Designing effective regulation for the "Dark side" of the Web

By Wolf Richter[1] and Ian Brown (University of Oxford)

1. Problem Statement

After the initial excitement about the Internet as a space outside the reach of governmental control has evaporated and courts in several states have applied national laws to 'Cyberspace', there is now a consensus among scholars and activists that the Internet is in principle subject to national regulation.[2] Nevertheless, the application of the 'traditional' tools of regulation to the challenges amplified by the Internet like unlicensed file-sharing, defamation, or email spam, has either shown to be ineffective or has produced severe side-effects.

For example, copyright law has undergone several major revisions during the past decade, which resulted in a substantial expansion of scope and term of copyright law, but has not been effective in curbing copyright violations facilitated by the Web.[3] At the same time, the transaction costs for enforcing copyright law have sky-rocketed since the diffusion of broadband Internet access to most private households as none of the various enforcement strategies – suing the creators of file-sharing tools, suing individual file sharers, partnering with Internet Service Providers to throttle or terminate the accounts of repeat infringers - has provided an efficient remedy.

The all-purpose PCs connected to the Internet is the dominating form of the digital workplace of today. Despite existing liability rules for emitting unsolicited email ("spam") or malicious software code ("Badware") and the ample availability of security software, security attacks and infections of PCs occur regularly and cause massive damage to businesses and private users.[4] The global span of the Internet makes it even more difficult to get hold of the creators of spam and malware, who can launch and control their attacks from remote places. Current protection technologies and liability laws have failed to solve the problem effectively.

While file sharing initially was a major driver of demand for broadband Internet access, it is also a major source of bandwidth.[5] Internet Service Providers (ISPs) who offer own media services struggle to compete with the "free" competition. ISPs are therefore motivated to impose limits or even restrict the use of certain applications entirely, but have no incentive to publicize their management practices. The FCC's decision in the US against the consumer ISP Comcast has confirmed the interest of the regulator to maintain the Internet open for new applications without prior approval by the network operators. But the decision has also shown that the line between acceptable network management and undue discrimination is a fine line, and violations, unless blunt ones like in the Comcast case, are very hard to detect.

In this paper we will investigate a different approach to addressing these central questions for the future of the Web, which builds on the self-regulatory tradition of the Internet: Entrepreneurial regulation.

2. Entrepreneurial regulation

The Web as a social space is not subject to physical constraints. The affordances of its users are created and constrained by the Web's architecture, which is engrained in software code.[6] The architecture is the result of many loosely coordinated and at times controversial decisions taken by the designers of the Internet and the creators of its many applications.[7] This fundamental difference between the regulation of physical space and the regulation of the Internet poses a challenge for traditional regulators: The regulated subjects can avoid regulation by "coding around it". This phenomenon can be nicely seen in the development of file-sharing applications and protocols: the coders of these tools were fast to "implement" the court orders by adjusting the design of their tools to escape liability.[8] The global nature of the Web makes avoiding regulation by changing the jurisdiction for web sites is as easy as a few clicks, a popular ploy among the senders of spam, and phishing emails.

[1] Acknowledgements: The author would like to thank the Web Science Research Initiative for their generous support and Prof. William Fisher for his supervision during an enjoyable stay at Harvard's Berkman Center for Internet and Society during summer 2008. Finally he would like to thank his interview partners at the Berkman and the EFF for their patience and their support of this research project.

[2] *See* J. Goldsmith and T. Wu: Who Controls the Internet? Illusions of a Borderless World, Oxford University Press (2006).

[3] *See* A. Webb and British Music Rights: Music Experience and Behaviour in Young People (2008) *available at* http://www.ukmusic.org/cms/uploads/files/UoH%20Reseach%202008.pdf (last visited 26/10/2008).

[4] *See* J. Kreutzer: Somebody Has to Pay: Products Liability for Spyware, American Business Law Journal Vol. 45 No. 1 (2008) 61 *at* 62f.

[5] *See* A. Asvanund et. Al: An Empirical Analysis of Network Externalities in Peer-To-Peer Music Sharing Networks (203) *available at* SSRN: http://ssrn.com/abstract=433780.

[6] *See* L. Lessig: The new Chicago School *in* Journal of Legal Studies Vol. 27 No. 2 (1998) p. 661–691.

