The innovative power of overlay networks

Willem Toorop
w.toorop@os3.nl

Abstract

This paper tries to illustrate how innovations on the Internet emerged alongside a continuously construction of overlay networks on the available infrastructure (the underlay networks) [1]. The incentive for those constructions: resource sharing. Stable, well defined and understood underly networks are identified as a condition for the construction of overlay networks. The question evaluated: can market competition in the underlay networks provide the stability needed for future innovations.

Contents

1 The incentive for resource sharing

In society goods are not equally dived among people. Some have more than they need, others have a shortage. To minimize shortage and maximize the availability of goods, is the social benefit that is pursued by governmental regulations that try to prevent anti-competitive practices. Monopolies or cartels could induce artificial scarcity for a certain good with the aim of asking a higher price and gain more profit from its sale. It is in societies best interest that fair competition between the suppliers of certain goods is in place so that its consumption can enjoyed by the widest possible group [2].

Although such competitive regulations make sure that the potential distribution of a good is maximized (by disabling artificial scarcity), they do not prevent a good being over-provisioned to some and under-provisioned to others. Other measurements of the welfare state towards an optimal equilibrium are: leveling of incoming, study financing, tax reductions, social security benefits, etc.

Some goods are delivered by resources. The holders of those resources might sometimes need the goods and sometimes not. With resource sharing, the holders of resources cooperate in such a way, that over- and under-provisioning of the goods they provide are minimized. The resources are coordinated by the holders in such a way, that the goods (provided by the resources) not needed by a holder can be used by others.

Many of the goods in information and communication technology (ICT) are provided by certain resources: Hard disks and memory provide persistent and temporary storage, processors provide computing power, cabling or Internet subscriptions provide bandwidth. Most users of those goods do not need the capacity available to them most of the time, but might other times feel a shortage of them. For example, when writing this text, I use almost nothing of the memory and processing power my computer has available. When I will later render this document to PDF format however, it will use all resources available to produce the file as quickly as possible.

Rendering of a LaTeX document to a PDF file is still bearable with the resources of a single computer, but this becomes very different when three-dimensional photo realistic animations need to be rendered, or when simulating the events during the big-bang or analyzing radio-telescopic data in search of extra-terrestrial life. In such computational intensive cases, one can't have enough processing power providing resources (CPU's). Cooperation with other holders of computing resources that sometimes need massive amounts of processing power is for such users very desirable.

The potential goods that come from computing resources are expressed by their capacity: The number of instructions or computations a processor can perform in a seconds (nowadays expressed in gigahertz) and the potential memory that can be used (expressed in number of bytes). With resource sharing, the capacity is the sum of all the participants. The pool of individual resources have become a single resource for the use of the common. The contribution of low capacity resources by all, delivers a high capacity resource for them in return.

It is essential though, that the usage of the common pool is well balanced. The cost of the goods is distributed over all participants in the pool and not felt by the individual user of the goods anymore. Also, with a big pool, the addition of a participants individual resources is marginal compared to the capacity of the common pool. There is no personal incentive that restricts unlimited usage of the goods provided by the resources anymore. The common capacity is however finite and this creates the danger of loosing all the benefits that a common pool would otherwise offer. This is known as "The tragedy of the commons" [3].

2 Overlay networks for resource sharing

The sharing of computing resources has been a major initial motivation for the ARPANET[1] and the Internet [4, 5]. In the late sixties and early seventies Computers were expensive and not commonplace commodity. For universities, research institutes and the military, to be able to use each others computers (when available) was a worthwhile aspiration.

In the design of the ARPANET, telephone lines were planned to be used to interconnect the computers on the different universities. Because of robust-

[1] The predecessor of the Internet

ness concerns[2] packet-switching was chosen as the communication-modus on the exclusive telephone-connections (leased lines) that would interconnect the computers. Packet-switching is a method in which multiple communication sessions can be shared on a single communication-line by cutting those sessions into small packets and transmitting them interlaced.

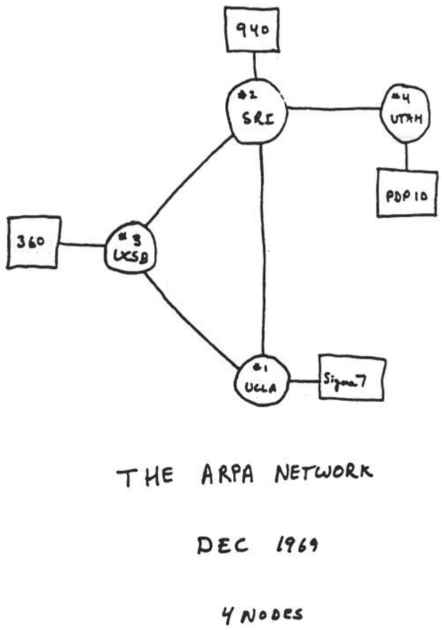

Fig. 1: © Alex McKenzie

In figure 1 a picture of the ARPANET in 1969 is shown. The robustness gain can be illustrated as follows: If the Stanford Research Institute (SRI) and - the University of California, Los Angeles - (UCLA) would want to communicate, but their direct leased line would malfunction, their session could be relayed by - the University of California, Santa Barbara - (UCSB). The packets of the session between SRI and UCLA could, by using the packet-switching communication modus, then be interlaced with the packets of the sessions between SRI and UCSB, and between UCLA and UCSB.

