

Ralf Bendrath

**Global technology trends and national regulation:
Explaining Variation in the Governance of Deep Packet Inspection**

**Paper prepared for the International Studies Annual Convention
New York City, 15-18 February 2009**

first draft

Updated versions will be available from the author upon request.

author contact information

Delft University of Technology
Faculty of Technology, Policy, and Management
Section ICT
P.O. Box 5015
2600 GA Delft
Netherlands
r.bendrath@tu-delft.nl
<http://bendrath.blogspot.com>

Abstract

Technological advances in routers and network monitoring equipment now allow internet service providers (ISPs) to monitor the content of data flows in real-time and make decisions accordingly about how to handle them. If rolled out widely, this technology known as deep packet inspection (DPI) would turn the internet into something completely new, departing from the “dumb pipe” principle which Lawrence Lessig has so nicely compared to a “daydreaming postal worker” who just moves packets around without caring about their content. The internet’s design, we can see here, is the outcome of political and technological decisions and trends.

The paper examines the deployment of DPI by internet service providers in different countries, as well as their different motives. In a second step, it will offer a first explanation of the varying cases by examining the different factors promoting as well as constraining the use of DPI, and how they play out in different circumstances. The paper uses and combines theoretical approaches from different strands of research: Sociology of technology, especially the concept of disruptive technologies; and interaction-oriented policy research, namely the approach of actor-centric institutionalism.

Table of Content

1. INTRODUCTION: THE INTERNET AND THE ACTORS WHO RUN IT	3
2. INTERNET GOVERNANCE THEORY	5
2.1. IG STUDIES BETWEEN TECHNO-DETERMINISM AND POLICY RESEARCH	5
2.2. THE SOCIAL CONSTRUCTION OF DISRUPTIVE TECHNOLOGY	6
2.3. TECHNOLOGY-AWARE POLICY ANALYSIS	7
3. DEEP PACKET INSPECTION AS A DISRUPTIVE TECHNOLOGY	9
3.1. THE END-TO-END-PRINCIPLE AS THE BASIS OF THE “GOOD OLD“ INTERNET	9
3.2. THE END OF THE END-TO-END-PRINCIPLE?	11
3.3. DPI FUNCTIONS AND TRENDS.....	12
3.4. DPI USE CASES AND DRIVERS.....	14
4. CASE STUDIES: USAGE AND GOVERNANCE OF DPI.....	16
4.1. NETWORK SECURITY	16
4.2. BANDWIDTH MANAGEMENT	17
4.3. AD INJECTION	20
4.4. COPYRIGHT CONTENT FILTERING	22
4.5. GOVERNMENT SURVEILLANCE	24
5. CONCLUSIONS: EXPLAINING VARIATION IN DPI GOVERNANCE	25
LITERATURE	27

1. Introduction: The Internet and the actors who run it¹

Academic and political debates as well as research about Internet governance have so far mostly focused on the topologically extreme parts of the network: On the one hand, they have addressed centralized functions, notably the domain name system and its management organization, ICANN. The struggles over the control of ICANN even lead to conflicts at the highest political levels, including the UN World Summit on the Information Society (WSIS).² On the other hand, much of the remaining scholarly literature as well as political debates on internet governance have addressed the people and the content providers that use the internet for the exchange of data, information, and communication. Here, the issues addressed are legion, ranging from the privacy of search engine usage to the taxation of electronic commerce. “Internet governance” so far has mostly meant governance of either the core or the end-points.

An important part of the internet, though, consists of the wide electronic landscape between the core and the end-points. What make the end-points reach each other, except for the availability of a functioning addressing system, are the cables, routers, and actors in between – the internet service providers (ISPs) and their hardware, as well as the protocols used by them and the governance arrangements among them. One could even say that the internet does not consist of the endpoints at all, and only to a very limited extend of the core naming infrastructure, but of the network of networks connected to each other by the TCP/IP protocol suite. There has been surprisingly little debate as well as research about the changes and governance processes going on here. Only in the last few years has this drawn more attention under the label “network neutrality”.

Most of the network neutrality debate, though, has focused normatively on the political and economic responsibilities of the ISPs towards the end-users or the content providers. An empirical analysis of the underlying technology and its changes and trends has been largely missing³. Ten years ago, Joel Reidenberg (1998) and Lawrence Lessig (1999) already made convincingly clear that the technological properties of the internet have as much a role in controlling the users as have laws, norms, and the market. The way the routers and access points are designed, and the way the network protocols handle the data packets, all influence the way internet traffic, and therefore user behavior, can be controlled. The argument that we should care about the political implications of technical structures as much as of social institutions has become general knowledge, but has inspired surprisingly little research in the internet governance community.⁴ Jonathan Zittrain’s recent and widely perceived book “The Future of the Internet and how to stop it” (2008) addresses the technological changes going on with the internet, and therefore is a valuable contribution in this tradition. But he again focuses on one extreme – the endpoints end devices attached to the network for various

¹ The author has spoken with a number of ISP, regulators, and content industry staff in the course of the project. Not all of them were willing to be quoted by name. Therefore, some of the data in this paper is not referenced.

² See <http://www.wsis.org> and the reports from the preparatory process at <http://www.worlsummit2005.org>.

³ A recent exception is Ohm (2008).

⁴ The political nature of technologies, and the fact that their design can influence and control social behavior, had already been widely discussed in the science, technology and society community (Bijker and Law 1992) (Winner 1986) which has shown that structures like road bumps, bridges, and even the design of locks for backyard doors (Latour 1994) can determine how people move and behave. But this has rarely been applied by researchers interested in internet governance. I myself have used this for a study on internet privacy governance (Bendrath 2008), but, like Zittrain, have only looked at the endpoints.

usages.⁵ The book does not really address the internet in itself – the network of networks connected by the TCP/IP protocol suite.⁶

This paper is about the internet in between the endpoints. It is about the possible *future* as well as the possible *end* of the internet as we have known it so far. I thereby also pick up a thread started by another author some years ago. John Walker, though not an academic researcher like Zittrain, in 2003 published a lengthy text under the title “The Digital Imprimatur” (Walker 2003), which was later published in German under the title “Ende des Internet” (Walker 2004). Walker listed the various attempts to put the internet as well as the devices attached to it under closer control through monitoring and filtering technologies, digital restriction management, and less open, “trusted computing”-restricted hardware for the end-users. More recently, Jeff Chester from the Center for Digital Democracy used “The end of the Internet?” as the title of an opinion piece in *The Nation* in which he explicitly warned that new network monitoring technologies may threaten free speech and the overall openness of the internet (Chester 2006).

A central element of this new technology is known as “Deep Packet Inspection” (DPI). DPI introduces “intelligence” into the routers and switches, which allows for discrimination of traffic and implies the end of the end-to-end principle. What will be its impact on the internet as we know it? Will it become a central part of the internet infrastructure, controlling most traffic flows, and therefore changing the network and the way we can use it dramatically? Or will it be domesticated and only used for limited tasks, according to the principles and norms of the good old open internet we all learned to love?

These questions also imply assumptions about the relationship between technology, society, and policy. We can conceive of DPI as a disruptive technology like the printing press, the internal combustion engine, or packet-switched networking. In this case, we would expect that it will have a structural, long-term impact on the internet and the social interactions based on it, no matter how much user, engineers or regulators try to fight this trend. Or we can see it as a socially constructed technology, being used only according to social norms and further engineered to oblige and even enforce these norms – one of them being that the internet should stay as neutral as possible. In this case, we would expect that the governance norms and structures that are developed around DPI will tame its impact, but also vary across different contexts and societies. The approach taken here is a combination of these two. It will use the “disruptive technology” concept together with elements from the “social construction of technology” theories. In order to study and explain the variance across cases, I will also add to these an interaction-oriented policy-analysis approach.

⁵ According to Zittrain, these endpoints become less and less open and “generative” and increasingly turn into special-purpose machines under the control of the manufacturers and vendors, thereby narrowing down the openness that was the basis for the very success of the internet. For a critique, see (Owen 2008).

⁶ Zittrain has addressed the issue of ISP filtering and internet control earlier (Zittrain 2004), but not as systematically, and also not widely perceived.

2. Internet Governance Theory

2.1. *IG studies between techno-determinism and policy research*

Internet governance research has so far mainly oscillated between techno-determinism and a fascinating neglect of technology. Many of the early publications on Internet governance have circled around the question: Can the internet be governed at all? This implicitly techno-deterministic view was less founded in well thought-through research questions, but reflected a dominant discourse in the 1990s. The decentralized architecture of the internet back then was widely seen as a paradigmatic break with the prevailing organizational principles of modern infrastructures. Unlike telephone networks, which were designed by hierarchical forms of coordination in and between nation states, the internet seemed to be immune to any form of central steering.⁷ Because the internet crosses and to some extent ignores national borders, it undermines territorial forms of control. However, national sovereignty and the authority of law are based on territories and intact borders between them (Drake 1993). In the view of many observers, this could only mean that cyberspace had to develop its own form of post-nation state control:

Global computer-based communications cut across territorial borders, creating a new realm of human activity and undermining the feasibility – and legitimacy – of applying laws based on geographic boundaries. (Johnson and Post 1997, 3)

Based on this notably techno-deterministic perspective, a large part of the emerging internet research community discussed various scenarios of “nongovernmental governance” for the net that ranged from Barlow’s famous “Declaration of the Independence of Cyberspace” (Barlow 1996) to visions of the “virtual state” that only plays a networking role and is not primarily based on its territory anymore (Rosecrance 1996).

This view of course was naïve, and it was quickly neglected by empirical studies, which showed that even the governance of the internet’s core resources is subject to heavy political struggles (Mueller 2002). The next generation of governance-related internet research therefore shifted towards the different forms and politics of regulating and controlling the internet. Here, we find the usual suspects of political science traditions well represented. While some still see the nation-state as the main actor (Goldsmith and Wu 2006), others insist on the relevance of an internationalist, regime-theory-informed view (Mathiason 2006). Some have found that internet governance forms depend on the specific policy field (Bendrath et al. 2007), others see it as the outcome of a realist power struggle between great powers (Drezner 2004). But more or less everybody concerned with internet governance now agrees that the internet or certain aspects of it can be governed at all (for an overview, see (Mayer-Schönberger 2003). But with this came a surprising neglect of the technological properties and effects of the internet that were so predominant in the debates of the 1990s.

The “internet” as an object of policy and governance is also not carved in stone. Changes in its underlying structure and functioning might weaken or strengthen the capacity of governments, private, or international actors to govern it. Internet governance research has done quite some work on how internet governance structures and processes affect the technological design of the internet and e.g. lead to new privacy standards (Cranor 2002) or security functions in the domain name system (Kuerbis 2009). But it has not really addressed

⁷ It is true that the core resources of the internet, like the root server for the domain-name system or the allocation of IP address ranges, are organized centrally. But normal usage is not affected by this, as routing and packet switching take place in a decentralized way.

how technological changes that emerged elsewhere (e.g. innovations in transmission speeds, in packet filters, or in video compression methods) have been reacted to and affected by governments and other actors. What is needed, therefore, is a perspective that integrates the techno-deterministic perspective and the governance-of-technology perspective, taking seriously the effects of technological changes on one hand as well as the governance structures and impacts on technology on the other hand.

