on a single processor die is rising faster than the amount at which power per transistor is falling. Consequently, each generation of processors requires more power, which means more cost. Perhaps as it is suggested in [13], the future figure of merit may no longer be the number of operations per second but instead the number of operations per second per watt. Fig. 7 [16] shows the trend of power consumption in the processors.

6.2.5. System Dynamics

Looking at different storage and processing factors affecting the price and technology trends of a computing system, we came up with a simple system dynamics diagram in Fig. 8.

The diagram explains affecting factors in system cost based on the capacity and processing attributes. It demonstrates major part of the discussion from the rest of this paper. From the upper right corner, it starts from the effect of increasing data on increasing capacity in storage systems, showing that the cheap cost of buying storage also affects the growth of data in a loop.

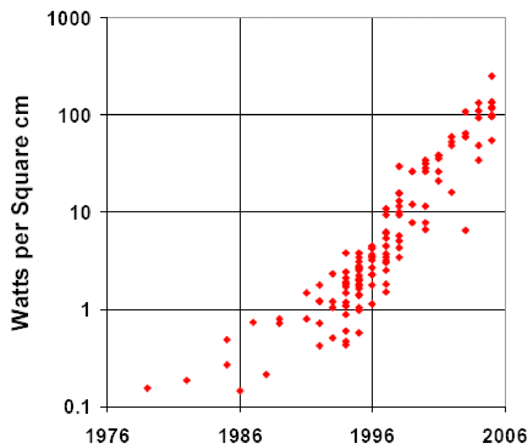

Fig. 7: Microprocessor power consumption

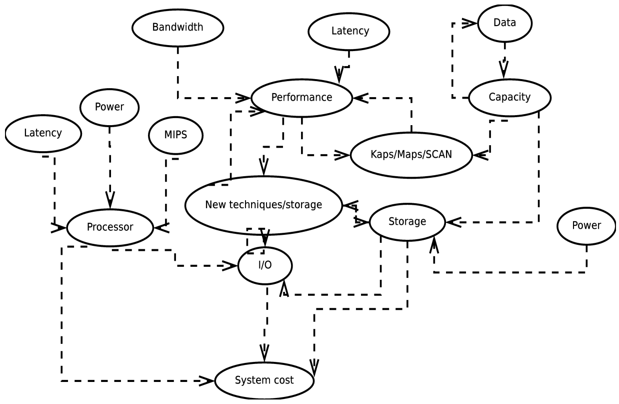

Fig. 8: System dynamics straw-man diagram

The diagram also shows bandwidth and latency are basic affecting factors in storage performance, which in interaction with capacity affects new performance metrics such as Kaps,Maps and SCAN. Then it is shown that all this metrics together can create new techniques or storage types which in fact in turn again affects the overall storage performance and cost.

On the left side of the diagram it is shown that latency and MIPS are affecting factors for processors. Storage and processor together affect the overall system cost. It is also shown that power consumptions should not be forgotten from the cost of both storage and processor. Also, one should remember the interaction of storage and processor is yet another important factor on overall system cost and performance.

7. Conclusions

The amount of digital data is increasing all the time, they are almost doubling every year. Storing and handling this huge amount of data, creates a new business case which should be addressed by the market. Industry tries to address the need with creating bigger capacity storages and faster processors to handle huge amount of data requests. But all the parameters don't follow the same trend, and some lags others. This affects both the design of the computing systems and administration decisions as well as the formation of the market.

Technology and price normally go hand in hand and have direct effects on each other; and of course the base price trend of the components should be considered when calculating the overall price of the system. In this paper, we have looked at the technology and price trends of different storages and data processing metrics. We have followed their improvement trends and noticed different ratios in their improvement curve.

Based on these ratios we have suggested what are the new business cases introduced to the market. We have also looked at the emergence effect of the other markets to the storage market, by introducing new kind of devices which need innovative new storage types. We suggest that while both capacity and performance of the computing systems are improving the actual bottleneck is in pipelining the data between disk and processor, which in turn emerges the new business cases such as introducing of the flash memories.

Finally, we conclude that hardware is driven by mass market economics and the ideal platforms are the one which lower capital costs, dramatically improve productivity and lower operational expenses, which means they are economically feasible.

References

[1] Nanotech memory, http://www.wired.com/gadgets/miscellaneous/news/2007/10/ion_memory#.

[2] Singularity.com statistics, http://singularity.com/charts/page62.html.

[3] Krste Asanovic, Ras Bodik, Bryan C. Catanzaro, Joseph J. Gebis, Parry Husbands, Kurt Keutzer, David A. Patterson, William L. Plishker, John Shalf, Samuel W. Williams, and Katherine A. Yelick, The landscape of parallel computing research: A view from berkeley, Tech. report, EECS Department, University of California at Berkeley, 2006.

[4] Jacek Becla and Daniel L. Wang, Lessons learned from managing a petabyte, 2005 CIDR Conference, Very Large Data Base Endowment, VLDB, 2005.

[5] Gordon Bell, Jim Gray, and Alex Szalay, Petascale computational systems: balanced cyberinfrastructure in a data-centric world, IEEE Computer 39 (2006).

[6] K. G. Coffman and A. M. Odlyzko, Internet growth: Is there a "moore's law" for data traffic, Handbook of Massive Data Sets, Kluwer Academic, 2001, pp. 47–93.

[7] P. Dubey, Recognition, mining and synthesis moves computers to the era of tera, Technology@Intel Magazine (2005).

[8] Jim Gray and Prashant Shenoy, Rules of thumb in data engineering, In Proc. 16th Internat. Conference on Data Engineering, 2000, pp. 3–12.

[9] E. Grochowski and R. D. Halem, Technological impact of magnetic hard disk drives on storage systems, IBM Syst. J. 42 (2003), no. 2, 338–346.

[10] John Hennessy and David Patterson, Computer architecture - a quantitative approach, Morgan Kaufmann, 2007.

[11] Sally A. McKee, Reflections on the memory wall, CF '04: Proceedings of the 1st conference on Computing frontiers (New York, NY, USA), ACM, 2004, p. 162.

[12] Andrew Odlyzko, Content is not king, http://www.dtc.umn.edu/odlyzko/doc/history.communications2.pdf.

[13] John Owens, Streaming architectures and technology trends, SIGGRAPH '05: ACM SIGGRAPH 2005 Courses (New York, NY, USA), ACM, 2005, p. 9.

[14] David A. Patterson, Latency lags bandwith, Commun. ACM 47 (2004), no. 10, 71–75.

[15] Philip E. Ross, Engineering: 5 commandments, IEEE Spectr. 40 (2003), no. 12, 30–35.

[16] Thomas Sterling, Peter Kogge, William J Dally, Steve Scott, William Gropp, David Keyes, and Pete Beckman, Multi-core for hpc: breakthrough or breakdown?, SC ’06: Proceedings of the 2006 ACM/IEEE conference on Supercomputing (New York, NY, USA), ACM, 2006, p. 73.

Transmission vs. Processing

András Zahemszky
andras.zahemszky@ericsson.com

Abstract

This survey investigates the important trade-off between transmission and processing in communication networks. It shows in a brief overview where this trade-off is clearly visible. Then the focus is mainly on Internet cores. We argue that it was the change in the price ratio between the cost of transmission lines and the cost of routers that was one clear motivation of packet switching, a clear success story of communications. Investigating the price change, as well as the rapid growth of Internet, we present some technological consequences appeared in the need of fast packet processing. It became clear that the relation between transmission and processing cannot be discussed without touching the question of the Memory Wall, as memory access speed may be a bottleneck when trying to build IP routers with 40 Gbps or faster interfaces.

1. Introduction

The trade-offs between transmission and processing appear in several aspects of computer architectures and communication networks. On a micro-level, they are visible in allocating the resources in multi-processors. In communications, on medium-level they can be investigated in the case of grid computing. Furthermore, in macro-level, they can be identified in the whole structure of the current Internet architecture, both in the edge networks and in the core.

As an example, in grid (distributed) computing, the goal is to solve a task as fast as possible, by utilizing the processing resources of several, remote CPUs. The rough idea behind this that computing power is everywhere and idle resources should be shared to others. This is a similar phenomenon that was identified in the 1950's in the case of expensive computers: they are usually idle waiting for human input, so why are they not shared among several users? Now, as computers became reasonably cheap and most of the times still underutilized, sharing resources by distributing tasks became feasible. For this, it is important to identify parts of the programs that can be run parallel, without too much communication between the different components. The jobs should be distributed so that minimal processing power is wasted (ideally no idle CPUs present in the grid) and the transmission (data transfer) between parts are limited (i.e, the computation is as parallel as possible). If there are frequent data transfers between the CPUs, CPUs may wait for each other, therefore computing power will be wasted. (Scheduling and load balancing are therefore crucial tasks here, but clearly out of scope of this paper).

Though the advent of grid computing is clearly related to the low price of CPUs, there are clearly other networking phenomena where the trade-off and the effect of the change in the price ratios are more visible, and resulted in more serious consequences with regards to the Internet architecture, as a whole. Consequently, the focus of the paper is more on those architectural consequences that made today's network to look like what they are today.

One of the most historically important result that proved to be the foundation of current Internet technologies is the *packet switching*. It replaced circuit switching with a new, radical way of thinking in the 1960's, because of efficiency and economic reasons (we will discuss this more deeply in Section 2).

The time cost of transmitting a message over a network has 3 distinct components: CPU (i.e. protocol operations), transmit delay (setup times and propagation delay) and link bandwidth (Gray, 1988). The delay caused by light, of course, cannot be eliminated, so the attention should be directed to the CPU performance and available bandwidth. This clearly shows the relevance of the prices of these in the evolution of telecommunication networks (see Section 2 for further details).

The price of transmission cannot be discussed without discussing the demand for bandwidth. The growing bandwidth demand, as well as the increased installed capacity in the core of the Internet, together with the transmission prices will be discussed in Section 3, where we will also give some numbers about historical processing power increase.

In a simple case study in Section 4, we will examine the case of Internet routers and provide some straw-man calculations about their processing requirements. Because of the hard constraints, processing-intensive tasks are usually in the edges of the networks. The core network is more like a large "bitpipe", transferring large amount of bits from one edge to another as fast as possible. Therefore we will consider only the most simple packet forwarding operations.

Finally, the high-level conclusions of this paper are collected in the form of a simple causal loop diagram in Section 5.

2. Historical case studies

In this section, we present two very interesting historical findings and analysis, which are shedding the light to the relevance of the trade-off between transmission and processing in communication networks. The first case study (Section 2.1) shows that the birth of packet switching technology for data communications was –among other reasons- due to the price trends of transmission and processing. The second case study (Section 2.2) overviews a simple cost model for transmitting messages over networks (Gray, 1988), providing another insight to the complex relations between transmission and processing in different sized networks.

2.1. Circuit switching vs. packet switching

Traditionally, telephony networks implement circuit switching: before any communication can take place between two remote endpoints, resources have to be allocated for the call, which will be after that uniquely used by the call. As a consequence, if the connection is still alive but at the moment there is not anything to transmit, resources will be wasted as they are cannot be used by other connections.

It was first Leonard Kleinrock who addressed that circuit switching does not fit the purposes of data communications. In his pioneering work (Kleinrock, 1962), Kleinrock introduced the concept of dividing messages into smaller pieces (packets) and sending these individually towards the remote end. On very simple terms, its efficiency is based on the bursty nature of data traffic: sources are idle and then suddenly they generate large amount of data to be transmitted; idle and bursty periods follow each other. The packet switching relies on sharing the resources, so other processes can use the network resources, while others are idle. Resources are dynamically allocated. Recall that by the preallocation manner of circuit switching, the whole bandwidth would have been wasted.

Lawrence Roberts, another founding father of the Internet phrased the very simple economical trade-off applied to data communication networks: if transmission lines are cheap, circuit switching is suitable, but if computing is cheap, packet switching is recommended (Roberts, 1978).

Roberts analyzed various trends in the Internet, first concentrating on the cost of communication and cost of computing, the two important aspects for packet switching. According to the analysis (Roberts, 2000b) communication costs halved only in every 79 months, from 1960 until 1995. Besides, cost of computation, i.e. the bits to process in a computer for one dollar doubled in every 18.6 months (Roberts, 1969). These results are similar to the statement of Moore's law. His updated calculations in the turn of the century (Roberts, 2000b) report the trend of doubling computer performance in every 21 months. He also confirms that the router performance also followed the same trends than computers in the period between 1969 and 2000.

Examining the price of transmission lines and processing (router cost), the turning point was around 1969. From this year, it was cheaper to perform some computations on a bit than to send it over a link. Together with the fact that the performance difference between packet switching and circuit switching for data is 15-fold (average-to-peak ratio of data traffic is 15:1), packet switching clearly became more and more advantageous, especially with the favourable change of prices. Though more computational operations were needed, like packetisation of messages and packet processing in intermediate nodes, huge amount of expensive bandwidth was saved. These relations are also shown in Figure 1, which is based on the findings of (Roberts, 1974), but taken from (Roberts, 2000a). Further trends inspected by Roberts will be examined in Section 3.

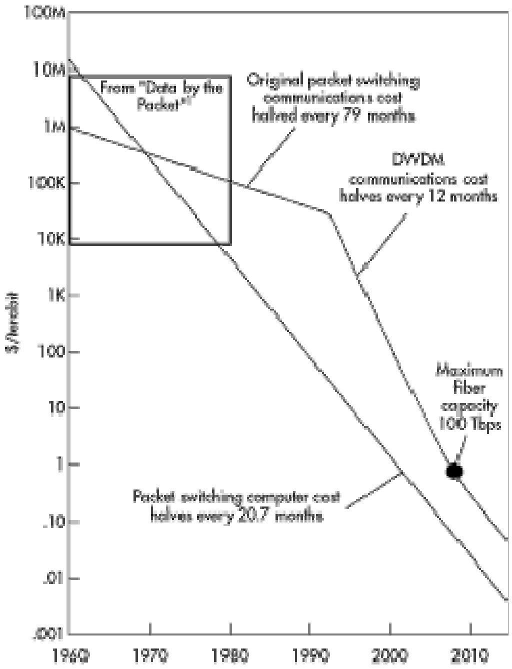

Figure 1. Cost of communications vs. cost of routers Source: (Roberts, 2000a)

2.2 The cost of messages

Jim Gray's work (Gray, 1988) gives another important insight, which helps understanding the trade-offs and their effects on networking technologies. Here, 20 years after its birth, we discuss his model in detail, emphasizing the selection of the parameters and commenting their values.

Gray investigated the message transport costs in 3 different settings (shared memory, local area networks and wide-area networks) and quantified the main differences for the past and made future predictions (1980, 1990 and 2000). For our paper's point of view, shared memory, where the communication processes are located in the same host, is out of scope and we focus instead of the implications related more closely to the networking.

The cost of sending a message from one process to another (in terms of delay) has 3 distinct components in this model: a) CPU, b) transmit delay, c) link bandwidth.

First, CPU reflects the protocol operations. Clearly, for LAN protocols the processing can be done with less instructions than in the case of wide-area protocols. Gray estimates 2500 instructions for sending and receiving a message in a LAN environment, while 12000 instructions for WAN-environment (based on implementations of the X.25 protocol). Besides, he estimated that CPU speed would be 100 millions of instructions per second (mips) in 2000. This estimation turned out to be quite conservative.[1] Meanwhile, 1 mips and 10 mips are used in the calculations for 1980 and 1990, respectively.

Second, transmit delay for long-distance connections are dominated by the propagation delay, i.e the speed of light. Certainly, it is the factor that cannot be eliminated (quantum-communications is again out of scope). Gray's usage of 10 ms for transmit delay in wide-area networks can also be considered conservative, when inter-continental connections are investigated (for comparison, delay in LANs are counted as 0.1 ms in the work).

Third, the (historical and projected) available link bandwidths for local area networks are 10 Mbits, 100 Mbits and 1 Gbps for 1980, 1990 and 2000, respectively.
The same numbers for wide area networks are 5 Kbits, 50 Kbits and 1 Mbits; remember that these are the access, and therefore maximum bandwidths for the processes communicating over long distances. These estimations, both in the case of LANs and WANs seem quite precise.

Now, we are briefly repeating Gray's results in Table 1 (omitting intermediate calculations).

Table 1: Cost (in ms) to send a 100 byte message

Network/Year	1980	1990	2000
Local area	2	0.27	0.036
Wide area	222	31	11

If one considers the dominating components, one can conclude that in the case of local area networks it is the CPU that contributes to the delay. However, in the case of wide-area networks, it is first the access bandwidth (1980 and 1990), but later the light starts to dominate.

Other important finding is that there is two orders of magnitude difference between LANs and WANs; furthermore the cost difference is apparently increasing (Table 1.).

Though it remarkably captured the characteristics of the message costs in distributed systems and presented insight to the differences between local area and wide area networks, Gray's precise, but simplified analysis may lack of some important parameters that are worth considering in current network scenarios. (We should not forget that the calculations and the conclusions were made more than two decades ago!)

In the model of wide area networks he did not consider intermediate network elements, i.e routers between the two ends in wide area networks. As one consequence, the processing delay of packets is not taken into account. Another consequence is that no queuing delays (i.e. time the packets spend in buffers of routers waiting to be processed/transmitted) are considered. This eventually means that the effect of background traffic and possible congestion is completely missing from the calculations.

Now, the ever increasing bandwidth demand should not be omitted from similar calculations. As the bandwidth demand increases, routers deep in the core networks need to perform rapid packet processing operations in line-speed, so that they can serve the increasing link capacities attached to them. These aspects will be thoroughly discussed later in Section 4.

2.3 Lessons learned from the case studies

Our case studies revealed two important facts. First, it was the different price curves of transmission and processing that made packet switching competitive. This emphasizes that price ratio changes do have crucial effects on the Internet architecture. Second, clear pictures can be obtained even without going deep into the models of economics. This emphasizes that performance ratio changes do have important effects to the Internet.

3. Towards understanding of price trends and performance changes

Now, we turn our attention towards different trends important for investigating backbone networks. First, we start with the trends related to transmission, i.e. trends in transmission prices, trends in traffic growth (demand), as well as trends in the network's capacity (supply). Then, we continue with processing performance, and finally summarize the lessons learned.

3.1 Backbone transmission prices

Collecting prices of transmission for backbone networks is a challenging task. The data and data

[1] As 4000 mips was reported for year 2000 in the website http://singularity.com/charts/page64.html

interpretations available are usually a bit controversial, and moreover, the data is hard to compare. This is mostly because of the fact that complex business situations play a big role in determining the prices, and they are changing from situation to situation. The attempt in this paper is still to find some trends and interpretation from well-referenced sources.

Roberts has investigated several trends affecting the Internet (Roberts, 2000). According to his statement, monopoly situations and government regulations kept backbone transmission prices high. Namely, the price of transmitting a bit halved only in every 79 months. According to the work, this trend continued until 1995, where a turning point came due to the advent of WDM technology with halving of prices in every 12 months. WDM technology cuts the prices because now more than one wavelength can be transmitted on a single wire, multiplying the available bandwidth in a cable. Deploying cables with higher capacities consequently became cheaper.

A study from Douglas Galbi (Galbi, 2000) investigates the change in the trend of bandwidth prices between 1990 and 2000. According to the data collected, there are some controversies, if we compare the data to Robert's conclusions. In Figure 2, we show the backbone prices found in that paper for AT&T between 1983 and 1998. The prices until 1993 more or less follow Roberts' results, but from 1993, a small increase was experienced.

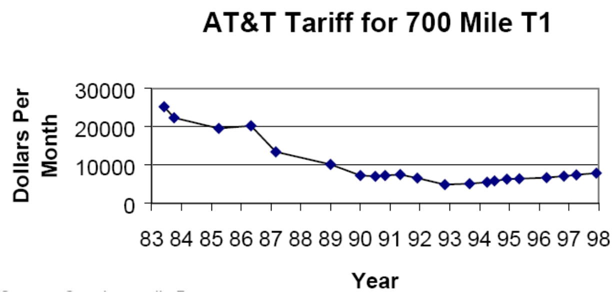

Figure 2. AT&T Prices for T1 connections (1.5 Mbps)

On the contrary to this the prices of cable deployment show more or less stable decrease trends over the decade (around 10-fold decrease in every 5 years). This can be examined in Table 2, where we summarize the investment costs of transatlantic cables in different years from 1988 until 2001.

More recent data can be found by looking into the data provided by Telegeography Inc., a company specialized in making thorough analysis on the market of telecommunications. Though access to their data is usually bound to subscriptions (priced thousands of dollars), some trends can be investigated from their public presentations.

Table 2. Cost of deploying transatlantic cables

Year	Cable	Cost ($mil/Gbps)
1988	TAT 8	744.0
1989	PTAT 1	367.4
1992	TAT 9	419.6
1992	TAT 10	206.7
1993	TAT 11	192.9
1995	TAT 12/TAT-13	32.6
1997	Gemini	22.4
1998	Atlantic Crossing (AC-1)	13.7
2000	TAT-14	3.0
2001	Level 3	2.4
2001	Hibernia	5.1

The data sets reveal that are some clearly visible characteristics in the price differences (Beckert, 2009) in different environments. One important aspect is that location does matter. For example, monthly prices per Mbps in a one Gbps bitpipe have a great variation: while it is 10$/Mbps in Los Angeles, around 40$/Mbps in Hong Kong, while almost 80$/Mbps in Rio de Janeiro (in the year of 2008). This emphasizes the fact that market conditions can cause huge variation in prices (competition, supply/demand, etc.). Another important phenomenon is that per bit prices are cheaper by higher capacities. For example, in Los Angeles, cost of 1 Mbits was $35 a month when 100 Mbits line is purchased, and around $7.5 if 10 Gigabits Ethernet connection is bought (year 2008). A further observation is that the annual price decline is higher by larger capacities (for instance, 20% compared to the 29% in the previous example). Last, but not least, prices vary from carrier to carrier. Therefore, in the next figures, always the median prices are shown.

From the plenty of tables and various figures in the presentations (Beckert, 2007), (Stronge, 2008) and (Beckert, 2009) we selected somewhat arbitrarily two to present the current trends of pricing in backbone networks.

On Figure 3, we show the price decline between 2002 and 2006 between major sites. The prices reflect the price of an STM-1 line, so per Mbps data can be derived by dividing the cost by 155 (STM-1 rate is 155 Mbps). The message of the figure is that after huge drops (around 50% per year in many locations), prices continue to decline in a slower pace. On Figure 4, the annual decline of transit prices is depicted. It shows an annual decline between 16%-26%, based on the location of the link. The data is derived by considering the prices of 1 Gbps lines.

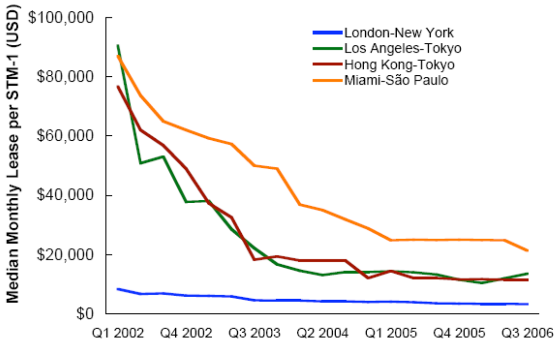

Figure 3. Price decline between 2002 and 2006 for STM-1 connections

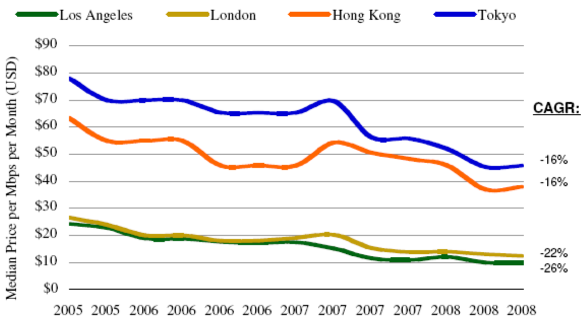

Figure 4. Price decline in various places between 2005 and 2008 (for Gigabit Ethernet connection)

3.2 Internet traffic

After considering the prices, we should turn our attention towards the increasing bandwidth demand in the Internet. Roberts states (Roberts, 2000) that between 1970 and 1982, the traffic doubled in every 21 months as new hosts were added to the network. Then, after the deployment of TCP/IP, the number changed to doubling every 9 months, while from 1998 (until the publishing of the paper) it doubled every 6 months.

The data provided by Odlyzko (Odlyzko, 2003) has more or less the same trends. He collected the total traffic on US backbones and investigated the trends. His data is briefly repeated in Table 3.

One can conclude from Table 3 that there was a dramatic increase in the demand in 1995 and 1996. This was a short period, where Internet traffic doubled approximately in every 3 months, meaning of 16x increase per year. Clearly, this did not last for long, though it was a common myth in the Internet business to claim that this increase still exists (Odlyzko, 2003). The growth numbers after the huge explosion are more conservative compared to the statement of Roberts (doubling every year compared to doubling every 6 months). According to Telegeography Inc. (Beckert, 2009) Internet traffic growth is between 50% and 100% annually, depending on the location, meaning that the traffic is doubling approximately in every 12-20 months.

One can conclude from Table 3 that there was a dramatic increase in the demand in 1995 and 1996. This was a short period, where Internet traffic doubled approximately in every 3 months, meaning of 16x increase per year. Clearly, this did not last for long, though it was a common myth in the Internet business to claim that this increase still exist (Odlyzko, 2003). The growth numbers after the huge explosion are more conservative compared to the statement of Roberts (doubling every year compared to doubling every 6 months). According to Telegeography Inc. (Beckert, 2009) Internet traffic growth is between 50% and 100% annually, depending on the location, meaning that the traffic is doubling approximately in every 12-20 months.

Table 3. Monthly traffic in US backbones

Year	TB/month
1990	1.0
1991	2.0
1992	4.4
1993	8.3
1994	16.3
1995	N/A
1996	1500
1997	2500-4000
1998	5000-8000
1999	10000-16000
2000	20000-35000
2001	40000-70000
2002	80000-140000

3.3 Capacity of wires

The speed of lines changed in the backbones from 50 Kbps (1969) to 10 Gbps lines until 2000. This is approximately a growth of doubling of speed in every 22 months (Roberts, 2000b).