[7] *See* J. Zittrain: The future of the Internet – and how to stop it, Allen Lane (2008) *at* 31.

[8] *See* T. Wu: When Code Isn't Law in Virginia Law Review Vol. 89 No. 4 (2003) 679 *at* 726ff.

In this paper, we will consider the malleability not as a challenge, but as a chance to innovate regulation. But this potential can only be realized if more appropriate models and tools to design innovative regulation are being developed and their effects are properly understood. A few examples of innovative models, which address the challenges posed by unlicensed file sharing, Internet filtering, and malware.

- *StopBadWare* is a web platform, which aggregates users' judgments about malicious software code ("Badware") spread over the Internet and enables users to take more educated choices when downloading and installing software
- Its sister platform *Herdict* aggregates users' successful or unsuccessful attempts to access web sites and generates a global map of Internet filtering and blocking, which allows users to assess if a service outage is the result of technical difficulties or systematic filtering
- The *Digital Media Exchange* project addresses the challenges to copyright law created by web-based, unlicensed file-sharing networks by providing network providers with a model to offer their users legal access to media files and at the same time remunerating creators
- *Switzerland* provides users with a set of easy to operate tools to run tests on their Internet Service Providers to detect interference with certain applications, like VoIP or file-sharing tools

The issues addressed by these projects are inherent in the design decisions underlying the architecture of the Web as we know it today: application neutrality, openness, global reach, decentralization, anonymity.[9] The common approach of these projects is to address the resulting public policy challenges not by relying on law or changing the architecture of the Web, but by building tools and platforms, which leverage the power of distributed information systems without impacting the unique features of the Web.

These "policy interventions through code" are different from traditional legal interventions in several aspects:

- They do not rely on the power of a national legislator or the courts to regulate user behavior
- Instead they compete with other options for user acceptance and therefore need to actively find means of becoming effective
- They are not legitimized by an institutionalized democratic process. Instead they derive their legitimacy from being supported by a critical mass of users.

We will call this approach to designing and implementing public policy interventions *entrepreneurial regulation*, because the resulting interventions are driven by private actors, not the state or regulators, have to compete with other options for gaining user acceptance, and are not legitimized by the formal legislative or administrative process.

3. Case studies on entrepreneurial regulation

This paper is based on six one-hour long interviews with the regulatory entrepreneurs at the Berkman Center and the EFF, which allowed us to explore the strengths and limitations of the entrepreneurial approach to regulation and documented the principles underlying the design of the tools and platforms. We put the focus of this paper on how these projects addressed the challenge of becoming effective, i.e. how to provide benefits for the individual user ("micro effect"), which translate into the production of a desired social benefit ("macro effect"), and the closely related question of legitimacy: without the legitimizing support of the state, the interventions have to find other means to convince users that they are pursuing a worthwhile cause.

In the following we will present four case studies and investigate the means employed to gain effectiveness and legitimacy. The four projects are StopBadWare.org, Herdict.org, the Digital Media Exchange project, and Switzerland.

- StopBadWare / Herdict

StopBadWare (SBW) started as an online database with the URLs of web sites that contain malware. SBW achieved immediate effectiveness through an agreement with one of the major points of control[10] of the Internet: search engine Google: When SBW classifies a site as "Badware", Google displays a warning next to its search result entry. When a user nevertheless decides to click on the link, a page with details about the detected infection is displayed. Evidence shows that only few users decide to ignore the warning or "nudge" [11], and as a result traffic to the affected pages drops significantly. SBW operates with a full team and a number of partners in the industry to classify sites suspected of containing bad- or malware and to work with the system administrators of these sites to remove the malicious code.

[9] *See* IETF RFC 1958: Architectural Principles of the Internet.

[10] *See* J. Zittrain: Points of Control – A history of Online Gatekeeping *in* Harv. J. Law & Tec Vol. 19 (2006) 253.

[11] *See* Thaler and C. Sunstein: Nudge – Improving decisions about wealth, health, and happiness (2008).