2.2. *The Social Construction of Disruptive Technology*

A similar shift in debates and perspectives has been taking place in the sociology of technology (for an overview, see (Dolata and Werle 2007)). The first generation of scholars mostly focused on technology assessments and therefore – implicitly or explicitly – had a technological-deterministic approach.⁸ Partly as a reaction to this, the second generation to the contrary has largely been interested in the development and social construction of technology, implying a socio-deterministic perspective. More recently, there have been approaches to find a middle ground. Technological changes are not seen as *determining* social interactions, but still have effects on society that are based on its material properties (Grunwald 2007). One attempt to combine these perspectives uses the concept of “disruptive technologies”. New technologies can “disrupt” social structures and practices:

Life is organized around technology. Despite our desire to maintain control over our lives, the bulk of what we do with our lives has been coordinated and adapted by and to the technology that surrounds us. Therefore it should come as no surprise that when existing technology evolves or old technology is made obsolete that the phase where new technology enters our lives could be seen as being disruptive. The disruption occurs when the technology, which is introduced, effects the social arrangements around which we build our lives. (Klang 2006, 7f)⁹

These disruptions, though, are not determining the effects new technologies have on society. This is where the social construction of technology, as well as its governance, becomes relevant. Similar arguments have been developed by scholars who examined the co-evolution of technological and social systems (Weingart 1989). They note that in order for a stable socio-technical system to develop, there must be a “fit” between the technological structures and machines on the one hand and the social and institutional structures they are embedded in (Dierkes et al. 1992). Otherwise, we can expect adaptive conflicts. The outcome of these conflicts either leads to the change and integration of technology into existing social structures, or to the change of social structures according to the technology. Neither way is pre-determined, but they are constrained by the characteristics of both technology and society.

⁸ This argument, of course, is not new. In 1620, Francis Bacon wrote in the “Novum Organum”: “printing, gunpowder, and the magnet (...) these three have changed the whole face and state of things throughout the world; the first in literature, the second in warfare, the third in navigation; whence have followed innumerable changes, insomuch that no empire, no sect, no star seems to have exerted greater power and influence in human affairs than these mechanical discoveries.” (Bacon 1858)

⁹ Matthias Klang notes that the concept has recently been narrowed down to industrial innovation and the survival of old and new products on the market: “In recent work the concept of disruption is being used to explain organizational change and innovation. Undoubtedly the most popular use of the term disruptive technology has been presented by Christensen in his book *The Innovator’s Dilemma*. Christensen defines a disruptive technology as a new technological innovation (product or service) that will eventually overturn the dominant technology in the market sector. (...) While this description has been effective in bringing the concept of disruptive technologies into the more general debate of the role of technology it has also changed our concept of technology, since it limits our general view of disruptive technologies to being one of a less economically viable technology.” (Klang 2006). I will follow Klang and use the more general meaning of “disruptive technology” in this paper, which includes all technologies that have more than just gradual responses as effects.

A lot of this research has focused on these adaptive conflicts that often occur around new technologies. Scholars have analyzed the different perceptions, images, frames and “leitbilder” of technology and how they are contested. Conflicting frames make the introduction of new technology difficult and contested (Orlikowski and Gash 1994). Or, in other words, different “leitbild” visions lead to conflicts about the adoption and use of technology (Dierkes et al. 1992). The exact paths of socio-technical adaptation and integration processes depend on the specific models and understandings of technology and organization, as well as on the outcomes of struggles and conflicts between actors involved in the creation as well as a usage and the governance of technology (Dierkes and Marz 1990). The specific conflicts in the end determine if technology has an impact on society, and how big it is; or if society and policy are able to control and govern new technologies. Examples that come to mind here of course are the variation in the use of nuclear energy, but also e.g. the different censorship regimes for the internet that have been established in different countries (Deibert et al. 2008).

2.3. Technology-Aware Policy Analysis

The analysis of public conflicts and struggles, of course, is the domain of political science. The challenge for technology-aware policy research is to combine the potentially disruptive powers of specific technologies with an analysis of concrete conflicts and interactions around its usage and governance. In principle, this is nothing new. Policy analysis has always dealt with the specifics of a policy field, be it social welfare, arms control or energy production. But political scientists do not seem to be very good at taking the specific technological properties of a techno-centric policy field (like internet governance) into account. Even many internet governance researchers tend to focus on the purely institutional, social or normative aspects of IG, typically fascinated with new governance models, global norms, or network analysis.¹⁰ Technology-aware policy analysis must combine the two sides of the coin: It must address how technological innovations impact the actors and their interactions; and it must address, in turn, how social norms and institutions have an effect on the governance of new technologies.

In technology-related policy fields, one factor therefore is the *technical properties that determine the interests as well as the governance capacities* of the actors involved. (Knill and Lehmkuhl 2002) The governance capacities can vary over how close the actors are to the core of the problem at hand. It makes a difference if they run technological systems, if they are end-users and customers, or if they are state actors that can only intervene from the outside. They also can vary over the scope of the problem. It makes a difference if the technological trends to be governed are happening within the reach of actors (locally or nationally), or if they are global and therefore more likely to be beyond the reach of national regulators. Lastly, the governance capacity of actors can vary across the type of conflict or policy problem at hand. If the changes are based on global technology shifts or transnational market pressures, it is harder to deal with them nationally or even locally.

When analyzing technology-related political conflicts and the eventual emergence of new governance structures around these, it is therefore important to connect the specific properties of new technology with the governance capacity of the actors involved. This can lead to different strategic constellations and therefore different types of interactions, even for different use-cases of a single new technology. Coordination games are easier to solve without hierarchical decision-making than games with defection or free-riding effects, or even re-distributive games (Scharpf 1997). Constitutive, regulative, and re-distributive policies will

¹⁰ The programs of the last symposiums of the Global Internet Governance Academic Network (Giganet) give a good illustration, see <http://giganet.igloogroups.org/>.

lead to different types of interaction. Lowi's famous hypothesis - "policies determine politics" (Lowi 1972) - can therefore here be translated into "technologies determine politics".

But this is, again, only one side of the story. A second factor is the *social and institutional context in which technology is embedded*. The political interactions around new technologies are as well shaped by institutional settings and legal restrictions, by the existence or non-existence of partisan groups and their degree of organization, and by the properties and legacies of the political system. These have an impact on the way the conflicts about technology are played out, how interactions are restrained or actors empowered. It also has an impact on the mode of interaction, which can range from highly conflictive to very cooperative. (Scharpf 1997) While this is normal political science, the technology aspect here is the social construction of the technology in these political interactions – technology's perceptions, the norms and governance structures that are developed around it, and the usage patterns that emerge. All of these factors play a role in relation to the social and institutional context. If e.g. the technology perceptions of specific actors can be linked to institutionally established norms, we can expect that they have a stronger position in the interactions that take place within these institutions. In this way, the institutions and social norms and perceptions will have an impact on how the games and conflicts around technologies are played out, and will therefore in the end also influence the technology itself. Looking at this perspective, we could say "polity determines technology".

It is this duality of the relation between technology and society that has to be dealt with. Sociologists of technology have used the "duality" concept developed by Giddens before, but only to analyze new technologies in organizations (Orlikowski 1992). On a larger, societal level, institutions, collective actors, laws and political structures of course play a much larger role. In the times of the internet, we can also see spontaneous attention and specialized public spheres emerge around contested issues. We can of course see this even more where the internet itself is affected and contested. This means that in no way is technology determining the *outcome* of politics, and equally, social and political factors do not determine the exact shape, usage and governance of technology. There is still considerable freedom for the actors involved to change their minds, to make mistakes, or to be very clever negotiators or campaigners. But their actions take place within the constraints of the technological and institutional possibilities. Such a theoretical framework as described here, which takes policy specifics, actors, institutions, and interactions seriously, is known as "actor-centric institutionalism" (Mayntz and Scharpf 1995). It helps guide the empirical analysis, while still being non-deterministic.¹¹

How do we apply this to new internet technologies? A first step is to analyze the technology itself, its trajectories and the way it is spreading, as well as the different use-cases it enables. Based on these use-cases, we can pinpoint the different actors and their interests in using the new technology. In the case of deep packet inspection, as shown in section 3, these are very diverse and include law enforcement agencies, content providers, politicians and groups that try to block illegal or harmful content, and the internet access providers themselves. The third step, provided in section 4, is to put the interests of the different actors into perspective by situating them in the institutional and social context. This includes the different kinds of

¹¹ Fritz Scharpf, who has developed this framework together with Renate Mayntz, puts a lot of effort into modeling the interactions in game-theoretical terms. While this is helpful in many cases, it can not be applied in cases where the actors are too many or where their interests and payoffs are not known yet. Especially the latter is what we often find in technology policy studies, where the researcher can not rest on an established and more or less stable basis of knowledge about the relevant actors, their interests, the institutions, and the modes of interactions, but where these instead are in flux because of technological innovations.

internet legislation, the organizational strength of industry and other associations, the existence of antagonistic groups, but also relevant societal norms and practices. These help understand the real interactions (not to say “games”) that in the end determine the outcome: The use and governance of new technology, in this case of DPI equipment.

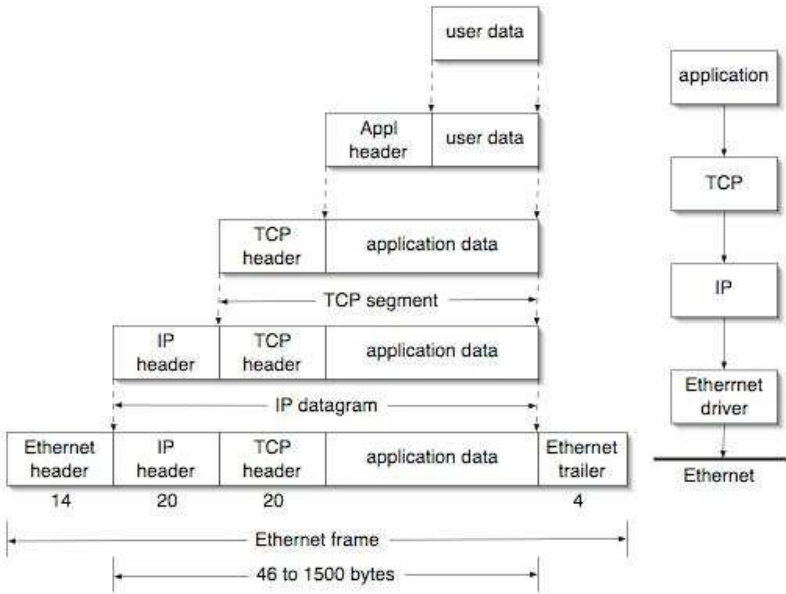
3. Deep Packet Inspection as a Disruptive Technology

In this section, I will briefly describe some central characteristics of the old internet (the one we’ve know so far), and then turn over to the changes enabled, if not caused, by Deep Packet Inspection. This chapter is thereby describing the disruptive potential of DPI technology.