The fibre technology seems up-to-date with the current and future demands. After the introduction of optical fibres introduction, the capacity of wires grew about 30% a year (Bux, 2001), which is a somewhat quicker upgrade than the one present in the previous decades. The limit of this approach seems to be at 100 Gbit/s, due to physical constraints. However, with WDM, it became possible in the mid-90’s to use several wavelengths on a single fibre, and it resulted in a radical annual growth rate. In 2001, 1.6 Tbits links were technologically possible; the limit of WDM is expected to be around 50 Tbits, again due to physical constraints.

According to AT&T's reports[2], it deployed 40 Gbits (OC-768) links into its backbone network in 2008, and besides, in lab environment, they made successful experiments with 17 Tbit/s links.

3.4 Processing power

The development of processing power is well characterized by Moore's law, which states that the number of transistors per integrated circuit doubles in every 18-24 months.

The conclusion of Roberts' (Roberts, 1969) early study was similar to Moore's law: the throughput per dollar of computers is doubling in every 18.6 months. Later, this doubling period was revisited to be 21 months (Roberts, 2000).

If we check the performance growth of general purpose processors, in the terms of million instructions per second, we can see the doubling of performance in every 1.8 years[3]. The performance reached the 4000 mips by 2000, which is 40x higher than Gray's estimation (see Section 2.2). Just as an interesting sidenote: released in 2008, the Intel Core i7 Extreme 965EE processor, which operates at 3.2 GHz, can perform more than 75000 million instructions per second.

3.5 Lessons learned from price changes and technological advance

Comparing the Internet traffic growth to the price data presented in Section 3.1, it seems that the price decline is slower than the overall traffic growth, giving place to some operator revenues. As it was pointed out in (Odlyzko, 2004), still, the major part of the revenue does not come from operating backbone networks. Furthermore, inducing traffic growth can drastically increase the revenue besides current pricing trends (Odlyzko, 2008). The technology advance in terms of fibre capacity seems to be well conform with the increasing traffic demand. Computing power is steadily increasing according to Moore's law, e.g. processor performance is 75000 mips in 2008.

[2] See for example the press release on the website http://www.networkworld.com/news/2008/102408-att-backbone-network-40gbps.html

[3] According to the data found on the website http://singularity.com/charts/page64.html

4. A case study: router performance

Another impact of traffic growth is the processing requirements of core routers. In a case study for examining the performance increase in the previous decades, we will now present their implications to core routers in this section. The approach presented here will be a straw-man view, as going deep into technological aspects is out of scope for this paper.

Because of its clear advantages, packet switching is truly the foundation of the Internet technology. After packetizing the messages to transmit in the sending host, all routers towards the destination host will operate on the packet headers, i.e. they must process the packets. In the simple case for IP routing, the routers must repeat the following steps for all incoming packets: checksum verification, TTL (Time To Live) check, determining the next hop by routing table lookup, and determining the next hop's Layer 2 address (either through a lookup or a query, see ARP for details).

As networks became more complex, sophisticated, time and memory-consuming packet processing operations appeared in different contexts. Some examples are: forwarding (IP, Ethernet, MPLS, ATM), gateway functionalities (IPv4-IPv6 translation, NAT, wireless TCP/IP), QoS (IntServ, DiffServ).

Let us continue by a simple calculation that may point out the incredible packet processing requirements current routers are facing. Table 4 shows that what is the inter-arrival time for 40 byte (small) packets by different line rates. Considering packets only consisting of TCP and IP headers means a pessimistic, worst-case estimate.

Table 4. Hints for requirements for core routers

Technology	Line speed (Gbps)	Million packets per second	Packet inter-arrival time (ns)
OC-12	0.622	1.95	576
OC-48	2.5	7.8	144
OC-192	10	31.25	36
OC-768	40	125	8

Table 4 shows that a router with for example 10 Gbps interfaces should be ready with all the processing on a packet in 144 ns, before the next packet arrives. Usually, parallelism is introduced to reduce the requirements: network processors are organized usually to be multi-core/pipeline architectures. In this way, if 16 cores are present, the budget for processing a packet is increased to be 16x144 ns = 2304 ns in the case of OC-48 interfaces. With OC-768 interfaces this number is still only 128 ns.

There are some data known about number of instructions required for packet processing (Wolf, 2000). It tells that a routing table lookup requires around 2.1 instructions/packet header byte. Another simple and frequent operation is fragmentation and defragmentation, which requires 7.7 instructions/packet header byte. More sophisticated operations (QoS, compression, encryption) may require 1-2 orders of magnitudes more operations.

Considering only table lookups, usually two has to be performed in a router: a) lookup in the IP forwarding table determining the next hop, b) determining the next hop's Layer 2 address. For a 40 byte packet, this is around 35-40 operations, which constitutes to 5000 mips, which is feasible by today's CPU rates.

It seems that the bottleneck of network processors is now the *Memory Wall* (McKeown, 2004). The following reads and writes have to be performed for each packet (Bux, 2001):

1) A packet is written to data store after arriving
2) Read header into processor complex
3) Write header back to memory
4) Read for outbound transmission

Considering again the worst case, the smallest packets of 40 bytes and the line rates of 40 Gbps, we will get that the memory access speed should be 160 Gbps, which speed cannot be reached by off-chip memories. More likely, ultra-wide, on-chip DRAMs are needed, as far as further parallelization by using multi-threadism, which can hide huge memory latencies.

McKeown's work on router architectures concludes that because of the slow memory access speeds, buffering of packets will be problematic for 40Gbps interfaces and at higher speeds. Furthermore, simple routers are needed – paradoxically, routers were always meant to be simpler. With our basic calculations we showed that simple packet processing is not a problem, but simply there will not be enough time for more sophisticated operations on a per-packet basis.

With the further emerge of optical technologies one can hardly predict the router architecture of future networks. McKeown predicts that in 10-15 years from now dynamic optical circuit switching or optical packet switching will be feasible (McKeown, 2004), which will surely initiate new questions and answers as well as a new technology, of which prices and performances should be considered: pure optical network elements.

5. Conclusions

Leaving the straw-man causal loop diagram to the final section was deliberate. In this case, it is a hard task to produce a loop diagram, as the real feedback loops are very complex, and was even hard to understand by the business actors involved in it. One *tentative* proof for that is the myth of that the Internet traffic is doubling in every 100 days. This belief was very commonly spread even years after this phenomenon was really true (in 1995 and 1996). This was likely to have effects on operator's network building strategies, and could easily lead to unnecessary upgrades, resulting in heavily underutilized networks. It might have had an effect on the fact that according to some reports, transmission prices stagnated in the 1990's in the US.

The currently developed causal loop diagram is strictly a straw-man approach and reflects the author's basic understanding of the technical/economic relationships of the different parameters and forces present in this document. The CLD is shown in Figure 5.

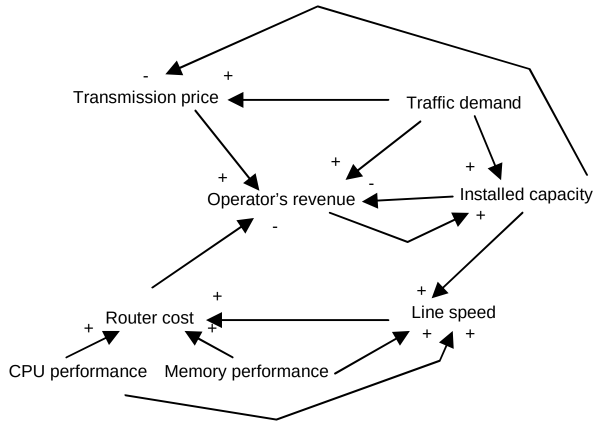

Figure 5. Causal loop diagram as a conclusion

Talking about backbone networks, the main output is the (Tier-1) operator's revenue. As Odlyzko pointed out (Odlyzko, 2008), it will be increased by increasing the traffic in the network (this is the reason for the direct connection). Furthermore, the growth of traffic demand increases the transmission prices (but may cut the transmission prices per Mbits as larger bitpipes are purchased and besides, competition is stimulated). Eventually the income of the operator comes from the amount of bandwidth it manages to sell, which depends on the demand and the prices.

Because of the increasing traffic demands, more capacity should be installed into the operator's network. This is either upgrade of current cables or deploying new ones. Surely, the cost of this decreases operator's revenue. Besides, if eventually the supply is higher than the demand, the transmission prices will decrease. If the capacity installations exceed a limit at some places of the networks, the line speed of routers eventually increases, leading to router equipment upgrade. Again, the cost of the routers (which mainly comes from the CPU and memory costs, which is related to their performance) operating in the backbone decreases operator's revenue.

Finally, as the operator's revenue is increasing it may decide to invest into the network by upgrading its network by installing additional capacity.

In this paper we investigated the trade-offs between transmission and processing. This trade-off had a crucial effect on current Internet: it formed the network to as it is today, simply because packet switching is turned out to be more economical than circuit switching at the end of 1960's.

40 years later, this needs some reconsideration. Because of a third actor, namely memory access, the growths will likely to be bounded. The future attention should be focused on technologies like dynamic optical circuit switching or optical packet switching.

References

Stephan Beckert, Eric Shoonover (2007), Global Traffic, Bandwidth, and Pricing Trends and Wholesale Market Outlook, in PTC 2007, 14 January 2007, http://www.ptc.org/past_events/ptc07/program/presentation/ST_BeckertSchoonover.pdf (accessed 1 March 2009)

Stephan Beckert, Alan Mauldin, Erik Kreifeldt (2009), Telegeography International Telecom Trends Seminar, in PTC 2009, 18 January 2009, http://www.ptc.org/ptc09/images/papers/PTC%2709_TeleGeography_Slides.pdf (accessed 1 March 2009)

Bux, W. Denzel, W.E. Engbersen, T. Herkersdorf, A. Luijten, R.P. (2001), Technologies and building blocks for fast packet forwarding, in Communications Magazine, IEEE,Volume: 39, Issue: 1, page(s): 70-77

Douglas A. Galbi (2000), Growth in the "New Economy": U.S. Bandwidth Use and Pricing Across the 1990s, Draft, http://www.galbithink.org/bandwidth.htm (accessed 1 March 2009)

Jim Gray (1988), The Cost of Messages, Proceedings of the seventh annual ACM Symposium on Principles of distributed computing, Toronto, Ontario, Canada, Pages: 1–7, 1988

Leonard Kleinrock (1962), Message Delay in Communication Nets With Storage, Ph.D. dissertation., MIT, 1962.

Nick McKeown (2004), Network Processors and their memory, Network Processor Workshop, NP-3. Madrid, Feb 2004. Keynote talk. http://tiny-tera.stanford.edu/~nickm/talks/NPU_2004_v1.pdf, (accessed 1 March 2009)

Nick McKeown (2005), How Emerging Optical Technologies Will Affect the Future Internet, talk at NSF FIND Meeting, Dec 2005, D.C., found in http://tiny-tera.stanford.edu/~nickm/talks/index.html (accessed 1 March 2009)

Andrew M. Odlyzko (2003), Internet traffic growth: Sources and implications, Optical Transmission Systems and Equipment for WDM Networking II, B. B. Dingel, W. Weiershausen, A. K. Dutta, and K.-I. Sato, eds., Proc. SPIE, vol. 5247, 2003, pp. 1-15

Andrew M. Odlyzko (2004), Pricing and architecture of the Internet: Historical perspectives from telecommunications and transportation, in TPRC 2004, http://www.dtc.umn.edu/~odlyzko/doc/pricing.architecture.pdf (accessed 1 March, 2009)

Andrew M. Odlyzko (2008), Threats to the Internet: Too much or too little growth?, Internet Evolution, February 25, 2008, http://www.dtc.umn.edu/~odlyzko/doc/too.little.growth.txt (accessed 1 March 2009)

Lawrence G. Roberts (1969), Data Processing Technology Forecast, http://www.packet.cc/files/Forcast69.htm (accessed 1 March 2009)

Lawrence G. Roberts (1974), Data by the Packet, *IEEE Spectrum*, Vol. 11, No. 2, Feb. 1974, pp. 46-51.

Lawrence G. Roberts (1978), Evolution of Packet Switching - Proceedings of the IEEE, Vol. 66, No. 1

Lawrence G. Roberts (2000a), Packet Switching or Optical Switching?, Internet Computing, IEEE Volume 4, Issue 1, Jan/Feb 2000 Page(s):50 – 51

Lawrence G. Roberts (2000b), Beyond Moore's Law: Internet Growth Trends, Computer, Volume 33, Issue 1 (January 2000), pp. 117-119.

Tim Stronge, Eric Schoonover (2008), Telegeography International Trends Workshop, in PTC 2008, http://www.ptc.org/ptc08/participants/speakers/papers/SchoonoverFinalSlides.pdf, (accessed 1 March 2009)

T. Wolf and J. Turner (2000), "Design Issues for High Performance Active Routers," *Proc. 2000 Int'l Zurich Sem. Broadband Commun.*, Zurich, Switzerland, Feb. 15–17, 2000, pp. 199–205.

Storage vs. Transmission

Matti Peltola

matti.peltola@mapelcon.fi

Abstract

This study reviews, if the seen and expected price development of the Internet network elements - like the memory devices and the transmission equipment - has influence to the architecture of the Internet. The new characteristics of the Internet traffic may effect to the optimum structure of the net. The scope of the review is 5…10 years.

The study shows that the prices of the memory elements and the transmission devices continue to decrease – the long range transmission systems less than the short range ones.

The caching will be a crucial element to keep the Internet operation costs low. The Peer-to-Peer (P2P) traffic and the streaming data will increase dramatically and the end-to-end latency time becomes more important. The increasing eCommerce will mean the enhanced customer data collection. The search engines will have more important role in the Internet services. The caching, which means storages close to the clients, have a solution to support all these changes.

Because the caching will offer economical benefits and the long-distance transmission costs don't decrease so fast than other cost elements in Internet system, the Internet architecture will favor local data storages in future.

1. Introduction

The Internet environment has been under continuous changes during its lifetime. The number of users has dramatically increased year after year, meaning often 100% annual growth. In the technology side Moore's law has caused significant performance improvements in network elements, including transmission speed, transmission capacity, storage capacity and memory access time.

The price development within the Internet network elements is continuously very fast; the price of the disk drives is going down to 1/100 part in a decade and the price of transmission to 1/5 in the same time. [Gray et al., 2000 and Nikander, 2009].

The increased P2P traffic, the dynamic web pages, the caching on search engines and the streaming multimedia will have influence to the cache hierarchy and the transmission capacity needs.

The cost changes within the network elements and the new characteristics of the traffic together with caching could lead to the changes in the storage and transmission architecture.

This study tries to understand the new balance between the storage and the transmission in Internet within next 5…10 years.

The chapter 2 reviews the Internet traffic statistics. The chapter 3 discusses of the price development of the Internet elements and defines the phenomena which may have influence to the total price of storage and transmission in Internet. The chapter 4 reviews the development of caching and possible changes in the caching strategy in case of the P2P traffic, caching within the search engines as well as caching with dynamic web pages. In the chapter 5 we will discuss shortly what the latency requirements mean to the network architecture.

2. Internet traffic statistics

The traffic type distribution has changed a lot in recent years. Schulze et al. (2008) have announced a study of the Internet traffic in the year 2007 in Germany and in Middle East. In Germany P2P is almost 75% of the total Internet traffic, the web browsing is "only" 10%, the audio and the video streaming about 8%. In Middle East the corresponding figures are 49%, 26% and 0.1%.

Figure 1a shows the traffic distribution in Germany based on the protocols in 2007; the measurements have been done in August and September.

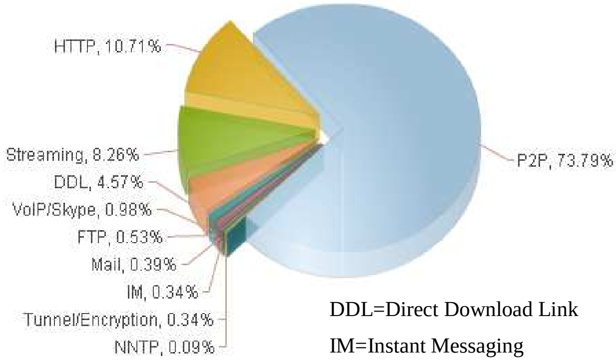

DDL=Direct Download Link

IM=Instant Messaging

NNTP=Network News TP

Figure 1a. Internet traffic distribution in Germany 2007 [ipoque, 2008]

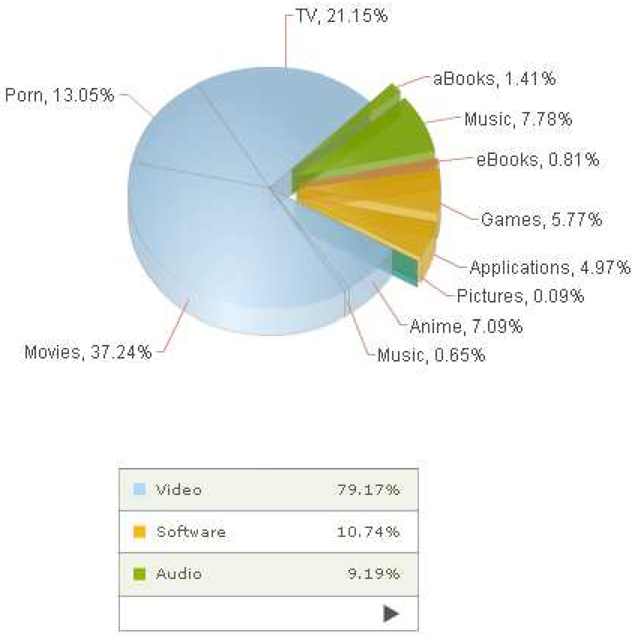

Figure 1b. Traffic (P2P) volume per Content type, BitTorrent, Germany 2007 [ipoque, 2008]

Figure 1b shows the more deep analysis of the P2P traffic [Schulze et al., 2008]. The main parts of P2P traffic are video (79%), audio (9%) and software (11%). The analysis is based on the follow-up of the anonymous BitTorrent P2P protocol traffic.

3. Price development of elements

The key parts in the storage vs. transmission techno-economical model are the development of transmission costs, the development of storage costs, the cache strategy, the latency requirements and the characteristics of the real-time streaming multimedia. In the figure 2 the straw-man systems thinking is used to understand the Internet storage and transmission costs, which are formed of different elements and their relationships. Except the actual device prices the characteristics of the new Internet traffic and the followed consequences have major influence to the network architecture and this way to the total costs of the network.

The figure 3 presents the transmission costs today and in 10 years [Nikander, 2009]. Transmission Tier-I transport price is today (2009) about \$25/Mbs/month, the DSL link costs are \$1…5/Mbs/month. The mentioned prices are going down 3…6x/decade. The technology development increases the available transmission bandwidths 4-fold in every decade.

The disk costs now about \$100/1TB, the prices go down to 1/100 part in a decade. The size of the biggest disks is today about 1TB. According to Gray (2000) we will have 100TB disks in 2018. The medium page size in

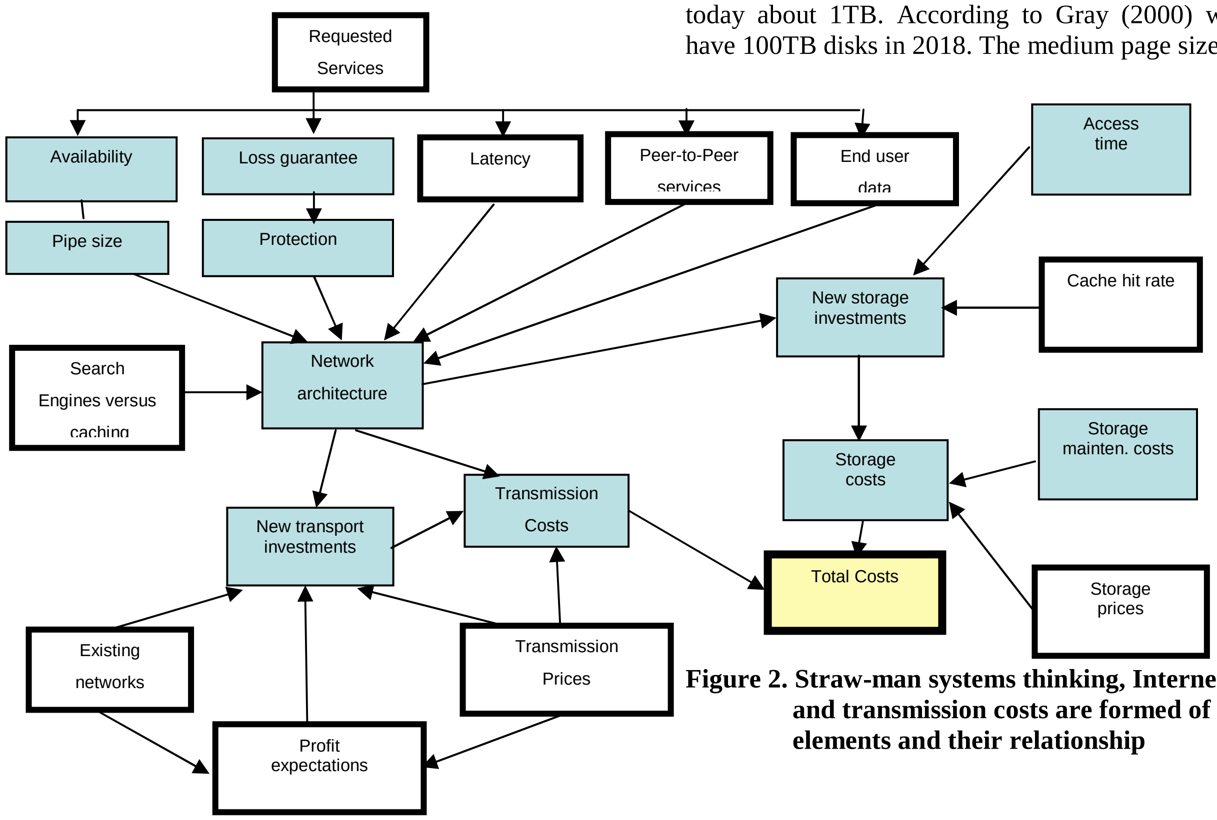

Figure 2. Straw-man systems thinking, Internet storage and transmission costs are formed of different elements and their relationship

Internet will be 5-fold in a decade. The memory access time will be improved to 10x in every decade.

The price of the RAM memory is today about $10000/1TB. The price will come down to 1/100 in a decade. The RAM size development continues - today the biggest device is 0.01TB. It is expected that in a decade the max RAM size is what the disk is today, e.g. the biggest RAM device in 2018 is 1TB.

The figure 3 demonstrates the listed price development. What can be seen is that the transmission and storage costs will go strongly down, however the Tier I transmission costs stay high comparing to the other element costs.

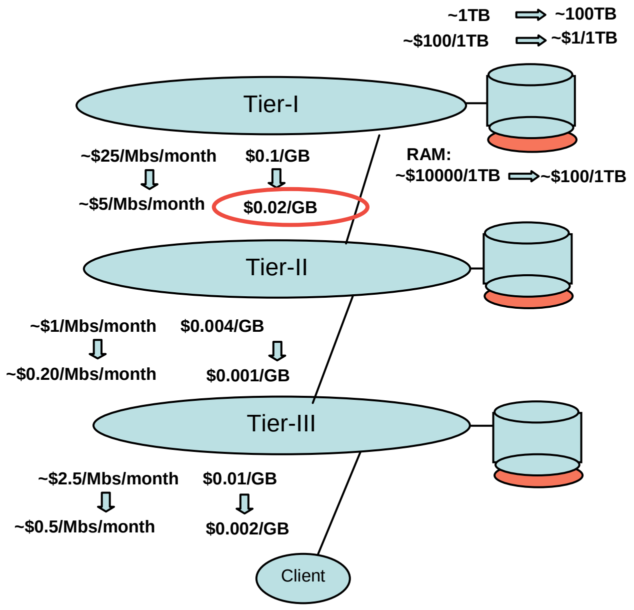

Figure 3. Internet caching hierarchy and price development 2009…2019

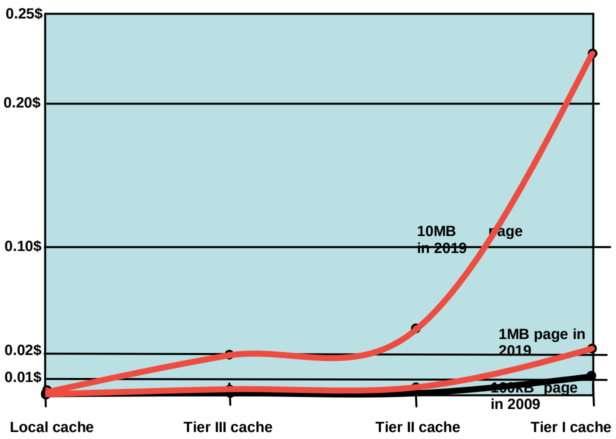

Figure 4. Average web page transmission cost in 2009 and 2019 depending on the cache site

The "relatively high" cost can be well noticed, if we expect that in 2019 the size of the typical web page is 10Mbytes – and we calculate what it costs to fetch that page from Tier level I – see figure 4. The result pressures the storage location towards to the lower Tier levels.

4. Caching

The caching saves bandwidth, improves the end-to-end latency times and prevents the purposely caused repeated queries to the same web page. When thinking the cultural unity of the society, it makes understandable the benefit of the caching, e.g. the people belonging to the same society use and are interested in the same type of information including the same language version of that information. The new traffic types and the new characteristics of the bit streams may have influence to how the caching can be implemented. In the following the essential cache cases will be analyzed.

4.1. Scale and performance of Web proxy caching

Typically the proxy cache sits in the front of an organization. The hit rate of the proxy cache grows with the client population (figure 5), however the hit rate doesn't increase much after 2500 clients [Wolman et al., 1999].

