The Structural Problems of the Internet for Cultural Policy

Christian Sandvig (csandvig@uiuc.edu)
University of Illinois at Urbana-Champaign

Author's Draft
October 2003

Sandvig, C. (forthcoming). The Structural Problems of the Internet for Cultural Policy. In D. Silver (ed.). *Critical Cyberculture Studies: Current Terrains, Future Directions*. (working title). New York: New York University Press.

The Structural Problems of the Internet for Cultural Policy

ABSTRACT

The technical features and underlying structure of communication systems are not the usual starting point for those concerned about content: a system's "plumbing" seems irrelevant to many analyses of culture and cultural policy. However, this essay demonstrates that on the Internet the link between the two can be direct and important. The case of broadcasting on the Internet (meaning one-to-many communication, multicasting, and Web content caching) shows that obscure changes in low-level Internet protocols dramatically shape who may participate in cultural production. Current design features distribute costs so that users are assumed to be passive receivers and that producers of popular content must be moderately well-capitalized. As the scholarly literature on technology predicts, "technical" questions about how a particular function should be realized mask answers about who uses it, what they do, and who pays. The egalitarian ideal of new communication technologies is then lost or gained in a series of early implementation decisions that are debated solely in technical terms, despite their political character and cultural import. This situation argues for an "infrastructural cultural policy:" one where structural, technical decisions about the system require direct involvement by public interest advocates charged with giving voice to the voiceless.

The Structural Problems of the Internet for Cultural Policy

Christian Sandvig (csandvig@uiuc.edu), University of Illinois at Urbana-Champaign

When we are concerned about cultural production, we think chiefly of content. We search for more money for our artists; we seek the preservation of artifacts; we organize exhibits, outreach, and public education programs; and we encourage treatises that explain whichever fading heritage or tradition worries us. This essay argues that at the present moment, the Internet sorely needs the attention of established advocates for culture and expression. But not their concern for content.

At first blush, the Internet seems a boon to culture—it doesn't need our help; it is a great gift. From the outset the network appeared to realize Brand's mantra that "information wants to be free" (Brand 1987). Anyone with an Internet connection can both send and receive: the depressing trend that each new form of communication from antiquity to the present has allowed a smaller number of people to speak to a greater number is now finally reversed.[1] We are free from the awkward prefix "mass" that plagued "communication" since the introduction of broadcasting (Peters 1996). In some respects, the costs of cultural production have fallen and the hoped-for diversity of voices has materialized. Any Internet user can start a blog, make a Web page, share unpublished writing, or distribute music that would not otherwise be heard. We even have before us the tantalizing prospect of digitizing the public domain and offering it in an instant, for free.[2] The Internet seems, at first blush, to deliver the whole of human creativity to us, and to open avenues of expression that were formerly chained.

Yet, the same structures of the Internet that grant these new ways to speak also ensure that no one will ever hear you. This new medium is not the participatory turn in the production of culture that it seems to be. It is true that most of the Internet's users prefer a few familiar

sources of information – that is, the most popular .01% of the Web accounts for about 50% of all traffic. But even more important is that we can't write a ratio that is very different from .01:50 without a very different Internet.[3] We will never hear those prophesied new voices, exercising their digital freedoms with new contributions to culture because the Internet is designed to keep them silent. Those who are concerned about cultural production are powerless, because by and large they cannot see inside the Internet's design. This is the problem that we must address, as this essay will explain by example.

A Problem: The Slashdot Effect

Let us consider a specific new voice. In a corner of the Web, nerds and geeks banded together to create what they thought might be a new form of journalism: the self-reported electronic newspaper (Baoill 2000). The newspaper existed before the professional reporter, they might have reasoned, and when reporting "news for nerds" (their motto) only the real nerds can be trusted to get the details just right.[4] They started *Slashdot* to report "news for nerds, stuff that matters;" a news service about technology, but a news service run like a bulletin board, where the reporters were readers and the readers were editors. Contributors to *Slashdot* receive "points" from other readers for having interesting things to say. First-hand accounts and diverse sources were encouraged. *Slashdot*'s users trawled the Web seeking news for nerds, but as the site increased in popularity, they encountered an unexpected problem.