3.1. The End-to-End-Principle as the Basis of the “good old“ Internet

A central design principle of the internet for the last 25 years has been the “end to end” argument (Saltzer et al. 1984). This basically meant that the data link, the network and the transport layers of the protocol stack should only care about sending data packets from the sender to the receiver, without worrying about their content, their security, or even the fact that they reached their destination. These issues should rather be addressed higher up the protocol stack, at the session or application layers. In a network designed according to this principle, the users at the endpoints then can use the data in manifold ways, according to the different applications they have running. At the same time, the devices and applications at the endpoints are also responsible for requesting a re-send in case of packet loss, for securing the communication, or for filtering out viruses. The network itself is dumb, and the intelligence is at the edges.

More technically speaking, application data (emails, music files, voice streams etc.) is encapsulated into TCP packets, which are in turn encapsulated into IP packets, which are in turn encapsulated into Ethernet frames, which are then sent over the wires or radio waves - like Russian dolls or envelopes within envelopes. Because of this encapsulation, the lower protocol layers do not have to care about what is in the packets or frames they transport.¹² Because of encapsulation, each layer only looks at the address and other information (header) that is relevant for him, without having to worry about the payload. Encapsulated within this payload are the headers and the payloads for the higher layers.



Encapsulation as it goes down the protocol stack. Source: Stevens (1993)

¹² For a common description of the networking protocol stack, see http://en.wikipedia.org/wiki/OSI_model.

Because of performance and other trade-offs, some applications like voice transmission (Voice over IP) do not need completely reliable transmission of all data packets. Instead of delays caused by frequent re-transmission requests, the users prefer higher transmission speeds and a constant flow of data packets, even if some of them get lost along the way.¹³ In the layered protocol stack, it is up to the applications to decide which way they deal with these issues. Simplicity at the lower protocol levels, which in our context also translates into simplicity at the network and transport level, therefore means *more* degrees of freedom for the application level. This is true even if the added functions in the network are considered improvements by most applications and users:

“[T]here may be some application that finds the enhancement is not worth the result, but it now has no choice in that matter.” (Saltzer et al. 1984)

Lawrence Lessig has used a very nice illustration for the way the end-to-end principle can be translated into the real world:

“Like a daydreaming postal worker, the network simply moves the data and leaves interpretation of the data to the applications at either end. This minimalism in design is intentional. It reflects both a political decision about disabling control and a technological decision about the optimal network design.” (Lessig 1999)¹⁴

In policy discourse, three arguments have been made to support the end-to-end principle. The first two are already provided by Lessig’s original argument. One is *political freedom* (disabling control), the other one is *technical flexibility* (optimal network design). In fact, the original inventors of the end-to-end argument did not have much of politics or freedom of the user in mind, but were concerned about lean, efficient, scalable and open network architecture that would allow any kind of application to run on top of it. Moral arguments like user freedom did not matter to them very much, because at that time, the internet was still largely an academic endeavor. But since then, the political motivation to have a “dumb” network has also gained a lot of support. A network that does not care about the content of the packets it moves is a network that can not easily be used for censorship or communications surveillance. Maximum freedom for the endpoints means maximum freedom for the users.

The third argument for such a simple and “dumb” network is one of *economic openness*. It was provided by Lessig only in one of his later books (Lessig 2002), because originally he was more concerned about the political control issues of the internet. It was also not even thought of by the inventors of the end-to-end-argument in 1984, because at that time, there were no commercial applications and providers on the internet.¹⁵ The economic argument is similar to the “disabling control” argument, but points to different motivations and potential drivers for control. A network that does not care about the content of the packets it transports is also one that enables innovation and competition at the transport *and* application layers, as long as they still provide a working interface to the TCP/IP layers which constitute the

¹³ The original authors of the end-to-end argument also point out the user-level error correction mechanisms that are available above the machinery level: “Excuse me, someone dropped a glass. Would you please say that again?” (Saltzer et al. 1984)

¹⁴ Paul Ohm (2008) has criticized the “envelope” metaphor, instead using the metaphor of a policeman watching road traffic pass by. While his analogy points at the important fact that the traffic speed (bandwidth) and the monitoring speed (DPI capabilities) are two distinct variables, it at the same time ignores the fact that on the internet, both traffic speed and monitoring speed are not fully independent, but controlled by one actor - the ISP.

¹⁵ The original end-to-end-argument refers to the way this principle more directly aligns costs with the applications and users causing the costs, but this was at that time and under its technical conditions (time-sharing in large computers was the paradigm back then) still seen as an engineering principle rather than an explicitly economic issue.

internet. In this model, no one needs permission from a gatekeeper to enter a service or content market. Unlike in specialized communication systems like the publicly switched telephone network or the German *Bildschirmtext*, anyone can come up with new totally products that run on top of the internet, or with innovative transmission methods that function underneath it, thereby contributing to more choice for the consumer as well as (in successful cases at least) to economic growth.

To summarize, the internet has so far been a loose network of interconnected data networks that share few central characteristics (see also Carpenter 1996):

1. *Technical Simplicity*: Because of the layered approach, they are only connected through the TCP/IP protocol suite and a shared address space. Therefore, they are highly open to new transportation methods (WiMax, UMTS etc.) as well as new applications (e.g. Twitter, Bittorrent, or XMPP/Jabber).
2. *Political Freedom*: Because the higher-layer payloads are encapsulated for the lower layers, the users have end-to-end communication channels at the application layer, which are normally not interfered with in transport.¹⁶
3. *Economic Openness*: Because of the openness for new applications, they do not discriminate traffic according to its source, therefore treating all innovations at the application layer equally and giving them a fair chance to succeed at the market.

3.2. *The End of the End-to-End-Principle?*

Deep Packet Inspection is a technology that changes this architecture of the internet by moving more intelligence into the network. Imagine a postal worker who is not just daydreaming and moving packets from one point to another in the transportation chain. Imagine the postal worker

- opens up *all* packets and letters,
- inspects and even reads the content,
- checks it against databases of illegal material and if finding a match, sends a copy to the police authorities,
- destroys letters he finds having prohibited or immoral content,
- sends packets with content from those mail-order companies which pay extra to the postal service to a special and very fast delivery truck, while the ones from the competitors go to an extra-slow and cheap sub-contractor.

Such a postal system would infringe on the values embodied by the internet as described above:

1. *Political Freedom*: The postal system would now invade the privacy of communications and introduce censorship, potentially leading to “lost” letters from trade unions or political dissidents.
2. *Technical Simplicity*: Such an inspection system would create an additional overhead that would slow down postal delivery and place a significant responsibility on the postal worker. The letters and packets would also be damaged when being opened. And, most importantly, the postal service would assume functions it never was founded for.
3. *Economic Openness*: The differential treatment of content from different senders and companies basically means blackmailing content companies like mail-order stores into

¹⁶ Even in the old network, man-in-the-middle attacks were possible, of course. But they were not by default built into the network design and the hardware and software implementing it.

signing additional and costly high-speed contracts. New business models that solely rely on innovative content being delivered through the normal postal system would have to get permission from the postal system to offer their products.

Now, imagine a postal worker could all do this without significant delays and damages to the packets and letters compared to his (former, now fired) daydreaming colleague. This is what deep packet inspection technology is designed for.

Deep Packet Inspection has the potential to change the basic nature of the internet. It provides a sharp and thoroughgoing departure from the end to end principle and introduces means for more centralized political and economical controls. It collapses the application, internet, and transport layers, and thereby gives ISPs a new role as potential gatekeepers of all their users' traffic. DPI also integrates the control functionalities of previous technologies for network surveillance, content filtering, bandwidth discrimination, or the tying of internet access to specific content subscription models. A DPI-controlled internet would not be the one we all have got used to, the one that gives maximum freedom for users and citizens to exchange information, communicate with each other, and invent new services and business models. Therefore, DPI has the potential to change the future of the Internet if introduced and used on a large scale.

But a potential does not necessarily, and rarely fully, translate into reality. In fact, it is inconceivable that someone “flips a switch” and we all end up with a new internet (Lee 2008a). Because DPI is a generic technology, the specific *real cases* of deep packet inspection equipment and usage do not have to implement *all* the above functions of the highly awake postal worker, though. It can just be used for filtering out content nobody wants to have (viruses and other malware), for managing bandwidth distribution according to the priority needs of each application and connection, for only triggering a police response when known illegal material is distributed, or for only discriminating against special content providers (like video on demand services) that are not part of the ISP's mother firm.

3.3. DPI Functions and Trends

DPI is a functionality embedded in the network that enables the network owner to analyze internet traffic (datagrams on the way through the network) in real-time and discriminate them according to their payload. Because of the high speeds of internet access even at the last mile that are widely available nowadays, DPI can not be done by software running on normal processors or switches. It has only become possible in the last few years through advances in computer engineering and in pattern matching algorithms.

Deep packet inspection is done using specialized hardware. In most cases, this is based on “application-specific integrated circuits” (ASICs), also called “system on a chip”. The problem with ASICs is the high costs and relatively long production times, because you basically have to re-arrange a whole semiconductor factory for each specific chip. In order to lower costs and save time, “field-programmable gate arrays (FPGAs) have also been used. These are integrated sets of large numbers of simple logic processors that can be freely configured and are also known under the general category of “configurable computing”. FPGAs are slower than the specialized ASCII circuits, cannot handle as complex a design, and

draw more energy. But their advantages include a shorter time to program, the ability to re-program them, and lower non-recurring engineering costs.¹⁷

Progress in pattern matching algorithms has also enabled DPI machines to match the payload of internet traffic against increasing sets of search patterns and discrimination rules without similar increases in the time needed for this. FPGA-based DPI based on the “Decoded partial CAM” (DpCAM) method can achieve a throughput between 2 and 8 Gbps with an ability to match about 2,200 patterns using a single machine. A more recent approach uses “Perfect Hashing memory” (PHmem) methods to determine a match. PHmem designs on the same devices can support 2 to 5.7 Gbps throughput. By this, it already reaches the speeds of comparable ASIC designs. (Sourdis et al. 2008)

Predecessors to DPI technology were shallow packet inspection (only the headers were inspected and the payload ignored) and packet pre-filtering (only a few packets are filtered out for thorough inspection, while the other, non-suspicious ones, are let through unchecked). These technologies can also be combined. Pre-filtering can e.g. be used to significantly reduce the workload of the actual inspection processors. Sourdis (2007) describes a method to avoid the DPI processing of over 98% of the packets, resulting in a throughput of 2.5 to 10 GBps in a single FPGA device. This approach is mainly aiming at filtering out mal-ware. Pre-filtering has its limits, though. It can be useful if there is a list of known terms and regular expressions to look for, like virus and mal-ware signatures, but is less useful when filtering is supposed to be based on behavioral information.