Using packet-switching introduced, besides sharing of computing power, another form of resource sharing. As can be seen in figure 1 there were (in december 1969) no leased line between the University of Utah (UTAH), and UCSB and UCLA, however because of packet-switching, the leased lines of SRI to those universities could be used for that purpose. Furthermore, each new participant connecting to a node in the ARPANET would not only be able to communicate with the resources offered at that node directly connected, but also with all other nodes in the network without having to buy leased-lines to all those other nodes as well.

It is clear, that this way of exploiting telephony lines was not in the interest of the telephony companies. For them it would have been better if each university would have to have a dedicated leased line to each remote university its computing resources it wished to use. Also, as the owner of the leased line, they had (and have) the potential power to restrict its usage. They could have tried to prevent packets flowing over the line not directly addressed to the other side of the line, thus motivating universities to acquire more leased lines from them.

However, the phone network was (and is) under substantial governmental regulatory control. Just like companies offering public transportation, by government regulation, had (and have) to offer there services in a non-discriminatory fashion, the telephone companies were (and are) obliged to offer common carriage access to all.

[2] There are rumors that the robustness came from the desire to build a nuclear war resistant network, but this is not the case. Even without a nuclear attack the links were unreliable.

Also, as can been seen in figure 1, on the early ARPANET, SRI provided more network resources then the other participants. It offered for example the potential communication with two extra universities to UTAH, were UTAH offered non extra. In the same fashion, it offered to UCSB and to UCLA (besides the backup path) also one extra resource. They could of course reflect that ratio in the division of the costs of the leased lines, but those costs are marginal compared to the main gainings, the capacity of the common computing resource, which usage provisioning can not be so easily divided. (The tragedy of the commons)

By introducing the packet-switching communication-modus, the ARPANET introduced new properties that were not found in the underlying telephony network. Besides the robustness, the resource sharing aspect added the possibility for two not directly connected parties to communicate. With packet-switching a new network emerged, providing new means for contacting and communication, on top of the telephony network. Such a network, on top of another network, is called an *Overlay network*.

The homogeneity of the underlying phone-network (common carrier access), stimulated the use of innovative new techniques (packet switching and forwarding) to create the desired properties that the phone-network itself did not provide. The overlay-network made the underlying phone-network more heterogeneous.

3 Towards layered networks

Packet switching introduced new desirable properties, but, because of the sharing of the communication line, at the cost of guaranteed bandwidth. The forwarding of packets was the responsibility of a smaller computer dedicated for this purpose: the Interface Message Processor (IMP, see figure 2).

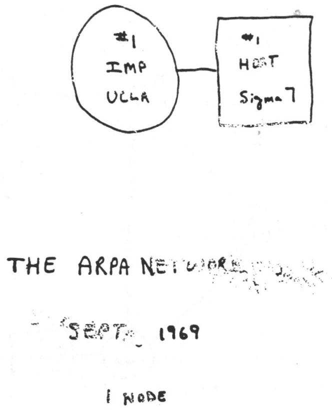

Fig. 2: © Alex McKenzie

It was the IMP's task to distribute the packets coming from the main computers (the hosts) over the leased lines. For this purpose it maintained first-in first-out queues. When a packet would arrive from the host or from a leased line, it would append the packet at the end of the queue for the destined communication line. When there was space on the line to sent a packet, it would remove the packet from the beginning of the queue and sent it over the line. When a queue for the most direct line would be full, the IMP would look if another line could be used to reach the destination and try to append it to that line's queue. If all would fail, it would simply drop the packet. Present-day IP-routers still operate in this way. It is generally referred to as "*Best effort delivery*".

Best effort delivery created a few problems for communication. When the network would be best utilized (the usage of the goods from the resources maximized), the packets that constituted a message might not arrive in the same order (because different paths were chosen) or might even not arrive at all. Even whole distinct messages might arrive in a different order (or partly or not at all).

In 1974 Vincent Cerf en Robert Kahn published a paper describing a method to overcome those issues [6]. By carefully numbering the packets and acknowledging the arrival of packets, endpoints in the network could construct a reliable connection in which bidirectional communication in the correct order was guaranteed. The method was described as the Transmission Control Program (TCP).

From the mid seventies on the awareness grew that certain aspects of computer networking could best be defined in layers. The International Organization for Standardization (ISO) presented its Open Systems Interconnection (OSI) model in 1980 [7]. For example IMP's forwarding packets did not have to be aware of the ordering and retransmission that was the responsibility of the hosts. The aspect of handling host-to-host connections, done by the hosts (transport layer in OSI), could be build *on top of* the handling of destination addresses, done by the IMP's (network layer in OSI).