When a *Slashdot* reader/contributor stumbled upon a juicy photo, story, or comment tucked away on a personal Web server (a bona fide unconventional news source) and then shared his prize on *Slashdot*, the clicks of *Slashdot* readers would overwhelm the bandwidth available to their target. That is, the act of promoting unusual content to even *Slashdot*'s modest audience caused that content to become instantly inaccessible. This phenomenon is common enough that they named it "the Slashdot effect."[5] The Web's eye, in the act of looking, destroys

the object of its gaze. But this doesn't happen for the more traditional Web destinations that we visit every day.

Two Solutions to "Delivering a Better Internet"

The New York Times never suffers from the Slashdot effect, partly because when you click on a link to nytimes.com, your bits don't really come from New York, as will be explained. The Slashdot effect cripples Web servers because of the way that communication from one to many—broadcasting—is implemented on the Internet. (Or rather, the way that it isn't.) In the terms of the Internet's original architecture, when five computers in the same house request the home page of The New York Times, five identical copies of the page are sent all the way from the Web server, wherever in the world that may be, to each machine. Since we are dealing with electrons and not broadsheet, it would make more sense to send just one copy. Copying the electrons is cheap while the bandwidth across the world is expensive, and so instead of sending five duplicates around the world a single copy could be quickly duplicated as close to the house as possible. Sending all five copies the whole distance only congests long Internet pipelines with five times more traffic than necessary.

With the New York Times homepage this effect might be trivial, but with the large files required by streaming multimedia the consequences are profound—today's Internet sends two (or ten, or a hundred) copies when one will do because the way it delivers traffic is typically not sophisticated enough to realize that each identical request is related. This is the reason why multimedia streaming on the Internet does not work very well and also the explanation of the Slashdot effect: the network near a source of very popular content becomes overwhelmed as duplicate requests proliferate.

So far, this sounds like an arcane technical problem: at best a tangent to cultural production. But for this problem there are solutions, and then there are solutions. One

approach proposes to modify the Internet's basic protocols to reduce the duplicate transmission of multiple streams. This solution is called "multicasting," and it is being advanced and refined in an open deliberative process in the standards bodies of the Internet world. These proposed changes to the Internet's protocols have not yet succeeded and the multicast backbone, or "MBONE," remains experimental (Eriksson 1994).[6] A second approach is called "content caching."[7] Like a guarded cache of pirate treasure, this approach involves employing a third party to store and copy your traffic at some intermediate point between source and destination, but as close as possible to the people that want it. The private company *Akamai* is the overwhelming leader in the obscure content caching market.[8] Content caching systems are both proprietary and expensive.

When you request a web page from the New York Times, *Akamai* (the Times' content caching supplier) intercepts your request. Your rough geographical location is traced—this is called geolocation. The page is then dispatched to you from a high-capacity *Akamai* data center as close to your computer as possible. The technology involved is a trade secret and this service is available only to *Akamai* subscribers (the *Akamai* motto: "Delivering a better Internet"). When *Slashdot* readers re-post links to mainstream news from the New York Times, these links never suffer from the Slashdot effect.

The Capital Requirements for Cultural Products to be Popular

In the example above, the most important difference between the proprietary, private solution to our arcane technical problem and an open deliberative solution is simple: who pays? Multicasting is a collective solution deep in the guts of the infrastructure: it requires each Internet user to pay for the cost of delivering content from anywhere to anywhere, though with flat-fee Internet pricing the user will never know it. Streaming and broadcasting from anywhere would simply *work*. Content caching—the solution we have today—requires the

provider of content to pay for an expensive add-on service, costs that they must recoup through advertising, subscription, or some other source of revenue. Streaming and broadcasting simply works, but only from the New York Times.

Therefore unknown to most, the present Internet requires those who produce popular content to be well-capitalized. You pay for your own Web hosting, and at anything other than small rates of traffic you pay for the bandwidth your visitors generate. The infrastructure for popularity is available, but it is expensive. The difference between these two solutions is nothing less than the decision between an Internet where only capitalized producers of culture can be popular and one where anyone can.