The early producers of DPI equipment were not the large router and network equipment companies like Cisco, Juniper or Nortel, but small start-up companies - though some of these have been spin-offs from larger companies. Today, there are around 30 DPI manufacturers. Most are based in the United States, but there are also advanced DPI manufacturers in Israel, Germany and elsewhere. DPI technology has been on the market since around 2002, and has increased in speed significantly since then. Currently, the fastest DPI equipment available can sustain a throughput of 80 Gigabit per second (Anderson 2008). This enables internet access providers to monitor and discriminate the traffic of 20,000 broadband subscribers with a bandwidth of 4 Megabits per second each – even if they all are online at the same time and fully use up their bandwidth.

DPI can be used for a variety of functions:

- *Network security:* Most DPI equipment is aimed at network operators who want to filter malware and other dangerous traffic before it reaches their customers or employees.
- *Network management:* DPI is also used for dealing with scarce bandwidth. It enables ISPs to throttle or even completely block unwanted traffic such as bandwidth-consuming peer-to-peer file-sharing. It can also be used for routing optimization based on the type of data transferred.(Bennett 2008)
- *Surveillance:* In the U.S., some large internet and phone companies were reported to have installed DPI equipment for real-time government monitoring and interception purposes in cooperation with the National Security Agency (Singel 2006). Other countries are also monitoring internet traffic, but it is not always clear if DPI technology is used here.
- *Content regulation:* A number of governments and attorneys have started initiatives that (would) mandate ISPs to block out any content that is seen illegal or “harmful”. These approaches range from filtering child-abuse websites (Dedman and Sullivan 2008; Kleinz

¹⁷ More recently, FPGAs have been enhanced with embedded microprocessors and other electronics, so you have a complete “system on a programmable chip”. Examples of such hybrid technologies are the Xilinx Virtex-II PRO and Virtex-4 devices.

2009) to censoring anything that is considered a threat to the government and public stability (Wagner 2008).

- *Copyright enforcement:* A number of big players in the content industry have also been pushing for mandatory filtering of copyrighted material that is shared on peer-to-peer platforms. They have tried to force ISPs to use filtering equipment that would automatically detect and block copyrighted music or video files, either through court cases (McIntyre 2008a) or through heavy lobbying (Krempf 2009). While it is still unclear if this is a technically functioning approach (McIntyre 2008b), this has certainly opened a new debate about ISPs' secondary liability and obligations related to this, which is far from an end yet.
- *Ad injection:* Another strand of DPI manufacturers has emerged more recently. Companies like NebuAd and Phorm offer advertisement packages for commercial ISPs that inject ads into websites that match the assumed interests of the users. The wishes of these potential consumers are inferred by a detailed analysis of their internet traffic (so-called behavioural targeting) (White 2007; Clayton 2008).

Deep Packet Inspection technology is also developing from just being an add-on to becoming embedded in the core internet infrastructure. Currently, an ISP who wants to do deep packet inspection has to buy a specialized box and insert it into its network, typically somewhere in the downstream between his backbone access and the last mile. In the mid-term, industry observers expect this functionality to migrate into the switches and routers themselves. Catch-phrases for this are „intelligent switching“, „layer7-switching“, or „application-based switching“ (Shafer 2002).

3.4. DPI Use Cases and Drivers

Many of the functions provided by DPI technology have been available before. The interesting and potentially paradigm-changing characteristic of DPI technology is the fact that it integrates these diverse functions into one piece of equipment. It also integrates the interests of a diverse set of actors:

- government agencies and content providers, who are interested in the monitoring and filtering of information flows (*political control*)
- network operating staff, who have to deal with more malware and bandwidth-hungry applications than ever before and who often have limitations for expanding bandwidth on the last mile (*technological efficiency*),
- ISPs that want to create additional revenues, or protect existing profit margins, achieved through vertical integration; e.g., by preventing independent services or applications from cannibalizing their voice telephony or video on demand revenues (*economic interests*).

These interests of course can overlap. For example, content providers may be interested in filtering out illegal content such as non-licensed distribution of their intellectual property. This is an economic interest as well as a political one. Network operators may want to filter malware and slow down “bandwidth hogs”, but they also may be interested in offering a “clean” internet for family use that saves the users from accidentally being confronted with disturbing content like pornography or hate speech. This is a technical, an economic and a political issue. And government agencies may be interested in censoring terrorist propaganda or instructions for building explosive devices, but they may also be interested in enforcing copyright law. The following tables illustrate the different drivers towards the interception, management, and discrimination of internet traffic, as well as the old technologies that can now be replaced by an integrated DPI system.

Table 1: use cases and drivers for DPI

Purpose	Old	New	Drivers
lawful interception, surveillance	TCPdump, Wireshark, dsniff etc. (store & analyze)	DPI (analyze packets and make decisions in real-time)	police, intelligence community
content regulation	blocking based on DNS, IP#, URL	hash-based blocking or surveillance	efforts against hate-speech, child-porn, political censorship
copyright enforcement	DRM Lawsuits	hash-based filtering	content industry
bandwidth management	TCP congestion management, QoS	application-based routing	ISPs: last mile over-subscription, P2P and video traffic
subscriber management	pay per minute, pay per volume	differentiated services and pricing	ISPs: heterogenous user behaviour and user needs in context of bandwidth scarcity
network security	stateful firewalls, asynchronous monitoring (TCPDump etc.)	content-based real-time monitoring	corporate network operators; anti-spam and malware efforts by ISPs
vertical integration	product tying	block or discriminate competing services	video on demand, integrated phone & internet providers, triple play.
behavioural-based advertising	cookies (website owners)	ad injection	ISPs, ad networks

Thus, we can see that DPI has a wide-ranging potential to change the nature of the internet by introducing means for political control and restricting its economic openness. But a potential does not necessarily translate into reality, and rarely does any technology's full potential become completely realized in one sector. Actual implementations of DPI do not have to implement *all* the above functions of the highly awake postal worker. Its use could be limited to filtering out content nobody wants anyway (viruses and other malware); restricted to managing bandwidth according to the priority needs of different applications; used only to trigger a police response when known illegal material is distributed; or used only for discriminating against content, applications or services that compete with the ISP's own offerings.

4. Case studies: Usage and Governance of DPI

Now that we have seen how DPI potentially might change the way the internet is run and operated, and identified the drivers for it, let us turn to the actual use cases and struggles around them. The empirical data in this section will show that DPI is not used as a general control technology, but in different ways by different ISPs in different countries. This can at least partially be explained by the specific interests around the use-cases, as well as by the governance structures that have been applied to or developed around DPI.

4.1. Network Security

The original use case DPI was developed for is network security. Traditional firewalls track which applications inside a local area network (LAN) have established which connections with which hosts on the internet. Thereby, they can control that there is no un-requested traffic coming from the outside, and they also can block ports that are not used by standard applications such as port 80 for http or port 25 for SMTP. There have been two challenges to this approach: First, because port 80 is almost always open, a number of applications have used this for their traffic, too. The internet telephony client from Skype for instance is famous for getting through almost every firewall, but BitTorrent clients and other applications are also using it. The other trend is the general move towards web services and cloud computing, which means there is no easy border between the inside of the LAN and the outside internet. This pushes network administrators towards completely inspecting and parsing the data that flows in and out, in order to see if it is what it claims to be. (Shafer 2002)

Many DPI vendors offer combined solutions now that provide intrusion detection, intrusion prevention and firewall capabilities together with a full inspection of the traffic. Their machines scan for known patterns of viruses, worms, and other malware, block their entering of the LAN, and compile reports and statistics. Some even offer additional privacy features for web services: The Sentry 1500 appliance from Forum Systems, for instance, “is designed to provide security in the world of Web services. Through packet inspection (...), the Sentry 1500 is able to add selective encryption to XML data in transit, thus ensuring the data is secure and cannot be read while it passes through the server.”(Shafer 2002) Network security is also the main field of university-based DPI research. (Artan 2007; Dharmapurikar et al. 2004) (Sourdis 2007). A recent achievement is to even be able to block Skype (Renals and Jacoby 2009).

This DPI use-case could potentially be contested and the subject of political conflicts. After all, even a DPI-based firewall looks into the users’ traffic data and could therefore be considered an infringement of telecommunications privacy. Why is this not happening? The answer is simple: This use-case is only found in corporate environments, where the employees have no reasonable expectation of privacy anyway. Companies use DPI to protect their corporate network, or universities their campus network. But DPI-based network security it is not used by internet access providers who sell to end-users, because the security problems of their customers can normally be externalized by the ISP – they simply don’t cost him money.(Eeten 2009) Administrators in the corporate world protect the company network and therefore are interested in utilizing the potential of DPI, but end-users just want cheap internet access. Even if one would consider content inspection with the aim of filtering malware a privacy breach: The EU data protection commissioners have already accepted that spam filters – which similarly read all email content – do not. So there is simply no leverage for this argument, and consequently no politicizing in sight.

4.2. Bandwidth Management

Bandwidth is and has always been a scarce resource, and therefore users can encounter congestion problems. This is a characteristic of any shared system that gives some freedom to the users, be it streets, mobile phone antennas or internet pipes. In fact, as early as the 1970s, network congestion was becoming an established research field. The internet protocols have a built-in way to deal with this (TCP congestion control), but new applications and the growth of the user base has for long kept up with an increase in bandwidth. Researchers spoke of "internet meltdown" in the 1980s, and the World Wide Web was often dubbed "World Wide Wait" in the 1990s. In the early days this affected the backbones and intercontinental lines in a variety of ways. One finding was that congestion only appeared on special routes, e.g. to New York City after 4 pm (MacKie-Mason and Varian 1996), others reported a complete jam of the undersea cable to New Zealand when the Apple Quicktime player was introduced in the early 1990s. While the argument that the "tubes"¹⁸ can be filled has also been made and exaggerated for political reasons¹⁹, just the case of video hosting service Youtube – now a subsidiary of Google – shows the orders of magnitude:

Each year the original content on the world's radio, cable and broadcast television channels adds up to about 75 petabytes of data -- or, 10 to the 15th power. If current estimates are correct, the two-year-old YouTube streams that much data in about three months. But a shift to high-definition video clips by YouTube users would flood the Internet with enough data to more than double the traffic of the entire cybersphere. And YouTube is just one company with one application that is itself only in its infancy. Given the growth of video cameras around the world, we could soon produce five exabytes of amateur video annually. Upgrades to high-definition will in time increase that number by another order of magnitude to some 50 exabytes or more, or 10 times the Internet's current yearly traffic. (Swanson 2007)

Today, there is still enough available bandwidth in the internet backbones for the foreseeable future (largely because of over-investment during the dot-com bubble). But the last mile has become a bottleneck due to the growth of these bandwidth-consuming applications like high-definition video streaming or peer-to-peer file sharing. This has especially an effect in two segments of the internet access market: Cable modem providers and mobile internet providers. Cable access suffers from the fact that the last mile is shared between a number of households in the same street, building, or neighborhood. If one or two heavy users are among them, they can easily consume all the available bandwidth and therefore degrade the experience even of those users who just want to download a few emails or surf the web. The problem here is to find a solution for sharing the available bandwidth between these users in a fair way. Mobile providers have a general scarcity problem, because the available frequencies are limited. Even if speeds are upgraded by new wireless technologies like HSDPA/UMTS-Broadband, the number of subscribers who can use these at the same time is limited per cell. In both cases, enlarging the capacity of the networks requires huge investments. A cable provider can either upgrade the cable modem termination systems (CMTS) or re-configure the physical cables so that less users share one CMTS; a mobile provider has to install more cells, which in turn also

¹⁸ The episode of the Daily Show built around the quote by Senator Ted Stevens ("the internet is a series of tubes") has become a classic reference in the internet community. Comedy Central, 12 July 2006, available at <http://www.thedailyshow.com/video/index.jhtml?videoId=126985>.

