

Not *The* Internet, but *This* Internet: How Othernets Illuminate Our Feudal Internet

Paul Dourish

Department of Informatics
University of California, Irvine
Irvine, CA 92697-3440, USA
jpd@ics.uci.edu

ABSTRACT

What is the Internet like, and how do we know? Less tendentiously, how can we make general statements about the Internet without reference to alternatives that help us to understand what the space of network design possibilities might be? This paper presents a series of cases of network alternatives which provide a vantage point from which to reflect upon the ways that the Internet does or does not uphold both its own design goals and our collective imaginings of what it does and how. The goal is to provide a framework for understanding how technologies embody promises, and how these both come to evolve.

Author Keywords

Network protocols; network topology; naming routing; media infrastructures.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

What does it mean to say, as many have, that “the Internet treats censorship as damage and routes around it” [14]? Or what does it mean to argue, as is also common, that the Internet is a democratizing force as a result of its decentralized architecture [8]? Or that it’s a platform for grass-roots community-building rather than mass media messaging [35]?

Statements like these have been critiqued for their treatment of topics such as democracy, censorship, broadcasting, or community – all complicated terms that are perhaps invoked and conjured more than they are critically examined [19]. However, the term that I want to focus on in these statements is “the Internet”. When we attribute characteristics to the Internet, specifically, what do we mean? Do we just mean “digital networks”? Do we mean digital networks that implement the Internet protocols? Or do we mean the one very specific network that we have

built – the Internet, this Internet, our Internet, the one to which I’m connected right now?

I ask these questions in the context of a burgeoning recent interest in examining digital technologies as materially, socially, historically and geographically specific [e.g. 13, 15, 36, 37]. There is no denying the central role that “the digital,” broadly construed, plays as part of contemporary everyday life. Wireless connectivity, broadband communications, and computational devices may be concentrated in the urban centers of economically privileged nations, but even in the most “remote” corners of the globe, much of everyday life is structured, organized, and governed by databases and algorithms, and “the digital” still operates even in the central fact of its occasional absence, the gap that renders particular spots “off grid.” Where digital technologies were once objects of primarily engineering attention, their pervasive presence has meant that other disciplines – anthropology, political science, communication, sociology, history, cultural studies, and more – have had to grapple with the question, “what are the cultural consequences of ‘the digital’?” The problem, however, has been to get a grip up on what ‘the digital’ itself is. Rather too often, ‘the digital’ is taken simply as a metaphor for regimentation and control, or it is used to name some amorphous and unexamined constellation of representational and computational practices. The somewhat casual assertion that “the Internet” has some particular properties runs this danger. It’s difficult to know how to read or assess these assertions without, for example, understanding something of the scope of what is being claimed through some kind of differential analysis of “the not-Internet.” We need to take an approach to “the Internet” that begins with an understanding of its technical reality, although not one that reduces it to simply that. In other words, we want to maintain a focus on the social and cultural reality of technologies such as Internet, but in a way that takes seriously its material specificities.

Two Transitions

To make this argument more concrete, let me begin by describing two cases from my own working history, two digital transitions that illustrate the complexity of talking about the Internet as a force for decentralization.

The first of these transitions occurred in the early 1990s when I worked for Xerox. Xerox had been a pioneer in

Copyright© 2015 is held by the author(s). Publication rights licensed to Aarhus University and ACM

5th Decennial Aarhus Conference on Critical Alternatives
August 17 – 21, 2015, Aarhus Denmark

DOI: <http://dx.doi.org/10.7146/aahcc.v1i1.21200>

digital networking. Early research on Ethernet and distributed systems constituted an important precursor to the development of the Internet Protocol (IP) suite, and the research systems such as PUP [4] and Grapevine [3] had subsequently given rise to a protocol suite called XNS (Xerox Network Services), which was the basis of a line of digital office systems products that Xerox sold in the marketplace. Xerox's own corporate network spanned the globe, linking thousands of workstations and servers together using the XNS protocols. In the 1990s, as Unix workstations began to dominate the professional workstation market, and as the arrival of distributed information services such as Gopher, WAIS, and WWW spurred the accelerated growth of the Internet, many inside Xerox became interested in using TCP/IP on internal networks too. Particular at research and development centers, various groups began to run TCP/IP networks locally and then increasingly looked for ways to connect them together using the same leased lines that carried XNS traffic between sites. What began as renegade or illicit actions slowly became organizationally known and tolerated, and then organizationally supported as TCP/IP became a recognized aspect of internal corporate networking within Xerox.