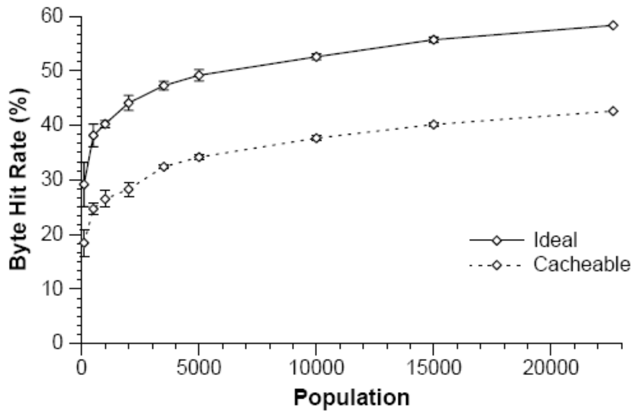

Figure 5. Proxy cache byte hit rate as a function of client population [Wolman et al., 1999]

The caching reduces bandwidth consumption compared no caching, but there is no benefit to increased client population trying to save bandwidth, figure 6.

Wolman et al. (1999) have come to the conclusion that a single proxy cache can provide the same benefits than the cooperative proxy in wide-area environment and do it with fewer resources and less overhead.

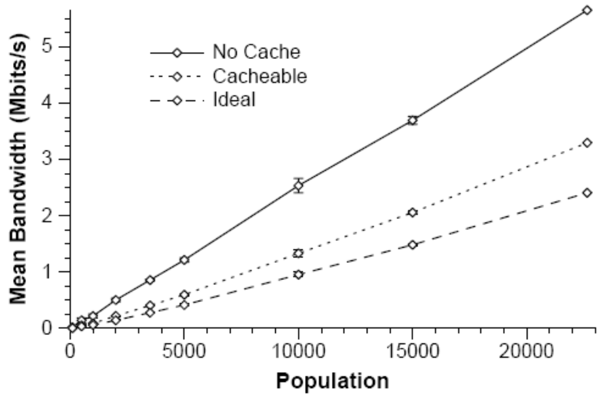

Figure 6. Used bandwidth as a function of client population [Wolman et al., 1999]

4.2. Caching of dynamic contents on the Web

Today many web pages are dynamic, e.g. their content changes upon every access, they include client-specific information or the content is the result of the query – however the caching mechanism would like to be used to improve the Internet performance.

Cao et al. (2000) have studied the caching of dynamic web pages – they propose to use the Active Cache protocol. When user requests the cached copy, the proxy has to invoke the cache applet which then decides if the cached document is returned as such or as a modified one.

The cache applets allow servers to obtain the benefit of proxy caching and without losing the capability to track user accesses and tailor the content presentation dynamically [Cao, 1998], which means the saving of bandwidth in Internet.

4.3. Search engines and caching

Baeza-Yates et al. (2007) have shown that "the caching is an effective technique in search engines for improving response time, reduce the load on query processors and save network bandwidth".

In their studies they have shown that the optimum caching strategy is to share the cache space between cache for answers and cache for posting lists.

Figure 7 shows the hit rate as a function of the cache size when caching posting lists. The cache size is a fraction of the total space required to store the posting lists of all query terms.

Also Lempel et al. (2003) have demonstrated the benefits of caching within search engines.

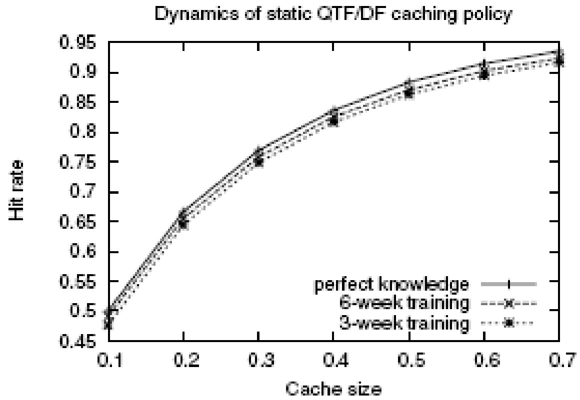

Figure 7. Impact of distribution changes on the static caching of posting lists [Baeza-Yates et al., 2007]

4.4. Peer-to-Peer and caching

Leibowitz et al. (2002) have studied the P2P traffic and the utilization of caching with it.

In their study they noticed that if you look the quantity of transmitted files, the major part of them are audio files – however among 100 most popular files the majority are video and application files. If you review the traffic from the volume (e.g. the number of bytes) point of view, the main part of P2P traffic is generated by the applications and the video files.

Leibowitz et al. (2002) noticed that the theoretical caching possibility is very high, even 67% of traffic could be cached – when the similar figure of HTTP caching is typically 30%, maximum 60%. The reason for this is that in video and application downloading there are typically relatively few files which are popular at the same time.

They also showed that 200GB cache is sufficient to achieve almost maximum caching, see figure 8.

The conclusion is that using caching we could save a lot of bandwidth within P2P traffic and the caching can be done quite close to the end-users.

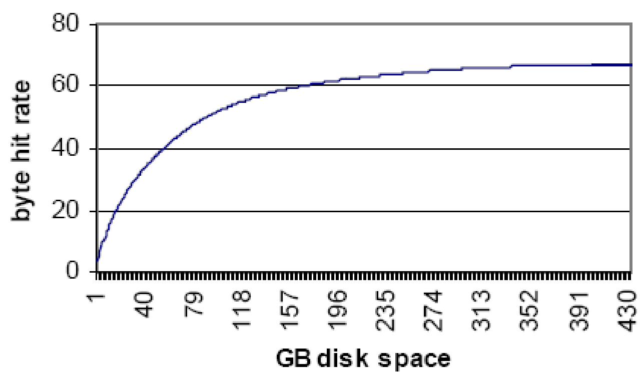

Figure 8. Theoretical Cache limits [Leibowitz, 2006]

5. New requirements for data network– low transaction latency

Internet traffic is bursty, it grows fast – often 100% annually. Odlyzko (2003) has studied the utilization rates of the Internet transmission in 1997, 1998 and found the similar results in 2003. Figure 9 shows the OC-48 link traffic in February 8-9, 2003 between Paris and Frankfurt. The average utilization of that link during the week was 13.8%. Odlyzko states that the backbones of the Internet are run at lower fractions (10…15%), comparing to the switched voice network at over 30% of capacity on average.

File transfers are likely to dominate in future [Schulze et al., 2008 and Odlyzko, 2004] – including P2P file transfers and the streaming multimedia. The streaming multimedia means to transfer - not huge, but a moderate amount of bits quickly through. The important new requirement what the streaming multimedia traffic will ask is low transaction latency, e.g. the low end-to-end delay.

Based on above the issue will be more the latency than the lack of transmission capacity. Odlyzko remarks that the networking industry has not taken this into account when trying to differentiate the pricing.

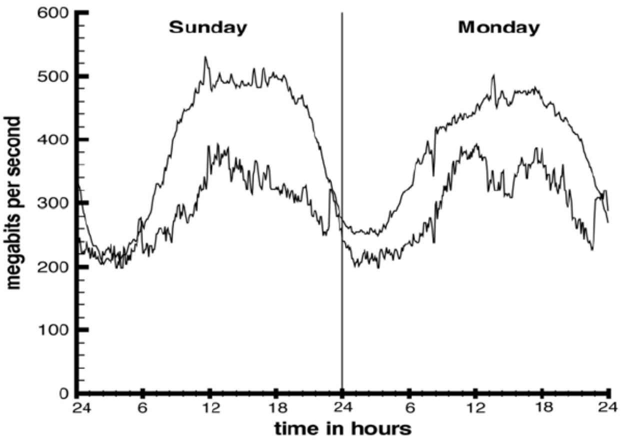

Figure 9. Traffic on AboveNet OC-48 link between Paris and Frankfurt, February 8 and 9, 2003. [Odlyzko, 2003]

6. Conclusions

The cost decrease of transmission and storage will continue with the same speed as earlier - the price reduction of the storage elements is faster than the transmission prices [Nikander, 2009].

The Internet network investments are ”once at a time” – type of investments, e.g. because of the size of the long distance network investments, the long distance network capacity has often grown in steps following the global economic trends. The present (2009) economical situation may slow down the building of the new transmission capacity – and by this way it has influence to the development of the Internet architecture.

The major part of the web traffic (in bytes) is now video, audio or software P2P traffic [Schulze H. et al., 2008]. The P2P traffic will benefit a lot of caching [Leibowitz et al., 2002]. Wolman et al. (1999) have shown that the caching can – even in the case of a relatively small number of customers (~2500) - achieve a good hit rate. The caching can be implemented also in case of the dynamic web pages [Cao et al., 2000]. The caching is an effective technique within the search engines – it will improve their performance [Baeza-Yates et al., 2007].

The growth of the streaming multimedia traffic will support the implementation of caching as close to the users as possible, because this will eliminate the end-to-end latency problems.

The increasing share of P2P traffic and the streaming multimedia together with the benefits of the caching technology will emphasize the benefits of the local storage solutions in the Internet architecture.

The study has shown that the transmission element and storage element prices are decreasing fast. In Internet the caching has fundamental influence to the amount of data what has to be transported in the network. Depending on the case the caching will decrease the data volume by 30…60%. This means that the emphasis in the Internet architecture will be more in the development of the caching solutions near to the clients than the solving the data transport by increasing transmission bandwidths in the Internet main routes.

References

Baeza-Yates R., Gionis A., Junqueira F., Murdock V., Plachouras V. and Silvestri F. (2007) The Impact of Caching on Search Engines. SIGIR ’07.

Cao P. (1998) Active Cache: Caching Dynamic Documents. http://www.cs.wisc.edu/~zj/active-cache/index.html (accessed 10 February 2009).

Cao P., Zhang J., and Beach K. (2000) Active Cache: Caching Dynamic Contents on the Web. Computer Science Department University of Wisconsin-Madison.

Chankhunthod A., Danzig P.B., Neerdaels C., Schwarz M.F. and Worell K.J. (1995) A Hierarchical Internet Object Cache. Computer Science Department University of Southern California, http://www.cs.princeton.edu/courses/archive/fall03/cs518/papers/harvest.pdf (accessed 10 February 2009).

Gray J., Shenoy P. (2000) Rules of Thumb in Data Engineering. MS-TR-99-100.

ipoque (2008) Internet Study 2007 report , company web page http://www.ipoque.com/resources/internet-studies/internet-study-2007 (accessed 26 February 2009)

Leibowitz N., Bergman A., Ben-Shaul R. and Shavit A. (2002) Are file swapping networks cacheable? Characterizing P2P

traffic. Proc. of the 7th Int. WWW Caching Workshop, http://2002.iwcw.org/papers/
18500253.pdf (accessed 10 February 2009).

Lempel R., Moran S. (2003) Predictive Caching and Prefetching of Query Results in Search Engines. 12th World Wide Web Conference (WWW2003).

Nikander P. (2009) Long Term Price Trens and Telco Industry. Presentation at S-38.4043.

Odlyzko, A. (2003) Data Networks are Lightly Utilized, and will Stay that Way. Review of Network Economics, Vol.2, Issue 3.

Odlyzko, A. (2004) Pricing and Architecture of the Internet: Historical Perspectives from Telecommunications and Transportation. Proceedings of the 32nd Research Conference on Communication.

Schulze H., Mochalski K. (2008) Internet Study 2007: The Impact of P2P File Sharing, Voice over IP, Skype, Joost, Instant Messaging. Tech. report, ipoque.

Wolman A., Voelker G.M., Sharma N., Cardwell N. Karlin A. and Levy H.M. (1999) On the Scale and Performance of Cooperative Web Proxy Caching.
17th ACM Symposium on Operating Systems Principles.

Transit and Access Prices – Evolution and Interdependencies

Tapio Levä
tapio.leva@tkk.fi

Abstract

Transport, and especially transit, costs, are a cost component in offering Internet access. The real importance of these costs is however not clear. In this paper we will study, how transit prices affect on ISPs costs and through that to access prices. The transit, peering and access cost structures are studied by using literature sources. This inspection is deepened by presenting transit and access price evolution during the last eight years based on gathered data from other sources and analyzing the factors of price formation after that. We will then depict the interdependency between the transit and access pricing by drawing a simple model presenting access cost and price formation. In addition, the effect of transit cost changes to ISPs cost structure is analyzed by presenting the effect of traffic growth and examples of found ways to save in transit costs.

1. Introduction

Prices of broadband Internet access have decreased during the whole 21st century while the amount of households having broadband access has increased. Decline of transit prices has been one of the enablers for lowered access prices, although not the only one. Other causes range from technological development to growing user base and tightening competition. But how much do transit prices really have effect on access prices compared to other factors?

Fundamental price ratio changes may have immense impact on both technical and industrial architecture of the Internet. The price ratio change between access and transit pricing has not been widely studied. Thus historical data and effects need to be analyzed to understand possible impacts for the future Internet evolution. The focus of this paper is on the access price formation from customer access –oriented ISP point of view and how the transit price changes affect that. Gathered access and transit prices for the last decade form the basis for the study. Due to historical view, only fixed access is considered, and wireless access is taken into account only in discussions concerning possible future development.

The paper is organized as followed. To understand how connectivity is arranged in the Internet, different connection methods are presented in the chapter 2. This covers both broadband access and ISP interconnectivity. On the access side especially DSL access is considered. ISP interconnectivity is examined from less technical perspective and the concentration is mainly on two types of interconnection agreements – namely paid transit and settlement-free peering. Chapters 3 and 4 go into transit and access pricing. Analysis of cost structure, price evolution and price formation is based on literature sources. Additionally the self-created model is used to explain the interdependency of transit and access prices. Chapter 5 discusses about the effects of transit price changes to ISPs costs and through that on access pricing. Finally, the findings and possible indications for future are summarized in the chapter 6.

2. Internet Interconnectivity

In the early days of the Internet, there was only one backbone network called NSFnet operated by National Science Foundation (Mills & Braun, 1987). After privatization of backbone and introduction of Border Gateway Protocol (RFC 1105, 1989) the market opened and the amount of backbone operators increased. Nowadays Internet consists of heterogeneous networks called autonomous systems (ASes), and the amount of them on February 2009 was over 30,000 (ASN, 2009). Autonomous systems are operated mostly by commercial Internet Service Providers (ISPs), but also by corporations and other enterprise providers, universities, government agencies, and content providers and other specialized service providers (Clark et. al, 2007). ISPs connect end users and businesses to the public Internet by selling Internet access. They compete over customers on price, performance, reliability etc. but they also must co-operate to offer universal end-to-end connectivity (Norton, 2001).

2.1. Interconnection agreements

The interconnectivity between ASes is arranged by two basic types of agreements – paid transit and settlement-free peering. The recursive combination of these standardized contracts resulting from complex and dynamic bargaining game between pairs of ASes creates the web of interconnections (Clark et. al, 2007).

Definition: A Transit Relationship is a business arrangement whereby an ISP provides (typically sells) access to the Global Internet (Norton, 2002b).

Transit relationship is clearly a customer (ISP A) – provider (ISP X) relationship. ISP X has responsibility of providing access to the global Internet – either by having itself the access or by purchasing it from another ISP. Through the transit relationship ISP A gets access to the global Internet and the Internet can access ISP A.

Definition: Peering is the business relationship whereby ISPs reciprocally provide access to each others' customers (Norton, 2002b).

A peering agreement is negotiated between equal peers. Parties share the costs and the traffic itself is exchanged directly between ISPs without any payment. Peering may either be arranged in public Internet exchange points (IXP) like Finnish FICIX[1] or then privately through private point-to-point connections. It is important to note that peering is not a direct substitute for transit. Transit provides access to the entire Internet whereas peering provides access only to each ISPs' customers.

The main motivation for peering is cost savings through reduced transit costs. If two ISPs, A and B, have significant amount of traffic between each others, and the traffic flows to both directions are symmetric[2], it is advantageous to both parties to change traffic directly without any payment and not through transit providers. This is illustrated in Figure 1. Cost savings are, however, not the only motivation for peering. Direct connection between peers reduces latency and enables thereby better performance for ISPs' customers (Norton, 2001).

2.2. Internet Ecosystem

The simplified structure of the Internet consists of Internet service providers, content and enterprise companies and end users. ISPs can be divided into two groups – Tier 1 ISPs and Tier 2 ISPs. Generalized Internet ecosystem is presented in the Figure 2.

Definition: A Tier 1 ISP is an ISP that has access to the global Internet routing table but doesn't purchase transit from anyone (Norton, 2001).[3]

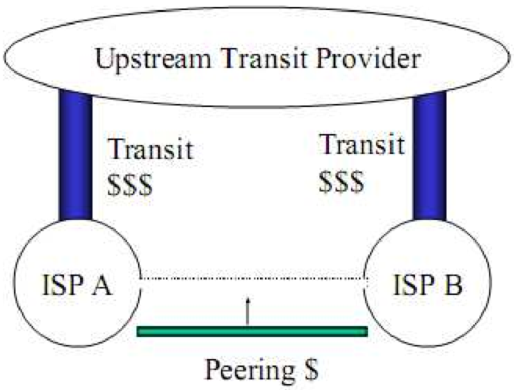

Figure 1: Cost savings through peering (adapted from Norton, 2001).

[1] http://ficix.fi/

[2] Typical traffic ratio requirements are either 2:1 or 1.5:1 (Clark et. al, 2007).

[3] Strict definition of Tier 1 ISP requires that ISP is not only transit-free but also all of its peering relationships need to be settlement-free. Technically there is no difference in settlement-free and paid peering, so the looser Tier 1 definition is used here.

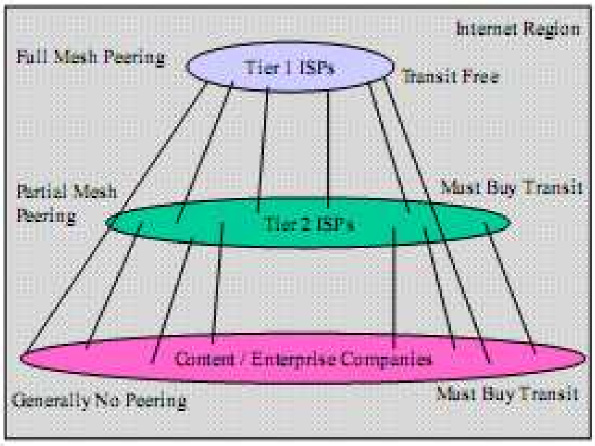

Figure 2: Generalized Internet ecosystem (adapted from Norton, 2003).

Since Tier 1 ISPs do not buy transit, they have to get access to the entire Internet solely through peering relationships. This means that every Tier 1 ISP must peer with all the other Tier 1 ISPs, and thus the amount of Tier 1 ISPs has stayed quite limited. According to Renesys Corporation (2009), there were 13 Tier 1 ISPs in January 2009.

Tier 2 ISPs are a heterogeneous group of ISPs that differ in geographical coverage, amount of customers and proportion of transit and peering traffic. Some small ISPs buy only transit and some large ISPs have vast amount of peering agreements. The common factor is that they still have to buy transit.

Content and enterprise companies are typically customers of ISPs. Their connectivity to the Internet is mostly based on transit agreements, and peering is rare. ISPs connect also end users (consumers) to the Internet by selling Internet access.

2.4. Consumer connectivity - Internet Access

In the Internet backbone and ISPs' core network the data is transmitted in optical fibers. End users can be connected to the Internet either using fixed or wireless access technologies. Due to the historical perspective of this paper the concentration is on fixed access, and wireless technologies, like WLAN, WiMAX or UMTS, are taken into account only in the concluding remarks concerning possible future development.

Obvious prerequisite for fixed Internet access is cabling from ISP premises to the customer. Commonly three types of cabling have been used to offer fixed Internet access: 1) copper cabling of PSTN, 2) coaxial cable of cable television and 3) optical fiber installed specially for Internet. Next chapter presents the most important technologies used in offering fixed broadband Internet access.

DSL (Digital Subsriber Line), and especially its asymmetric variant ADSL, is a dominant technology using copper cabling. It was standardized in 1999 (ITU,

1999) and later updates (ADSL 2, ADSL 2+) have increased possible transmission speeds up to 24 Mbps (LVM, 2004a). However, transmission speeds over 20 Mbps require short, under 1.5 km subscriber lines, which restricts the availability in rural areas. Because ADSL uses subscriber lines, the transmission speed can be guaranteed from the home to ISP premises.

Cable modem technology is standardized by CableLabs[4] and called Docsis (European version is EuroDocsis). The newest version of standard is 3.0, which enables both downlink and uplink transmission speeds over 100 Mbps. Cable modem offers better uplink transfer rates than ADSL, but the availability of cable is more limited and concentrated mostly in urban areas (LVM, 2007). Additionally, in cable modem technology the offered capacity is shared with all the users inside a network node.

FTTH (Fiber to the home) is technology in which optical fiber is brought to the homes. Optical fiber offers by far the fastest transmission rates (theoretically in Tbps) over huge distances (KPMG, 2005). The largest difference to the two other options is that in most countries there is no existing cabling. Bringing fiber to the homes is very expensive due to high installation costs, which together with unused potential of DSL and cable systems and lacking services requiring very fast transmission speeds have hindered the adoption of FTTH. Thus, for example Finnish operators, have developed broadband networks by bringing fiber closer to the customer (FTTB), especially in large cities (LVM, 2007).

Figure 3 presents the evolution of the number of broadband subscribers by access technology from 2004 to 2008. ADSL and cable modem have been the most important technologies. Other options, especially FTTH and wireless alternatives have constantly increased their market share (TeleGeography, 2009). Same kinds of figures are reported by FICORA (2008a) and OECD (2008).

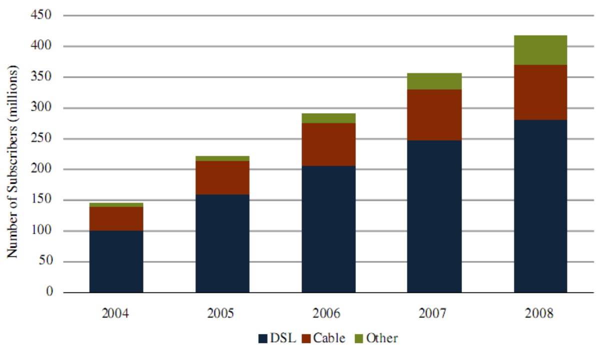

Figure 3: Broadband subscribers by access technology. Adapted from TeleGeography, 2009.

[4] http://www.cablelabs.com/ (accessed 3 March 2009)

3. Transit and peering pricing

In this chapter, the costs of connecting networks together are examined. Both transit and peering prices are covered, although the concentration and numerical data is presented only about transit. Understanding peering costs is anyhow relevant for understanding possible changes in Internet interconnection.

3.1. Transit and peering costs

Transit is normally priced usage based. The customer commits to some level of traffic and pays for that. The actual usage is measured[5] and compared to committed level of traffic, and if that is exceeded, additional usage is charged extra. Additionally there is typically an initial start-up cost and the price per Mbps declines stepwise when more traffic is committed (Norton, 2002b).

As a bought service, a transit customer does not require any infrastructure outside its core network and capital expenses are handled by the transit provider. Also the amount of operational costs relating to network maintenance and support are low, since transit contracts typically include service level agreements (SLAs) that make providers responsible for quick troubleshooting in problem situations (Clark et. al, 2007). As for transaction costs, purchasing transit is low cost compared to peering, since networks selling transit are interested in attracting customers and they carry most transaction costs.

Peering requires higher initial costs than transit. The key cost components of public peering are presented below in the Figure 4. Especially capital costs form important cost factor. These include transport hardware (routers, switches, etc.), rack space and switch port.

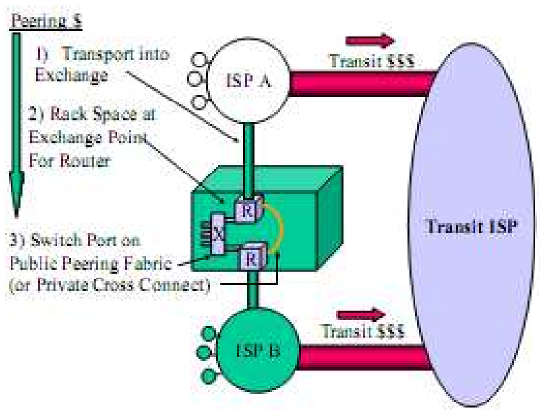

Figure 4: Cost components of public peering (adapted from Norton, 2002b).

[5] The common implementation is known as 95-5 rule, meaning that the traffic is measured in short intervals (e.g. five minutes,), which are then ranked in descending order to provide a measure of capacity utilization over time in Mbps (Clark et. al, 2007). The 95^{th} percentile is chosen to present the actual capacity utilization.

Also operational costs are higher than in transit, because operators have to deal themselves in problem situations, and the level of expertise needed is much higher than in buying transit. Additionally, transaction costs needed to enter into peering contract are high, since finding, contacting and negotiating with potential peer takes time and money, as well as configuration of network to match peer's peering policy.

As told in the chapter 2.1, main motivation for peering (especially from ISP's point of view) is possibility to decrease transit costs. Norton (2002b) has modeled operator's problem situation of choosing between peering and transit. The model is presented below in the generalized peering break-even graph (Figure 5). The main conclusion derivable from the graph is that after the amount of traffic exchanged between two ISPs exceeds the break-even point, it becomes sensible to peer instead of buying transit.

3.2. Transit price evolution

To get understanding of IP transit price evolution, actual price data is important. The problem with transit prices is that there are no list prices available but the prices are negotiated between ISPs and protected by non-disclosure agreements. Thus gathering data is difficult, and the amount of available data is petty. Telecom statistics and analysis company TeleGeography Research collects transit prices in the major cities worldwide, but unfortunately their data is not public. They have however presented some of their data in PTC conferences and these presentations are available (TeleGeography, 2007, 2008, 2009). Although the methods and details of data collection are not revealed and the data sets of different years' presentations are not comparable, the data is still used because of lacking alternatives.