Even leaving diversity of expression aside, other benefits weigh on the side of multicasting. Changing the Internet's fundamental protocols is now accomplished through a public process negotiated in International standards bodies, and the resulting solution will be published and available for free to anyone, as is the specification for the Internet's protocols.[9] This means simply that challenging "technical" problems for the production of culture, or any human value, can be unearthed in these public documents and addressed. In contrast, proprietary solutions like content caching may function in ways that raise privacy concerns (as geolocation might) but we must guess at how these systems function by snooping: peeking at traffic, interpreting sales brochures, and reading annual reports of the companies involved.[10]

This does not mean we should value open processes in and for themselves. There is no reason to assume that an open solution to a knotty problem of Internet architecture will be more likely to lead to freedom (or any positive normative value) in the long run than a secret solution. In fact, the present dilemma of multicasting vs. content caching presents us with a case where we have an open process and this has produced a flawed solution. That is, our

Internet architecture debates are currently open, but multicasting has not (yet) been incorporated into the Internet protocols.

Why Not the MBONE? General Problems for the Public Interest

The preceding summary may suggest a conspiracy to the conspiracy-minded. If one solution (multicasting) is so clearly beneficial, why isn't it yet the standard? One set of answers is clearly related to process: even if we prefer open deliberation and public results, open deliberation about complicated issues is hard work. As each new group of interested parties joins the debate about the future of the Internet, consensus becomes even more elusive. In addition, the solution becomes ever more complex to accommodate each new set of interests. Worse, even the long and contentious debates we have today may involve the wrong people. No group of advocates has stepped forward to agitate that multicasting deployment is crucial for freedom of expression; this group of advocates does not exist because those that concern themselves with freedom of expression typically don't participate in Internet standards debates.

The second set of answers is clearly related to incentives: Some multicast functionality has been available in the off-the-shelf equipment used by your Internet Service Provider for the last few years, but your provider likely has no interest in using these features. Simply put, "Receivers do not care whether they receive their…streams from unicast or multicast" (Diot, Levine, Lyles, Kassem, and Balensiefen 2000: 81). It doesn't matter to you how the New York Times arrives on your screen. It *does* matter to the sender who wishes to provide multimedia content and has no access to the New York Times' expensive private content caching supplier. And it does matter to society that successful Internet speech carries a requirement that the speaker be well-capitalized, but these concerns are not immediately felt by the users that click on Web links to the Times. So to implement a collective solution to broadcasting on the Internet that would let anyone be a sender, we would need to secure the cooperation of the Internet

Service Providers that gain their revenue from receivers: the people that click on these links. Implementing multicasting is an additional cost for your provider, and yet there will be no demand from you.[11]

If we had the answers to these normative problems of Internet architecture in our pocket, it would still be unclear what exactly we should do with them. These problems of process and incentive arise in a context where the Internet is thought to be free of anyone's control. In reality, the Internet is at most an uneven anarchy—aside from pockets of ungovernability the inability of governments to regulate cyberspace has been greatly exaggerated. For instance, the Domain Name System is effectively under centralized control and this control is a direct delegation of authority from the US government, although the US government hopes that this situation is perceived as international cooperation (see Froomkin 2000).[12] Yet no such governmental relationship exists with Internet architecture, where a jumble of overlapping standards bodies continue to operate under David Clark's credo that, "We don't believe in kings, presidents, or voting. We believe in rough consensus and running code."[13]