¹⁹ AT&T's vice president Jim Cicconi claimed in April 2008: "In three years' time, 20 typical households will generate more traffic than the entire Internet today" (Donoghue 2008), but the original comparison had been made to the Internet more than ten years ago: "By 2010, the average household will be using 1.1 terabytes (roughly equal to 1,000 copies of the Encyclopedia Britannica) of bandwidth a month, according to an estimate by the Internet Innovation Alliance in Washington, D.C. At that level, it says, 20 homes would generate more traffic than the entire Internet did in 1995." (Cauley 2008)

require more uplinks. This kind of scarcity is not prevalent with DSL internet access, because here, the users do not share the last mile. For DSL providers, what matters is only the sum of all traffic that they have to transport to and from the next internet exchange point or backbone provider. But for DSL, the increase of traffic can bring problems, as well. Normal end user internet providers have much less bandwidth available than the sum of all their customers have subscribed to (over-subscription).²⁰ This makes economic sense, because not all subscribers are online at the same time, and not all of them need the full bandwidth on top of that. But now, file-sharing programs often run in the background and even over night, and video downloads have added to the problem. Normal TCP/IP congestion control mechanisms can not deal with these problems, because they only work if the applications behave in a fair way and e.g. do not establish several TCP connections simultaneously or ignore the IETF congestion control standards.(Floyd 2000)

A number of ISPs have therefore started what is now called “network management” or “bandwidth management”. The first generation of this approach throttled down applications that were perceived as illegitimate anyway, especially peer-to-peer (P2P) file-sharing platforms like Bittorrent. While these are used for many legitimate purposes, including the distribution of free content or open source software, it is safe to assume that much of the P2P traffic consists of illegal copies of movies and music (Mateus and Peha 2008). On top of that, cable networks are designed to have more download than upload capacity, therefore they have problems dealing with P2P users who seed or upload large files. Either way, the P2P users were (and still are) considered “bandwidth hogs” who degrade the experience of other users. Some ISPs blocked Bittorrent traffic completely; others only slowed down the P2P traffic.

Interestingly, information about these practices had been collaboratively collected by users for quite a while. Collaborative P2P websites like Azureuswiki.com, named after a popular Bittorrent client, have long lists of ISPs in a number of countries which inform about the details of traffic shaping practices, as well as how to work around these issues. The data available shows that mostly cable ISPs use this kind of traffic shaping²¹, which conforms to the economic incentives discussed above. A number of such tools for detecting traffic shaping practices are available already. They mostly detect the transfer times and available bandwidth for different applications and protocols by technical means, other focus especially on peer-to-peer traffic. So far, they have required advanced computer skills (compiling a source code, setting up a server) beyond the grasp of a normal user. More recently, internet giant Google has joined forces with the New America Foundation and Planet Lab and set up a platform (“Measurement Lab” or “M-Lab”) for internet measuring tools.²² So far, it only contains a list of tools smaller than the one already provided by the Electronic Frontier Foundation (EFF)²³, but given Google’s resources and its experience in usability as well as aggregating and leveraging user data, we may see more readily usable tools for research here, as well as an informative and comprehensive collection of internet management all over the world. M-Lab’s stated future plans include cooperation with the research community, including support for “a wide variety of internet research & measurement tools” as well as making “all data publicly accessible” (M-Lab 2009).

²⁰ Most ISPs therefore offer internet access “up to” a certain bandwidth in their marketing material and customer contracts.

²¹ For example, the only German ISP slowing down P2P traffic seems to be Kabel Deutschland. In the U.S., cable ISPs tend to either slow down P2P traffic or prevent seeding (uploading), see. http://azureuswiki.com/index.php/Bad_ISPs.

²² <http://www.measurementlab.net>.

²³ <http://www.eff.org/testyourisp>.

The type of interaction among the actors involved one could expect here would be mutual adjustment through the market. With such an internet-enabled market transparency and the growing popularity of P2P among users, the threat for ISPs that their customers switch to a competitor that does not engage in traffic shaping could prevent ISPs from doing it in the first place. Indeed, in most European countries, where there is unbundling of the last mile and competition between cable and DSL providers, there are no or very few reports of P2P throttling by ISPs (Anzureuswiki.com 2009).²⁴

On the other hand, in large parts of the U.S. outside urban areas, there is only one internet provider available, so this market interaction could not take place. The issue therefore was dealt with on a political level. In August 2007, a blog post (Ernesto 2007) about an especially intrusive practice used by U.S. cable provider Comcast sparked a public debate. Comcast not just slowed down P2P traffic, but injected false “RST” (restore) packets into users’ uploads, giving the other end of the TCP connection the impression that the connection was terminated. For detecting P2P traffic, deep packet inspection equipment from Sandvine was used, as the company revealed later.²⁵ Similar practices were also reported from Canadian ISPs Cogeco and Rogers.(Ernesto 2007) This made public interest groups working on the “Net Neutrality” debate pick up the issue, place it on the public agenda, and kick off several political processes. The NGOs Electronic Frontier Foundation and Public Knowledge as well as the Bittorrent software vendor Vuze filed complaints and petitions with the Federal Communications Commission (FCC). The FCC had already established principles for broadband providers in 2005:

*consumers are entitled to access the lawful Internet content of their choice (...).
consumers are entitled to run applications and use services of their choice (...).
consumers are entitled to connect their choice of legal devices that do not harm the
network (...).
consumers are entitled to competition among network providers, application and
service providers, and content providers. (FCC 2005)*

After lots of unwelcoming media coverage and two FCC hearings, at the first of which Comcast even was caught engaging in questionable practices²⁶, the company tried to correct its public perception and avoid a ruling. In March 2008, it announced a switch to a new network management technique by the end of the year for dealing with heavy bandwidth use and congestion. For this, it was even partnering with BitTorrent Inc., in order to develop an application-neutral approach to traffic management. (Kravets 2008) (Lee 2008b) Still, in August 2008, the FCC decided that Comcast had “unduly interfered with Internet users’ rights” and ordered an end of this practice as well as a full disclosure of the methods used and replacement plans (FCC 2008).²⁷

As a result of this FCC ruling as well as negative press coverage, customer complaints and other problems²⁸, ISPs have started to switch to application-agnostic network management. While the FCC case against Comcast was still pending, competitor Time Warner Cable announced tests with customers in June 2008 with metered internet access. Under these

²⁴ Only the U.K. is an exception, which may be due to some informal agreement of major ISPs.

²⁵ See the Comcast filing to the FCC from 19 September 2008, [http://www.eff.org/files/Complete_Comcast_NM_Filing -- Date-Stamped 9 19 2008.pdf](http://www.eff.org/files/Complete_Comcast_NM_Filing_-_Date-Stamped_9_19_2008.pdf).

²⁶ It was disclosed that a number of Comcast’s supporters in the audience during the first hearing were random people that got money for showing up. As a consequence of this embarrassment, Comcast representatives did not appear at the second hearing.

²⁷ Comcast has appealed the ruling in September 2008 (Kravets 2008).

²⁸ Canadian civil liberties groups filed a complaint with the Federal Privacy Commissioner against Bell Canada because of its use of DPI for traffic management, especially P2P throttling. The case is still pending.

subscription plans, heavier users would pay more. (Kravets 2008) Other ISPs since then have experimented with volume caps, where a customer e.g. gets 250 GB of download volume per month, after which either the speed is reduced or every additional Gigabyte will cost extra. Similar volume-based approaches to data traffic have been in use by mobile internet providers. In this field, users had always been used to paying per minute or volume, so mobile providers did not seem to see a need for more intrusive network management practices.

While this subscriber-centric approach still uses DPI technology and may be considered intrusive, the FCC and other regulators' policies seem to allow this as long as the customers have enough transparency over what they subscribe to. (Karpinski 2009) Because DPI equipment is so flexible, the vendors still profit from this trend. A recent industry report therefore issued "strong buy" and "buy" recommendations for leading DPI vendors ProCera Networks and SandVine Corp, based on the prediction that

2009 will be a breakout year for the traffic management and deep packet inspection subsector of the networking industry. (...) Over the past few months, many, but not all, of the competing factions have compromised on a set of policies that will allow certain less aggressive network management policies to be enacted by ISPs. Cox Cable, the country's third-largest ISP, recently issued a policy statement discussing how it will implement an application based prioritization policy that will delay some types of traffic, such as large file transfers, during periods of peak demand. While some parties are still not happy with this policy, the ISP industry nevertheless feels it now has a set of parameters on which sound traffic management policies can be set. We believe this signals the beginning of a period of much more aggressive rollouts of traffic management systems. We believe this will be beneficial to several vendors within this subsector of the networking industry.(Noel 2009)

4.3. Ad Injection

Online marketing and advertising has been one of the growth markets in the marketing industry over the last years. The biggest player in this field nowadays is Google, with its own program AdSense and the acquisition of DoubleClick as the main components. One approach that has gained momentum is to serve ads in websites not based on the context (i.e. the content of the website in which the ad will be inserted), but on the profile of the user. This profile is based on tracking where users are going on the net, which websites they visit and what they do there. If, based on this, an advertising company knows e.g. that a user is mainly interested in video-games and new technology gadgets, it can serve him gaming or technology ads even if he is just looking for a train connection or for the weather forecast. This approach is known as behavioral advertising.

Companies like Google aggregate knowledge about users' online behavior by feedback from cookies and scripts that are placed on a number of websites, e.g. because the website owner uses Google AdSense to earn some money, or because he runs a Google Analytics script on his blog to see where the visitors are coming from. The problem, of course, with this approach is that even Google does not know everything about all internet users, because not all websites in the world incorporate Google's services.²⁹

²⁹ German data protection commissioners even have made clear that the use of Google Analytics is illegal for German website owners, because it conflicts with German data protection laws, and that even Google Inc. in Mountain View is bound by this through the EU-USA Safe Harbor Agreement (ULD 2009).