SBW explicitly does not leverage the "wisdom of the crowd" to gather and process data. Ideas to let the community vote on what they consider malicious and what not are regularly discussed, but have not yet been implemented for fears of abuse. The cooperation between Google and SBW is beneficial for both sides: Google does not have to defend itself against allegations of discriminating against individual sites. SBW as an independent and not-for-profit initiative provides for a transparent review process and handles complaints as Google's "Appellate Court".[12] SBW benefits, because the partnership extend its reach to the millions of Google users, which helps SBW to achieve its stated goal to reduce the spread of malware across the Internet. SBW was started as the result of Jonathan Zittrain's work on generativity and his concerns that the spread of malware over the Internet would motivate people to abandon general purpose PCs and to switch to "locked-down" appliances to access the Internet. The project was started without previous demand analysis and the smart integration into Google's search engine has made SBW independent from the need to actively seek or generate user demand. On the other side, SBW now depends on the centrality of Google as the dominant provider of search results to become effective.

Alternatives to the SBW solution could include stricter liability rules for the creators and disseminators of malware. Since many infections originate in other countries, which do not necessarily have similar liability rules, a liability solution would be set up to fail. Stricter vicarious liability for system administrators to regularly screen servers for infection would make it easier to identify a responsible person, but again only for servers in a country with the respective liability rules. The liability rules could be easily avoided by moving servers to countries with less strict rules, Only few Web users will be able to determine where a server is located and moving servers will hence most likely go unnoticed by the majority of users. Another solution would be to see consumers move away from the open PC architecture to "locked-down" devices, which are immune against infections from the Internet, but may also restrict users from installing new applications. This trend is of concern to Zittrain and several others because it could severely hamper the spread of new applications on the Internet.[13] Finally, imposing an obligation on ISPs to protect users from malware could be an effective remedy, although several questions, in particular who decides what is malware and how to protect users from an abuse of this system to stiffle competition, remain unanswered.[14]

StopBadWare's sister project *Herdict* was also started by Zittrain at the Berkman Center as an extension of his work on Internet filtering with the OpenNet Initiative. Herdict's purpose is to map Internet filtering in different regions of the world. In constrast to SBW, Herdict leverages distributed information collection or crowdsourcing techniques to monitor the accessibility of web sites. The aggregation of these results allows Herdict to produce a map of blocked or filtered web sites by country and potentially even by Internet Service Provider. To facilitate the reporting of blocked web sites, contributors can volunteer to install a plug-in on their PCs web browser, which reports successful and unsuccessful attempts to access web sites to Herdict's servers. When a user finds that a website is inaccessible, the user can check if other users on his network have a similar experience, and draw conclusions if the web site is affected by technical problems or blocking by her Internet Service Provider.

User demand for the Herdict service had not formally been established before starting the project. Instead the team had been thinking about ways how to provide the users with a benefit, which would entice users to contribute and enable the project to reach its goal of mapping interference with Internet traffic. In this respect, the team relied on the "build it and they will come" approach and the hope that a community of users would find the Herdict platform useful or otherwise worthwhile to be supported. In designing the platform, the team could rely on the experience gathered during the work on the OpenNet initiative and tap into the pool of knowledge gathered StopBadWare. Like many other crowdsourcing projects, Herdict requires a critical number of contributors to become effective. At the time of this writing, the number of contributors outside of the US does not seem to have reached the critical mass with about 30 reported sites (both accessible and inaccessible) in China, Canada, and Spain in the past 30 days, and less in all other countries.[15] Internet filtering is a highly sensitive subject. While in some countries filtering or blocking is based on law, e.g. the laws against children pornography in the European Union, the majority of filtering happens because of higher security concerns and lacks both legal foundation and independent review, An independent review process therefore has to circumvent the institutions of the national regulators. In contrast to StopBadWare, Herdict's data gathering is decentralized and the initiative does not rely on a central node like Google to become effective. Herdict aims to attract individual users by providing with a tangible benefit, by lowering the effort required to contribute to the platform, and by stressing that users ("reporters") are contributing to a good cause.

[12] *See* http://newsfeedresearcher.com/data/articles_t6/google-site-search.html (last visited 25/02/2009).

[13] *See* J. Zittrain: The future of the Internet – and how to stop it. Allen Lane (2008).

[14] *See* Fifth Report from the House of Lords Science and Technology Committee Session 2006-07. HL Paper 165. Personal Internet Security *at* §3.41ff.

[15] Status: 24th of February/02/ 2009, the day before the official announcement of the launch of Herdict.org.