Also, the concept of interlacing packets on a medium, such as a telephone line, doesn't really involve routing. On a medium such as radio or COAX cables, all communication is heard by all. The communicating parties can all address each other without any packet needing forwarding. The OSI model therefor defines two types of addressing. Path determination and potentially forwarding needing **logical addressing** on the Network layer *on top of* the **physical addressing** on the Data Link layer. To make the distinction with the Network layer, packets on the Data link layer are called frames. The frames can be transmitted on a diversity of different physical media: the Physical layer.

The incentive for creating those layers was to provide more potential heterogeneity in computer networks. The ARPANET community shared the vision of a layered network, although not as rigorously as ISO, and later specifications of TCP did not define it as a monolithic system, but divided the connection maintaining properties from the internetworking properties (Internet Protocol). The combination called the Internet Protocol Suite (or informally TCP/IP).

It was not until January the first 1983 that the TCP/IP protocol was, with some struggle, deployed on the ARPANET. One of the main motivations was the growth of the network. The protocol that was used until then, Network Control Program (NCP), could only address 256 hosts.

The layered architecture was a success however. It stimulated initiatives for creating applications on the endpoints (hosts) on top of the Internet Protocol Suite using its diverse communicational modes in a uniform way. It also made it possible to easily connect Local Area Networks (LANS) and other types of networks below the IP layer. More importantly, the separation of the routing from the connection handling created a new underlay network (or middle) just for routing (See figure 3). The IMP's became routers. The new IP-network could be extended by just adding routers without adding new endpoints. The possibility for a market just for the routing underlay IP-network emerged.

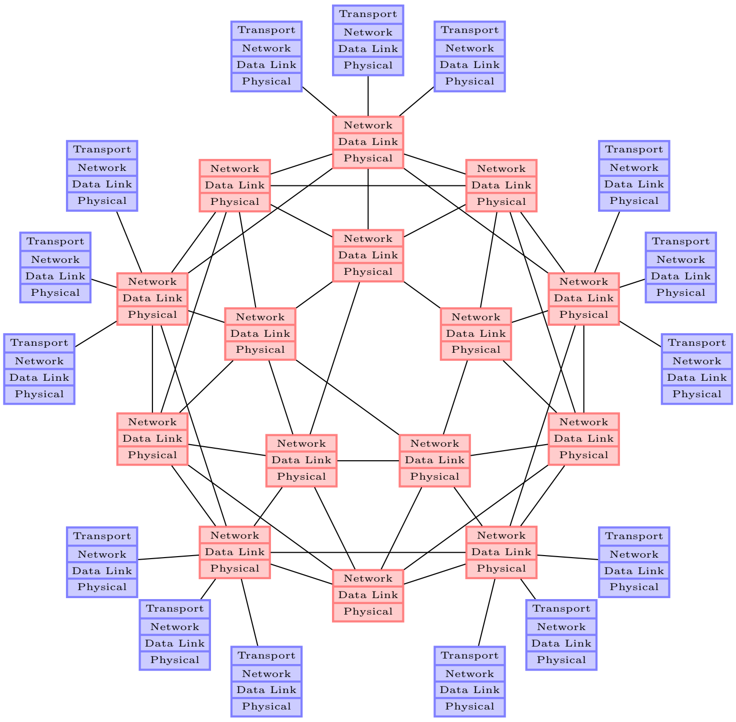

Fig. 3: A three layered middle of routers.

4 The middle market (Network layer and below)

In 1981 CSNET was initiated to interconnect the Computer Science departments at universities, stimulating the growth of the ARPANET. CSNET was funded by the United States government agency: National Science Foundation (NSF). In 1985 NSF itself began a program to create a national network for the higher education community regardless of discipline: NSFNET. Other educational computer networks emerged all around the globe around the same time. Also commercial network initiatives emerged.

NSFNET stimulated universities to seek commercial, non-academic clients, providing the means for locally initiated commercial networks to exchange traffic. Internet Exchanges (IX) emerged as the points where different network-providers could interconnect. The costs of those exchanges shared by the participants. NSFNET did not allow the commercial clients to use the long-haul interconnections of NSFNET with the successful intention to stimulate the emergence of commercial long-haul network providers. The Internet Exchanges are were commercial and educational networks met and still meet.

With all universities moved to NSFNET, and an increasing number of commercial networks interconnecting with NSFNET as well, in 1990 ARPANET retired. The internetwork communicating IP was from then on called the Internet.

Network providers have two types of relations with other providers on IX's, peering and transit. They buy transit from other (upstream) providers for

Internet access. Transit can also be sold to other (downstream) providers or to clients in their own network. Peering means that providers exchange access to their customers (downstream networks) freely for the mutual benefit [8].

Since 1994, providers exchanged what access they offer via the Border Gateway Protocol (BGP) version 4[3]. BGP let the providers exchange how far the different networks on the Internet are via their router interfaces. It is common in the IX market that providers try to demotivate upstream providers (for which they have to pay) to route traffic to networks through them, and to try to route traffic via downstream providers (that pay them) whenever and as much as possible.