Internet service providers, on the other hand, have access to all their subscribers' web surfing data, because they are the one who transport all their internet traffic. If they start looking into the internet traffic in real-time using DPI, they can establish a pretty comprehensive picture of the users' interests and profiles. Based on this, a number of companies have sprung up over the last few years that offer exactly these services for ISPs. They track users' behavior, aggregate it into a user profile, and then inject ads into the websites the users visit based on their profiles. This approach is known as "ad injection". It has mainly been used in the U.S. and the UK so far.

Since 2007, it had seemed like there was a general move towards ad injection usage by ISPs. U.S. ISP Knology publicly acknowledged that it had been testing deep packet inspection with a few hundred customers. Other internet access providers such as Embarq or Wide Open West (WOW) changed their terms of service in order to allow tracking of user websurfing habits (Krempel 2008). CenturyTel was also testing NebuAd's service (Anderson 2009). Similar tests were conducted or planned by British ISPs. But in 2008, ad injection and related DPI usage were pushed onto the public agenda and met heavy scrutiny and criticism. Interestingly, the results were markedly different in the United States and in the UK.

In the United States, ISP Charter Communications announced its own plan to partner with an ad injection company called NebuAd in May 2008. Charter's Senior Vice President sent a letter to customers informing them of the plan and giving them instructions on how to opt out. Like its industry peers, Charter was criticized following its announcement. The public advocacy groups Free Press and Public Knowledge hired a technical consultant to produce a report dissecting NebuAd's methods (Topolski 2008). Congressmen Edward Markey and Joe Barton wrote a letter to Charter's CEO arguing that the plan might violate federal law and urging the company not to act until it had consulted with Congress. The Senate Subcommittee on Interstate Commerce, Trade, and Tourism held a hearing about interactive advertising prompted by the controversy. Connecticut's Attorney General also released a letter urging Charter not to implement the program. In the face of this criticism, about a month after announcing the plan, Charter abandoned it. (Ohm 2008) A number of other ISPs who had been in talks with NebuAd or similar companies also suspended or cancelled trial runs of ad injection systems. As a result, the market for DPI-based ad injection in the United States collapsed. In September 2008, NebuAd had to announce that it had ended its behavior-based marketing activities, and CEO and founder Bob Dykes resigned. (Nakashima 2008) Adzilla, another US-based ad-injection company, even took its website offline in late summer 2008 and since then has replaced it with a one-page notice saying it plans to "develop new solutions and services to delight users and enhance the online ecosystem experience". The story is not even over yet: in late 2008, NebuAd and the partnering ISPs were sued in a California federal court for intercepting, copying, and otherwise interfering with users' privacy (Anderson 2009).

British ad injection provider Phorm was hit by similar public scrutiny in early 2008 when it became known that major ISP British Telecom (BT) had conducted a trial of its "Webwise" system with 10 000 customers in 2006 and 2007. The public interest think tank Foundation for Information Policy Research (FIPR) published two lengthy and critical analyses of Webwise. While the technical report (Clayton 2008) made clear to the interested public how intrusive and deceptive the system works, the legal analysis (Bohm 2008) drew drastic conclusions:

(...) that deployment by an ISP of the Phorm architecture will involve the following illegalities (for which ISPs will be primarily liable and for which Phorm Inc will be liable as an inciter):

- *interception of communications, an offence contrary to section 1 of the Regulation of Investigatory Powers Act 2000*
- *fraud, an offence contrary to section 1 of the Fraud Act 2006*
- *unlawful processing of sensitive personal data, contrary to the Data Protection Act 1998*
- *risks of committing civil wrongs actionable at the suit of website owners such as the Bank of England.*

But while Phorm's stock price dropped by 85% (Mitchell 2008), the company has continued its business. One reason for this is the soft reaction of the UK government and the national information and privacy commissioner. After questions were raised by the EU Commission about the conformity with European telecommunications privacy law, the British government stated that Phorm's ad injection system did not breach EU data laws nor according UK regulations. It only required that any future versions of the system have to be done with user consent and make it easy for ISP customers to opt out. The information commissioner's office stated that the trials were a breach of the EU Telecommunications Privacy Directive of 2003, but they would not take action against BT, because they had only conducted purely "technical tests". (Hanff 2008) As a reaction, the two other ISPs initially also debating to test Phorm's system announced ongoing but skeptical interest. Virgin Media said it was "still evaluating the system", and Carphone Warehouse announced it would only run Phorm on an opt-in basis (anonymous 2008). BT is still planning to deploy the Phorm's Webwise system, and Phorm announced in January 2009 that it is getting first overseas customers (Andrews 2009). An interesting factor in the Phorm case was the division of privacy advocacy groups. While FIPR spoke out loudly against Phorm, and other activists set up a campaign website (www.nodpi.org) and started criminal procedures against Phorm and BT, the well-known London-based watchdog group Privacy International (PI) was curiously quiet in the whole conflict. The reason was that the core members of PI are also running a privacy consulting firm, which had been contracted by Phorm to advise them on their system. In fact, the architecture used by Phorm does not give away any customer information to the advertisement buyers, and the identifiers for tracking the customers are pseudonymized. This is also the reason the British information commissioner has not taken serious measures against Phorm or BT.

In the U.S., on the other hand, the "shadow of hierarchy" could be felt clearly by ISPs who thought about using NebuAd's services. This may be because the company was not as openly trying to picture its products as privacy-friendly, or because the ad injection ideas fell on fertile soil for protests by network neutrality activists and politicians ready to start procedures. Either way, the technological specifics of each case aligned with the set of the actors involved, and led to the different outcomes.

4.4. Copyright Content Filtering

As deep packet inspection looks into the payload of the internet traffic, it has always been an option to filter and block traffic based on what is in it (not just based on the protocol, like in the bandwidth management case). There are basically two versions of this story: One is around filtering demands from copyright owners, the other is around censoring illegal or harmful content based on political pressure. Let us look at copyright filtering in Europe here.

The filtering of illegitimately distributed copyrighted material is a DPI use case that is especially interesting, because while it had its ups and downs, it now seems to lack most chances for survival in Europe. Since 2004, the European music industry has tried to use the

courts to establish a secondary liability for ISPs whose customers illegally share copyrighted material. The aim was to force ISPs to set up filtering technology that would detect and block copyrighted music automatically. This was seen as a technical alternative to identifying and suing each user who was found participating in a peer-to-peer network and distributing copyrighted material. One product that has been specifically marketed by the music industry is sold by a company called Audible Magic. It uses a fingerprinting technology to recognize copyrighted music files in the data stream. It is said to be used in about 75 universities in the U.S. to monitor their networks for peer-to-peer sharing of copyright music (anonymous 2007) and has been heavily marketed by the recording industry.

In June 2007, this strategy had what looked like a quick success in Belgium. The music industry association (SABAM) demanded in court that ISP Scarlet installs such a filtering technology. An injunction by the Court of First Instance in Brussels in fact established exactly that obligation.(EDRi 2007) After this initial success, the music industry moved on to Ireland in 2008 and sued the largest internet provider, Eircom. EMI, Sony, Warner and Universal sought an injunction from the Dublin High Court which would have required Eircom to establish the same filtering system as in Belgium.(McIntyre 2008a) Here, we can clearly see how different national markets are affected by the same technology trends as well as transnational interest groups.

But in October 2008, the Belgian case turned out different than expected. ISP Scarlet convincingly demonstrated to the court that the technology suggested by SABAM as well as in Ireland (Audible Magic) - did not work and that the music industry even had deceived the court by falsely claiming it had already been used elsewhere. Therefore, the trial court in Belgium lifted the injunction against Scarlet.(McIntyre 2008b) This, in turn, created a problem for the music industry in Ireland. The obvious reaction was to prevent a precedent in Ireland that would ban any copyright filtering obligations for the future, and instead try to reach the second best outcome. Just two weeks into the proceedings of the Irish case, the Belgian case was settled out of court. The parties agreed that Eircom implements a „three strikes, you’re out” policy: disconnecting users from the internet after they have been identified as illegally distributing music on file-sharing platforms.(McIntyre 2009)

The main reason for the outcome in this case was not just the technological feasibility – there have been many cases where politicians or judges have decided to use non-functional technology. It was the open opposition of the ISP to simply complying with the ruling and installing a filtering system. While the decision in the end was made by hierarchical order (the court), and important part was the negotiations over the truth claims here. Because Scarlet had control over the network and the exact technical set-up, it was in a position to convince the judges of its version of the truth claims, and in the end convince the Belgian court that Audible Magic does not deliver what its promises. The out-of-court settlement in Ireland a bit later shows even more signs of as negotiated agreement. The content industry could not win the current case anymore, but by settling also prevented a precedent and thereby left the legal space open for technological improvements of copyright filtering software.

While the policies, actor interests and interactions are really different, a similar observation can be made in the discussions about mandatory child-porn and “harmful content” filtering. ISPs have a central veto position here and openly exploit it. For instance, a major ISP in Australia announced openly that it would only participate in government-run field tests of a planned filtering infrastructure to generate more proof that these technologies do not work.(Moses 2008)

4.5. *Government Surveillance*

Deep packet inspection can also be used for monitoring internet communication in real-time., without changing the content or discriminating different streams. Law enforcement and intelligence agencies have always been trying to catch up with the latest advances in communications technology. While the rise and spread of broadband internet access made it hard to catch even for well-funded agencies such as the National Security Agency, deep packet inspection now makes full-scale surveillance of internet traffic possible again. In the United States, an inside whistle-blower in December 2005 reported that the NSA was massively intercepting domestic phone calls and internet communications without any court order or oversight. It seemed that president George W. Bush, based on a theory of broad executive power in times of war, had reauthorized this surveillance a number of times, even after the Department of Justice found the program to violate criminal laws. (Singel 2006) A bit later, former AT&T technician Mark Klein revealed that the telecommunications giant had installed a fiber-optical “splitter” at its facility in San Francisco, which sent copies of phone and internet traffic to a secret room controlled by the NSA (Klein 2006). Later reports showed that most of the nation’s telecommunications companies had set up similar interception and surveillance facilities under control of the NSA. The technology used was reportedly based on DPI equipment by Israeli-American vendor Narus (Ewert 2006). It captured all the traffic, inspected the content, and stored a smaller subset for detailed analysis and combination with other data available to the intelligence community.