XNS was a protocol suite designed for corporate environments. Although technically decentralized, it depended on an administratively centralized or managed model. The effective use of XNS was tied together by a distributed database service known as the Clearinghouse, which was responsible for device naming, address resolution, user authentication, access control, and related functions. Users, servers, workstations, printers, email lists, organizational units, and other network-relevant objects were all registered in the Clearinghouse, which was implemented as a distributed network of database servers linked via a so-called "epidemic" algorithm by which they would keep their database records up to date. Access control mechanisms distinguished administrators, who could update Clearinghouse databases, from regular users, who could look up names but couldn't introduce new ones. The Clearinghouse service was central enough to the operation of XNS services that this administrative access was needed for all sorts of operations, from adding new users to installing new workstations.

By contrast, the TCP/IP network, and the Unix workstations that it generally linked, was much less centrally administered. For the Unix machines, users could be defined locally for each computer, and similarly, workstations could maintain their own machine addressing and network routing information. Even when systems were interconnected, much less coordination was required to get machines connected together effectiveness in the IP network than in the XNS network. As a result, the rise of the IP network provided a route by which people could to some extent become more independent of the corporate IT management structure through which the XNS network was

operated. Since XNS was the dominant technology for organizational communication, it wasn't entirely possible for people using TCP/IP to "route around" the corporate network, but it started to provide some independence.

The second transition was also a transition to IP, but in a very different context. This transition was going on while I briefly worked at Apple in the late 1990s. As at Xerox, the rise of the Internet in general was reflected in the increasing use of TCP/IP in a network that had originally been put together through a different network protocol – in this case, AppleTalk. AppleTalk was a proprietary network suite that Apple developed to connect Macintosh computers; it had evolved over time to operate over the Ethernet networks commonly deployed in corporate settings, although it had originally been developed for linking computers together in relatively small networks. One important feature of the Appletalk networking protocols is their so-called "plug-and-play" approach, which allows a network to be deployed with minimal manual configuration. For example, Appletalk does not require that network addresses be pre-assigned or that a server be available for network resource discovery; these features are managed directly by the networked computers themselves. Accordingly, setting up Appletalk networks requires little or no administrative intervention. TCP/IP networks, on the other hand, do require some services to be set up – DHCP servers to allocate addresses, name servers to resolve network addresses, and so on. (In fact, the contemporary networking technologies known as Bonjour or Zeroconf are mechanisms designed to re-introduce Appletalk's plug-and-play functionality into TCP/IP networking.) So, where the transition from XNS to TCP/IP was a decentralizing transition at Xerox, one that increased people's independence from corporate network management, the transition from Appletalk to TCP/IP at Apple moved in the other direction, creating more reliance on network infrastructure and hence on network administration.

These examples illustrate two important concerns that animate this paper. The first is that statements about "the Internet" and its political and institutional character suffer for a lack of contrast classes ("critical alternatives," if you will). The Internet may well be decentralized – but compared to what, exactly? More decentralized internetworks could be imagined and have existed. We might be able to make some more fruitful observations if the "Internet" that we are trying to characterize weren't such a singular phenomenon. In this paper I will briefly sketch some cases that might serve as points of comparison that provide for more specific statements about contemporary phenomena by showing how they might be, and have been, otherwise. The first concern, then, is to provide a framework within which the characterizations can take on more meaning. The second concern is that the singular nature of the Internet makes it hard for us to distinguish the conceptual object of our attention, the object that we are characterizing. Given that the Internet – the

specific Internet to which we can buy or obtain access today – is generally the only Internet we have known, it’s hard to be able to pin down just what object it is that we are characterizing when we talk about, say, decentralization. Do we mean that a world-spanning network-of-networks is inherently decentralized? Or is decentralization a characteristic of the specific protocols and software that we might use to operate that network (the Internet protocols)? Or is it rather a characteristic of the specific network that we have build, which doesn’t just use those protocols, but implements them in a specific network made of particular connections, an amalgam of undersea fiber-optic cables, domestic WiFi connections, commercial service providers, and so on? Is it our Internet that’s decentralized, while we could still imagine a centralized one being built? Or is our Internet actually less decentralized than it might be, failing to achieve its own promise (if that’s something we want)?