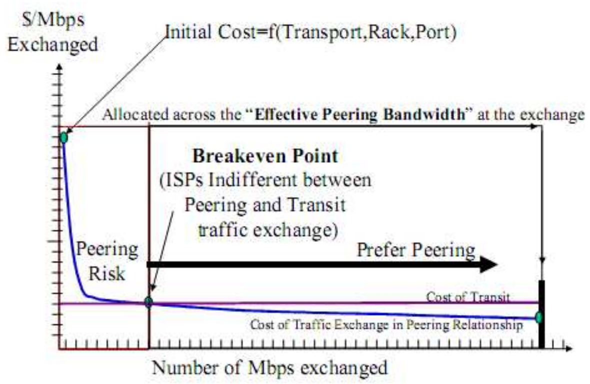

Figure 5: Generalized peering break-even graph (adapted from Norton, 2002b)

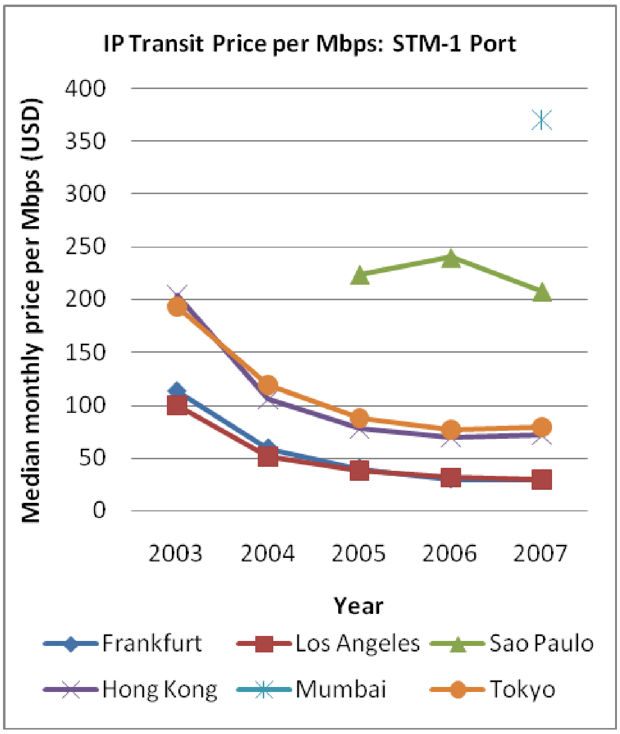

Figure 6: IP Transit Price per Mbps: STM-1 Port (adapted from TeleGeography, 2008).

Figure 6 present the transit prices per Mbps of STM-1 port in four cities (Frankfurt, Los Angeles, Hong Kong and Tokyo) from 2003-2007. Additionally a couple of data points from two other cities (Sao Paulo and Mumbai) is available. The main observation from the figure is that prices declined from 2003 to 2006 but stabilized after that. Interestingly the price development curves remind significantly each other.

The geographical difference in prices is clear. Prices in Asia are twofold and in developing markets like Brazil and India even over tenfold compared to Europe and US. The newest PTC presentation (TeleGeography, 2009) has comparable price data from 2005 to 2008 concerning Gigabit Ethernet Port. Also this data shows price stabilization during the year 2006, but the prices started to decline again in the summer of 2007. The compound annual price decline is still not very fast, only 16-26 % depending on the geographical location.

3.3. Factors affecting transit prices

Telegeography (2009) graphs show that transit prices vary not only by committed traffic level and geographical location but also by capacity and carrier (TeleGeography, 2009). Generally speaking price per Mbps decreases with increasing capacity (e.g. $36/Mbps in 10 Mbps port compared to $8/Mbps in 10 GigE port) and the prices for higher capacities decline faster. Carriers also price their offering differently depending on for example served market segments and SLAs. Also bargaining position of an ISP has significance on transit prices.

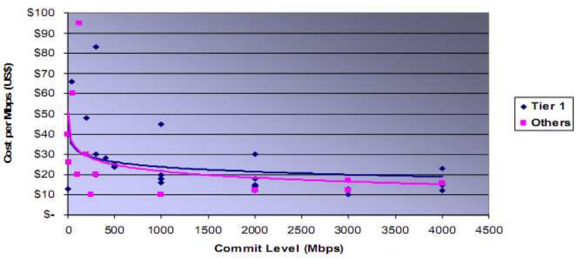

Figure 7: NANOG 2006 Transit Survey (adapted from Norton, 2006).

The impact of all the factors affecting transit prices can be observed in the Figure 7, which presents the paid transit prices of 42 ISPs in 2006. The data was gathered by Norton in 2006 NANOG meeting and presents the costs per Mbps in the function of commit level. Average cost per Mbps was $25, but it varied highly, especially in the small commit levels. The main observation is the importance of commit level, but the variation inside the same commit level can be explained by the other factors: geographical location, capacity, carrier and bargaining position of an ISP.

4. Access pricing

In the 90s, when analog modems and ISDN where used to connect to the Internet, pricing was usage based, but the broadband brought simpler flat rate pricing with. This development follows Odlyzko's (2001) observation of traditional evolution of communications that prices decrease and become simpler while quality rises and increased usage increases total revenues. Consumers (and providers) like flat rate due to its predictability. Additionally ISP saves money, because extensive accounting is not needed. The drawback of flat rate is the possibility to uneven resource usage.

Because Internet is best effort network and attempts to insert quality of service guarantees have mostly failed, the product differentiation of access subscriptions is done based on differing transmission rates. Due to asymmetric traffic pattern of most Internet applications, the concentration is on downlink speeds. Typical ADSL subscriptions available during the last ten years have had downlink transmission speeds from 128 kbps to 24 Mbps whereas uplink speeds have varied between 128 kbps and 2 Mbps (LVM, 2001-2008). Although pricing has been flat, ISPs in some markets have used traffic limits to restrict excess use of resources (LVM, 2004b & 2006).

4.1. Costs of offering Internet access

McKinley and JPM analysis presented by Newman (2003) (Figure 8) illustrates costs of offering DSL and cable subscription. Prices are for U.S. market and DSL costs are calculated for incumbent local exchange carrier (ILEC). Although the data is little bit old and estimated, the cost components and their coarse proportions provide interesting information.

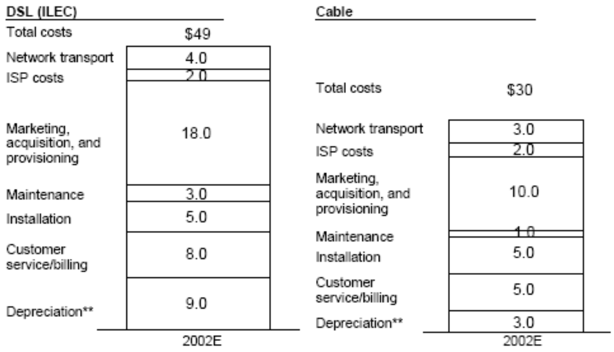

Figure 8: Average broadband costs per subscriber per month in 2002 (Adapted, from Newman, 2003).

Importantly, the costs relating to technology employed form only a small fraction. Network transport forms approximately only 8-10 % of costs but unfortunately transit and peering costs are not specified. Additionally some capital expenses relating to network transport are included in depreciation. Regardless of that, the proportion of costs relating to network transport are small when compared to the marketing and customer service costs, which cover about half of the costs.

In Newman case, the local loop cabling was neglected due to existing and provider owned subscriber lines. However, if the subscriber line is not owned by an ISP, leasing existing local loops forms relevant cost component for this study. Copper cables are normally owned by incumbent telco and coaxial cables by cable TV provider. Because installing parallel cables is not reasonable and monopolistic ownership would form barrier for market entry, local loop unbundling regulation has been introduced in many countries to foster competition (Bourreau & Dogan, 2005). New entrants can lease subscriber lines from incumbent ISPs which enables market entrance without high initial costs. These costs naturally affect access prices. When new access technologies, like FTTH, are considered, there is not existing cabling. In the case that an ISP wants to offer FTTH access, huge investment on cabling will show up in access prices for long period of time.

4.2. Access price evolution

Ministry of Transportation and Communications Finland has gathered broadband access prices from the year 2001 (LVM, 2001-2008). The prices of ADSL access of different transfer rates are presented in the Figure 9. The prices are in euro and they present the mean value. It can be seen that the prices have declined significantly in seven years. The faster the connection, the more the prices have dropped. Especially rapid change happened during the years 2003 and 2004, but after that the price level has stabilized. In recent years providers have concentrated on bringing faster downlink transfer rates available for lower price, for example all the large operators[6] provide 24/1 Mbps transfer rate for approximately 50 euro. At the same time the price of the slowest connection, 256/256 kbps, has actually increased. Besides, most operators do not sell transfer rates below 1 Mbps anymore. Transfer rate growth, however, is limited to downlink, which increases the asymmetry. For most users and applications this is irrelevant, but for example peer-to-peer applications would enjoy faster uplink rates.

Ministry of transportation and communications Finland has also gathered ADSL access prices for 512 kbps, 1 Mbps and 2 Mbps connections from all EU countries during the years 2003 to 2005 (LVM, 2004b & 2006). Figure 10 presents the mean monthly prices of 512 kbps subscription in those countries, for which there was data available for every year from 2003 to 2005. The evolution in other countries follows Finnish markets in declining price trend, although in many countries the prices dropped much more from 2003 to 2004 than 2004 to 2005. Actually only Austria and Sweden differ from the normal pattern. Moderate price declines in Sweden can be explained by the well-developed markets, since the operators concentrated on higher speed subscriptions, the trend that started in Finland just in 2006.

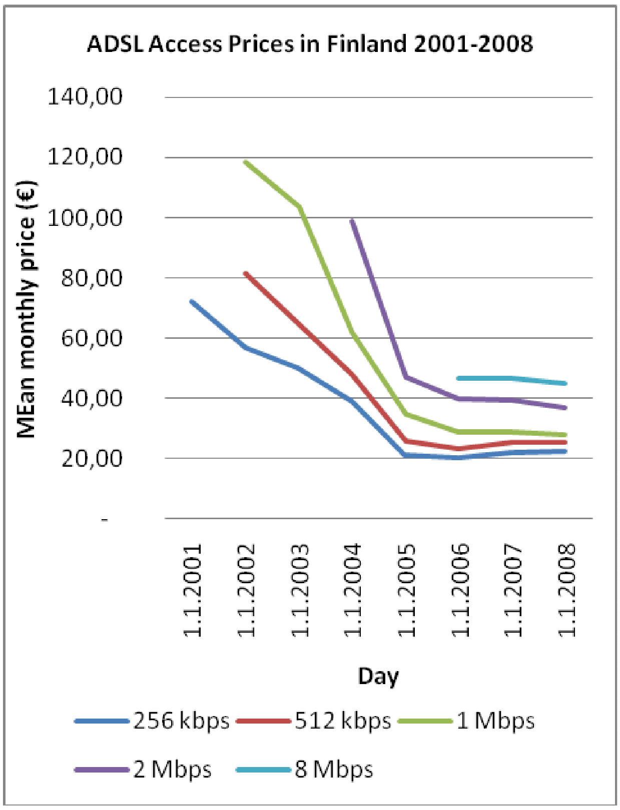

Figure 9: ADSL Access Pricing 2001-2008 in Finland (LVM, 2001-2008)

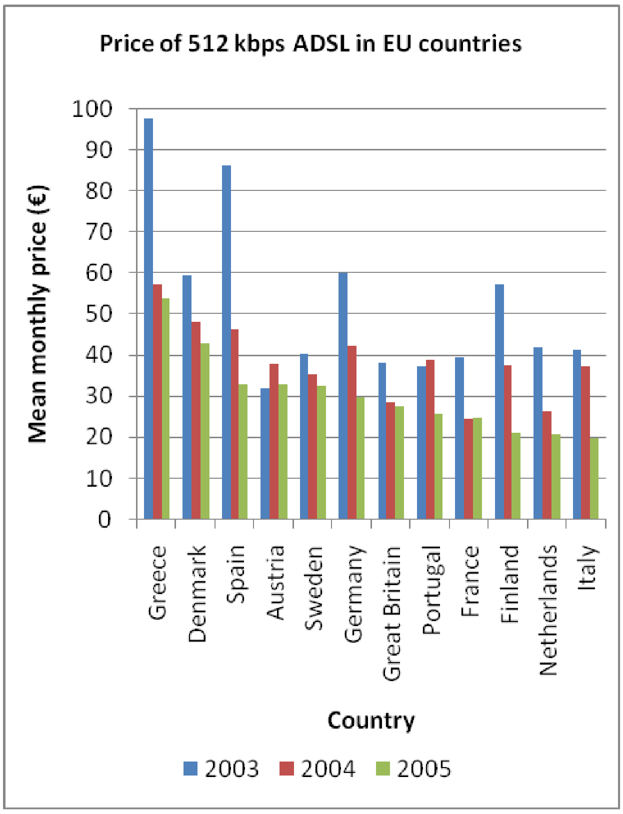

Figure 10: Price of 512 kbps ADSL in EU countries (LVM, 2004b & 2006)

If the access prices are compared in €/Mbps, there is a clear difference between different speed classes. Figure 11 presents the same data as Figure 9, but now the prices are made comparable by using euro per Mbps as the y-axis. Overall the price difference between different transmission speeds has remained in the same level. The slowest connection, 256 kbps has been for last three years over 15 times more expensive than the 8 Mbps connection.

Although the price comparison of access prices in EU countries (Figure 10) showed quite a similar price level in all countries, the OECD level comparison from year 2003 Figure 12 reports huge differences in per Mbps prices[7]. But when the actual data source for the figure (ITU, 2003) is inspected, it is revealed that large part of the differences can be explained by the different transfer rates. In the comparison the connection speeds range from 256 kbps (Denmark) to 10 Mbps (Japan). Even though the price difference between these two countries is over 100-fold when price per Mbps is considered, the monthly payment is only twofold ($48.97 vs. $21.29). Generally the monthly prices vary between 20 and 55

[6] Operator web sites (Elisa, Saunalahti, Sonera, DNA) on 2 March 2009.

[7] The prices are purchasing power parity adjusted ADSL prices (except US Comcast that is cable) from single operators.

dollars, which are all under 1 % of the monthly household income. This leads to the conclusion that the acceptable monthly broadband price from the customer perspective is in the same level around the world and the offered transmission speeds depend on other issues.

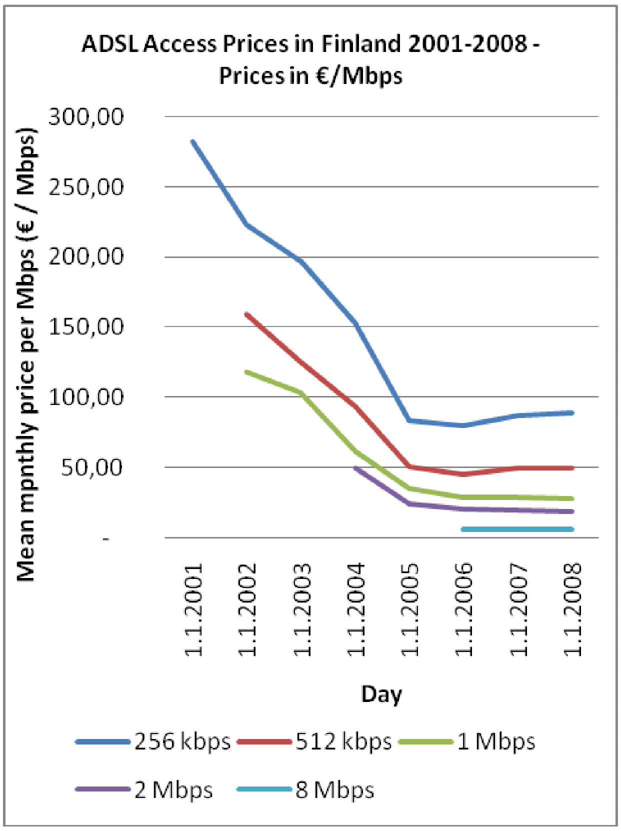

Figure 11: ADSL Access Prices 2001-2008 in Finland - Prices scaled to 1 Mbits/s (LVM, 2001-2008)

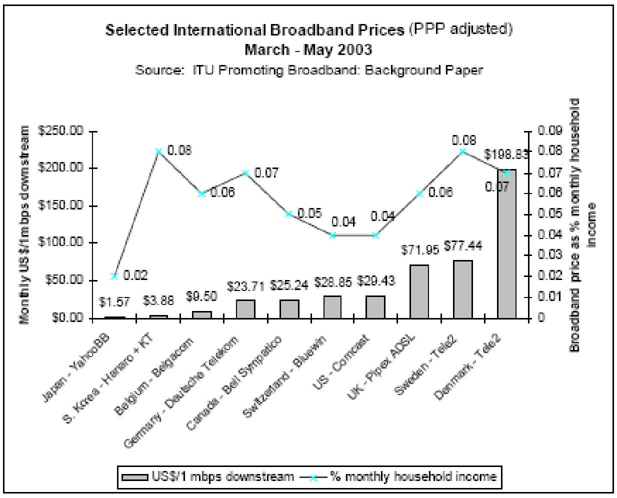

Figure 12: International comparison of broadband prices (adapted from Ismail & Wu, 2003).

4.3. Factors affecting access prices

Based on the Figure 11 and Figure 12 it is clear that access price does not follow transfer rate in a linear fashion, but fast subscriptions cost far less per Mbps than slow ones. This can be examined from both customer and ISP sides. From the customer perspective there is an interesting question, how are the connectivity and better transmission speeds valued. It can be argued that the connectivity is more valuable than the transfer rate. It can also be questioned, if every extra-Mbps is equally valuable to the customer.

The inspection from ISP side is more straightforward. The fixed costs for offering access form substantial part of the cost structure (Newman, 2003) and of access price. Besides, it is not clear, how much more traffic, and thereby costs, the customers having faster subscriptions cause. Data networks are typically lightly utilized (Odlyzko, 2003). The Internet backbone utilization (Figure 13) remained below 55 % in the year 2008 and utilization decreases when moved towards edges of the network (Odlyzko, 2003). Due to the light utilization, faster transmission speeds do not necessarily increase require updates to the ISPs network. One anecdotal thing is also that the customers with slower connections possibly subsidize the customers with faster connections.

When access price evolution analysis is brought further, it can be claimed that even though the same technology is available around the world, the market situation has huge importance in the transmission speeds and thus in price per Mbps. Market maturity is one aspect. The prices decrease faster when the markets are growing, and after saturation (in Finland 2007) the price level stabilizes. At least in Finland, the ISPs have in recent years constantly kept the prices in the level that customers have accepted but increased transmission speeds and made their offering more lucrative that way (Kilpailuvirasto, 2008). Competition is another market-related phenomenon that has important effects. Increasing competition drops prices and encourages ISPs to provide better services. The combination of high growth market and fierce competition can be seen in Finnish ADSL prices (Figure 11) during the year 2003, when the prices declines approximately 50 %.

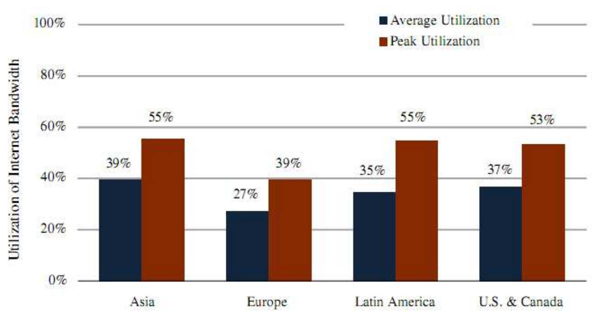

Figure 13: Average and peak backbone utilization by region in 2008. Adapted from Telegeography, 2009.

Regulator can also affect pricing. Means to increase competition, like local loop unbundling, is an example of indirect ways. Requiring uniform pricing is an example of more direct ways. In sparsely populated countries costs to offer access in urban and in rural areas differ clearly. If pricing is left to be defined by the markets, the access prices will probably be higher in the countryside or the services are not even offered in some regions. Regulator may intervene by requiring uniform pricing and/or higher geographical coverage (Foros & Kind, 2003).

4.4. Interdependency model

Different factors of ISP's costs of offering Internet access and the Internet access price formation are combined to interdependency model presented in Figure 14. The study concentrates on total costs, and per Mbps prices are showed only for transit and peering. The relationships between factors are showed by arrows. Plus-sign in the end of arrow indicates that change in the value of the start of the line creates change to same direction in the end of the line. For example, increase in the number of customers increases amount of traffic. Respectively, minus-sign means reverse relationship, e.g. increasing level of competition decreases access prices. Additionally, an arrow without sign indicates important factor whose impact cannot straightforwardly be seen direct or inverse.

The main idea behind the figure is following. Number of customers and their actual network utilization increases the amount of traffic. Transmitting the traffic creates indirectly large part of total costs. Increasing amount of traffic requires ISP network upgrading. It also increases transit and peering costs. If higher percentage of traffic is transit, transit price per Mbps declines but the overall transit costs increase. Same applies to peering. Marketing, customer service and billing costs are added up to total cost. The final access price per customer results from total costs. Larger number of customers results in lower price per customer. Level of competition and regulator's actions has also effect on access price.

The first thing to note is the importance of traffic amount. More traffic means more costs, although the dependence is not linear and actually the price per Mbps decreases. Almost as important question is how the traffic is transported. There are three possible options, where the end point of the customer traffic is located: 1) inside the ISP network, 2) in the network with which ISP has peering agreement, 3) in the network that is reachable only through a transit link. ISP network upgrading is needed in every case, but it can be possible to save in peering and transit costs.

The other aspect that requires underlining is the combination of access rates and actual network utilization. As expressed in the chapter 4.3 an ISP can offer higher transmission speeds without substantial extra costs, if the utilization of access connections remains in the same level. The users experience lower latency, but the amount of traffic does not increase significantly. This is supported by Odlyzko (2004), who has come to conclusion that it is getting moderate amount of bits quickly what matters to users, not huge amount of data. However, faster transmission speeds enable usage of new, bandwidth-intensive applications. Thus the overall effect of increasing transmission rates to the amount of traffic is unclear and depends on usage patterns.

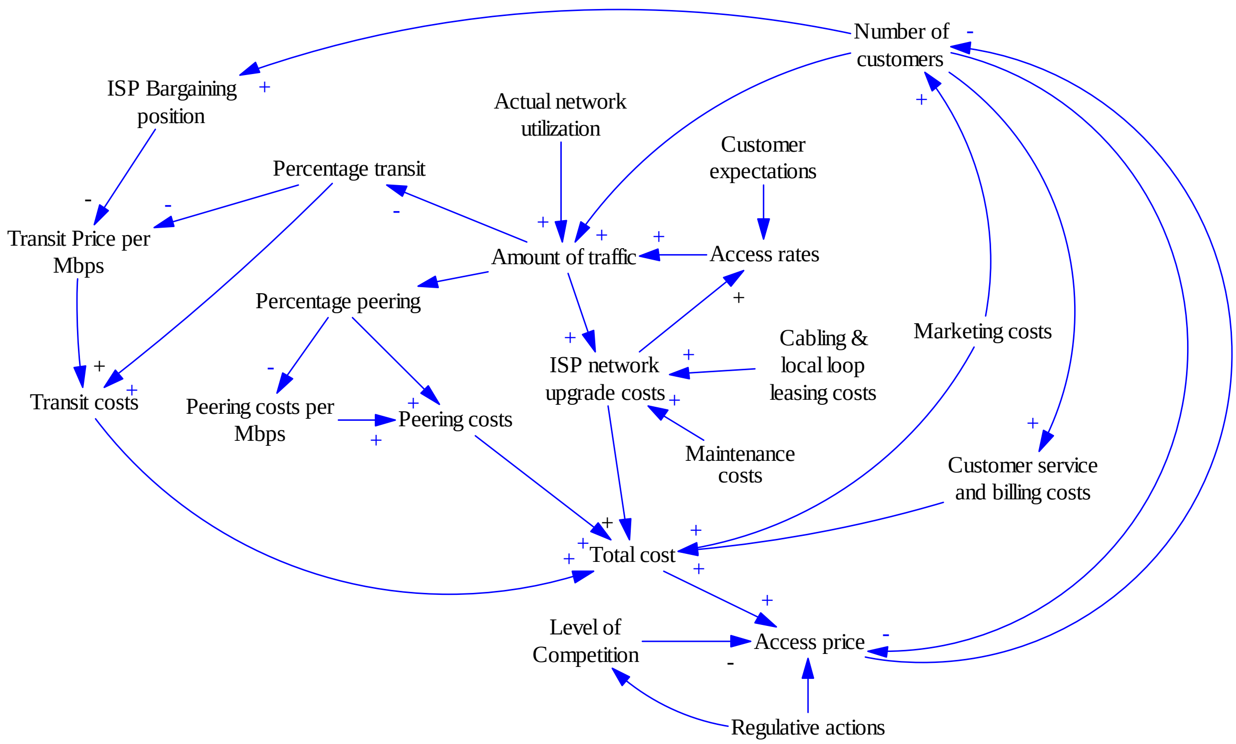

Figure 14: Interdependency model of Internet access costs and pricing

5. Effects of transit cost changes

Short time span, petty amount of transit price data and the complexity of access price formation together influence so that direct comparison of access and transit price evolution is not sensible. The key finding is that although transit costs are only one part of an ISP's cost structure, their role is important since the ISPs have streamlined their other costs and cost savings are started to be sought also from transport costs. Thereby the analysis is focused on how transit cost changes affect on an ISP's costs and what can be done to save in transit costs.

5.1. Effect of traffic growth

As explained in the chapter four, the impact of transit costs, and thus the impact of transit price evolution, is only indirect to the access pricing. Since the market maturing and economies of scale have dropped ISPs' other costs, the costs relating to the amount of traffic are growing their role in ISPs' cost structure. The annual growth rate of Internet traffic was around 60 % in 2005-2007 (Figure 15), meaning that traffic amount has doubled every 18 months. At the same time, annual transit price decline was between 16 and 26 %, as told in chapter 3.2. This means that the traffic growth has outpaced transit price declines, and thus the transit costs of ISPs have increased although the transit prices have declined. This is demonstrated in Figure 16, which presents the transit price declines from 2006 to 2007 compared to the traffic growth. Increasing traffic also increases other costs, especially those related to upgrading ISPs' core network. At the same time access prices have declined, meaning lower revenues per customer. Increasing customer base has helped and brought scale advantages, but also increased costs through increased traffic.