This does not necessarily lead to the *most* optimal paths, but to the cheapest for the providers in between the communicating parties. However, because equally sized providers are likely to peer and because bigger providers have the long-haul lines, the competitive modus-operandi creates a hierarchy in which the closer to each other smaller providers are at the bottom and the bigger, providing access to more distant networks, higher in the hierarchy. Because the cheapest paths are likely to tend towards the shortest paths (because the costs of those communication lines is cheaper), the market incentives consequently converges automatically to a near optimal routing policy, and a fair division of the shared bandwidth on the communication lines. The IP market releases the users from the burden of fair division of bandwidth on communication lines, effectively eliminating the "Tragedy of the Commons" problems the universities on the initial ARPANET had.

5 End-point applications (Above the network layer)

Accessing the remote computers remotely from a network terminal was one of the first ARPANET applications[4]. The transfer of files from one host to the other another[5]. Those two applications were a direct simple endpoint to endpoint (end-to-end) usage of the new communicative services provided.

Email predates the Internet. The introduction of time-sharing in computer systems allowed multiple users to use a computer at the same time. Email on such hosts meant appending a message to a file that served as another users mailbox. The initial ARPANET email system was simply an extension to FTP [13]. Sending a message meant connecting to the computer on which the mailbox of the recipient resided and then appending the message to that mailbox.

This was also a simple end-to-end style usage of the early ARPANET. The communicating entities were the computers. The network presence of the people (their netizenship[14]) an account on those computers.

In the early eighties, computers became commodity hardware. It became feasible to use computers for a dedicated task. This changed the way in which the ARPANET applications worked from a end-to-end fashion to a client server approach.

In 1982 a new standard for email was proposed: the Simple Mail Transfer Protocol (SMTP). Its objective: "to transfer mail reliably and efficiently"

[3] Initially defined in 1994 in [9] but because of a number of errors and ambiguities and the legal clarity required by industry, its final definite description only appeared in 2006 [10]

[4] TErminaL NETwork (Telnet) [11]

[5] The File Transfer Protocol (FTP) [12]

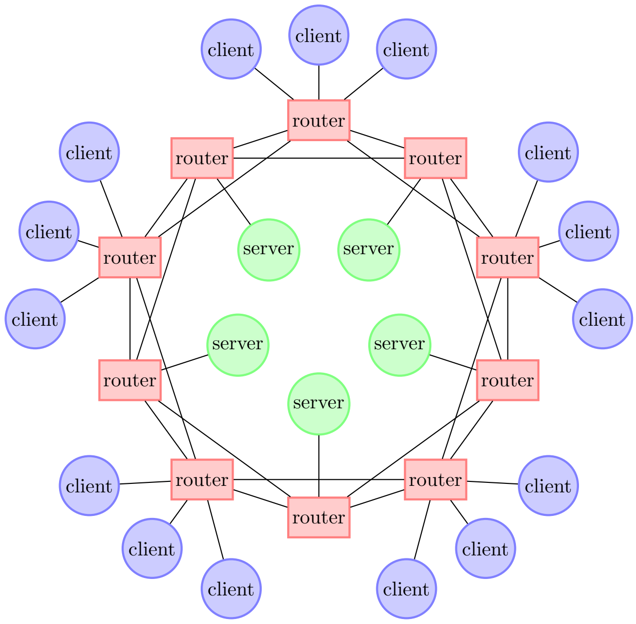

Fig. 4: Servers as the new middle

[15]. Users were not merely accounts on computers, but had their own email address. SMTP made sure that email send to an email address would be send possibly through multiple SMTP servers to eventually arrive on the host were the mailbox of the user would reside. In 1984, the Post Office Protocol (POP) [16] allowed for remote access of the mailbox. Client programs could be run from any computer and access the mailbox via POP and send email via SMTP. SMTP and POP introduced a new coherent overlay network allowing persistent messaging. SMTP servers acting as routers for messages using people (or organizational) bound addresses and POP servers acting as the mechanism to access persistent mailbox storage *in* the network. The servers acting as a new middle; the client email programs as new endpoints.

Another mayor innovation realized around the same time was the Domain Name System (DNS) [17]. With NCP the maximum number of hosts on the ARPANET was 256. With NCP all hosts on the ARPANET where enumerated in a file, `HOSTS.TXT` that could be retrieved from SRI. The transition to TCP/IP increased the theoretical potential number of addressable hosts to 4 billion. Clearly a new naming mechanism was needed.

DNS introduced a hierarchy in the naming scheme. DNS servers are authoritative for a position in the hierarchy and everything below that point (the zone). Any DNS server can be queried for a domain name. The answer can be the IP-address of a host, or the redirection to another "nameserver" lower in the hierarchy. DNS servers authoritative for a zone can as such delegate authority of names lower in the DNS-namespace to other DNS servers. The load of the lookups is as such distributed and delegated to the organizations serving certain parts in the namespace. Furthermore additional caching nameservers were introduced to further mitigate the burdening of name-lookups from clients.