In order to provide the resources to think about these questions fruitfully, I will approach the topic from two perspectives. The first is to briefly catalog some alternatives to “the” Internet. Some of these are entirely alternative networks; some are small components of the broader Internet that do not always operate the same way as the whole. The second is to take in turn some key aspects of network function – routing, naming, and so on – and examine their contemporary specificities, with particular focus on the relationship between specific commercial and technical arrangements and the openness or range of possibilities encapsulated by the network design. Taken together, these allow us to develop an understanding of the landscape of potential network arrangements within which our current arrangements take their place, and perhaps more accurately then target or assess statements about what “the Internet” is or does.

A CATALOG OF OTHERNETS

Our starting point are what I have been calling “othernets” – non-Internet internets, if you will. Some of these are networks of a similar style but which happen to use different protocols; some of them are radically different arrangements. Some of them are global networks, and some more localized; some are specific networks and some are ways of thinking about or approaching network design. The collection listed here is far from comprehensive.¹ What they do for us here, though, is to flesh out an internetworking space of possibilities, one that helps to place “the Internet” in some context.

Fidonet

A Bulletin Board System (BBS) hosts messages and discussions, generally on a quite simple technology platform such as a conventional home PC with a modem that allows people to call into it over regular phone lines to

read and post messages. In the US, where local phone calls were generally free, BBSs flourish in the late 1970s and early 1980s, as the home computer market grew. With very simple software, they allowed people to communicate by dialing in to the BBS at different times and posting messages that others would read later. While one certainly could make a long-distance call to connect to a BBS, most BBS use was local to take advantage of toll-free calling, so that most BBS activity was highly regionalized.

Fido was the name of BBS software first written by Tom Jennings in San Francisco in 1984 and then adopted by others elsewhere in the US. Before long, Fido was updated with code that would allow different Fido BBS systems to call each other to exchange messages; Fidonet is the name both of this software and of the network of BBS systems that exchanged messages through this mechanism. Fidonet’s growth was explosive; from its start in 1985, Fidonet had around 500 nodes by the end of 1985, almost 2,000 by 1987, 12,000 by 1991 and over 35,000 by 1995. Each of these nodes was a BBS that server ten to hundreds of users, who could exchange email messages, files, and discussions on group lists.

Fidonet’s design was (with one critical exception) radically decentralized. Based as it was on a dial-up model rather than an infrastructure of fixed connections, it employed a model of direct, peer-to-peer communication. The Fidonet software was originally designed with a flat list of up to 250 nodes; system of regions, zones, and networks was introduced within a year of the original software when it became clear that the system would very soon grow beyond that capacity. This structure, which mapped the topology of the network geographically, provided a message routing structure which reduced costs (by maximizing local calling and by pooling messages for long-distance transfer) but with a minimum of fixed structure; direct communication between nodes was always central to the Fidonet model. The structure essentially exhibited a two-level architecture; one of conventional structures (that is, the conventional pattern of daily or weekly connections between sites) and an immediate structure, made up of those nodes communicating with each other right now. (The Internet – being made up in its current form primarily of fixed infrastructure and broadband connectivity – largely conflates these two.)

Within a year, Fidonet was in in danger of hitting an initial limit of 250 nodes, and so the network protocols and software were redesigned around the concepts of regions and nets. Until this point, Fidonet had used a single, “flat” list of nodes, which was directly maintained by Jennings. The introduction of regions and nets allowed for a more decentralized structure. This was simultaneously an addressing structure and a routing structure, linked together – the structure of the network was also the structure that determined how messages would make their way from one place to another.

¹ Most problematically, local and community “DIY” networks usefully illuminate the issues in question but had to be omitted for space [e.g. 1, 16, 21, 30].

Fidonet was built around a file transfer mechanism that would allow files to move from one node to another. Other facilities could be built on top of this mechanism, such as electronic mail. One of the major features of Fidonet from the user perspective was the discussion groups known as “echoes”. Echoes allowed users to post messages that would be distributed to all users on the system interested in a topic. A “moderation” mechanism allowed echo managers to discourage off-topic posts, but this was a post-hoc mechanism (rather than an approach which required explicit approval before a message was sent out). As in systems such as The Well [39], the topically-oriented discussion forums provided by echoes were the primary site of interaction and community building across Fidonet (although unlike The Well, Fido echoes were networked rather than centralized.)