While this essay has focused on this problem of *broadcasting* on the Internet in order to illustrate the structural problems of the Internet for cultural policy, it is important to interject that this problem for freedom of expression is only one problem from an infinite series. Even in the realm of content distribution there are many more: widespread deployment of filters and caches in firewalls by Internet Service Providers have other worrying structural effects on content.[14] Problematic biases also exist in search and directory services (Introna and Nissenbaum 2000; Rogers 2000). The phenomenal success of *Google's* revolutionary *PageRank* algorithm makes it much easier to locate the most popular Web content and much harder to locate unpopular content that uses the same words as popular content. Preferred placement services on search engines and portals also reward capitalized content providers. Concern

about the Internet's structure must also encompass complicated overlaps between the legal and the technical.[15] The lack of any legal nondiscrimination requirement for Internet Service Providers (unlike, say, telephone companies), means that they can turn away content they disagree with. The Internet's vulnerability to some very specific forms of hacking combine with the nondiscrimination problem to make unpopular content a pariah: Aljazeera, the controversial media network of the Arab world, initially could not find a hosting provider for its English-language Internet site in part because this site would attract hackers to any provider that hosted it.[16] The legal and technical conspire to make unpopular views a "poisoned chalice" for skittish Internet Service Providers.[17] This essay, then, is not meant highlight the specific problem of broadcasting on the Internet, but to raise the problem of a continuing series of technical decisions that need attention from those concerned about the public interest and the role of communication systems in society.

Pragmatic Steps toward the Techno-Socio-Legal

After outlining a number of obstacles that presently exist between us and the Internet we hope for, we must turn to strategy. Yes, the situation is difficult, but not hopeless. If we can articulate our normative goals for the Internet's development, a course of action suggests itself. Our situation is far from unprecedented, and indeed not a surprise. As the scholarly literature on technology predicts, "technical" questions about how a particular function should be realized in technology mask assumptions, interests, and political bargains.[18] The egalitarian potential of new communication technologies like the Internet is then lost or gained in a series of early implementation decisions that are debated solely in technical terms, despite their political character and cultural import.

It is true that we are presently faced with a messy, semi-anarchic Internet that we want to change (but we aren't sure how). But the birth of the Internet was just as messy, and the

early stages of other communication technologies may have been even be messier.[19] Even those who pine for the more tranquil framework of telecommunications are remembering an imagined past.

First, the easy answer. To advance normative goals such as freedom of expression, a straightforward strategy is to continue the successful government policy of funding applied networking research projects. The same sort of projects that produced the Internet provide a development environment that (while not neutral) at least develops standards and software under pressures that are orthogonal to the insistent and competing factions of commerce. The keystone of this policy was public ownership of the resulting software and standards, and this needs to continue. Research investments in the Internet's protocols and software should be a priority for technology policy in the US and elsewhere.

Second, and most important, this situation argues for an "infrastructural cultural policy:" one where structural, technical decisions about the development of society's communication system require direct involvement by public interest advocates charged to give voice to the voiceless. To succeed this requires an awkward combination of technical, social, and legal expertise. It requires skilled engineers whose technical acumen is matched by an understanding of the place of technology in society, and the normative issues raised by engineering work. It requires scholars who study the political economy of communication systems but who their research agendas to better embrace the minefield of the technical. The best current researchers have recognized this need and are moving toward a techno-socio-legal convergence, from both inside and outside of engineering (e.g., Clark, Wroclawski, Sollins, and Braden 2002; Shah and Kesan 2003).

Institutional structures still need to change course to encourage the fusion of this necessary triad. Foundation programs that deliver more money for our artists, the preservation

of cultural artifacts, and public education programs need to be reconsidered to include the structural problems of the Internet. A successful program of giving to change these Internet fundamentals will provide a ten- or hundred-fold return on philanthropic investment when compared to more traditional, narrowly-defined giving for the purpose of cultural diversity. Similarly, re-thinking educational programs as fusions of what is now found in communication, science and technology studies, law, and computer science will directly prepare a new generation of student to navigate the entangled current landscape of communication technology. Finally, academic institutions and national research councils need to encourage and recognize inter- and multidisciplinary work—not as a general good, but as collaboration across this specific gulf to face these problems.

Toward an Infrastructural Cultural Policy

The future of broadcasting the Internet may yet be resolved in favor of diversity in cultural production: we can hope for the smooth introduction and adoption of multicasting in the Internet's core. This will render *Akamai's* private pay-to-speak services unnecessary, and Internet popularity will become affordable. Still, the resolution of this example is no redress for the larger problem for cultural policy: a communication system's "plumbing" presently seems irrelevant to those concerned about cultural production, and this must change.