This interception was obviously violating legal protections and safeguards like the Foreign Intelligence Surveillance Act (FISA) and the Fourth Amendment to the U.S. constitution. It also pointed at illegal behavior by both the NSA *and* the telecommunications companies. The Electronic Frontier Foundation (EFF) sued AT&A and tried to hold the company accountable for its participation in illegal dragnet spying.³⁰ As a reaction, the U.S. Congress passed the FISA Amendments Act which granted immunity to telecom firms that participated in the wiretapping program in June 2008. The decision was bipartisan. Even Barack Obama, who was still campaigning to become president at that time and risked losing his liberal followers, voted for retro-active immunity of ISPs. The EFF is currently challenging this law, and is also representing victims of the surveillance program in a lawsuit against the government since September 2008.³¹

In this case, the post-9/11 situation in which the distinction between domestic and foreign threats was blurring had clearly shifted the interest of the intelligence towards monitoring domestic traffic, too. The emergence of DPI technology for high-speed internet traffic monitoring at the same time made this possible. The specialties of the “war on terror“ then lead the telecommunications companies to fulfill NSA’s (and indirectly the president’s) demands even if in clear breach of the law. The government and the legislature were in turn willing to grant post-hoc immunity to them after the facts of this practice were published. This shows that technological opportunities can be used by governments to breach previously established and long stabilized norms and practices. In special situations with extremely high importance on the public agenda, the ISPs are willing to cooperate and use the technologies against established law, but in turn, they expect protection from the government. While this is a general pattern, individual deviations are still possible. Telecommunications provider Qwest did not obey to the NSA demands, insisting on due legal procedures and FISA warrants. In turn, it reportedly got under heavy pressure because the U.S. government withdrew contract options in the dimension of hundreds of millions of dollars. (Nakashima and Eggen 2007)

³⁰ See <http://www.eff.org/nsa/hepting>.

³¹ See <http://www.eff.org/cases/jewel>.

5. Conclusions: Explaining variation in DPI Governance

Comparing the cases, we can see different ways DPI is or can be embedded into or break the social context, and how it impacts established interests. If we again focus on the role of the ISPs as the operators of the internet, we can build a draft typology, based on their different roles, interests, and the following interactions.

Autistic practices: This is what we see in the case of network security. The owner of a LAN or corporate network is running DPI equipment on the edge between his own network the internet, using it as an intrusion-detection system. This is purely self-centered behavior, and in fact it does not affect other actors on the internet or elsewhere in any way (employees are expected to accept this). So there are no conflicts and no politics developing around these cases. In game-theoretic terms, the usage of DPI here is a unilateral move that does not interfere with the interests of other actors. The outcome (no interaction, not politicizing) is as expected. New technology just gets adopted and used if it fits the interests of actors, as long as no-one else's interests are affected.

Unilateral market behavior in regulated markets: This is what we find in the cases of bandwidth management and ad injection. The ISPs have an economic incentive to create additional revenues or manage available resources more effectively, and DPI allowed them new ways for doing this. As expected, DPI usage varies over the type of ISP (mostly cable and mobile). Here as well, the ISPs used DPI for their own network and their own benefit. The difference to the autistic practices, though, is that in these cases, the ISPs are market players. They sell connectivity to end-customers, they have other ISPs competing with them, and they have to comply with some rules set by regulatory or other authorities. This led to political conflicts with other actors, mainly consumer and net neutrality groups, who openly opposed traffic discrimination and ad injection. While the cases so far are pretty similar and the technology more or less the same in each use-case, the inputs as well as the outcomes vary across countries. On the input side, it seems that bandwidth management is mainly an issue in the U.S. and other regions where there is less competition on the last mile because of a lack of unbundling regulations. European ISPs except for some cable providers (and the outlier U.K.) tend to not resort to bandwidth management because of functioning competition. In the U.S., a specific way of application-agnostic bandwidth management is now on the rise, after the FCC has specified the rules for the market and thereby set the acceptable use-cases. In the cases of ad injection, the outcome was remarkably different in the U.S: and the U.K., though the technology used was more or less the same. The fact that ad injection now seems to be allowed in the U.K. can be explained by the actor constellation (Privacy International staff actually helping Phorm) and the institutional setting (the Information Commissioner allowing ad injection if opt-out is provided).

Negotiated agreement: This is what we find in the cases of content filtering. The use of DPI here was not in the ISPs' own interest, but was demanded by others, such as the copyright holders of films and music. These needed an agreement with the ISPs, who did not want to become "copyright cops" and therefore had a veto position. In the U.S., the Audible Magic technology is mainly used at universities, but not by end-user access providers.³² In Europe, the ISPs were successful in even demonstrating to the courts that the filtering technology does not work. The negotiations therefore were taking place on the level of truth claims, and the ISPs as operators of the networks again were able to leverage their expertise. While not dealt with in depth in this paper, the same pattern seems to be the case for child-porn and "harmful content" filtering as demanded by governments and other actors without explicit legal

³² In the university cases, we can refer to the "autistic practices" pattern again.

requirements to do it. When the filtering lists get too broad, the ISPs openly refuse to cooperate, as is the case in Australia. When it is very narrow about real child porn material, they are willing to cooperate. In both cases, they have a veto position, unless they are not forced by law.

Hierarchical decision-making: This is what happened in the NSA surveillance case in the United States. Under pressure by the president and the National Security Agency in the post-9/11 situation, most ISPs complied with the government's demand to install spying hardware, even though it was not required by the law. The protests and lawsuits by the public and by NGOs did not change the situation. While one could argue that the deviation of Qwest shows that we also have a *negotiated agreement* here, there was a clear hierarchy in the actors and the discourse: National security and the agencies dealing with it beat other arguments. The special hierarchical role of the state here is also illustrated by the fact that the ISPs' compliance was legally covered by government and congress decisions after the fact.

The combination of theories and approaches from the sociology of technology and from interaction-oriented policy research used here has proven helpful for understanding variations in DPI use. It also helped shedding some light on the impact of technology and society respectively:

On the input-side, technology-oriented policy analysis helped explain the *variation across use-cases of DPI*: The different use-cases provide for varying interests, motivations, and capabilities of ISPs and other actors. This also structured the initial strategic interaction situation. Roughly said: If ISPs have a self-interest in DPI usage, they can and often will go ahead and just do it. If DPI is in favor of the interests of third parties, like copyright holders or government agencies, these have to put significant efforts into making the ISPs cooperate and not play their veto position. Again: Technology determined politics.

On the outcome side, the analysis of the norms and institutions in which the concrete interactions took place helped explain the *variation within use-cases of DPI*. Roughly said: If there are regulatory norms or oversight bodies to which the ISPs can refer, this works in their favor, as shown in the ad injection case in the U.K. Strong moral norms and/or hierarchical decision-making by government authorities, on the other hand, play an important role in convincing ISPs to do things they would not normally do, as shown in the NSA surveillance case and in the filtering of child-porn.

Still, there is neither techno-determinism nor social-determinism here. The framework of actor-centric institutionalism is helpful in nailing down the initial strategic setting of the actors involved, and the added technology analysis can adapt this to techno-centric policy fields. But actor-centrism also means that there is still a lot of freedom for the parties involved to behave differently, to make mistakes, or even to still act in favor of the most open internet in the face of a disruptive technology like deep packet inspection.

Literature

- Anderson, Nate. 2008. "Throttle 5 million P2P users with \$800K DPI monster." *Ars Technica*, 12 May 2008.
- . 2009. "ISPs: don't blame us; NebuAd did all the dirty work!" *Ars Technica*, 6 February 2009.
- Andrews, Robert. 2009. "Phorm Denies It May Pay Users To Opt In." *paidcontent.uk*, 8 January 2009.
- anonymous. 2007. "Audible Magic emerging as top copyright cop in digital revolution." *International Herald Tribune*, 26 March 2007.
- . 2008. "Anti-Spyware Coalition untersucht Praktiken für personalisierte Anzeigen." *Heise News*, 28 April 2008.
- Anzureuswiki.com. *Bad ISPs* 2009 [cited 13 February 2009. Available from http://azureuswiki.com/index.php/Bad_ISPs].
- Artan, N. Sertaç. 2007. *High-speed network intrusion detection and prevention (PhD Dissertation)*. Brooklyn, NY: Politechnic University.
- Bacon, Francis. 1858. *Instauratio Magna, Part II: Novum Organum, Book I*. Translated by James Spedding et al. Edited by J. Spedding, R. L. Ellis and D. D. Heath. London: Longmans.
- Barlow, John Perry. 1996. "A Declaration of the Independence of Cyberspace." Davos, Switzerland.
- Bendrath, Ralf. 2008. "The Social and Technical Self-Governance of Privacy." In *Responsible Business? Self-Governance and the Law in Transnational Economic Transactions*, ed. O. Dilling, M. Herberg and G. Winter. Oxford: Hart.
- Bendrath, Ralf, Jeanette Hofmann, Volker Leib, Peter Mayer, and Michael Zürn. 2007. "Governing the Internet: Legitimate and Effective Rules for a Globalized Medium?" In *Transforming the Golden Age Nation State*, ed. A. Hurrelmann, K. Martens and P. Mayer. Houndmills: Palgrave.
- Bennett, Richard. 2008. "In neutrality debate, carriers get blamed for Net's weaknesses." *San Jose Mercury News*, 17 April 2008.
- Bijker, Wiebe E., and John Law, eds. 1992. *Shaping Technology, Building Society: Studies in Sociotechnical Change*. Cambridge, MA: MIT Press.
- Bohm, Nicholas. 2008. "The Phorm "Webwise" System - a Legal Analysis." Cambridge: Foundation for Information Policy Research.
- Carpenter, Brian E. 1996. *RFC 1958: Architectural Principles of the Internet*. Fremont/CA: IAB Network Working Group.
- Cauley, Leslie. 2008. "Avoiding Net traffic tie-ups could cost you in future." *USA Today*, 20 April 2008.
- Chester, Jeff. 2006. "The End of the Internet? ." *The Nation*, 1 February
- Clayton, Richard. 2008. "The Phorm "Webwise" System." Cambridge: Cambridge University Computer Lab.
- Cranor, Lorrie Faith. 2002. "The role of privacy advocates and data protection authorities in the design and deployment of the platform for privacy preferences." In *Computers, Freedom and Privacy 2002*. San Francisco: ACM.
- Dedman, Bill, and Bob Sullivan. 2008. "ISPs are pressed to become child porn cops." *MSNBC.com*, 16 October 2008.
- Deibert, Ronald, John Palfrey, Rafal Rohozinski, and Jonathan Zittrain. 2008. *Access Denied: The Practice and Policy of Global Internet Filtering*. Cambridge, MA: MIT Press.
- Dharmapurikar, Sarang, Praveen Krishnamurthy, Todd S. Sproull, and John W. Lockwood. 2004. "Deep Packet Inspection using Parallel Bloom Filters." *IEEE Micro* 24 (1):52-61.