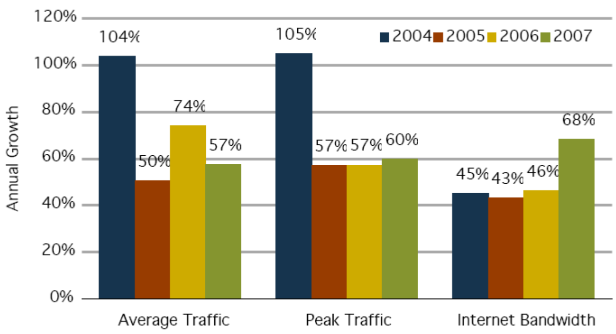

Figure 15: International Internet traffic and bandwidth growth 2004-2007 (adapted from TeleGeography, 2008)

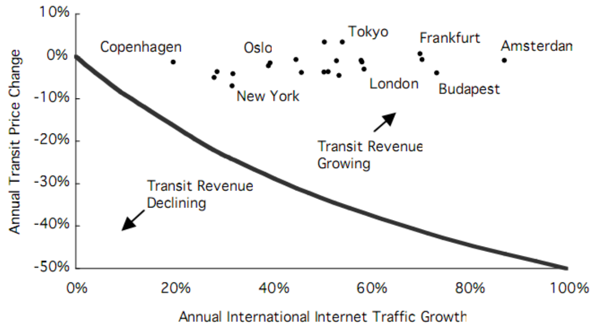

Figure 16: IP Transit Price Declines versus Traffic Growth in Selected Cities, Q2 2006-Q2 2007 (adapted from Telegeography, 2008).

5.2. Ways to save in transit costs

After all, transit and peering costs – even though they form only small part of costs – are the ones, in which ISPs may try to save money. There are two broad options available. Firstly, ISPs can try to restrict excess use. Most users behave well from ISPs' point of view and use only moderate amount of bandwidth. Nevertheless, there are also bad-behaving users who hoard bandwidth and thus produce high costs. Second option is to try to affect on the proportion of the most expensive traffic type – transit traffic.

The ISPs have used couple of means to restrict excess use. Bandwidth hoarding has raised a question, if unlimited flat rate is reasonable. A monthly bandwidth cap is a simple way to limit the excessive use and in some markets ISPs have used them successfully (LVM, 2004b & 2006). Interestingly, TeleGeography (2009) reports that many ISPs in the U.S. are imposing monthly bandwidth caps, which supports the assumption that proportion of transport costs is increasing. Since P2P file sharing produces high percentage of traffic (25 % according to TeleGeography (2009) and 43-70 % (depending on region) according to ipoque (2009)), any means to restrict its cost effects are taken into inspection. Different kinds of P2P management and deep packet inspection systems are available and some operators have used them to block or restrict P2P file sharing networks. Comcast is one of those ISPs who have been accused (and proved) to be using them (MSNBC, 2007). These P2P traffic restrictions have created passionate discussion about net neutrality and using them may affect negatively in the image of an ISP.

An ISP can decrease the proportion of transit traffic either by keeping the traffic inside ISP's core network or negotiating peering agreements, as mentioned briefly in the chapter 4.4. The most low-cost option is, if the traffic can be kept inside ISP's network. Thereby caching most popular content is highly interesting option to an ISP. Even storing some popular and bandwidth-intensive content, like YouTube[8] videos, in the servers inside ISP's network may be an affordable solution. The large amount of P2P traffic has lead to searching ways to optimize cost effects of peer-to-peer traffic also other ways than restricting it. P4P concept tries to keep traffic inside ISPs core network by favoring peers within the same ISP and in close geographical proximity instead of typical way of P2P where peers are connected randomly. The technology has not yet been widely adopted, but field tests of Telefonica and Verizon have proven promising (TeleGeography, 2009).

Increasing transit costs make peering more lucrative option. Especially interesting change is happening because of networks are specializing to access-selling ISPs and large content providers (like Yahoo and Google). Norton (2002a) and Clark et al. (2007) have named these providers to eyeball networks and content networks, and both of them have highly asymmetric traffic patterns, which typically has prevented peering. Earlier eyeball-heavy access operators have valuated themselves higher than content-heavy networks, but this situation is changing since the amount of content increases. Due to business reasons, changing traffic reciprocally with content-heavy networks starts to look sensible for ISP's, even though the traffic flows would be highly asymmetric. If all the content would be transmitted through the transit links, only party benefiting would the transit provider. This change in peering incentives affects not only on ISPs transit costs but also increases the complexity of Internet interconnection.

6. Conclusions

Both transit and access prices have declined in recent years. Access price depends on transfer rate but the relationship is not straightforward. Fast subscriptions cost far less per Mbps than slow ones, which underlines the high fixed costs of offering Internet access. Access price change has stabilized, but the ISPs are bringing faster transmission speeds available, which decreases the prices in €/Mbps.

Based on presented price data and access price formation, it can be said that the relationship between transit and access pricing is definitely not linear. This can be reasoned since the transport costs form only one, small part of ISPs costs, and actually marketing and customer service are far larger cost factors. This situation is changing since the ISPs have optimized their other costs and the proportion of transport cost in ISP's cost structure is increasing.

The importance of relationship between traffic growth and transit price evolution is one of the key findings. In recent years the traffic growth has outpaced transit price declines meaning higher transit costs to ISPs. Thus the amount of traffic is as important issue as transit pricing. Increased transit costs (and other costs to offer Internet access) can be seen in many ways. In last couple of years the access price evolution has stabilized, especially in the slowest transmission speeds. The ISPs have fought against price erosion by offering faster subscriptions for the same price (Kilpailuvirasto, 2008). Offering faster transmission speeds is not a problem to ISPs, if that does not change consumers' usage patterns and increase traffic amount. Recent introduction of bandwidth caps in the U.S. supports this finding.

The ISPs fight against increasing traffic amount by trying to limit bandwidth hoarding usages like P2P file sharing. Also means to decrease the most expensive traffic, namely transit, gain popularity. Increased peering is one option, caching the most popular content or even placing the servers of most popular content providers are the other ones. Altogether keeping a larger deal of traffic inside the ISP's network is a good way to save costs.

This paper leaves still some questions open. If we consider basic access technologies, wireless options start to disturb and compete against fixed options. Another interesting issue is, will FTTH be widely adopted and how fast that happens. Fiber would drop per Mbps prices to totally new level and enable huge increase in bandwidth and traffic amount. To improve the analysis, the peering cost evolution could be included. Another interesting question is how increase in transmission rate increases the amount of data. Also understanding costs of ISP network upgrading would be relevant in the future.

References

ASN (2009) The 32-bit AS Number Report. http://www.potaroo.net/tools/asn32/ (accessed 25 February 2009).

Bourreau M., Dogan P. (2005) Unbundling the local loop. European Economic Review, Vol. 49, pp. 173-199.

Clark, D., Faratin, P., Gilmore, P., Bauer, S., Berger A., and Lehr W. (2008) Complexity of Internet Interconnections: Technology, Incentives and Implications for Policy. http://web.si.umich.edu/tprc/papers/2007/797/Clark%20Lehr%20Faratin%20Complexity%20Interconnection%20TPRC%202007.pdf (accessed 25 February 2009).

FICORA (2008a) Telepalveluiden käyttötutkimus 2007. http://www.ficora.fi/attachments/suomiry/5vnqZFclE/Tutkimusraportti2007Telepalveluidenkayttotutkimus.pdf (accessed 2 March 2009).

Foros, Ø., Kind H. J. (2003) The broadband access market: competition, uniform pricing and geographical coverage. Journal of Regulatory Economics, Vol. 23, pp. 215-235.

ipoque (2009) Internet Study 2008/2009. http://www.ipoque.com/resources/internet-studies/internet-study-2008_2009 (accessed 3 March 2009).

[8] http://www.youtube.com

Ismail S., Wu I. (2003) Broadband Internet access in OECD countries: A comparative analysis http://www.caltelassn.com/reports/BroadbandOECD7Oct03.pdf (accessed 1 March 2009).

ITU (1999) G.992.1 : Asymmetric digital subscriber line (ADSL) transceivers. http://www.itu.int/rec/T-REC-G.992.1/en (accessed 2 March 2009).

ITU (2003) Promoting Broadband: Background Paper http://www.itu.int/osg/spu/ni/promotebroadband/PB03-PromotingBroadband.pdf (accessed 10 March 2009).

Kilpailuvirasto (2008) Kilpailukatsaus 24.11.2008. http://www.kilpailuvirasto.fi/tiedostot/kilpailukatsaus-24-11-2008.pdf (accessed 25 February 2009).

KPMG (2005) Europe’s Broadband Future. http://www.kpmg.fi/Binary.aspx?Section=174&Item=2579 (accessed 3 March 2009).

LVM (2001-2008) Suomen telemaksujen hintataso – selvitykset 2001-2008, LVM:n julkaisut 18/2001, 15/2002, 15/2003, 21/2004, 27/2005, 19/2006, 17/2007 ja 15/2008.

LVM (2004a) Laajakaistatekniikoiden kehitys 1995–2010. http://www.lvm.fi/fileserver/53_2004.pdf (accessed 25 February 2009).

LVM (2004b) EU-maiden laajakaistahinnat http://www.lvm.fi/fileserver/Julkaisuja%2076_2004.pdf (accessed 3 March 2009).

LVM (2006) EU-maiden laajakaistahinnat 2005 http://www.lvm.fi/fileserver/Julkaisuja_2_2006.pdf (accessed 3 March 2009).

LVM (2007) Kansallinen laajakaistastrategia - loppuraportti http://www.lvm.fi/fileserver/Julkaisuja%203_2007.pdf (accessed 3 March 2009).

MSNBC (2007) Comcast blocks some Internet traffic. http://www.msnbc.msn.com/id/21376597/ (accessed 3 March 2009).

Mills, D. L., Braun H. (1987) The NSFNET backbone network. Proceedings of the ACM workshop on Frontiers in computer communications technology, pp. 191 – 196.

Newman, S. (2003) Broadband Access Platforms for the Mass Market: An Assessment. http://web.si.umich.edu/tprc/papers/2003/254/BbandAccessPlatforms.pdf (accessed 25 February 2009).

Norton, W. B. (2001) Internet Service Providers and Peering. http://www.cs.cmu.edu/~srini/15-744/readings/Nor01.pdf (accessed 25 February 2009).

Norton, W. B. (2002a) The Art of Peering: The Peering Playbook. http://arneill-py.sacramento.ca.us/ipv6mh/playbook.pdf (accessed 25 February 2009).

Norton, W. B. (2002b) A Business Case for ISP Peering. http://pages.cs.wisc.edu/~akella/CS640/F06/readings/norton.pdf (accessed 25 February 2009).

Norton, W. B. (2003) Evolution of the U.S. Peering Ecosystem. http://www.nanog.org/mtg-0405/pdf/norton.pdf (accessed 25 February 2009)

Norton, W. B. (2006) Transit Cost Survey, Data collected at NANOG 36. http://www.nanog.org/mtg-0606/pdf/bill.norton.2.pdf (accessed 25 February 2009).

Odlyzko, A. (2001) Internet pricing and the history of communications. http://www.dtc.umn.edu/~odlyzko/doc/history.communications1b.pdf (accessed 25 February 2009).

Odlyzko, A. (2003) Data Networks are Lightly Utilized, and will Stay that Way. Review of Network Economics, Vol. 2, pp. 210-237.

Odlyzko, A. (2004) Pricing and Architecture of the Internet: Historical Perspectives from Telecommunications and Transportation. http://www.dtc.umn.edu/ odlyzko/doc/pricing.architecture.pdf (accessed 25 February 2009).

OECD (2008) OECD broadband statistics - OECD Broadband subscriptions, by technology, June 2008. http://www.oecd.org/sti/ict/broadband (accessed 2 March 2009).

Renesys Corp. (2009) Internet Captivity and the De-peering Menace. http://www.renesys.com/tech/presentations/pdf/nanog-45-Internet-Peering.pdf (accessed 25 February 2009).

RFC 1105 (1989) A Border Gateway Protocol (BGP). http://www.ietf.org/rfc/rfc1105.txt (accessed 3 March 2009).

Telegeography (2007) Global Traffic, Bandwidth, and Pricing Trends and Wholesale Market Outlook, PTC 2007. http://www.ptc.org/past_events/ptc07/program/presentation/ST_BeckertSchoonover.pdf (accessed 25 February 2009).

Telegeography (2008) TeleGeography International Telecom Trends Seminar, PTC 2008. http://www.ptc.org/ptc08/participants/speakers/papers/SchoonoverFinalSlides.pdf (accessed 25 February 2009).

TeleGeography (2009) TeleGeography International Telecom Trends Seminar, PTC 2009. http://www.ptc.org/ptc09/images/papers/PTC%2709_TeleGeography_Slides.pdf (accessed 25 February 2009).

Price Ratio Evolution of Cellular Data and WiFi

Thomas Casey and Juuso Karikoski
{thomas.casey, juuso.karikoski}@tkk.fi

Abstract

In this paper we try to model the price ratio evolution of Cellular Data and WiFi in the past and speculate how it could evolve 5 years into the future. We use system dynamics to get an overall understanding of how the price and value ratio of Cellular Data and WiFi could evolve and what could be the critical issues affecting the evolution.

1 Introduction

Wireless data usage is growing in a rapid rate. Both WiFi (Wireless Fidelity) and Cellular Networks have become common data connection enablers with prices for both steadily decreasing. The increase has been fueled for example by simple flat rate pricing and growing coverage of 3G (Third Generation) network services and WiFi access point footprint.

Although Cellular Networks clearly beat WiFi in terms of absolute geographical coverage, the coverage of WiFi for typical nomadic users (i.e. users that roam around but remain static while using the wireless data connectivity) in an urban densely populated environment can be quite good.

In a situation where both connectivity types are available, Cellular Data and WiFi could be considered as substitutes to each other. And since the largest factor for a consumer is arguably price, it is interesting to examine how Cellular Data and WiFi prices and price ratio have evolved, how they could evolve in the future and what could be the resulting impact on their usage.

Will ubiquitous, easy to use, vertically integrated Cellular Data dominate in the future or will WiFi serve a significant share of the demand for wireless connectivity? This depends on many factors and we try to model and elaborate those in this study with system dynamics.

The rest of the paper is organized as follows. In chapter 2 we will define the scope and assumptions regarding our study. In chapter 3 we conduct an analysis of Cellular Data and WiFi historical price evolution, model the pricing structure of both and try to project the trends and uncertainties affecting the price. In chapter 4 we first build a simple system dynamics model and describe two possible extreme scenarios, one where Cellular Data clearly dominates and one where WiFi becomes a considerable substitute for Cellular Data. In chapter 4 we also try to build a more comprehensive model with more variables and speculate what could be the actual outcome between the two extreme scenarios. Finally in chapter 5 we draw conclusions.

2 Scope and assumptions

In order to be able to study the price ratio in a meaningful way, we need to narrow down the scope of our model and make some assumptions. We focus on studying typical nomadic laptop users in a large European city. We also assume that no separate Quality of Service (QoS) is provided and that the users are served with a Best Effort service class. We focus only on connectivity and the applications or content that are used won't be within the scope. The city is assumed to have both ubiquitous Cellular Data coverage and a considerable WiFi footprint.

Cellular Data refers to all forms of Cellular Data connectivity, i.e. GPRS (General Packet Radio Service), EDGE (Enhanced Data rates for Global Evolution), UMTS (Universal Mobile Telecommunications System), HSPA (High-Speed Packet Access) and LTE (Long Term Evolution) and its evolved versions (possibly featuring user deployed and operated Femtocells in the future). The WiFi footprint includes all forms of WiFi connectivity available to a nomadic user, from private home WiFi access, to public WiFi (e.g. city operated, or hotspot based), enterprise WiFi, operator WiFi and WiFi community access.

The price for the consumer is the target for our study, but the operator costs are also considered as part of the pricing structure. Our study period for price and price ratio evolution is based on available historic data from the near past and range 5 years into the future.

For simplicity, in this study Cellular Data and WiFi are considered as substitutes to each other for the nomadic users, i.e. either Cellular Data or WiFi is used exclusively. In reality, especially technically competent users might be using both Cellular data and WiFi for example to save on costs or to increase capacity. Nevertheless, provided that flat rate prices are used and that sufficient QoS can be supplied, Cellular Data could be the only source of connectivity for many users (especially for the ones in need of vertical integration).

3 Price evolution analysis

This chapter gives insight to and analysis of the historical price data for mobile data. The historical price data presented in this chapter is collected from Finland and used to evaluate the possible effects the changes in the price ratio between Cellular Data and WiFi might have on the Internet technology and business architectures.

3.1 Analysis of and insights to historical prices for mobile data

The historical price data has been divided to WiFi and Cellular Data prices. The prices have been collected from the operators' web pages using the Internet archive[1] and from various price comparisons conducted by the Finnish Communications Regulatory Authority[2] and the Finnish computing magazine Tietokone[3] for example.

The WiFi prices have been divided to operator provided monthly 1 Mbps WiFi subscriptions for private customers and public hotspots. The 1 Mbps bit rate was selected, because it was the subscription that all of the operators offered. The monthly prices have been gathered from various operators like Haminetti[4], MSOYNET[5] and WiFinet[6] for example and the hotspot prices from operators like Netikka[7], SparkNet[8] and Zonet[9].

Figure 1 below shows the evolution of the monthly WiFi subscription prices between 2003 and 2009. As one can clearly see, the price trend is decreasing. There is a 43 percent decrease in the prices since 2003.

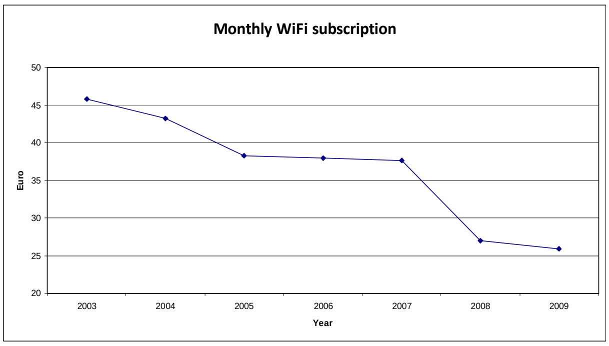

Figure 1. WiFi monthly subscription prices

Regarding the hotspot prices such a decreasing trend could not be observed. All the operators offering public WiFi hotspots studied in this report had exactly the same prices in the year 2009 as they had in 2005 or 2006. The average price for 10 hours of WiFi hotspot usage is approximately 10 Euros. This might be result of a local monopoly that the hotspot providers have in their premises.

The mobile data prices in the cellular networks were divided to prices per used megabyte and prices for a monthly flat rate subscription of up to 384 kbps downlink capacity. This bit rate was selected because it was the only bit rate offered by all of the operators during the whole reference period.

The Cellular Data prices have been collected from Finnish operators including Elisa[10], DNA[11] and Sonera[12] and their virtual operators Saunalahti[13], Kolumbus[14] and TeleFinland[15].

The Cellular Data prices per megabyte had the same price trend as WiFi hotspots. All the operators had the same prices in 2009 for one megabyte of usage as they had in 2005, except for Sonera. Sonera changed their pricing scheme from megabyte based to time based usage in 2007, but nonetheless their time based prices also stayed the same during 2007-2009.

Regarding the flat rate subscriptions similar behavior as in monthly WiFi subscriptions can be observed. As Figure 2 shows the flat rate prices have gone down 84% from 64 Euros a month to less than 10 Euros a month between 2005 and 2009. Cellular Data flat rate subscriptions have only been available since 2006 and thus the price for 2005 is an approximation for 100MB of monthly data usage for subscriptions that had a data block included in the monthly fee.

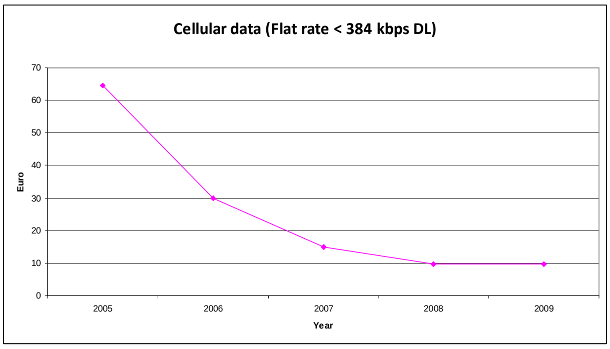

Figure 2. Cellular flat rate subscription prices

Because the prices have been constant in WiFi hotspots and usage based cellular data, the flat rate Cellular Data

[1] www.archive.org/ [Accessed 16.02.2009]

[2] http://www.ficora.fi/ [Accessed 25.02.2009]

[3] http://www.tietokone.fi/ [Accessed 25.02.2009]

[4] www.haminetti.net/ [Accessed 16.02.2009]

[5] www.msoynet.fi/ [Accessed 16.02.2009]

[6] www.wlannet.com/ [Accessed 16.02.2009]

[7] www.netikka.fi/ [Accessed 16.02.2009]

[8] www.sparknet.fi/ [Accessed 16.02.2009]

[9] www.zonet.fi/ [Accessed 16.02.2009]

[10] www.elisa.fi/ [Accessed 16.02.2009]

[11] www.dna.fi/ [Accessed 16.02.2009]

[12] www.sonera.fi/ [Accessed 16.02.2009]

[13] www.saunalahti.fi/ [Accessed 16.02.2009]

[14] www.kolumbus.com/ [Accessed 16.02.2009]

[15] www.tele.fi/ [Accessed 16.02.2009]

and monthly WiFi prices are only used in describing the price ratio (Figure 3).

Figure 3. Price ratio of Cellular Data and WiFi

During the life time of the Cellular Data flat rate subscriptions WiFi has been the more expensive alternative, but before the year 2006 one can assume that WiFi was the cheaper alternative mainly because Cellular Data was still in an early phase and was priced accordingly. One must notice that the prices are not for subscriptions that had the same bit rate and thus the price ratio has to be interpreted accordingly.

Although the scope of this report only includes mobile data analysis, the historical data regarding mobile voice prices might give insight to what the evolution of mobile data prices could be. Figure 4 shows the evolution of mobile voice prices in Finland between 1995 and 2007. The prices have been decreasing steadily between 1995 and 2003 mainly because of the GSM (Global System for Mobile Communications) technology development and network externalities. In the year 2004 there is a steeper drop in the prices. This is probably result of the obligatory mobile number portability which was enforced on operators in the year 2003 (Smura, 2004) and the obligatory leasing of capacity to mobile virtual operators which was introduced in 2004 (Kiiski and Hämmäinen, 2004). In 2006 3G bundling was made legal (MinTC, 2003) which could explain the minor increase in prices between 2006 and 2007.

Analyzing the historical price data of mobile voice leads to a question whether regulators or other external influences could have similar effects on mobile data prices as they had on mobile voice prices. These effects are speculated in the coming chapters.

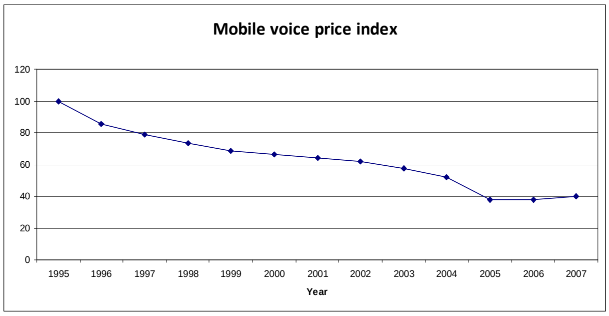

Figure 4. Mobile voice price index (MinTC, 2008)

3.2 Pricing structure

In order to understand the basic structure and dynamics of prices for Cellular Data and WiFi we will take a look at some fundamental concepts of economics (Hämmäinen, 2008) (Figure 5).

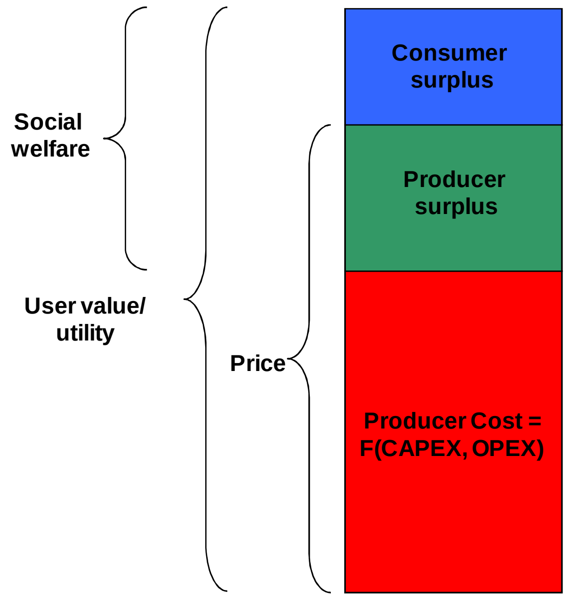

Figure 5. Simple pricing structure

The first element to observe is the overall value of the service for the user, i.e. utility. In terms of wireless data connectivity, the utility of the service is dependent for example on how much bandwidth can be supplied and also on the geographical availability of the service. Utility can also be dependent on how easy to use the service is (i.e. the degree of vertical integration) and on what is the need for wireless bandwidth which is in turn dependent on the application characteristics. The utility of a product can also depend on the availability and price of its substitutes and in the case of Cellular Data and WiFi they can be viewed as a closed market of two services where both are substitutes to each other.

The price of the service is a sum of the cost to produce the service and producer surplus (PS). The cost of providing the service is a function of both capital (CAPEX) and operational costs (OPEX) where capital costs are usually divided over a period of time for a particular investment. Technology development and network externalities decrease costs over time which, depending on the competitive situation in the markets, might or might not be reflected into the prices. When competition is low (i.e. a monopoly or an oligopoly of a few significant market power operators) the decreased cost might only increase producer surplus and not decrease prices. Thus if the market is somehow regulated, the regulator tries to maximize the social welfare and balance it for both the consumer and the producer.

Cellular markets have traditionally been highly regulated, whereas WiFi regulation has been rather low leading to the possibility of a local monopoly and somewhat high prices in a particular area.