It is true that attention to these problems requires an unusual combination of expertise: they beg those who care about culture, society, and the law to care about technology at a level of detail where few outside of computer labs are comfortable. But if the history of other media is any guide, the structural decisions made in the early decades of the Internet—these decisions, made these days—have the potential to endure for years to come. After these decisions are made, it will be very costly to change our mind once we realize we are unhappy with the Internet we have built. To usefully care about the freedom of expression and the production of

culture on the Internet, we must care about the "plumbing." While an attention to infrastructural cultural policy may be a lot to ask, it is the least that is required. The unsettled character of today's advanced communication systems is not our burden, it is our chance to act.

ACKNOWLEDGEMENT

This research was kindly supported by a visiting fellowship from the Oxford Internet Institute at Oxford University. The author would like to thank Rajiv Shah for his helpful comments on an earlier draft of this paper.

SOURCES CITED

Abbate, J. (1999) *Inventing the Internet*, Cambridge: MIT Press.

Almeroth, K. C. (2000) 'The Evolution of multicast: From the MBONE to Inter-Domain Multicast to Internet2 Deployment' *IEEE Network*, 14.1: 10-20.

Baoill, A. Ó. (2000) 'Slashdot and the Public Sphere', *First Monday*, 5.9.

Barabási, A., Albert R., and Jeong H. (2002) 'Scale-free characteristics of random networks.' *Physica A*, 281: 69-77.

Bijker, W. E., Hughes, T. P. and Pinch, T. J. (eds) (1987) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology,* Cambridge: MIT Press.

Brand, S. (1987) *The Media Lab: Inventing the Future at MIT*, New York: Viking.

Clark, D. D., Wroclawski, J., Sollins, K., and Braden, R. (2002) 'Tussle in Cyberspace: Defining Tomorrow's Internet', *Proc. ACM SIGCOMM*, 222.

David, P. A. (2001) 'The Evolving Accidental Information Super-Highway', *Oxford Review of Economic Policy*, 17.2.

Diot, C., Levine, B. N., Lyles, B., Kassem, H., and Balensiefen, D. (2000) 'Deployment Issues for the IP Multicast Service and Architecture', *IEEE Network*, 14.1: 78-88.

Eriksson, H. (1994) 'MBONE: The Multicast Backbone', *Communications of the ACM*, 37.8: 54-60.

Froomkin, M. (2000) 'Wrong turn in cyberspace: Using ICANN to route around the APA and the Constitution,' *Duke Law Journal*, 50.17: 17-184.

Hindman, M., Tsioutsiouliklis, K. and Johnson J. A. (2003, July 28) '"Googlearchy:" How a few heavily-linked sites dominate politics on the Web,' unpublished manuscript, Princeton University. Available HTTP: <http://www.princeton.edu/~mhindman/googlearchy--hindman.pdf> (accessed 5 Oct 2003).

Information Technology Association of America (2001) *Alexa Research Finds the 50 Leading Sites Account for More Than 25% of Traffic*, Arlington: ITAA. Available HTTP: <http://www.itaa.org/isec/pubs/e20013-13.pdf> (accessed 6 Oct 2003).

Innis, H. A. (1964) *The Bias of Communication*, Toronto: University of Toronto Press.

Introna, L. D., and Nissenbaum, H. (2000) 'Shaping the Web: Why the Politics of Search Engines Matters', *The Information Society,* 16.3: 169-85.

Lettice, J. (2003, March 27) 'Al Jazeera's Web Site: Ddosed or Unplugged?' *The Register,* London.

McChesney, R. (1993) *Telecommunications, Mass Media, and Democracy: The Battle for the Control of US Broadcasting, 1928-1935*. New York: Oxford University Press.

Online Computer Library Center. *Web Characterization Project*, Dublin, Ohio: OCLC. Available HTTP: <http://wcp.oclc.org/> (accessed 7 Oct 2003).