- Dierkes, Meinolf, Ute Hoffmann, and Lutz Marz. 1992. *Leitbild und Technik. Zur Entstehung und Steuerung technischer Innovationen*. Berlin: edition sigma.
- Dierkes, Meinolf, and Lutz Marz. 1990. "Technikakzeptanz, Technikfolgen, Technikgenese. Zur Weiterentwicklung konzeptioneller Grundlagen der sozialwissenschaftlichen Technikforschung." Berlin.
- Dolata, Ulrich, and Raymund Werle, eds. 2007. *Gesellschaft und die Macht der Technik. Sozioökonomischer und institutioneller Wandel durch Technisierung*. Frankfurt a.M.: Campus.
- Donoghue, Andrew. 2008. "AT&T: Internet to hit full capacity by 2010." *CNET News.com*, 18 April 2008.
- Drake, William J. 1993. "Territoriality and Intangibility." In *Beyond National Sovereignty: International Communications in the 1990s*, ed. K. Nordenstreng and H. I. Schiller. New York: Ablex.
- Drezner, Daniel W. 2004. "The Global Governance of the Internet: Bringing the State Back In." *Political Science Quarterly* 119 (3):477-98.
- EDRi. 2007. "Belgium ISP ordered by the court to filter illicit content." *EDRi-Gram*, 18 July 2007.
- Eeten, Michel van. 2009. "The Broken Model of Internet Security: Dealing with the Security Externalities of Key Market Players (lecture)." In *"Focus on Internet" Seminar*. TU Delft.
- Ernesto. 2007. "Comcast Throttles BitTorrent Traffic, Seeding Impossible." *Torrentfreak.com*, 17 August 2007
- Ewert, Bruce. 2006. "All About NSA's and AT&T's Big Brother Machine, the Narus 6400 " *Daily Kos*, 7 April 2006.
- FCC. 2005. *Policy Statement*. Washington DC.
- . 2008. *Press Release: Commission orders Comcast to end discriminatory network management practices*. Washington DC: Federal Communications Commission.
- Floyd, Sally. 2000. *RFC2914: Congestion Control Principles*. Fremont/CA: IETF Network Working Group.
- Gabel, David, and David F. Weiman, eds. 1998. *Opening Networks to Competition. The Regulation and Pricing of Access*. Boston, Dordrecht, London: Kluwer Academic Publishers.
- Goldsmith, Jack L., and Tim Wu. 2006. *Who Controls the Internet? Illusions of a Borderless World*. New York: Oxford University Press.
- Grunwald, Armin. 2007. "Technikdeterminismus oder Sozialdeterminismus: Zeitbezüge und Kausalverhältnisse aus der Sicht des "Technology Assessment"." In *Gesellschaft und die Macht der Technik. Sozioökonomischer und institutioneller Wandel durch Technisierung*, ed. U. Dolata and R. Werle. Frankfurt a.M.: Campus.
- Hanff, Alexander. 2008. "ICO admit BT 2007 Trials breached PECR 2003 but refuse to act!" *nodpi.org*, 31 May 2008.
- Johnson, David R., and David G. Post. 1997. "The Rise of Law on the Global Network." In *Borders in Cyberspace. Information Policy and the Global Information Infrastructure*, ed. B. Kahin and C. Nesson. Cambridge/MA: MIT Press.
- Karpinski, Rich. 2009. "Comcast's Congestion Catch-22." *Telephony Online*, 23 January 2009.
- Klang, Mathias. 2006. *Disruptive Technology. Effects of Technology Regulation on Democracy*. Göteborg: University Department of Applied Information Technology.
- Klein, Mark. 2006. "Declaration in Support of Plaintiff's Motion for Preliminary Injunction." U.S. District Court of Northern California.
- Klein, Torsten. 2009. "Bundesregierung treibt Netzblockaden gegen Kinderpornografie voran." *heise news*, 13 January 2009.

- Knill, Christoph, and Dirk Lehmkuhl. 2002. "Private Actors and the State: Internationalization and Changing Patterns of Governance." *Governance: An International Journal of Policy, Administration, and Institutions* 15 (1):41-63.
- Kravets, David. 2008. "Comcast Beginning 'Net Neutrality' Testing." *Threat Level*, 3 June 2008.
- Krempf, Stefan. 2008. "Provider sollen Kunden umfassend ausgespäht haben." *Heise News*, 6 April 2008.
- . 2009. "Internetsperren und Filter erneut auf der Brüsseler Agenda." *heise news*, 27 January 2009.
- Kuerbis, Brenden. 2009. "Securing critical Internet resources: Influence and control of Internet standards through delegation and social networks." In *International Studies Association 50th Annual Convention*. New York City.
- Latour, Bruno. 1994. *Der Berliner Schlüssel. WZB Working Paper FS II 94-508*. Berlin: Wissenschaftszentrum Berlin für Sozialforschung.
- Lee, Timothy. 2008a. "Changing The Internet's Architecture Isn't So Easy." *Techdirt.com*, 24 April 2008.
- Lee, Tom. 2008b. "What Comcast/Bittorrent Actually Means: Bittorrent Selling Hardware." *TechDirt*, 4 April 2008.
- Lessig, Lawrence. 1999. *Code and other Laws of Cyberspace*. New York City: Basic Books.
- . 2002. *The Future of Ideas. The Fate of the Commons in a Connected World*. London: Vintage Books.
- Lowi, Theodore. 1972. "Four Systems of Policy, Politics, and Choice." *Public Administrative Review* 32 (July/August):298-310.
- M-Lab. 2009. *Get Involved* 2009 [cited 30 January 2009 2009]. Available from <http://www.measurementlab.net/getinvolved.html>.
- MacKie-Mason, Jeffrey K., and Hal R. Varian. 1996. "Some Economics of the Internet." In *Networks, Infrastructure, and the New Task for Regulation*, ed. W. Sichel and D. L. Alexander. Ann Arbor: University of Michigan Press.
- Mateus, Alexandre, and Jon M. Peha. 2008. "Dimensions of P2P and Digital Piracy in a University Campus." In *TPRC*. Arlington VA.
- Mathiason, John. 2006. "Internet Governance Wars, Episode II: the Realists Strike Back: A review of Goldsmith and Wu's 'Who Controls the Internet? Illusions of a Borderless World'." Syracuse: Internet Governance Project.
- Mayer-Schönberger, Viktor. 2003. "The Shape of Governance. Analyzing the World of Internet Regulation." *Virginia Journal of International Law* 43:605-73.
- Mayntz, Renate, and Fritz W. Scharpf. 1995. "Der Ansatz des akteurszentrierten Institutionalismus." In *Gesellschaftliche Selbstregulung und politische Steuerung*, ed. R. Mayntz and F. W. Scharpf. Frankfurt a.M.: Campus.
- McIntyre, T.J. 2009. ""Three strikes" for Ireland - Eircom, music industry settle filtering case." *IT Law in Ireland*, 29 January 2009.
- McIntyre, T.J. 2008a. "Filter or Else! Music Industry Sues Irish ISP." *Computer & Law*, April - May 2008.
- . 2008b. "SABAM v. Scarlet: Belgian ISP released from obligation to filter network for illegal downloads." *IT Law in Ireland (Blog)*, 26 October 2008.
- Mitchell, Jordan. 2008. "Behavioral Targeting and Privacy - Game Over for ISP Level Profiling?" *MetaMuse (blog)*, 8 September 2008.
- Moses, Asher. 2008. "Net censorship plan backlash." *The Age*, 11 November 2008.
- Mueller, Milton. 2002. *Ruling the Root. Internet Governance and the Taming of Cyberspace*. Cambridge MA: MIT Press.
- Nakashima, Ellen. 2008. "NebuAd Halts Plans For Web Tracking." *Washington Post*, 4 September 2008, D02.

- Nakashima, Ellen, and Dan Eggen. 2007. "Former CEO Says U.S. Punished Phone Firm." *Washington Post*, 13 October 2007.
- Noel, Joseph. 2009. *Breaking the Traffic Management Deployment Logjam. As the Rules Have Become Clearer Deployments Are Beginning at a Rapid Pace*. San Francisco/CA: Emerging Growth Research, LLP.
- Ohm, Paul. 2008. *The Rise and Fall of Invasive ISP Surveillance*: SSRN.
- Orlikowski, Wanda J. 1992. "The duality of technology: rethinking the concept of technology in organizations." *Organizational Science* 3 (3):398-427.
- Orlikowski, Wanda J., and Debra C. Gash. 1994. "Technological frames: making sense of information technology in organizations." *ACM Transactions on Information Systems* 12 (2):174-207.
- Owen, Bruce M. 2008. "As long as flexibility has value to users, suppliers will have incentives to offer it." *Boston Review*, March / April 2008.
- Renals, Peter, and Grant A. Jacoby. 2009. Blocking Skype through Deep Packet Inspection. Paper read at 42nd International Conference on System Sciences, at Hawaii.
- Rosecrance, Richard. 1996. "The Rise of the Virtual State." *Foreign Affairs* (4):45-61.
- Saltzer, Jerome H., David P. Reed, and David D. Clark. 1984. "End-to-end arguments in system design." *ACM Transactions on Computer Systems* 2 (4):277-88.
- Scharpf, Fritz W. 1997. *Games Real Actors Play*. Boulder/CO: Westview.
- Shafer, Scott Tyler. 2002. "Switching made smarter." *Infoworld*, 8 May 2002, 34-6.
- Singel, Ryan. 2006. "Whistle-Blower Outs NSA Spy Room." *Wired News*, 4 July.
- Sourdis, Ioannis. 2007. *Designs and algorithms for packet and content inspection*. Delft: TU Delft.
- Sourdis, Ioannis, Dionisios Pnevmatikatos, and Stamatis Vassiliadis. 2008. "Scalable Multi-Gigabit Pattern Matching for Packet Inspection." *IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Special Section on Configurable Computing Design—II: Hardware Level Reconfiguration* 16 (2):156-66.
- Stevens, W. Richard. 1993. *TCP/IP Illustrated, Volume 1: The Protocols*. Upper Saddle River, NJ: Addison-Wesley.
- Swanson, Bret. 2007. "The Coming Exaflood." *Wall Street Journal*, 20 January 2007.
- Topolski, Robert M. 2008. "NebuAd and Partner ISPs: Wiretapping, Forgery and Browser Hijacking." Washington DC: FreePress.
- ULD. 2009. "Datenschutzrechtliche Bewertung des Einsatzes von Google Analytics." ed. U. L. f. Datenschutz. Kiel: Unabhängiges Landeszentrum für Datenschutz.
- Wagner, Ben. 2008. "Modifying the Data Stream: Deep Packet Inspection and Internet Censorship." In *3rd Annual GigaNet Symposium*. Hyderabad, India.
- Walker, John. *The Digital Imprimatur. How big brother and big media can put the Internet genie back in the bottle*, 4 November 2003 [cited 27 August 2008. Available from <http://www.fourmilab.ch/documents/digital-imprimatur/>].
- . 2004. "Ende des Internet?" *Telepolis*, 2 February 2004.
- Weingart, Peter. 1989. "'Großtechnische Systeme' - ein Paradigma der Verknüpfung von Technikentwicklung und sozialem Wandel?" In *Technik als sozialer Prozeß*, ed. P. Weingart. Frankfurt a.M.: Suhrkamp.
- White, Bobby. 2007. "Watching What You See on the Web." *Wall Street Journal*, 6 December 2007, B1.
- Winner, Langdon. 1986. *The Whale and the Reactor: A Search for Limits in an Age of High Technology*. Chicago: University of Chicago Press.
- Zittrain, Jonathan. 2004. "Internet Points of Control." In *The Emergent Information Policy Regime*, ed. S. Braman. Houndmills, Basingstoke: Pargrave Macmillan.
- . 2008. *The Future of the Internet and how to stop it*. New Haven & London: Yale University Press.