¶ Digital Media Exchange project

The Digital Media Exchange (DMX) is the offspring of a whitepaper authored by the Berkman Center's Digital Media Project.[16] It aims to provide ISPs with the tools and licenses to offer legal file sharing to their customers. The model is based on the blanket license approach suggested by William Fisher and Neil Netanel.[17] Instead of trying to change users current file sharing behavior, it aims to monetize file sharing. Every user pays a monthly content fee together with the normal subscription fee to her ISP. The tools provided by the DMX project monitor the popularity of songs, videos, and pictures on the file-sharing network and distribute the collected fees accordingly to the copyright holders, which have licensed their content to the non-profit coop operating the royalty collection and distribution mechanism. The benefits to the user are obvious: In exchange for a small monthly fee they gain protection from prosecution for copyright infringement. Copyright owners would benefit from the revenue generated by the so far not monetized file-sharing activity. The cannibalization effect on current record sales is still unclear. ISPs should be interested in the model to differentiate themselves from their competitors and to be better able to manage the file-sharing traffic on their networks. The cooperation with ISPs makes the DMX model gain immediate effectiveness with end users. But the project has to convince ISPs and copyright holders to accept the model. Attempts to implement the model in China and in the UK have (for now) not been successful for lack of support by copyright owners and ISPs.

The enforcement of copyright law against users of file sharing networks has proven to be expensive and little effective. As a result, the music industry has decided to change their strategy again and will now focus on collaborating with ISPs to throttle or terminate copyright infringer's Internet connection.[18] This solution raises many concerns about the proportionality of the termination and, in the age of ubiquitous Internet access, about the effectiveness.[19]

¶ Switzerland

The Electronic Frontier Foundation is a group of civil liberty advocates based in San Francisco. The EFF played a major role in detecting the blocking of the Bittorrent file-sharing protocol by the ISP Comcast. After the publication of their report, the EFF started receiving regular calls from concerned citizens, whose Internet connections were slow and who wondered if the EFF had information about mal-conduct of their ISP. As a result, the EFF started developing a suite of applications to detect interferences by Internet service providers for use by "normal users" as the tools used by the EFF technicians require substantial expertise. The name "Switzerland" alludes to the Swiss neutrality principle and the purpose of the tool to detect violations against network neutrality. The tools will enable users to run tests on their ISPs, and the EFF will assist users and developers of file sharing applications to analyze the results. Users benefit from the knowledge about their ISPs network management practices. Switzerland by itself does not solve the issues of discriminatory packet management, but it provides transparency about its existence, the methods used, and the extent to which it interferes with traffic. In the absence of competition between ISPs the results of the Switzerland tests do not enable consumers to make better choices. Instead, further action like bringing a case in the court of public opinion or traditional regulatory intervention is required to stop the detected discriminatory practices.

4. Findings

¶ Gaining effectiveness

We identified two main strategies of making the interventions effective: Providing users with compelling benefits that outweigh the costs of using the service or partnering with central gatekeepers of the Internet like ISPs or search portals. The benefits to the users are both tangible like the knowledge about an ISP's Quality of Service provided by the Switzerland project and intangible like the warm fuzzy feeling of supporting a good cause.

Surprisingly, none of the projects employed principles of participatory design in the solution design phase. Instead all relied on the "Build it and they will come" approach currently dominating in the "Web 2.0" arena. The majority of the projects were devised as a result of internal discussions at the Berkman Center, but not with prospective users or beneficiaries. The anticipated demand of users and the likely benefits were devised as the result of thorough analysis or experience. In this respect, the solution design followed the traditional top-down approach and relied on feedback

[16] *See* Copyright and Digital Media in a Post-Napster World *available at* http://cyber.law.harvard.edu/media/files/wp2005.pdf

[17] *See* W. Fisher: Promises to Keep. Technology, Law, and the Future of Entertainment, Stanford University Press (2004); *see* N. Netanel: Impose a Noncommercial Use Levy to Allow Free Peer-to-Peer File Sharing in Harv. J. Law & Tec Vol. 17 (2003).

[18] *See* Eliot E. Van Buskirk: RIAA to Stop Suing Music Fans, Cut Them Off Instead *in* WIRED December 19, 2008 available at http://blog.wired.com/business/2008/12/riaa-says-it-pl.html (last visited 25/02/2009).