Peters, J. D. (1996) 'The Uncanniness of Mass Communication in Interwar Social Thought', *Journal of Communication*, 46.3: 108-23.

Rogers, R. (ed.) (2000) *Preferred Placement: Knowledge Politics on the Web,* Maastricht, The Netherlands: Jan van Eyck Akademie.

Sandvig, C. (forthcoming) 'Shaping Communication Infrastructure and Innovation: The End-to-End Network that Isn't', in D. Guston and D. Sarewitz (eds) *Shaping Science and Technology Policy: The Next Generation of Research*, Madison, University of Wisconsin Press.

Schudson, M. (1981). *Discovering the News: A Social History of American Newspapers*. New York: Basic Books.

Shah, R. and Kesan, J. (2003) 'Manipulating the governance characteristics of code', *Info* 5.4: 3-9.

Vichare, R. (2002) 'Content Caching Vendor Market Share', research report, International Data Corporation.

Winner, L. (1980) 'Do Artifacts Have Politics?' *Daedalus,* 109.1: 121-35.

NOTES

---

[1] This compelling though debatable trend was expressed by Innis (1964).
[2] Similar to the promise now offered by *Project Gutenberg* after the concept of charging for content was dropped. See: http://gutenberg.net/
[3] Statistics on the concentration of Web use among a few sites are notoriously unstable. This estimate is based on data from 2001, combined from market research and academic data (Information Technology Association of America 2001; Online Computer Library Center 2001). While the numbers given in other estimates vary, the general trend of concentrated attention is clear across a wide variety of studies and methods (e.g., Barabási, Albert, and Jeong 2002; Hindman, Tsioutsiouliklis, and Johnson 2003).
[4] For the history of newspapers, objectivity, and reporters, see Schudson (1981).
[5] An excellent overview of the Slashdot effect can be found in the Wikipedia. See: http://www.wikipedia.org/wiki/Slashdot_effect
[6] For a history and overview of multicasting, see Almeroth (2000).
[7] The content caching market was worth about $430 million in 2001, see Vichare (2002). Participants include *Akamai, Inktomi,* and *Cable & Wireless*.
[8] See: http://www.akamai.com/
[9] This openness at the basic level of architecture is an often-overlooked benefit that came from the Internet's origins as a government project. The original contractor, BBN, was ordered by DARPA to make its technical specifications freely available, and the home of the early Internet in the academic culture of computer science departments promoted an openness that led to the present public "Request for Comments" system that explains the core of the Internet. For more, see Abbate (1999).
[10] Closed systems also raise technical concerns. For example, how do engineers design and plan the future Internet if they cannot determine how applications and protocols will behave? However, in this section I mean to highlight the problems for human values.
[11] This is not to say that multicasting carries a fee to your provider, but only that the configuration and support for a new feature like multicasting entails some effort. For a discussion of this infrastructure migration problem in the context of innovation, see David (2001) or more generally the economic literature on network externalities.
[12] This comment does not mean to endorse the governance of the Domain Name System, but simply to point out that control is centralized and effectively under the control of one government.
[13] David Clark is now a Senior Research Scientist at the MIT Computer Science and Artificial Intelligence Laboratory and he is an undisputed leader of the development of the Internet's architecture. Several versions of this comment have been attributed to him at various times—most famously in an address to the Internet Engineering Task Force.
[14] This argument has been framed by others in terms of the end-to-end argument in Internet system design, however the end-to-end argument is a technical gloss obscuring more the familiar debates about mediation that occur in all communication systems, as I argue elsewhere (Sandvig forthcoming).
[15] Although to be fair, *every* example in this essay can be conceptualized as a complicated overlap between the legal and the technical. The legal components of a problem are, however, less obvious for some of the more arcane examples.
[16] For instance, the distributed denial of service (DDoS) attack.
[17] The phrase "poisoned chalice" has been used in exactly this way in the trade press catering to Internet Service Providers (Lettice 2003).
[18] See, e.g., Winner (1980), Bijker, Hughes, and Pinch (1987).
[19] See, e.g., McChesney (1993).