DNS also made it possible to lookup other records than just host addresses. It could also be used, for example, to lookup the SMTP server for a zone. Over the years many new record types have been introduced to direct client to information associated with a certain zone. The added communicative properties provided by the DNS overlay network: a generic distributed way for looking up information.

Of course the most prominent exploitation of the client-server architecture is the World Wide Web (WWW) [18]. It introduced hypertext, which content is distributed over the Internet.

6 The consumer market

In the early nineties commercial networks existed and emerged that practiced two-sided marketing models. Besides offering networking applications themselves, such as their own propriety email, file-sharing and bulletin board services, they offered their consumers information services to their other customers: the information providers. CompuServe was one of the biggest. It offered restricted access to websites from 1990.

However the attractiveness of the externalities provided by the Internet made them all transit to the role of Internet provider. The distributed way applications evolved by protocols on the Internet provided larger and more diverse information services and communication methods than the propriety systems developed by the initial commercial network providers. They could no longer ask information providing customers for money, when the biggest audience was on the Internet anyway.

The entrance of consumers on the Internet however dramatically altered its architecture. No longer were all hosts equal endpoints. The homogeneity offered by IP did not extend into the consumers houses. Initially Internet access was provided by dial-up telephony links which made the deployment of servers in the home problematic. Later, broadband access and cable modems provided a permanent Internet connection, but isolated the home network of the consumers with Network Address Translation (NAT) to alleviate the IP address exhaustion that was already apparent to manifest.[6] The middle in the overlay networks became the de-facto middle of the Internet, at least for the consumers. The dominance of the World Wide Web as the platform for information services and Content Providers (CP's), changed the roles of the participants as show in figure 5.

Until the bursting of the dot-com bubble in 2001, the Internet was perceived by the market as a way to push content. Content and information services providers tried to create a two-sided market, generating attractive services and content for the consumers to attract customers: advertisers. After 2001, successes were made by initiatives that exploited the World Wide Web as another platform to build an overlay network upon. Social media and Web 2.0 applications are designed to harness collective intelligence [19]. The human mind being the resource tried to be maximally exploited by stimulating collaboration. The customers are, again advertisers or parties interested in all sorts of demographic

[6] With Network Address Translation the addresses of the internal (home) network are translated by the modem to one single public Internet IP-address. It is thus for other hosts on the Internet impossible to address such a host on the internal (home) network directly.

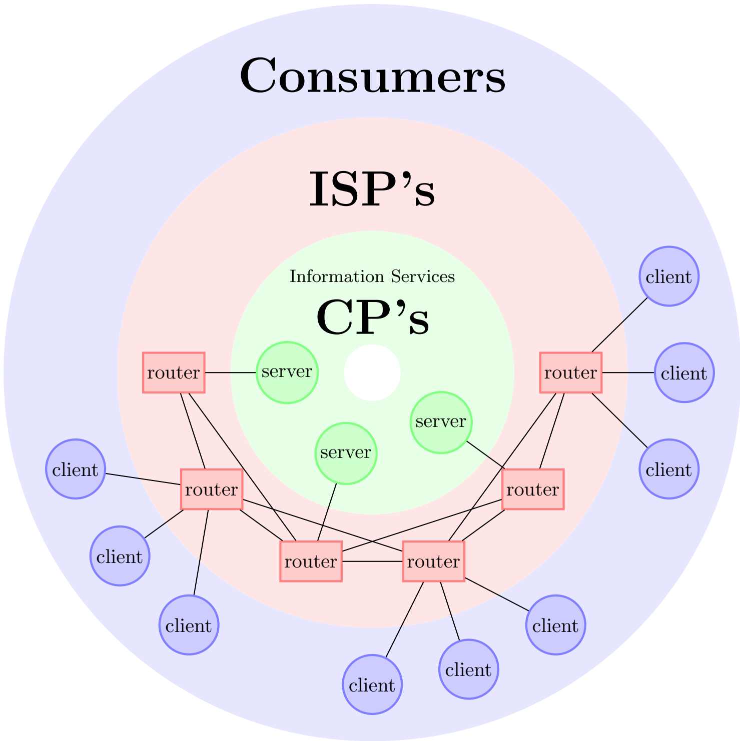

Fig. 5: The consumer-Internet network architecture

data that can be collected by analyzing the by the consumers provided data and their behavior.

Many of the Web 2.0 applications are provided by centrally coordinated businesses, but not necessarily for profit (Wikipedia).

In Augustus 2005 the FCC lessened the restriction of the regulatory framework for broadband and cable providers by classifying them as providers of information services in stead of telecommunications services [20]. This means that IP doesn't have to be provided as a "common carrier" service in the U.S. anymore and opened up the possibilities for price discrimination. This initiated new activity in the *network neutrality* debate [21]. A famous comment from Ed Whitacre, the CEO of AT&T (then SBC) in November that year on the subject of content providers: [22]

> How do you think they're going to get to customers? Through a broadband pipe. Cable companies have them. We have them. Now what they would like to do is use my pipes free, but I ain't going to let them do that because we have spent this capital and we have to have a return on it. So there's going to have to be some mechanism for these people who use these pipes to pay for the portion they're using. Why should they be allowed to use my pipes?