What is left after the price is subtracted from the utility of the service is the consumer surplus (CS). It can be used to roughly indicate how many users there are for the service and how much they are using the service.

Cellular Data and WiFi have different characteristic regarding these elements. For example if we look at the above mentioned elements for utility, we can see that the availability for Cellular Data is currently substantially higher than for WiFi in all of its forms, due to large coverage and simple SIM authentication of Cellular Data versus a small WiFi footprint with fragmented authentication. On the other hand the capacity of WiFi is usually larger than for Cellular Data provided that there is a reasonable number of other users.

The cost structure of Cellular Data and WiFi differs also. To showcase this, both costs are roughly estimated in terms of Capital Expenditures (CAPEX) and Operating Expenditures (OPEX) in equations (1) and (2).

WiFi cost = $CAPEX_{WiFi}$ + $OPEX_{WiFi}$ =

F(equipment investment) +

F(ADSL/cable monthly fee, authentication server OAM) (1)

Cellular Data cost = $CAPEX_{Cell}$ + $OPEX_{Cell}$ =

F(frequency cost, NW equipment cost, NW planning/dimensioning costs) +

F(NW&customer maintenance costs, QoS&Mobility management, site rentals) (2)

The cost structure for WiFi is rather light. In terms of capital expenses, investments to simple consumer grade equipment, such as an access point and possibly an authentication server, are only needed. A monthly fee for wired access and authentication server maintenance costs constitute the operational expenses.

Cellular Data cost is considerably heavier with initial investment costs stemming from spectrum fees, complex network equipment and the careful planning of the network to avoid interference between the Base Stations (most of which are not needed with WiFi). The operational costs are also higher for Cellular Data with e.g. customer and network maintenance and site rental costs.

3.3 Trends and uncertainties affecting price

There are many factors that might have an effect on the price ratio between WiFi and Cellular Data. The growing number of mobile data users and high bandwidth consuming applications are factors that are likely to increase the overall demand for mobile networks. Furthermore, with the rapid diffusion of social networking and user created and distributed content, the demand for Uplink (UL) capacity could grow substantially faster than the demand for Downlink (DL) capacity leading to challenges especially for Cellular Data. On the other hand storage costs are becoming smaller and smaller in comparison to transmission costs (Gray and Shenoy, 2000) which could lead to more preloaded content and thus decrease the demand for bandwidth.

Regarding the pricing flat rate is probably going to be a likely pricing scheme, because the general trend is towards simple pricing (Odlyzko, 2004) which in turn is likely to further increase demand.

On the supply side there is always a possibility of running out on spectrum as the demand increases. Ficora (2007) has estimated that the allocated unlicensed spectrum for WiFi (2,4 GHz) should be sufficient until 2015. After that the increasing demand has to be covered somehow. One solution might be the usage of the 5GHz band.

Some sources also go as far as saying that the mobile data traffic will grow 300 fold until the year 2015[16]. The emerging new technologies such as UMTS900, Long Term Evolution, IEEE 802.11n, IMT (International Mobile Telecommunications) Advanced, adaptive antennas and flexible spectrum usage naturally increase the supply side when deployed and there is also a decreasing trend in cell sizes which increases bandwidth. The downside is that the smaller cell sizes also increase the number of sites and thus especially the operational costs.

The regulator actions can have a big impact on the prices as happened with mobile voice price (Figure 4) for example. The regulators want to maximize the social welfare arising from the use of radio spectrum and balance it equally between the producers and consumers and thus can affect the prices directly or indirectly by enabling competition. The regulators also might introduce some kind of vertical bundling (Hämmäinen,

16 http://www.digitoday.fi/bisnes/2009/02/18/nsnn-beresford-wylie-mobiilidata-kasvaa-300-kertaiseksi/20094537/66 [Accessed 20.02.2009]

2008) in order to increase the utility of the technology as happened in Finland in 2006 with the 3G subscriptions.

One interesting issue is international roaming. Currently the data prices for cellular usage in foreign countries are very high and thus there is an incentive to use WiFi when being abroad. The regulators can, however, introduce for example flat rate data roaming fees for operators that could increase international mobile data roaming. There is already a proposal made by the EU commission for a price cap of one euro per megabyte and also commercial launches of data roaming exist (Hämmäinen, 2009), but it remains to be seen what happens.

Other quite recent phenomena that might have an impact on the prices are user operated access points (AP), i.e. WiFi communities such as Wippies[17] and FON[18] that substantially extend the availability of WiFi to the community members and user deployed base stations (BS), i.e. home base stations and Femtocells that can increase the coverage and capacity of a cellular network in a convenient and cost efficient way. Both of these leverage a social effect and as a result increase the value of the network and reduce the operating costs and thus might have a considerable effect on the prices and price ratio of WiFi and Cellular Data. However, the operators offering fixed broadband subscriptions and carrying the transit traffic are probably not going to be pleased that the WiFi communities are emerging, because these social APs load the core network of the operator more as the WiFi is being used by more people.

In addition, the evolution of vertical bundling and integration (i.e. ease of use and seamless service) of WiFi through more intelligent devices is also an important factor.

4 Conceptual modeling

The conceptual modeling in this seminar paper is carried out with system dynamics. System dynamics is a system engineering method for complex systems. The basis lies in system's thinking, i.e. being able to see the world as a complex system where everything is connected to everything and one cannot assume that a change in one variable wouldn't affect anything else (Sterman, 2000). The software used is Vensim PLE[19], which is free for personal and educational use.

4.1 Basic System dynamics model

The basic system dynamics model can be seen in Figure 6. The core of the model consists of three stock variables, namely potential users, Cellular Data users and WiFi users. Stock variables are containers that can increase or decrease over time. The increase or decrease in stocks happens through flows, which in this model represent the adoption of Cellular Data and WiFi and the substitutions between these networks. The model is divided to two domains, namely Cellular Data and WiFi,

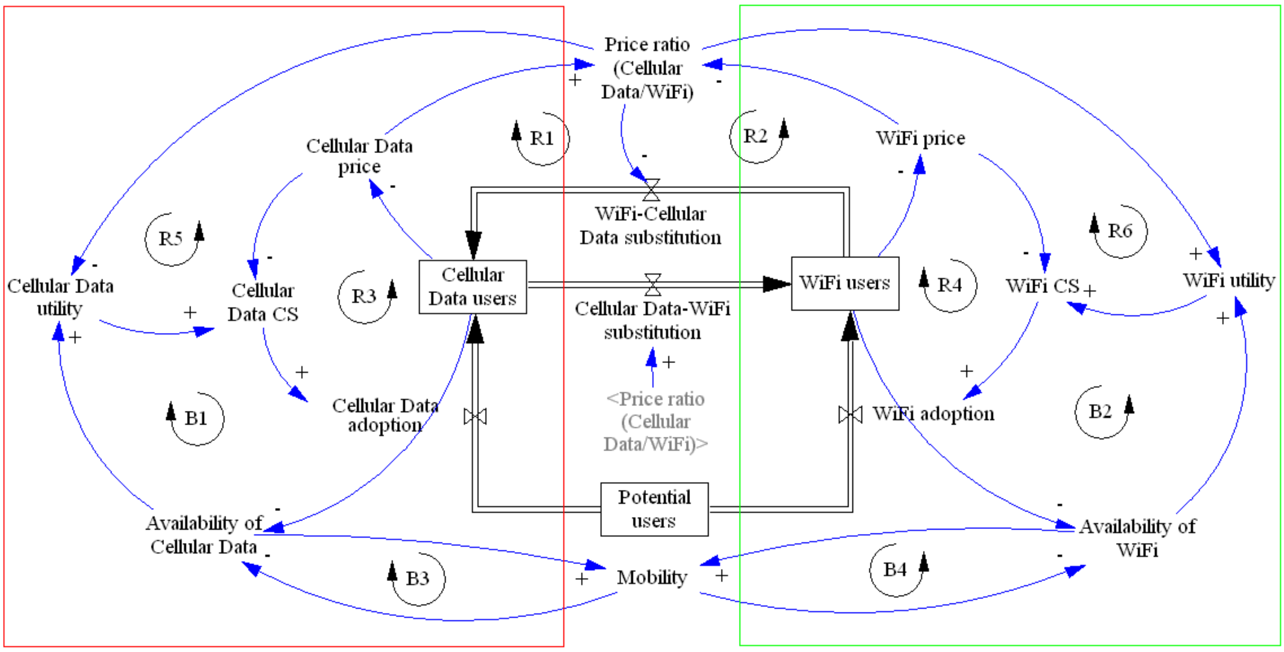

Figure 6. Basic system dynamic model

17 www.wippies.com/ [Accessed 18.02.2009]

18 www.fon.com/ [Accessed 18.02.2009]

19 http://www.vensim.com/venple.html [Accessed 24.02.2009]

which are represented by the red box on the left and the green box on the right, respectively. The units for the different variables are money for price, CS and utility, users for persons, bit rate for availability and a dimensionless factor for mobility.

The other parts of the model consist of four balancing loops and six reinforcing loops. The first balancing loops (B1 and B2) represent the effect the number of users has on available bandwidth for both network types. The more bandwidth there is available, the more value and thus consumer surplus the network has for the end user. This in turn increases the number of users and interference, which reduces available bandwidth and thus has a balancing effect.

The other two balancing loops (B3 and B4) show the interaction between mobility and availability. As the users are mainly spending their time at work or home (Statistics Finland, 2001), the possibility that there will be a network available is bigger than when they are moving in a larger area including parks, restaurants and other households for instance. Thus increased mobility decreases availability, but on the other hand increased availability can increase mobility of the nomadic users, because there are more networks available. The negative effect of mobility is larger for WiFi than Cellular Data.

The first two reinforcing loops (R1 and R2) represent the substitution effects between Cellular Data and WiFi. The better the price ratio for a given network is, the bigger the substitution effect to that network is.

The next two reinforcing loops (R3 and R4) show the effect that the number of users has on price. The more users there are the smaller the price gets, because of network externalities and economies of scale, and the smaller the price gets, the more users will adopt the network in question. The last two reinforcing loops (R5 and R6) represent the effect the price ratio has on the network utility, i.e. the value the network has for the end user. As discussed in section 3.2 Cellular Data and WiFi can be viewed as forming a closed market where the increased price of one increases the demand and value of the other.

4.2 Two basic scenarios

Based on the introduced model two possible extreme scenarios (Figure 7) can be identified: one where Cellular Data clearly dominates and another where WiFi becomes a considerable substitute for Cellular Data. Although these scenarios are not necessarily very likely they serve as a good starting point for the analysis.

Figure 7. Two extreme scenarios

4.2.1 Scenario 1: Cellular Data dominates over WiFi

In the first scenario the easy to use and ubiquitous Cellular Data is clearly preferred by users and as a result it substitutes WiFi. In this case, in the system dynamics model above, most of the users would flow to the Cellular Data users stock.

The cellular networks would be able to cater to the bandwidth needs of the growing number of users and evolving applications or even surpass them and as a result, on a larger scale, the availability of Cellular Data would remain sufficient for all users desiring connectivity (i.e. the balancing loop B1 would not be that strong). This would be positively reflected in the utility and Consumer Surplus of Cellular Data which would in turn increase its adoption and the number of users. The increase in number of users would reinforce positive feedback with network externalities and economies of scale resulting overall in lower and simple flat free prices for Cellular Data (i.e. reinforcing loops R1, R3 and R5 would be very strong).

If at the same time the availability of WiFi for nomadic users would remain low for example due to a small WiFi footprint and increased mobile use of wireless data (which Cellular Data can cater to more efficiently), it would lead to a decrease in WiFi utility, Consumer surplus and WiFi adoption (i.e. balancing loops B2 and B4 would be strong and the reinforcing loops R2, R4 and R6 would be quite weak). This would eventually lead to a situation where WiFi prices would not be competitive enough with Cellular Data fueling the users urge to not to adapt WiFi and substitute it with Cellular Data.

4.2.2 Scenario 2: WiFi a major substitute for Cellular Data

In the other extreme scenario the cheap and easily operated WiFi would become a considerable substitute for Cellular Data. In this scenario the availability of WiFi would improve substantially with for example a vast growth of WiFi communities and public WiFis (i.e. weak balancing loop B2). Another possible reason for WiFi becoming the dominating source for wireless connectivity could be the fact that people spend most of their time either at home or at work (Statistics Finland, 2001) as was already mentioned. If the mobility of nomadic users would remain rather low the negative impact the on availability of WiFi would not be that dramatic (i.e. weak balancing loop B4).

With an abundant supply of WiFi connectivity the utility and consumer surplus of WiFi would increase and as a result so would the WiFi user base. This would lead to a positive feedback cycle that would drive down the overall costs of WiFi (i.e. strong re-enforcing loops R2, R4 and R6).

Furthermore if the cellular networks could not supply enough data for the high bandwidth consuming applications their utility for the consumer would drop (i.e. weak re-enforcing loops R1, R3 and R5). This could happen for example due to fierce competition between cellular operators and the resulting aggressive usage of flat rate pricing. This could result (and has already resulted) in a tremendous growth of Cellular Data usage and in heavy congestion (especially during peak hours) which would in turn decrease the experienced quality for the user (i.e. balancing loop B1 very strong). The outcome of this would be a lower number of Cellular Data users and non-competitive prices with WiFi.

4.3 Enhanced model

The modeling conducted above is very elementary and should be enhanced. Thus an enhanced model (Figure 8) is built which is based on the other trends and uncertainties introduced in chapter 3.3 and excluded from the basic model.

The enhanced model includes a variable for average bandwidth demand per user to model the possible increase in demand that for example the high consuming applications bring along. The bigger the storage transmission ratio gets the more BW is demanded. This in turn increases Cellular Data usage and the number of operator Picocells and possibly user-deployed Femtocells.

Another important new variable is the regulator actions to maximize and balance social surplus. Regulators might intervene if the degree of local competition (loop B6) in the Cellular Data side is too low or if vertical integration is needed to increase the utility of the

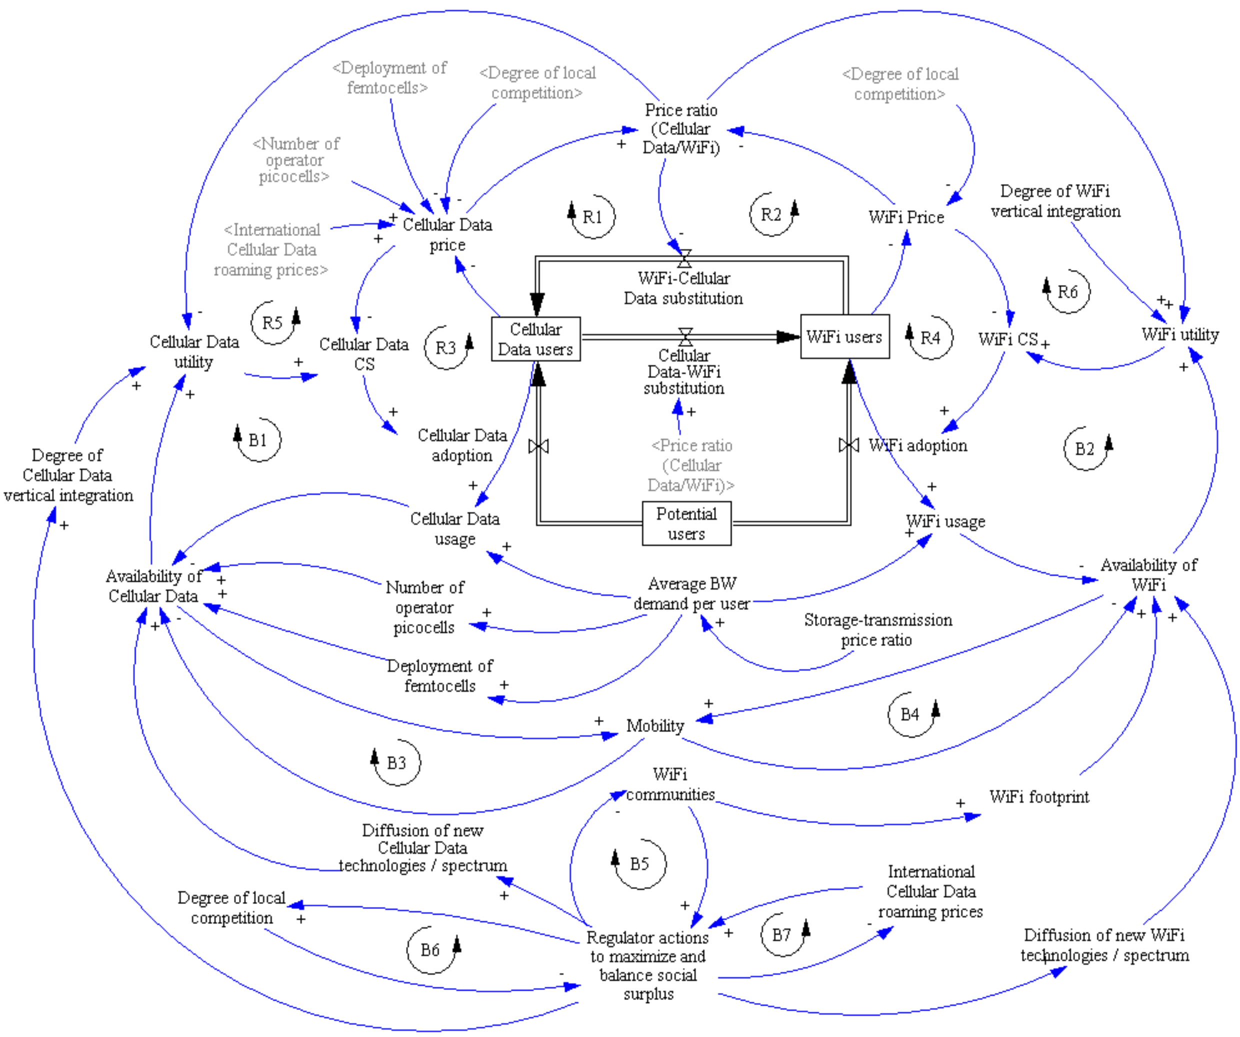

Figure 8. Enhanced model

technology (as was done with 3G bundling). Regulators also might support and accelerate the diffusion of new technologies or be more liberal with frequency spectrum to increase the availability.

The WiFi communities might increase the costs of the operators carrying the transit traffic (possibly reducing their Producer Surplus) and thus regulators could for example limit the diffusion of WiFi communities to balance the overall social surplus (loop B5). The international roaming prices might also be affected by regulator actions as discussed in chapter 3.3.

These variables affect the prices of the networks. The more there is competition the lower the prices get. On the Cellular Data side the need to build more operator Picocells increases the operational costs and thus consumer prices, but on the other hand the more user-deployed Femtocells there are the more the operational costs are outsourced to the users and the cheaper prices get.

In general, for some of the concepts in the model it can be rather hard to find a good unit to describe the change (e.g. for mobility or the degree of vertical integration), and for those some kind of an index could be used based on a formula comprising other constituting elements and an expert opinion.

4.4 Speculation of other possible evolution scenarios

Looking at the two extreme scenarios discussed in section 4.2 it is evident that in reality the outcome will quite likely be something in between the two extremes. Here, based on the enhanced model introduced above, we will conduct some speculation on other possible evolution scenarios for the price and user ratio of Cellular Data and WiFi. We will try to identify critical issues that could tilt the advantage towards either of the two access types.

4.4.1 Issues supporting Cellular Data

With the strong trend of decreasing cell sizes and increasing number of small base stations (i.e. Picocells) one critical element for keeping the costs of Cellular Data reasonable is the success of Femtocell deployments and related outsourcing of operational costs to users ("hidden OPEX"). There are many factors supporting the success of Femtocells ranging from the users urge to "help themselves" in order to provide better coverage and capacity at home to subvented equipment from the operator.

Another issue that might support Cellular data is if the vertical integration and ease of use of WiFi remains low. There is a threat that WiFi networks remain fragmented, and that the separate credentials and different authentication mechanisms cannot be efficiently integrated e.g. with intelligent clients and by WiFi aggregators (e.g. boingo[20]) and WiFi communities that could enable seamless roaming between single hotspots.

What could hinder the expansion of WiFi communities and other forms of private connectivity sharing with WiFi (e.g. Joiku[21]) is regulator intervention. If the traffic volumes for this kind of sharing grow considerably, the increased costs for wired and wireless operators providing the private connectivity might force the regulators to restrict such sharing.

Another factor contributing for the case of Cellular Data is the possibility of strong enforcement by the regulator to enable a fast diffusion of the new cellular technologies (e.g. UMTS900 in Finland, and LTE etc.). Furthermore, there have been ongoing discussions about the possibility of more liberal frequency spectrum policies by the regulator for example in terms of technical neutrality of the allocated frequency blocks and reselling which might increase the availability of Cellular Data.

4.4.2 Issues supporting WiFi

One contributor for the case of WiFi (that can already now be observed) is the fact that the high international data roaming prices are driving users to utilize WiFi while they are abroad. If no reasonable harmonized cross-country Cellular Data price regulation is not achieved, such usage of WiFi might remain high or even grow substantially.

Another possibility that could support WiFi is a scenario where Femtocells are not widely deployed leading to expensive maintenance of many Picocells (due to e.g. many site visits etc.) which would be reflected in increased Cellular Data prices.

The main reason for the low deployment rate of Femtocells could be its inability to provide carrier grade service. Many issues, such as network planning and interference, access control of users, the network management of possibly millions of customer access points, the support for emergency calls and lawful interception, are still unresolved. On the other hand the WiFi communities that are free from these restrictions and provide only consumer grade connectivity could spread more rapidly leveraging the same "hidden OPEX" phenomena.

Furthermore if intelligent self-configuring devices and networks emerge, become common and enable easy, vertically integrated WiFi service, it could accelerate the uptake of WiFi for the technically incompetent users (early/late majority) and could substantially increase both WiFi user base and WiFi sharing and hence the utility of WiFi.

---

[20] www.boingo.com [Accessed 27.02.2009]

[21] www.joiku.com[Accessed 27.02.2009]

The resulting growing variety of ways of sharing and accessing WiFi could increase the degree of competition within WiFi providers and drive down the prices.

5 Conclusions

In this paper we have tried to analyze the historical price and price ratio evolution of both Cellular Data and WiFi and model how they could evolve in the future. Since both technologies are rather young, and since the availability of price data for both access types is rather limited, most of the focus in the study was put on understanding different issues affecting the future of their price evolution and on what could be the degree of substitution between the two.

First, we conducted an overview analysis of the price evolution including the study of historical data and the structure of prices for both and took a general look at the different trends and uncertainties that could affect the evolution in the future. Then we built a basic system dynamics model and derived two extreme scenarios to get a good starting point for the analysis. An enhanced model was then constructed to get a better understanding of the different factors affecting the evolution and of what could be the actual outcome. Looking at the current situation one might argue that Cellular Data is likely to become the dominant access method for wireless data, but as could be seen from the system dynamics figures, powerful positive feedback loops and re-enforcing cycles can also emerge for WiFi. With a large number of variables it is very challenging to forecast what would be the actual outcome of the battle of the two.

Nevertheless, system dynamics seems to be a powerful tool to get an overall understanding of the dynamics of the wireless data market and to comprehend the possible underlying reasons.

Especially regarding the enhanced model, the analysis was rather light and could be elaborated more in the future coupled possibly with a wider scope than what was used here. Quantitative modeling based on the introduced system dynamics model(s) could also be the target of future work.

References

Ficora, 2007. Finnish Communications Regulatory Authority Report: Future demand of radio frequencies (In Finnish). [Online] Available at: http://www.ficora.fi/attachments/suomi_R_Y/5n8HwrsDL/Files/CurrentFile/Radiotaajuusraportti2007.pdf [Accessed 18.02.2009].

Gray, J. and Shenoy, P., 2000. Rules of Thumb in Data Engineering. In: *16th International Conference on Data Engineering* (ICDE'00), pp.3.

Hämmäinen, H,. 2008. Operator Business Course, Helsinki University of Technology. Lecture 1, Lecture 4, and Lecture 8. [Online] Available at http://www.netlab.hut.fi/opetus/s383041/k08/Lectures/

Hämmäinen, H., 2009. Flat-rate Data Roaming. Presentation at the COST605 Econ@tel meeting in Limassol (3. February 2009).

Kiiski, A. and Hämmäinen, H., 2004. Mobile Virtual Network Operators: Case Finland. [Online] Available at: http://www.netlab.hut.fi/tutkimus/lead/leaddocs/Kiiski_ITS_MVNO.pdf [Accessed 24.02.2009].

MinTC, 2003. Ministry of Transport and Communications Report: Communications market act 393/2003, section 25. [Online] Available at: http://www.finlex.fi/en/laki/kaannokset/2003/en20030393.pdf [Accessed 24.02.2009].

MinTC, 2008. Ministry of Transport and Communications Report: Price level of telecommunications charges in 2007 (In Finnish). [Online] Available at: http://www.lvm.fi/fileserver/LVM15_2008.pdf [Accessed 17.02.2009].

Odlyzko, A., 2004. Pricing and Architecture of the Internet: Historical Perspectives from Telecommunications and Transportation. [Online] Available at: http://dtc.umn.edu/~odlyzko/doc/pricing.architecture.pdf [Accessed 18.02.2009].

Smura, T., 2004. Mobile Number Portability: Case Finland. [Online] Available at: http://www.netlab.hut.fi/opetus/s38042/s04/Presentations/06102004_Smura/Smura_slides.pdf [Accessed 24.02.2009].

Statistics Finland, 2001. Time use survey. [Online] Available at: http://www.stat.fi/meta/til/akay_en.html [Accessed 25.02.2009]

Sterman, J.D., 2000. Business Dynamics: systems thinking and modeling for a complex world. McGraw-Hill, Boston.

CAPEX vs. OPEX from the Perspective of a Mobile Operator

Heidi-Maria Rissanen
heidi-maria.rissanen@ericsson.com

Abstract

Increasing data traffic and mobile technology evolution among others are driving the operators to upgrade their networks. This has given rise to the cost-effectiveness of the mobile operators.