[19] *For example the European Parliament, see Michael Geist: "Three Strikes and You're Out" Policy Strikes Out* http://www.michaelgeist.ca/content/view/2850/159/ *(last visited 25/02/2009See* A. Adams and I. Brown: Keep looking, the answer to the machine is elsewhere *in* Computer and Law 20 (1)).

loops for further refinement. The design of the three Berkman initiatives SWB, Herdict, and DMX rely on the academic work of Fisher and Zittrain, who are leading experts in their fields and have studied the subject matter of their interventions for many years. The knowledge involved in the EFF project stems from years of advocacy work on network management and the coding experience of the EFF's network analysts. Further research is required to identify comparable projects, which employ a more participatory solution approach, which becomes increasingly common in public policy design in the form of citizen round tables or open hearings, to evaluate differences in performance.

¶ Gaining legitimacy

Traditional forms of regulation like legislation or adjudication are governed by a formal process, which despite its admitted deficits in everyday practice fosters the involvement of relevant parties, transparency, due process, and provides for independent review by the courts. In the absence of these established institutions in Cyberspace, the interventions require a different form of legitimization. In this respect they face similar challenges like Social Enterprises innovating in areas like environmental protection to gain legitimacy.[20] Stressing the cultural validity of the social values they espouse, associating themselves with universal goods, acquiring a reputation for specialist knowledge and representing important constituencies are some of the tools employed by successful NGOs.

Projects like Switzerland and Herdict provide the users with information from which they are free to draw conclusions and adjust their behavior accordingly. But users can also decide to ignore the information provided and stop using the service. The critical mass of users required to become effective could serve as an indicator for the validity and importance of their cause. But support by a critical mass of users is not sufficient to establish legitimacy. The fact that a strong minority or even a silent majority in society supports recreational drug use or engages frequently in binge drinking does not by itself legitimize these activities. According to Jepson NGOs possess several legitimating assets. One is regulatory legitimacy, i.e. operating empowered by a law or on behalf of a regulator. Since the initiatives presented here don't have a public mandate, they to manage other cognitive assets to gain legitimacy. Normative legitimacy is the relationship with wider social or political ideas. One strategy we found at all four initiatives is the alignment of the initiative's brand identity with trusted academic institutions like Harvard University or strong advocacy brand names like the EFF. The non-profit character of the initiatives signals independence from profit-thinking and determined pursuit of the intended goal. Switzerland and Herdict position themselves as serving as watch-dogs of the public interest of maintaining the Internet as a free and open space of innovation. If the public opinion buys into these legitimating efforts depends on how skillfully the initiatives will be able to manage the cognitive assets of the public. Further research is needed to determine how successful they have been to position their interventions as legitimized.

Partnering with the central gatekeepers of the Internet makes SBW's and DMX's model gain immediate effectiveness. Although users are free to ignore the "Badware" warnings displayed by Google, the architecture of the warning makes most people follow them. This centralized form of enforcement raises more severe questions of legitimacy. It is interesting to observe that both projects have established transparent review and complaints processes, which allow those who are negatively affected, e.g. by the Badware warnings to file complaints. Additional voluntary commitments to transparency, e.g. to publish the revenue distribution key for the royalties collected by the DMX project or releasing the list with web sites that are affected by Badware warnings are designed to mirror the due process institutions governing a "traditional" regulator. Integrating democratic governance principles can be observed at the larger self-regulatory bodies of the Internet like ICANN, who employ another legitimizing asset: structural legitimacy.[21] Again, the non-profit nature of the initiatives and their brand alignment with an academic institution are targeted at providing them with the required level of legitimacy in public perception.

5. Conclusion

Common across all projects is the finding that an "intervention through code" is capable of achieving effects, which would not have been possible using the traditional tools of regulation. Either the required laws would prove ineffective because they are too easy to circumvent, or they would produce unintended side-effects, which could even run counter to the desired regulatory goal. This finding is a strong indicator that the approach taken by the regulatory entrepreneurs at the Berkman Center and the EFF could be a model for designing future regulatory regimes for the Web and hence deserves further investigation. Learning from successful Social Enterprises could mean designing better-adjusted and even more effective regulation for the Web to ensure its social benefit.

[20] *See* P. Jepson: Governance and accountability of environmental NGOs *in* Env. Science & Policy Vol. 8 (2005) 515 *at* 519.

[21] *See* R. Weber and M. Grosz: Legitimate Governing of the Internet (2008).