Despite the bold remarks from Whitacre, this has not yet been made effective. It is not completely clear that price discrimination, in the form of providing better access to paying content providers, would benefit the ISP's. More specifically, it is not clear if selling investments in better bandwidth to CP's would be more beneficial than selling it to the consumers [23].

In Europe regulations for communication services are technology neutral. It tries not to make distinctions between the type of technology used. In contrast with the U.S., anti-competitive practices can be addressed in communication markets after-the-fact (ex post). In the U.S. the communication regulation are the solitary before-the-fact (ex ante) measures. No competition law can be applied with communication markets [24].

7 The Grid and Cloud Computing

The market changed the way the Internet looked like for consumers, but not for universities, research institutes, information services and content providers. They still had and have the end-to-end connectivity.

For universities and research institutes there was, and still is, the same old incentive to share their computing resources. Ongoing research in the subject let to an umbrella term for "standard, open, general-purpose protocols and interfaces to coordinate resource that are not subject to centralized control to deliver nontrivial qualities of service": *the Grid* [25]. Working on *the Grid* universities and research institutes have developed Grid middleware that allow universities to provide access to computing resources in a homogeneous way, with ways to locate distributed resources and authorization mechanisms for virtual organizations in which authorization can be re-delegated by the authorized parties [26]. With Grid middleware universities can engage in Grid partnerships or sell access to their computing resource for a fixed-rate price.

Information services and content providers have build their own pools of computing resources to alleviate peaks in demand for certain services or content or from certain locations. Recently those enterprises have started to use technologies originally developed for the Grid, to sell usage of their own pool of resources to customers [27]. The umbrella term for those commercially provided computing resources is: *Cloud computing*.

In contrast to the Grid, Cloud Computing is offered to anyone paying on a pay-per-use basis. Their primary targeted markets are the information services and content providers. Using "the Cloud", providers of content and information services do not need their own server-infrastructure anymore. In stead, they acquire servers from the Cloud provider on demand and release them when not needed anymore. Server usage is offered by Cloud-providers at prices from as low as €0.10 per hour.

Cloud Computing can in that sense be seen as a new underlay network. The role of the server in the client-server way of communication is split between the software offering the services (or content) and the Cloud underlay network providing the computing resources needed to run that software.

Cloud Computing significantly lowers the threshold for new market initiatives in the information services and content provider segment. Newcomers do not need to make big investments or plan on potential usage to try out new ideas for new applications. The Cloud underlay solves all that for them and in doing so stimulates new innovations.

8 Peer-to-peer systems

A convincing demonstration of the innovative power of overlay networks are peer-to-peer systems [28]. Despite all the frustrations on the network layer introduced by the consumer market (primarily by NAT), peer-to-peer systems have managed to create true distributed end-to-end overlay networks, each endpoint acting as a client and a server simultaneously[7].

The incentive for the emergence of peer-to-peer systems is resource sharing. BitTorrent provides an overlay network that effectively eliminates the necessity of having a large amount of bandwidth when providing content to many downloaders [33]. It does this by letting the downloader also become the server for the file. The next downloader will retrieve only parts of the file from the initial provider, and parts from the first downloader. The more popular a file will be, the more participants a downloader will use to fetch the file and the bigger the bandwidth capacity available for downloading the file.

In a similar fashion peer-to-peer can be used for broadcasting live video streams. The more viewers a certain channel will have, the better the bandwidth capacity for the channel. In China, peer-to-peer TV is very popular. QQLive provides more that 100 TV channels.

Peer-to-peer systems can overcome the ossification of IP. Certain communicative properties are high on the wish-list of the Internet Engineering Task Force (IETF), but if the deployment of such a property involves the participation of the middle-boxes (the routers), it is unlikely to happen for legacy reasons.

Multicasting is one of such communicative properties. Multicasting is a publish subscribe modus of communicating. A subscriber joins a multicast group. It will then receive all traffic from that group from the publisher. The publisher has to sent its information only once. The network then makes sure the information is delivered to all subscribers. The first specification of multicasting is from 1989 [34]. Still, the current Internet does not provide reliable multicasting. However gossiping peer-to-peer systems can provide it and a multiplicity of other communicative services [35, 36].

Skype is a Voice Over Ip (VOIP) service providing telephony services on the Internet in a peer-to-peer style. Skype is also a commercial organization. It offers free dialing to other Skype users, but charges for dialing to the regular telephony-network. It also charges for incoming calls from the regular telephony network.

[7] NAT was introduced to alleviate the IP address exhaustion. The impossibility to address hosts on the internal (home) network was not the motivation for NAT. NAT boxes are not designed to provide security. Because of that, NAT boxes often behave in such a way that often "holes" can be "punched" into them, enabling outside hosts (on the internet) to directly connect to internal (home) endpoints.