In this paper we will study how the total cost of ownership (TCO) of a mobile operator is like. We will describe the different cost types and show some reference data from previous studies. We will also discuss, what kind of mechanisms there are to decrease both capital and operational expenditures. We will study how heterogeneous networks affect the costs of a mobile operator. In addition, femtocells and site sharing are presented as techniques for reducing TCO.

1. Introduction

New data services are a major driver for the mobile operators to invest in higher bandwidth. Higher bandwidth requires investments both in access and in core network. This has given rise to the cost-effectiveness of the mobile operators, as they are facing a challenge to provide high bandwidth services and still maximize their profits. The TCO (total cost of ownership) of a mobile operator consists of capital expenditures (CAPEX) and operational expenditures (OPEX). Capital expenditures are for example costs of the network infrastructure. OPEX, in turn, consists for example of personnel costs and monitoring the network.

A recent study shows that OPEX is 70-80 % of TCO for mobile operators (ABI Research, 2008). Thus, operators try to minimize their TCO through savings rather in OPEX than CAPEX.

Even though the management of currently deployed 3G networks is mainly done remotely, operators still spend a lot of money for configuring and managing the network. Instead of expensive personnel performing these tasks, it would be more cost-efficient to have the network to operate and maintain itself automatically. In next generation networks beyond 3G, like in Long Term Evolution (LTE), more attention is already being paid to the automatic operation and maintenance.

There is interdependency between capital and operational expenditures. According to a recent study (ABI Research, 2008) CAPEX directly affects 70-80 % of OPEX. As the costs of the hardware are going down, CAPEX is directly decreased. This implies that the operators need to understand the importance of the network architecture. In this paper, we discuss the historical and possible future trends of CAPEX/OPEX ratios for a mobile operator. We will also study the affects price changes have had and possibly will have on the network architecture in the near future.

The structure of this paper is as follows. In Chapter 2 we first define the different cost types used in this paper and describe the overall cost structure of a mobile operator. In Chapter 3, we discuss and present data about the infrastructure costs of the operators. We will also discuss how the costs of different equipment are affected by the used technology. In Chapter 4, we will study different techniques and technologies to decrease the TCO. These techniques include heterogeneous networks, self-organizing networks and network sharing. Finally, we will conclude and discuss some issues for future work the paper in Chapter 5.

2. Cost Structure of Mobile Operators

As described in Chapter 1, TCO is divided into capital and operational expenditures. In the following, the definitions of capital and operational expenditures are given. An example of the cost structure for a mobile operator is also discussed.

2.1. Definitions of OPEX and CAPEX

There are several different ways to model capital and operational expenditures. In this paper, we follow the model used for example in Loizillon (2002) and Smura (2007). In comparison to the model used in this paper, for example Nokia uses a model where implementation expenditures consisting of the cost to build a network are calculated separately.

Capital expenditures for a mobile operator are defined as the costs that are related to the network infrastructure of the operator. CAPEX consists of the costs to acquire network infrastructure and software. In the case of a mobile operator this means for example equipment for radio and transmission. Spectrum and license fees, as well as costs to purchase land and buildings are included in CAPEX. In addition, costs to upgrade existing equipment are counted in CAPEX. Capital expenditures are subject to depreciation.

Operational expenditures are defined as costs related to operating the network (Verbrugge et al., 2005). Costs that are needed to keep the network operational in a failure free situation are counted in OPEX. This includes the costs for power, air-conditioning and leased network equipment. Costs for transmission is also considered as part of the costs to keep the network operational. Costs for maintaining a network operational are also counted in OPEX. This includes actions like network monitoring and software management. Also, repairing faulty equipment, pricing and billing, and marketing are other costs counted into OPEX. Operational expenditures are not depreciated over time as they do not contribute to the fixed assets of the operator.

2.2. Cost Structure

The cost structure of an operator has been studied for example in Loizillon (2002) and Smura (2007). The overall cost structure of a mobile operator varies a lot depending on the country and also operator. Thus, for example in Loizillon (2002) scenarios with different population densities and the size of the country are studied.

In the first scenario, the mobile operator is an operator in a large, densely populated country. In the second scenario, the operator operates in a small, sparsely populated country. Both of the operators provide WLAN infrastructure as well.

The total capital expenditures of the operators in these two scenarios are shown in Table 1. In addition to the costs listed in Table 1, the cost for license is 1 % in the small country. In the large country scenario the cost of the license fee is much higher, around 6 billion Euros. However, this is excluded from the cost structure for a large country shown in Table 1.

Table 1 shows that in both scenarios access network costs are bigger than the costs of the core network. The cost of the core network in small country is even less than 10%. This is also aligned with for example Odlyzko (2004). Odlyzko (2004) shows that the costs of the Internet are migrated to the edges of the network. Core is not expensive.

Table 2 shows the operational costs of the two scenarios. Transmission costs in small country are significantly higher than in large country. This is because the distance between the access network and the core network is longer. The personnel costs are in both scenarios the biggest part of the operational costs.

Other costs not related to operating to network like marketing and terminal subsidies are also significant in both cases. Subsidies are needed to lower the terminal prices to attract customers. It is noteworthy that the costs for WLAN infrastructure and operation are very small.

Table 1. CAPEX for a 3G network (Loizillon et al., 2002)

Cost type	Large country	Small country
WLAN	4%	1%
Access network	30%	38%
Sites	38%	53%
Core	28%	7%

Table 2. OPEX for a 3G network (Loizillon et al., 2002)

Cost type	Large country	Small country
Transmission	5%	25%
Terminal subsidies	16%	13%
Personnel costs	49%	39%
Marketing	23%	18%
Site rental	3%	0%
WLAN	4%	5%

3. Drivers affecting CAPEX

Even though the capital expenditures are smaller component of the TCO than the operational expenditures, there is more data available about the capital expenditure. However, there is interdependency between capital and operational expenditures. This applies that the choice of the network technology and infrastructure will eventually affect also operational expenditures. Thus, in this Chapter we will concentrate on studying drivers for new investments from mobile operator' perspective.

Factors affecting the capital expenditures of the operator depend mainly on the market it is operating in. We will concentrate in this paper on the situation in Western Europe, which can be considered as a developed market. The main drivers in such a market are new data services, which require higher bandwidth. Operators also want to provide seamless access to the users meaning that there will be connectivity to the end user across heterogeneous access networks and technologies (Kellerer et al., 2003). Minimal input can be required from the end user. Mobile network technology evolution is of course one of the drivers to invest into new network infrastructure. New technologies enable the operators to do for example faster network rollouts. Also, as in 3G network processes are not automated, operators would rather invest in such technology that would help them to reduce both capital and operational expenditures (Brown, 2008).

3.1. Trends of CAPEX components

Figure 1 shows the total capital expenditures for four different network equipment types in Western Europe. Data is anonymised from ABI Research (2008). Figure 1 shows that for network Equipment A the CAPEX has been and is expected to stay same over the years. For Equipment C, in turn, the cost is increasing dramatically. This can be explained with the new 4G network technology, Long Term Evolution (LTE), taken into use. As new data services are provided to the users, traffic load is also increased and operators need to invest in the backhaul, which means the connection between base stations and the base station controller, which is connected to the core network. Equipment B is needed

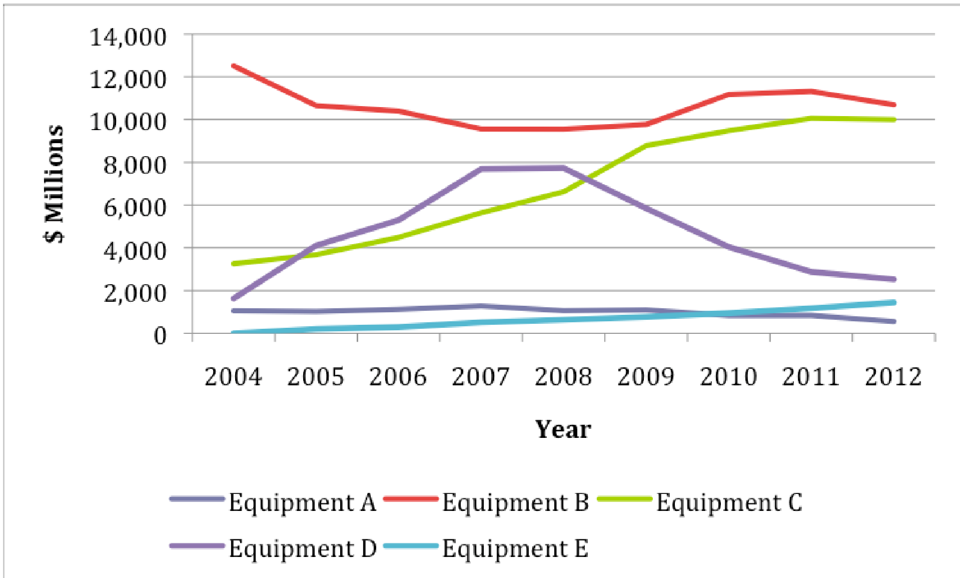

Figure 1. Costs of four different network equipment types in Western Europe, 2004-2012 (ABI Research, 2008)

independently of the technology. Price per unit is expected to drop as hardware costs are decreasing. However, the number of installations is expected to grow and thus the total cost is increasing. Equipment of types D and E are related to investments enabling access to the Internet from different kinds of access networks. Equipment D is related to operators starting to take the IP Multimedia Subsystem (IMS) into use. Even if the initial investments for IMS are substantial, the IMS is expected to reduce the total cost of capital expenditures. The IMS is assumed to help to achieve savings through economy of scope. Economy of scope means that users in a wide area are reachable through a single system. Equipment E is related to operators starting to offer other types of access to the end users. This is a result of cheaper technologies such as Wireless LAN (WLAN) that can be used to build access points in buildings to provide high-speed data services. WLAN can be considered as a complement to an operator's service offering. This means that networks in future will be heterogeneous, as suggested for example in Niebert (2005).

Figure 1 shows that for some cost components the used technology does not make a difference, but the cost remains almost the same. For some components the cost is tightly coupled with the used technology. As the price per unit for some of the equipment is going down, the number of installations is growing at the same time.

4. Mechanisms for Minimizing TCO

To be able to respond to the demands of higher traffic loads in the future it can be seen that the operators will need to be able to provide higher capacity in their networks. The problem in wireless networks has been that the cost per bit has not decreased as much as the demand from the users. This means that the users do not want to pay for the mobile service more than they already pay. This drives the operators to look for cheaper technologies.

As mentioned in Chapter 3, future networks are likely to be heterogeneous. There are two levels of heterogeneous networks (Johansson et al., 2007). On one level heterogeneous networks refer to multiaccess networks. In a multiaccess network there are several radio access technologies based on different standards that can be accessed with a multiradio capable terminal On the other level, there are hierarchical cell structures using the same radio access standard with different types of base stations. In the following, we will first describe the infrastructure and operation costs for different types of base stations.

We will also discuss self-organizing networks and femtocells, which are very small cellular base stations.

4.1. Hierarchical Cell Structures

As we discussed in Chapter 2, the cost of the access network is significant independently of the type of the market. Access network means the Radio Access Network (RAN), which consists of base stations, sites, radio network controllers and last-mile-transmission (Johansson et al., 2004). Selecting the characteristics of the base stations correctly can minimize the costs of base stations. This means that base stations for example with different cell ranges can be used.

Sample costs related to 3G base stations with different cell ranges are shown in **Table 3** (Johansson et al., 2004). An urban environment is assumed. The cell range of a macro base station is at maximum 1 km. For a micro base station the maximum cell range is 0.25 km, and for a pico base station the maximum cell range is 0.1 km. Macro base station is much more expensive to build and

Table 3. Base station costs (Johansson et al., 2004)

Cost/Base station (k€)	Radio	Sites	Transmission	Total
Macro	70	168	28	266
Micro	27	35	28	90
Pico	12	10	28	50

operate than the other two. This is because it has higher output power and coverage

Also, as a macro base station has more users connected to the base station, the reliability of the base station has to be better. This affects the operation and maintenance (O&M), installation and site costs. Micro base stations are in this example much cheaper than macro base stations. Pico base stations, on the other hand are not that cheaper than micro base stations. Johansson et al. (2004) also shows that it is possible to minimize the cost of infrastructure. All three base station types are feasible in some scenarios. Assuming that the capacity per user is 1kbps, which corresponds to a speech service, the more there are users per km2 the smaller base station should be. Macro base stations are the most cost-efficient solution if there are less than 4000 users per km2. If the user density is higher, micro base stations should be used. Pico cells are the cheapest solution if there is a very high user density (20 000 users/km2). Figure 2 shows the cost per user as the function of user density. The figure tells that the denser the deployment is, the cheaper the cost per user is. This implies that there is some kind of scale of economics.

4.2. Multiaccess Networks

As different technologies are designed with different objectives and for different purposes, by combining different technologies it might be possible achieve savings. Mobile protocols, like 3G, provide reliable service with high coverage but more limited data speed. For WLAN, in turn, the objective is to provide high-speed data connectivity but with a very small coverage. As shown earlier in Section 4.1, the costs for macro base stations are high. CAPEX costs per WLAN Access Point are according to Björkdahl et al. (2004) 16 600 €. Table

Table 4. Cost structure of a WLAN network (Björkdahl et al., 2004)

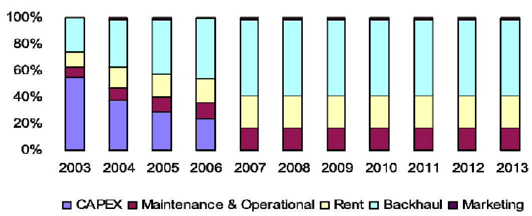

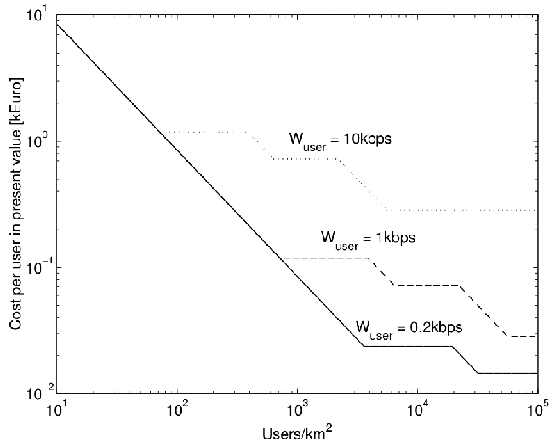

Figure 2. Minimum cost per user for different capacity (Johansson et al., 2004)

4 shows the deployment costs of a WLAN network. CAPEX costs are significant in the beginning, but after that the costs are mainly operational like marketing, rental and backhaul. Depending on the backhaul type, the costs for backhaul per site can be as high as 1800 € per month (Björkdahl et al. 2004).

Johansson & Furuskär (2005) show that it is feasible to use WLAN Access Points (APs) in case there is a sudden need for increased capacity. Also, deploying WLAN APs is reasonable if the market is such that there is a lot of variation in the traffic density geographically. As the prices of WLAN APs are going down, deploying WLAN becomes more attracting option for the operators.

4.3. Self-Organizing Networks

For the fourth generation mobile systems, LTE, reduction of cost, complexity and automated maintenance have become key drivers (NGMN Alliance, 2007 and Brown, 2008). Self-organizing networks are currently being studied in 3GPP.

As we have discussed earlier in this paper, costs related to actions performed by the personnel are a big part of the operational expenditures. High complexity for configuring current networks is one of the drivers for automating the mobile network operation and configuration. Self-organizing networks are trying to minimize the human interaction needed in operating a network. Some human action will be needed for example to verify results. An autonomous activity is fully taken care by a machine. One further driver for self-organizing networks is that when LTE is taken into use, operators are expected to run several different networks simultaneously. This should not cause any additional operational expenditures to the operators.

NGMN Alliance (2007) defines two types of NodeBs, which are basically base stations. The first type, Macro

Table 5. Comparison between macro cells and femtocells

Total Demand (Mbps)	Macro cells			Femtocells		
	Nof sites	CAPEX	OPEX	Nof APs	CAPEX	OPEX
200	34	3,4M€	1,02M€	100	0,1 M€	0,05M€
1000	167	16,7M€	5,01M€	500	0,5 M€	0,25M€

NodeB is maintained and operated by the operator. Home NodeB is installed by the end user. Even though some manual work would be needed to install a new node, configuration work is reduced by utilizing self-configuring functionality.

Next, we will describe femtocells, which are an implementation of self-organizing networks.

4.4. Femtocells

Femtocells, also known as home base stations, are a further technique to lower operators' TCO. They are very small cellular base stations that are designed to operate inside building. One femtocell can support 5-100 users. The idea is that the femtocell connects to the backhaul using a DSL line. Femtocells have very short range, low cost and low power.

Advantages of the femtocells include for example (Chandrasekhar et al., 2008):

- Thanks to the short range, femtocells can transmit with low power and achieve better signal-to- interference-plus-noise ratio (SINR)
- As femtocells are used for serving the traffic originating indoors, macro base stations can better serve mobile users
- In-building coverage will be better and customers will be more satisfied with the operator.
- As many operators have leased E1/T1 lines, reduced backhaul traffic will reduce operational expenses significantly. The cost of backhaul will move to the customer as he pays for the DSL subscription.

From the perspective of this paper, probably the most interesting advantage of the femtocells is the cost benefit. Deploying femtocells will lower both the operational and capital expenditures of a mobile operator. Annual operating expenses will be less than $200 dollar per femtocell where as the operating costs for a macro base station are tens of thousands of dollars (Chandrasekhar et al., 2008). The operator will save in site lease costs, backhaul and electricity costs.

As a new expenditure to the operators will be the subsidization of hardware installed in customers' homes. Operators will need to provide a scalable architecture to transport data over IP, which will be a new operational expenditure. As subsidies can be considered as costs for network infrastructure, they are counted into capital expenditures (Chandrasekhar et al., 2008). Possible hardware upgrades will also cause operational costs.

One of the challenges of the femtocells is that the cost for hardware has to be low enough. Hardware for a femtocell is forecasted to be something between $120-$200 (ABI Research, 2007). However, to be able to achieve a thriving market, the price should be less than $100. In addition, there are unsolved issues related to for example interference with macro network (Meissner, 2008).

Table 4 shows a comparison done between macro cells and femtocells both with a 2 Mbps capacity (Markendahl et al., 2008). First row shows the estimations with a low level demand with 200 Mbps in the network. The second case is with a high level demand, 1000 Mbps, in the network. Data shows that the CAPEX costs for building a femtocell network infrastructure is much less than building a macro cell network. The analysis assumes ideal conditions, meaning that there is no wall attenuation and that a femtocell can provide the coverage needed. Also, operational costs are much smaller for femtocells. This is in line with Section 4.1 where the cost differences between different kinds of base stations where discussed.

4.5. Network Sharing

Network sharing is said to be very efficient way to reduce especially network rollout costs. For example in the United Kingdom, 66 % of GSM sites are shared (Parker, 2004). Network can be shared on several levels. On Radio Access Network (RAN) level elements like Radio Base Stations, Radio Network Controllers and sites can be shared. In the core, for example Mobile Switching Centers or backbone network can be shared.

Network sharing can for example improve the network coverage of an operator. Sharing a network is also feasible in areas where user density is low. In such cases, national roaming is used. In addition, time to market (TTM) will be faster and market share can increase.

The biggest savings will be in capital expenditures, but depending on the sharing agreement, there might be savings in operational expenditures as well. By sharing a new mast site with one operator, the operator saves 25 % of the costs. If there are more operators, saving will be higher. Site sharing will save costs for acquiring sites, and will result in more efficient use of resources (Giles, 2004).

The disadvantage of network sharing is that the operator might loose its competitive advantage. In addition, there might be more complexity in shared networks. There are also problems related to trust and compatibility (Hultell, 2004).

5. Conclusions and Future Work

In this paper we discussed the different costs of a mobile operator. We defined, which kinds of costs are considered as capital expenditures and which are operational. We also showed some data from earlier studies describing the cost structure of a mobile operator. The affect of the used technology on the cost of equipment was also described.

From the data can be seen that the biggest capital expenditures are related to access network. The biggest part of the operational expenditures goes to actions performed by personnel. As operational expenditures are today 70-80 % of the TCO, operators pay more attention on finding ways to decrease OPEX.

We provided a brief study on some of techniques the operators and equipment vendors are currently investigating to decrease the total cost of ownership. Heterogeneous networks have two levels: hierarchical cells and multiaccess networks. Hierarchical cells using the same radio technology with different types of cells is one of the techniques to decrease both OPEX and CAPEX. Multiaccess networks, in turn, combine different radio access technologies. Combining for example 3G and WLAN accesses is feasible especially as the hardware cost per a WLAN AP is very low (Loizillon, 2002). Self-organizing networks are one of the areas for a mobile operator to save in operational expenditures. We describe femtocells, which are small base stations deployed indoors. One of the advantages of the femto cells is the backhaul cost moved from the operator to the end user. In addition, the hardware of a femtocell is cheap. Site sharing is one popular way to reduce capital expenditures. However, there are some problems related to site sharing like complex agreements between the parties.

Based on this analysis it is impossible to say, which of the techniques would be the most cost-efficient. There are mechanisms that reduce mainly CAPEX whereas other techniques can be reduced OPEX. It is anyway clear, that the operators will continue to pay attention to reducing the TCO.

In the context of this seminar, not much material or data was found about the operating costs of the mobile operators. For example, analyzing operating costs per a mobile user in different types of mobile networks would be interesting. This would be an interesting topic for possible future studies. Data about the trends would have been interesting to find as well.

References

ABI Research (2007) "Femtocell Market Challenges and Opportunities". http://www.abiresearch.com/products/market_research/Femtocell_Market_Challenges_and_Opportunities (accessed 20.2.2009)

ABI Research (2008) "Mobile Operator Capital Expenditures Analysis". http://www.abiresearch.com/products/market_research/Wireless_Infrastructure_CAPEX_and_Base_Station_Spending (accessed 20.2.2009)

Björkdahl J. et al. (2004) "Financial Assessment of Fourth Generation Mobile Technologies", CPR 2004 Conference

Brown G. (2008) Light Reading: Self-Organizing Networks & LTE. http://www.lightreading.com/document.asp?doc_id=158441 (accessed 20.2.2009)

Chandrasekhar V. et al. (2008) "Femtocell Networks: A Survey"

Giles T. et al. (2004) "Cost Drivers and Deployment Scenarios for Future Broadband Wireless Networks – Key research problems and directions for research"

Hultell J. et al. (2004) "Business models and resource management for shared wireless networks"

Johansson K. et al. (2004) "Relation between base station characteristics and cost structure in cellular systems".

Johansson K. and Furuskär A. (2005) "Cost efficient capacity expansion strategies using multi-access networks".

Johansson K. et al. (2007) "Modelling the cost of heterogeneous wireless access networks"

Kellerer, W. et al. (2003) "The operators 'vision of systems beyond 3G: Operators' key drivers for systems beyond 3G". Eurescom report.

Loizillon F. et al. (2002) "Final results on seamless mobile IP service provision economics", IST-2000-25172 TONIC. Deliverable number 11

Markendahl J. et al. (2008) "Analysis of Cost Structure and Business Model options for Wireless Access Provisioning using Femtocell solutions, 19th European Regional ITS Conference, Rome, Italy, September 2008.

Meissner P. (2008) "Zukunft der Mobilfunknetze – Perspektiven für den Mittelstand". http://www.synergie-fachtagung.de/fileadmin/files/synergie-fachtagung/schloss_vollrads_08/Meissner_09_09_2008.pdf (accessed 20.2.2009)

NGMN Alliance (2007) Use Cases related to Self Organising Network. Overall Description.

Niebert, N. et al. (2004) “Ambient networks: an architecture for communication networks beyond 3G”, IEEE Wireless Communications Magazine, Vol. 11, No. 2, pp.14–22.

Odlyzko A. (2004) “Pricing and Architecture of the Internet: Historical Perspectives from Telecommunications and Transportation”.

Parker J. (2004) “Planning a Cost Effective 3G Base Station Strategy”. http://www.coleago.co.uk/uploads///Downloads/Planning%20a%20Cost%20Effective%203G%20Base%20Station%20Strategy%20JP%20Jul%2004.pdf (accessed 20.2.2009)

Smura T. et al. (2007) “Final results on economics of converged network and service environments “. ECOSYS Deliverable number 22.

Verbrugge S. et al. (2005) “Modeling Operational Expenditures for Telecom Operators”. Proceedings of ONDM2005, Milan, Italy, February 2005 pp.455-466.

System Dynamics of Client-Server and Peer-to-Peer Content Distribution

Mikko Heikkinen, mikko.heikkinen@tkk.fi
Mikko Särelä, mikko.sarela@ericsson.com

Abstract

We develop a system dynamics model to study the differences and relations between client-server and peer-to-peer content distribution modes. Our model preliminarily confirms that client-server is the preferred distribution mode of ISP but peer-to-peer saves costs significantly for content provider. Further work is required for definitive conclusions.

1. Introduction

The aim of this paper is to develop a preliminary system dynamics model to investigate the differences and optimizations between the client-server (CS) and peer-to-peer (P2P) modes of content distribution.

We seek answers to the following research questions by investigating the model and literature:

- Which is the most efficient way to distribute information?
- Is there a point of transition where peer-to-peer becomes a more important mode for information distribution than client-server?

Our paper is structured as follows: first, we introduce the reader to the CS and P2P distribution modes. Then, we summarize the system dynamics modeling method. Based on the introductory chapters, we present our preliminary system dynamics model. We finish by some discussion and conclusions where we provide partial answers to our research questions and issues for further research.

2. Content distribution modes

Currently the two main modes of content distribution are client-server (CS) and peer-to-peer (P2P). We introduce both briefly in the following subsections, and compare them in the third subsection.

2.1. Client-server

Client-server (CS) is the most commonly used content distribution mode in commercial Internet services and applications. The basic idea of CS is that several clients request information from a server or a group of servers. Figure 1 has an example of a CS system where multiple Clients A(1), A(2), etc. located in the network of Internet Service Provider (ISP) A access the Server B located in the network of ISP B. Server B has the only distributable copy of Data which each client has to fetch independently.