Usenet

Usenet is a somewhat informal term for a worldwide network of computers linked by a series of mechanisms built on top of a facility provided by versions of the Unix operating system [18]. The facility was known as uucp, which stands for “Unix-to-Unix copy.” In Unix’s command-line environment, “uucp” was a command that users could use to copy files between computers. It was designed by analogy with the standard “cp” command for copying files; just as a user might use “cp” to copy a file from a source to a destination filename, they might also use “uucp” to copy a file from a source to a destination, where either the source or destination file location was in fact on a different computer.

Uucp was developed as a basic user file-copy facility, but the core idea of file-based interconnections between minicomputers was general enough that many other facilities could be built on top of it. For instance, uucp could be used to support email messaging between sites, exchanging individual messages as files that the remote system would recognize as drafts to be delivered by email locally. The same mechanisms that named remote files, then, might also be used to name remote users.

Since it had been designed initially simply to provide a user interface to a dial-up mechanism for file exchange, Usenet provided no global naming mechanism to identify sites, files, objects, or users. Rather, its naming mechanism was a sequence of identifiers (separated by exclamation points or “bangs”) that explained how a message should be routed. So, for instance, the path “seismo!mcvax!ukc!itspna!jpd” directs a computer to deliver the file first to a computer called “seismo”, at which point the path will be rewritten to “mcvax!ukc!itspna!jpd”; subsequently, it will be delivered to a computer called “mcvax”, then to one called “ukc,” and so on. “jpd” is the user, whose account is on the computer “itspna”. To send a message correctly, then, required that one knew not only the destination, but the route that the message should take – the series of peer-to-peer connections that must be made. Each path describes a direct

connection; our example bang path only works if “mcvax” is one of the computers that “seismo” regularly connects to directly. One couldn’t, for instance, route a message along the path “seismo!ukc!itspna!jpd” because seismo only dials up to certain other computers, and ukc isn’t one of them.

Two aspects of this are worth noting. The first concerns the dynamics of network structure in the presence of this route-based addressing mechanism. Route-based addressing via bang paths means not just that you need to understand how your computer is connected to the network; everybody needs to understand how your computer is connected to the network, if they want to reach you, and they need to understand how all the intermediate computers are connected to the network too. This arrangement does not allow, then, for frequent reconfiguration. Should seismo stop talking to mcvax, then people using that connection as part of their routing process will find their routes break.

The second noteworthy aspect, a partial consequence of the first, is that within the open, pairwise connection structure afforded by Usenet, a backbone hierarchy of major sites soon arose, at first through conventional practice and later through explicit design. These were major sites that engaged in significant data transfer with each other or with other groups, including ihnp4 at AT&T’s Indian Hill site, seismo at the Center for Seismic Studies in northern Virginia and mcvax at the Mathematics Centrum in Amsterdam, which effectively became the primary trans-Atlantic gateways, and national sites such as ukc at the University of Kent at Canterbury, which effectively became the gateway to the United Kingdom. Significantly, some of these also served as “well-known nodes” for routing purposes; one might quote one’s email address as “...!seismo!mcvax!itspna!jpd”, with the “...” essentially standing for the directive “use whatever path you use to regularly use to reach here.” The design of the network, then, is unstructured but the practice requires a commitment to some form of well-understood centralization.

Uucp mail, with its explicit use of bang paths, was not the primary or most visible aspect of Usenet, however. A distributed messaging service which came to be known as Usenet news was first deployed in 1980, using the same underlying uucp mechanism to share files between sites. Unlike email messages directed to specific users, however, articles in Usenet news were open posts organized into topically-organized “newsgroups.” Via uucp, these would propagate between sites. Usenet connected many of the same academic and industrial research sites that came to be incorporated into ARPAnet or its successors, and so over time, internet protocols became a more effective way for messages to be exchanged, at which point, Usenet newsgroups were distributed over TCP/IP rather than uucp.

DECnet

Digital Equipment Corporation (DEC, or just “Digital”) was a leading designer and vendor of minicomputers through the 1970s and 1980s. Indeed, many of the

computers connected to the early ARPAnet/Internet were DEC machines, especially systems in the DEC-10, DEC-20, PDP, and VAX ranges, which were widely used in the academic research community. At the same time, DEC had its own networking system, delivered as part of its RSX and VMS operating systems. This network system, known as DECnet or the Digital Network Architecture (DNA), was initially introduced in the mid-1970s as simple point-to-point connection between two PDP-11 minicomputers. Subsequently, the network architecture evolved to incorporate new technologies and new capabilities. The design and development of DECnet was, then, largely contemporaneous with the development of the Internet protocols. The last fully proprietary versions were DECnet Phases IV and Phase IV+, in the early 1980's [10]; DECnet Phases V and V+ maintained compatibility with the proprietary DECnet protocols but moved more in the direction of support for the ISO-defined OSI protocol stack.