The Internet Engineering Task Force (IETF) has explored several techniques to "punch holes" in NAT and developed protocols and standards for endpoints to find out how they might achieve a direct connection, or otherwise could use a third party [29, 30, 31]. A detailed exploration of those methods can be found in [32].

9 Conclusion

The Internet has a tradition in the emergence of overlay networks providing innovative new ways of communicating on the underlying network. In this paper I have just scraped the surface of the innovations that have taken place on the Internet until today. It would be impossible to address all ideas that have emerged. In stead, I have tried to classify the different appearances of the overlay networks in a historical context.

1. Initially, the Internet (ARPANET) itself was an overlay network on top of the telephony network.

2. The end-to-end connectivity provided the means for overlays based on client-server architecture.

3. With the commercialization of the Internet, the original end-to-end architectural properties were not delivered to the new consumers. This has lead to the simultaneous emergence of the

   human collaboration oriented class of overlay networks on top of the World Wide Web provided by social media and Web 2.0 application, and

   peer-to-peer networks reestablishing the end-to-end properties and realizing functionalities that are nowadays impossible to establish in the ossified Internet.

Alongside the identification of those classes of overlay networks, I have tried to illustrate how layering in general, and more specifically new underlay networks have created new markets. The IP-network market has without doubt let to cheap Internet access available to all. Cloud Computing already stimulates innovations in the information services and content provider markets. It has the potential to become a generalized computing and storage resource in the middle for all.

But the most empowering force is not offered by those markets. The possibility to construct overlay networks to utilize resource sharing is still, despite frustrations from the market, available to all. Peer-to-peer systems can offer impressive services that do not need huge investments in computing and storage resources, simply because the users of those systems contribute their own resources when using it. Therefor peer-to-peer systems are my bet for the next disruptive innovations. Still, a few issues need to be resolved.

Market driven resource sharing systems have the advantage that the goods cost money. Their usage felt personally. The seller is trusted by the buyers to make sure the price is being paid. Peer-to-peer systems are typically community collaborations. There is no central party that can check that the cost (contributing resources) for the usage is being paid. The evaluation of risk and trust of other peers has to be resolved to be able to effectively handle the "Tragedy of the Commons" problem [37]. Also, the user-friendly killer application, harnessing the collective intelligence, has to be found [38].

As a concluding remark: Not only do new ideas for communicating and resource sharing lead to the construction of overlay networks, but the overlay networks themselves stimulate innovation in the area of communication and resource sharing leading again to new overlay networks. Hence the title of this paper.

References

[1] Dave Clark, Bill Lehr, Steve Bauer, Peyman Faratin, Rahul Sami, and John Wroclawski. Overlay networks and future of the internet. In *In Journal of Communications and Strategies*, pages 1–21, 2006. URL http://people.csail.mit.edu/wlehr/Lehr-Papers_files/Clark%20Lehr%20et%20al%20Overlays%20CandS%20Sept2006.pdf.

[2] Adam Smith. *The Wealth of Nations (Bantam Classics)*. Bantam Classics, March 2003. ISBN 0553585975. URL http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0553585975.

[3] G. Hardin. The tragedy of the commons. *Science*, xx:1243–47, 1968. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.124.3859&rep=rep1&type=pdf.

[4] Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff. A brief history of the internet. *SIGCOMM Comput. Commun. Rev.*, 39(5):22–31, 2009. ISSN 0146-4833. URL http://doi.acm.org/10.1145/1629607.1629613.

[5] Steven King. Computer networks: Heralds of the resource sharing. [FILM], 1972. URL http://www.archive.org/details/ComputerNetworks_TheHeraldsOfResourceSharing.

[6] Vinton G. Cerf and Robert E. Kahn. A protocol for packet network intercommunication. *IEEE Transactions on Communications*, 22:637–648, 1974. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.113.7384&rep=rep1&type=pdf.

[7] H. Zimmermann. Osi reference model–the iso model of architecture for open systems interconnection. *Communications, IEEE Transactions on*, 28(4):425 – 432, apr 1980. ISSN 0090-6778. URL http://www.cs.cmu.edu/~srini/15-744/F02/readings/Zim80.pdf.

[8] Geoff Huston. Interconnection, peering and settlements, 1999. URL http://au.net/papers/1999-6-peer/peering.pdf.

[9] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 1654 (Proposed Standard), July 1994. URL http://www.ietf.org/rfc/rfc1654.txt. Obsoleted by RFC 1771.

[10] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Draft Standard), January 2006. URL http://www.ietf.org/rfc/rfc4271.txt.

[11] C.S. Carr. Network subsystem for time sharing hosts. RFC 15, September 1969. URL http://www.ietf.org/rfc/rfc15.txt.

[12] A.K. Bhushan. File Transfer Protocol. RFC 114, April 1971. URL http://www.ietf.org/rfc/rfc114.txt. Updated by RFCs 133, 141, 171, 172.

[13] A.K. Bhushan. Comments on the File Transfer Protocol. RFC 385, August 1972. URL http://www.ietf.org/rfc/rfc385.txt. Updated by RFC 414.