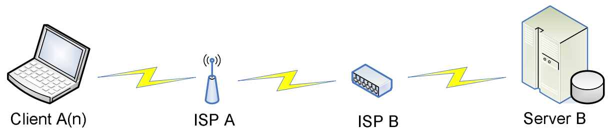

Figure 1. Client-server distribution mode

2.2. Peer-to-peer

Currently 30-80 % of Internet traffic is estimated to be peer-to-peer (P2P) based file sharing traffic. The proportion of the traffic varies significantly depending on the source, see Haßlinger (2005), Sandvine (2008), and Schulze and Mochalski (2009) for some studies. P2P is most commonly used in file sharing type of content distribution, but it has also been applied for instance to communications (e.g. Skype and Google Talk) and streaming (e.g. Coolstreaming, PPlive, PPstream, TVAnts, Sopcast, QQlive, among others).

No established definition for P2P exists. Nevertheless, according to Androutsellis-Theotokis and Spinellis (2004):

"Peer-to-peer systems are distributed systems consisting of interconnected nodes able to self-organize into network topologies with the purpose of sharing resources such as content, CPU cycles, storage and bandwidth, capable of adapting to failures and accommodating transient populations of nodes while maintaining acceptable connectivity and performance, without requiring the intermediation or support of a global centralized server or authority."

The same authors classify P2P content distribution systems into multiple subcategories based on several properties such as centralization and structure. Many peer-to-peer systems employ centralized elements for bootstrapping and authorization. They also differ significantly among their internal structure.

In this paper we use a simple model illustrated in Figure 2 for P2P content distribution mode. The basic idea is that Peer A(1) gets the Data from a source such as a Server, and starts to distribute the Data in a P2P network which for simplification is located completely within ISP A. To another peer downloading the Data from Peer A(1), such as Peer A(2), Peer A(1) acts as a Seed, i.e. a peer having a complete distributable copy of the Data. A third peer, Peer A(3), may acquire the Data from Peers A(1) and A(2) simultaneously.

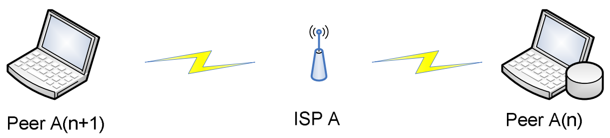

Figure 2. Peer-to-peer distribution mode

2.3. Comparison

There are significant differences between the client-server and peer-to-peer distribution modes. In CS, the server administrator is responsible for most of the costs related to the content distribution such as maintaining and securing the server, transit charges and ensuring sufficient data transfer, storage and processing capacity. The client typically only pays the access cost, unless the service has an additional cost, such as a subscription fee. In P2P, the cost of distributing the content is divided among peers. As the peers are typically behind residential access links, the ISPs providing consumer access to the Internet pay a significant amount of the transfer cost in form of increased transit fees. This has lead to efforts to control and restrict P2P traffic, see for example CRTC (2008) and FCC (2008).

In CS, the server administrator has complete control over who may access the server. The control also leads to the possibility of requesting fees to access the content. In P2P, the possibilities for controlling the distribution and usage of the content are very limited, or non-existing. Some experiments have been made with the distribution of Digital Rights Management (DRM) controlled content in a P2P network but with little success so far (TorrentFreak, 2008).

The main benefit of P2P systems is that they typically scale well to demand. They have been experimented in delivering video streams from high-interest public events such as the inauguration ceremony for President Obama (Roettgers, 2009).

However, the performance of P2P systems is typically more difficult to predict than of CS systems. It depends on the number of seeds and peers distributing the content in question. In CS systems, the performance can be adjusted within certain limits by traffic engineering.

Table 1 summarizes our comparison.

Table 1: Comparison between CS and P2P

Item	Client-server	Peer-to-peer
Cost	Server-centric	Distributed
Access control	Server, DRM	Limited DRM
Scalability	Limited	Unlimited
Performance	Predictable	Unpredictable

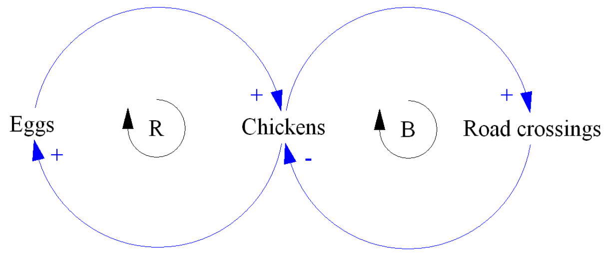

Figure 3. Causal diagram example

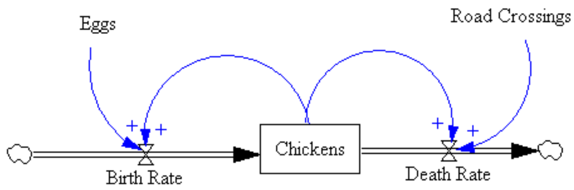

Figure 4. Flow diagram example

3. System dynamics method

System dynamics investigates different organizations, markets and other systems with the aim of understanding the reasons for the dynamic behavior of the systems while taking into account their holistic nature. Compared to non-structured time series models, the system dynamics models commonly predict better long time scale behavior and especially junctures at the expense of levels of single variables. Limited growth, instability and oscillations are commonly built within system dynamics models.

Two types of diagrams exist within system dynamics. Causal diagrams contain reinforcing and balancing loops, which affect the state of the system accordingly. Flow diagrams contain stocks and flows integrated with causalities affecting the volume of the flows. Figures 3 and 4 have examples of causal and flow diagrams, respectively, applied to a system where chickens are born from eggs and dying in road crossings.

The reader is encouraged to refer to e.g. Sterman (2000) for a review on system dynamics.

4. System dynamics model

We developed a preliminary system dynamics model to investigate the costs in transmission, storage and processing in both client-server and peer-to-peer content distribution modes. In this section, we present our cost distribution axioms model, the actual model, the input parameters to the model, and the results from the model.

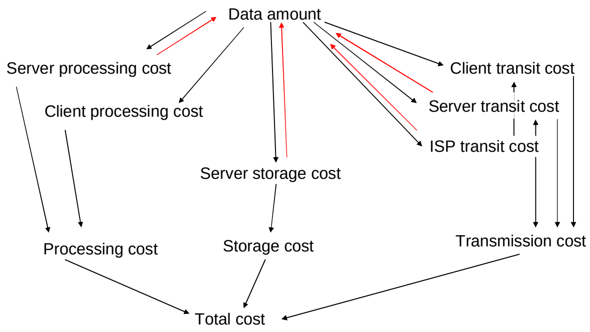

Figure 5. CS dynamic cost model

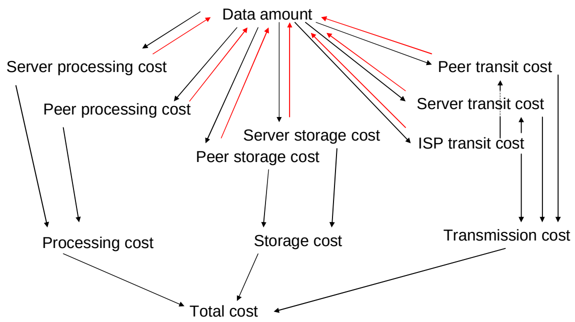

Figure 6. P2P dynamic cost model

4.1. Cost distribution axioms

The main cost variables we are interested in are processing, storage and transmission. Processing is the cost of processing effort required for successful content distribution by the processing elements in servers, peers and clients. Storage is the cost of storing the content in servers and peers. Transmission is the cost of transmission which occurs among servers, peers and clients, and their respective ISPs.

The dynamic cost models for CS and P2P are sketched in Figures 5 and 6, respectively. In the CS model, client is assumed to have no feedback options to the distributed data amount, whereas in the P2P model peers have the possibility to affect the amount of data distributed. Also, in the P2P model, the server is assumed to be the original source, or seed, for the data.

Our cost distribution axioms are summarized in Table 2. Server is the most expensive element as it is the centralized service element facing the heaviest requirements. Both processing and storage costs are significantly lower for clients than peers because clients act as passive participants in the CS system. Transmission cost depends on the intensity of peer actions and is not necessarily significantly higher for a peer than a client if the peer does not engage heavily in redistributing the content it has acquired.

Table 2: Cost distribution axioms

Cost item	Axiom
Processing	Server >> Peer >> Client
Storage	Server >> Peer >> Client
Transmission	Server >> Peer > Client

4.2. Model

In the actual system dynamics model both CS and P2P domains are combined, see Figure 7. The model consists of two main parts: stakeholder dynamics and user role dynamics. The stakeholder dynamics depicts the current status of capacities of a consumer, an Internet Service Provider (ISP) and a Content Provider (CP). The capacity growth rate is determined by a general technology development rate and an upgrade rate determining the actual upgrade from the potential development. For a consumer's growth rate, the growth rate of an ISP has a restricting effect. The upgrade rates of an ISP and a CP are dependent on the number of potential users, clients and peers from the user role dynamics part.

The user role dynamics part models the change of user role from a potential user to a client, and from a client to a peer. The transformation from a potential user into a client is controlled by a simplified Bass (1969) diffusion model (see Sterman, 2000 for details about the model). The transformation from a client into a peer is controlled by a peer threshold parameter related to consumer's capacity and a freeloading rate.

The model was developed using Vensim PLE 5.8c which is a system dynamics simulation software available free to academic use. The flow diagram structure described in Section 3 is in use in the model.

The model equations are listed in Table 3 using Vensim format. Measures taken to avoid division by zero errors have been omitted from the equations. Variables acting as input parameters are described in Section 4.3.

The size of the stocks is simply determined by the respective inflow and outflow rates. The adoption rate follows the Bass diffusion model. The client to peer rate is determined by the inverse proportion of freeloaders if the threshold criterion is met. The ISP upgrade rate is designed to increase as the number of clients and peers grows, emphasizing peer growth. The CP upgrade rate increases as the number of clients grows but decreases as the number of peers grow. Stakeholder capacity growth is determined by technology development and upgrade rates. Regarding consumer's capacity growth rate, the ISP's capacity growth rate acts as a constraint.

Table 3: Model equations

Variable	Equation
Adoption rate	Potential*Effect + (Client+Peer)*Contact*Fraction*Potential / Total
C rate	MIN(C tech dev*C upgrade, ISP rate)
Client	INTEG(Adoption rate - Client to Peer rate, 0)
Client to Peer rate	IF THEN ELSE(Consumer>Peer threshold, (1-Freeloading rate)*Client , 0)
Consumer	INTEG(C rate, C0)
Content provider	INTEG(CP rate, CP0)
CP rate	CP tech dev*CP upgrade
CP upgrade	MIN(MAX(Client/Potential, Client/(Client+Peer)), 1)
ISP	INTEG(ISP rate, ISP0)
ISP rate	ISP tech dev*ISP upgrade
ISP upgrade	MIN(MAX((Client+Peer)/Potential, Peer/Client), 1)
Peer	INTEG(Client to Peer rate, 0)
Potential	INTEG(-Adoption rate, P0)

Table 4: Common input parameters

Variable	Value (Unit)
C upgrade	0.5
C0	1
Contact	100
CP0	1
Effect	0.01
FINAL TIME	10 (Year)
Fraction	0.015
Freeloading rate	0.5
INITIAL TIME	0 (Year)
ISP0	1
Peer threshold	1
P0	10 000
TIME STEP	1 (Year)
Total	50 000

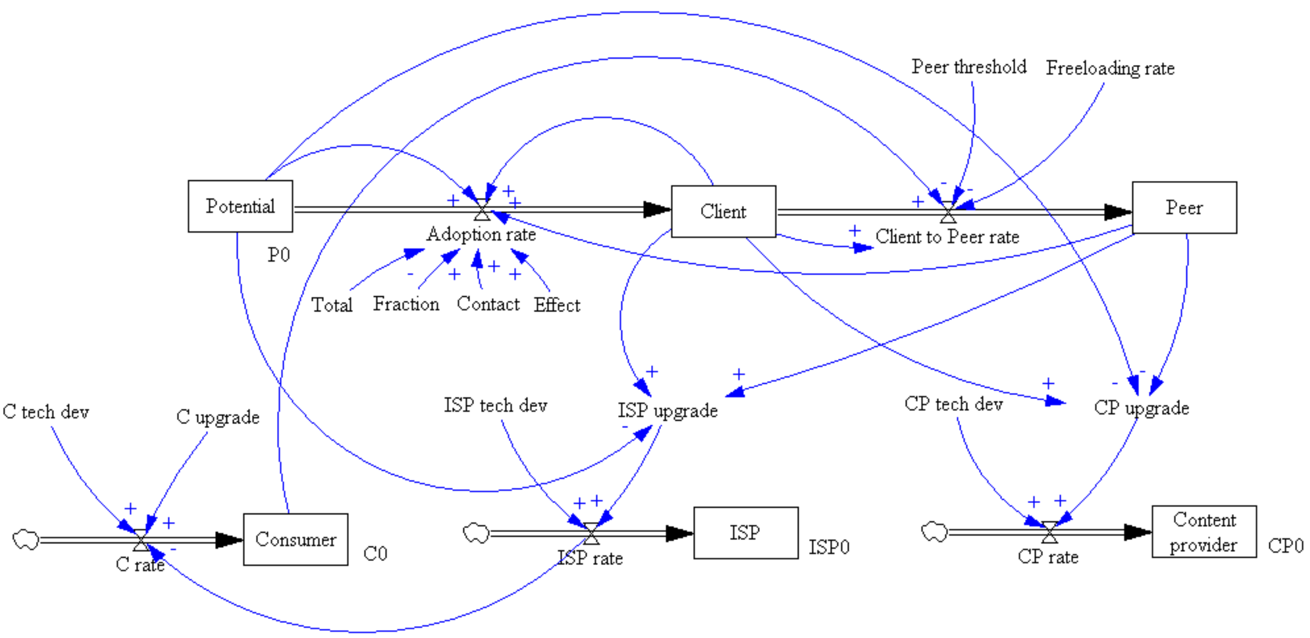

Figure 7. System dynamics model

4.3. Input parameters

Common input parameters to the model are listed in Table 4. Also simulation related time parameters are included: the simulation is run for 10 years with a time step of 1 year. We estimated the common input parameters to match consistent behavior. The initial technology growth for all stakeholders is set to 1. We note especially that we do not model the growth in absolute measurement units such as Mbps, MIPS or GB; instead we use a relative growth index.

The input parameters specific to transmission, processing and storage capacity scenarios are listed in Tables 5 and 6. Technology development parameters are based on Nikander (2009). Transmission technology is expected to develop 10 times better in ten years (i.e. 1.26^10≈10), processing and storage 100 times better in ten years (i.e. 1.58^10≈100).

Table 5: Transmission input parameters

Variable	Value (Unit)
C tech dev	1.26
CP tech dev	1.26
ISP tech dev	1.26

Table 6: Processing and storage input parameters

Variable	Value (Unit)
C tech dev	1.58
CP tech dev	1.58
ISP tech dev	1.58

4.4. Results

We illustrate the development of our stock variables by varying one of our key parameters, the Freeloading rate, in three different values (0.10, 0.50, 0.90). Figures 8 and 9 illustrate the effect on the number of clients. As expected, the number of clients increases as the freeloading increases (i.e. people are not willing to contribute by turning into peers), and the number of peers decreases as freeloading increases.

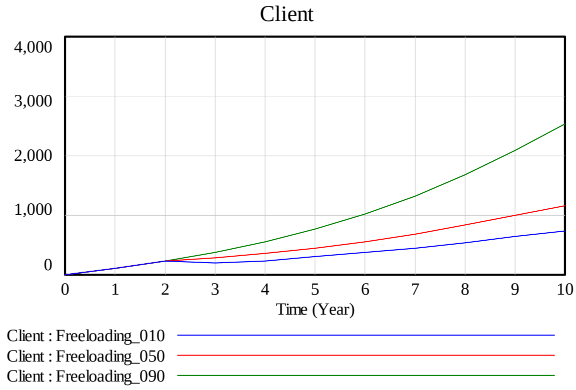

Figure 8. Clients (Transmission, Freeloading)

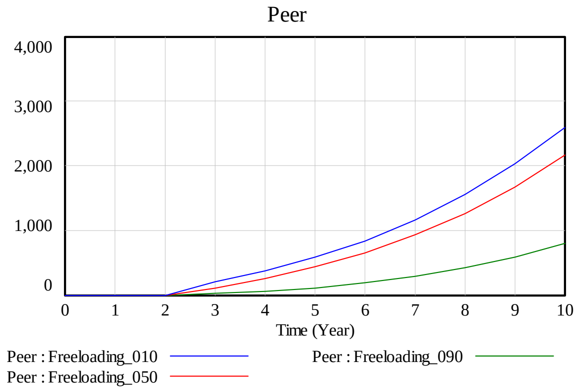

Figure 9. Peers (Transmission, Freeloading)

Figures 10-12 illustrate the effect of freeloading rate on consumer's, ISP's and content provider's transmission capacity. Intuitively rising freeloading decreases both consumer's and ISP's capacity demand but increases content provider's need for capacity. This is due to the axiom that servers and peers require more transmission capacity than clients: when there are less peers in the system, there is a growing need for client capacity (i.e. content provider's capacity). The similarity in consumer's and ISP's behavior is because of the consumer's dependence on ISP's capacity upgrades.

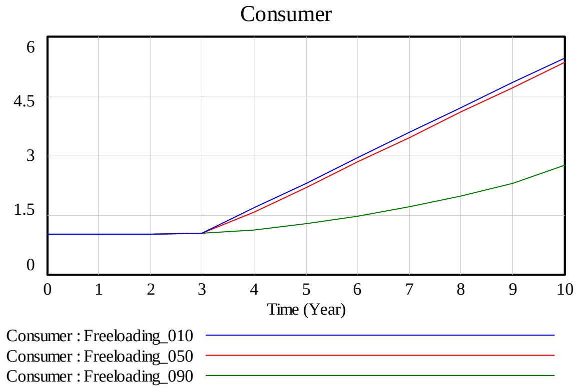

Figure 10. Consumer (Transmission, Freeloading)

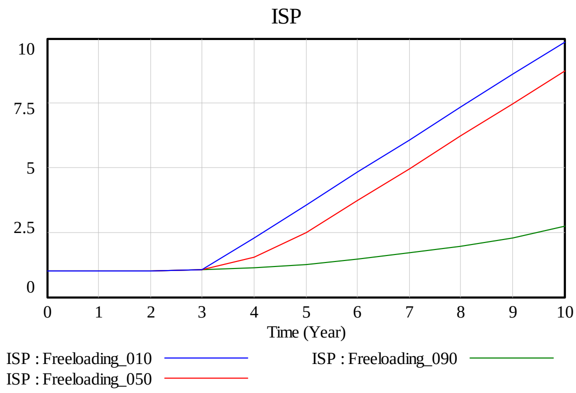

Figure 11. ISP (Transmission, Freeloading)

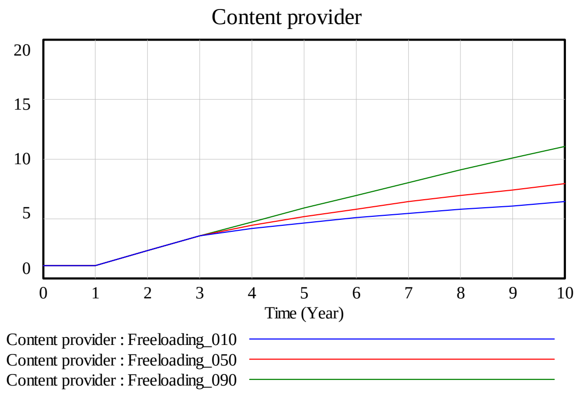

Figure 12. Content provider (Transmission, Freeloading)

Considering the capacity growth of different stakeholders, we notice that content provider reaches its maximum theoretical capacity when freeloading is very common (i.e. the number of peers is small); ISP acts vice versa. Freeloading has relatively larger effect between its upper values than its lower values.

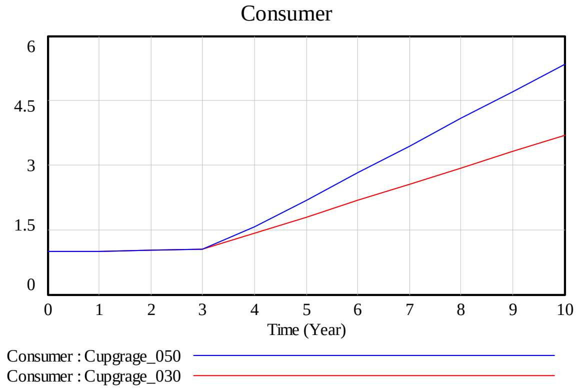

Figure 13. Consumer (Transmission, C upgrade)

Another interesting key variable is the Consumer upgrade (C upgrade) which has a significant impact on consumer's capacity growth as seen from Figure 13 with values 0.30 and 0.50. Intuitively, decreasing the C upgrade results in lower capacity growth for consumer.

Modeling processing and storage ends in similar results, only the order of magnitude in stakeholder capacity results is different.

5. Discussion

We acknowledge several limitations in the model and action points for further development. Some of them are discussed in this section.

The most significant flaw in the model is that it does not have return flows for users (i.e. flows from Client to Potential, Peer to Client and Peer to Potential). We omitted the flows for simplicity. Including the flows requires implementing an advanced user role dynamics with criteria for users to return to their previous roles.

Another flaw in the model is that it does not differentiate between different types of content. Implementing content differentiation would require the simulation of absolute capacities with additional variables in stakeholder dynamics.

The model does not differentiate between downlink and uplink transmission capacity. At its current stage the model allows the inclusion of them only through simulating different scenarios with different input parameters. The key parameters in modeling the difference are Consumer capacity and Peer threshold. Currently most ISPs market asymmetric Internet access to consumers where uplink is significantly slower than downlink. Also, P2P systems are sensitive to uplink capacity. Therefore, in an uplink scenario, both parameters should be scaled down.

According to Isaac and Walker (1988), freeloading increases with group size. In our model, for simplicity and variability, the Freeloading rate is constant. It could be implemented for instance using a lookup function to reflect changes in group size (i.e. the number of peers).

Furthermore, the model could be expanded by implementing technical variables behind capacity development in more detail for example by including ISP's potential efforts to restrict P2P traffic, and by applying sociological theory to the user role dynamics for example according to Fehr and Fischbacher (2003).

Considering the Sterman (2000) criteria for testing system dynamics models, we discuss some aspects. The model was limitedly tested for boundary adequacy: the minimum concepts for addressing the problem should be present. Regarding structure assessment: the model structure lacks detail and feedback loops; the model conforms to basic physical laws, and the model captures partially the behavior of actors relevant to the system. The model is dimensionally and parametrically consistent. The model was limitedly tested in extreme conditions and limited sensitivity analysis was done. Behavior reproduction is according to expectations: no behavior anomaly or surprise behavior was noticed. Integration error and family membership were not tested.

6. Conclusions

We did not find concise answers to our research questions. Our model does not give any definitive answer to the relative efficiency between client-server and peer-to-peer systems. Reaching a definitive answer would require calculating the total social cost in different content distribution scenarios.

According to our model, content provider benefits from the P2P model as its capacity requirement (and thus cost) increases slower with P2P compared to CS. On the other hand, ISP benefits significantly if P2P does not become a dominant distribution mode as its capacity development (and also cost) rises significantly if P2P gains popularity. Also consumer has the opportunity to increase its capacity more if P2P is popular and ISP's capacity increases, but it can be seen mainly positive from the viewpoint of the consumer.

In our model the consumer willingness to share resources is the most important impact factor for P2P popularity. To gain a realistic view on the factors affecting the relative importance of the two distribution modes, more parameters such as consumer willingness to adopt new P2P-based technologies, actual capacity measurements, and ISP peering and transit economics should be implemented in the model.

In conclusion, we were able to demonstrate the validity of our axioms with our model, but the model requires significant further work to give a realistic view on the differences and relations of the two distribution modes.

Acknowledgements

We gratefully acknowledge our colleague Juuso Karikoski as the author of Figure 3 and the discoverer of the following literature: Isaac and Walker (1988).

References

Androutsellis-Theotokis, S. and Spinellis, D. (2004) A survey of peer-to-peer content distribution technologies. ACM Computing Surveys, Vol. 36, pp. 335-371.

Bass, F. (1969) A new product growth model for consumer durables. Management Science, Vol. 15, pp. 215-227.

CRTC (2008) Review of the Internet traffic management practices of Internet service providers. http://www.crtc.gc.ca/eng/archive/2008/pt2008-19.htm (accessed 16 February 2009).

FCC (2008) Memorandum opinion and order 08-183. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-183A1.pdf (accessed 16 February 2009).

Fehr, E. and Fischbacher, U. (2003) The nature of human altruism. Nature, Vol. 425, pp. 785-791.

Haßlinger, G. (2005) ISP platforms under a heavy peer-to-peer workload. In P2P Systems and Applications, R. Steinmetz and K. Wehrle, Eds. Springer, Berlin, pp. 369-381.

Isaac, R.M. and Walker, J.M. (1988) Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism. The Quarterly Journal of Economics, Vol. 103, pp. 180-199.

Nikander, P. (2009) Long term price trends and telco industry. https://noppa.tkk.fi/noppa/kurssi/s-38.4043/materiaali/dr._nikander_s_presentation.pdf (accessed 17 February 2009).

Roettgers, J. (2009) Inauguration P2P stream a success, despite backlash. http://newteevee.com/2009/02/07/cnn-inauguration-p2p-stream-a-success-despite-backlash/ (accessed 16 February 2009).

Sandvine (2008) Sandvine releases global Internet traffic trends report. http://www.sandvine.com/news/pr_detail.asp?ID=203 (accessed 16 February 2009).

Schulze, H. and Mochalski, K. (2009) Ipoque Internet study 2008/2009. http://www.ipoque.com/study/ipoque-Internet-Study-08-09.pdf (accessed 19 February 2009).

Sterman, J.D. (2000) Business Dynamics: systems thinking and modeling for a complex world. McGraw-Hill, Boston.

TorrentFreak (2008) BitTorrent shuts down video store, brings back search. http://torrentfreak.com/bittorrent-shuts-down-video-store-brings-back-search-081207/ (accessed 16 February 2009).