[14] Michael Hauben. *Netizens: On the History and Impact of UseNet and the Internet*. Peer-to-peer Communi, 1996. ISBN 1573980315. URL http://www.columbia.edu/~hauben/netbook/.

[15] J. Postel. Simple Mail Transfer Protocol. RFC 821 (Standard), August 1982. URL http://www.ietf.org/rfc/rfc821.txt. Obsoleted by RFC 2821.

[16] J.K. Reynolds. Post Office Protocol. RFC 918, October 1984. URL http://www.ietf.org/rfc/rfc918.txt. Obsoleted by RFC 937.

[17] P.V. Mockapetris. Domain names: Concepts and facilities. RFC 882, November 1983. URL http://www.ietf.org/rfc/rfc882.txt. Obsoleted by RFCs 1034, 1035, updated by RFC 973.

[18] Tim Berners Lee. Information management: A proposal, 1989. URL http://cdsweb.cern.ch/record/369245/files/dd-89-001.pdf?version=1.

[19] tim Oreilly. What is Web 2.0: Design Patterns and Business Models for the Next Generation of Software. *Communications & Strategies, No. 1, p. 17, First Quarter 2007*. URL http://ssrn.com/paper=1008839.

[20] Kevin J. Martin. Comments on adoption of wireline broadband internet access order, 2005. URL http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-260433A2.pdf.

[21] Marjory S. Blumenthal and David D. Clark. Rethinking the design of the internet: the end-to-end arguments vs. the brave new world. *ACM Trans. Internet Technol.*, 1(1):70–109, 2001. ISSN 1533-5399. URL http://doi.acm.org/10.1145/383034.383037.

[22] Patricia O'Connell. At sbc, it's all about "scale and scope", 2005. URL http://www.businessweek.com/magazine/content/05_45/b3958092.htm.

[23] Jay Pil Choi and Byung-Cheol Kim. Net neutrality and investment incentives. Working Papers 08-03, NET Institute, September 2008. URL http://ideas.repec.org/p/net/wpaper/0803.html.

[24] J. S. Marcus, Christian Wernick, and Kenneth R. Carter. Network Neutrality: Implications for Europe. *SSRN eLibrary*, 2008. URL http://ssrn.com/paper=1522039.

[25] Ian Foster. What is the grid? - a three point checklist. *GRIDtoday*, 1 (6), July 2002. URL http://www-fp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf.

[26] Buyya R. and Venugopal S. A gentle introduction to grid computing and technologies. *CSI Communications*, 29(1):9–19, 2005. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.67.4243&rep=rep1&type=pdf.

[27] Ian Foster, Yong Zhao, Ioan Raicu, and Shiyong Lu. Cloud computing and grid computing 360-degree compared. 2008. URL http://arxiv.org/abs/0901.0131.

[28] Dejan S. Milojicic, Vana Kalogeraki, Rajan Lukose, Kiran Nagaraja, Jim Pruyne, Bruno Richard, Sami Rollins, and Zhichen Xu. Peer-to-peer computing. Technical report, 2002. URL http://www.hpl.hp.com/techreports/2002/HPL-2002-57R1.pdf.

[29] J. Rosenberg, R. Mahy, P. Matthews, and D. Wing. Session Traversal Utilities for NAT (STUN). RFC 5389 (Proposed Standard), October 2008. URL http://www.ietf.org/rfc/rfc5389.txt.

[30] R. Mahy, P. Matthews, and J. Rosenberg. Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN). RFC 5766 (Proposed Standard), April 2010. URL http://www.ietf.org/rfc/rfc5766.txt.

[31] J. Rosenberg. Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols. RFC 5245 (Proposed Standard), April 2010. URL http://www.ietf.org/rfc/rfc5245.txt.

[32] Willem Toorop. Desktop sharing with the session initiation protocol, 2009. URL http://www.delaat.net/~cees/sne-2008-2009/p29/report.pdf.

[33] Bram Cohen. Incentives build robustness in bittorrent, 2003. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.14.1911&rep=rep1&type=pdf.

[34] S.E. Deering. Host extensions for IP multicasting. RFC 1112 (Standard), August 1989. URL http://www.ietf.org/rfc/rfc1112.txt. Updated by RFC 2236.

[35] Anne marie Kermarrec and Maarten Van Steen. Gossiping in distributed systems, 2007. URL http://www.cs.vu.nl/~steen/papers/2007.osr.pdf.

[36] Arthur van Kleef and Willem Toorop. Gossiping software distribution, 2010. URL https://www.os3.nl/_media/2009-2010/students/willem_toorop/gossiping-software-distribution.pdf.

[37] Michal Feldman and John Chuang. Overcoming free-riding behavior in peer-to-peer systems. *SIGecom Exch.*, 5(4):41–50, 2005. URL http://doi.acm.org/10.1145/1120717.1120723.

[38] Panayotis Antoniadis and Bénédicte Le Grand. Incentives for resource sharing in self-organized communities: From economics to social psychology. In *ICDIM*, pages 756–761, 2007. URL http://dx.doi.org/10.1109/ICDIM.2007.4444315.