Given that DECnet was designed at roughly the same time as the Internet protocol suite, and given that it connected many of the same computer system types as the Internet protocols, it is again a useful point of comparison. DECnet was based on much the same "layered" protocol model that was the contemporary state of the art, and its basic architecture – a point to point connection layer, a routing layer, a layer for reliable sequenced delivery, and so on – is similar to that of systems like PUP, XNS, and TCP/IP. However, some key differences reveal the distinct character to the contexts in which DECnet was expected to operate.

One of these is that DECnet incorporated a sophisticated management interface, and indeed, that facilities for network management were designed into the protocol stack from an early stage. That is, DECnet was entirely imagined to be deployed in a managed environment. TCP/IP has to be managed too, of course, but the management of TCP/IP networks is not a network function in itself. (The Internet protocol suite includes a protocol, SNMP, for network management uses, but network-wide, management is not a key consideration.)

A second suggestive distinction lies within the sets of services standardized within DECnet. These included services, like network terminal access, similar to TCP/IP, but also some that the Internet protocol suite did not natively attempt to support, such as seamless remote filesystem access, in which disks attached to one computer would appear to be virtually available to users of other, connected computers. Remote file access of this sort (which was also a feature that had been part of the Xerox network system) goes beyond simply file transfer by providing users with the illusion of seamless access to both local and remote files. (Email, on the other hand, was not one of the standardized protocols, although networked email services were available through operating system applications.)

A third – although relatively trivial – distinction was that DECnet addresses were just 16 bits long. Since each

computer on a network needs to have a different address, the size of the address is a limit upon the size of the network. With 16-bit addresses, DECnet network implementations were limited to 64449 hosts.

These three of these features of the DECnet design speak to a particular context of use. They highlight the expectation that DECnet deployments would be uniform, well-regulated and actively managed. This makes perfect sense in the context of DEC's sales in corporate settings, where network implementation can be phased, planned, and centrally directed. Effective use of the shared file facilities, for instance, require a coordinated approach to the layout and conventions of filesystems across machines, while the network management infrastructure suggests that this was a key consideration in the settings for which DECnet was designed. An odd little implementation quirk in DECnet Phase IV similarly supports this. To make routing and discover easier, the network software running on a DECnet computer operating over an Ethernet network would actually reset the hardware network address of the computer to an address that conformed to the DECnet host address. This would cause considerable difficulty when DECnet was running in the same environment as other protocols, but in a highly managed environment where uniformity could be guaranteed, it was less troublesome.

In sum, although DECnet was based on the same decentralized, peer-to-peer approach to network connectivity that characterizes the Internet protocol suite, its specific configuration of that approach is one that was in practice designed for the highly managed, highly controlled setting of corporate IT.

CSNET

The conventional historical account of the emergence of today's Internet traces its beginnings to the ARPANET. ARPANET was both a research project and a facility; that is, the ARPANET project developed the networking technologies that are the underpinnings of the contemporary Internet and it also operated a facility – a network – based on those underpinnings. So this was not simply a research project funded by ARPA; it was also a facility that supported ARPA research. ARPA is the research arm of the US Department of Defense (DOD), and so in order to qualify as a site to be connected to ARPANET, one had to be a DOD site or a DOD contractor. At the time, this included many prominent research universities and computer science departments, but by no means all, and not even most.

Recognizing the value that the DOD-contracting universities were deriving from their participation in the ARPANET effort, wanting to expand beyond the smaller-scale network arrangements upon which they were already depending [6], and concerned that 'the ARPANET experiment had produced a split between the "haves" of the ARPANET and the "have-nots" of the rest of the computer science community' [9], a consortium of US computer

science departments partnered with the National Science Foundation and other groups to put together a proposal for the Computer Science Research Network, or CSNET. CSNET integrated multiple different network technologies, with a store-and-forward email system over regular phone lines as the base level of participation, but leased line and X.25 networking available for higher performance connections.

While the development of CSNET produced many important technical contributions, CSNET's most significant legacies might be historical and institutional, in that, first, CSNET represented a significant expansion in the spread of TCP/IP in the academic research community, and second, it was designed in collaboration with and in order to support the mission of the National Science Foundation (rather than, say, the military backers of ARPA and ARPANET). The CSNET effort laid the foundation for a later NFS-funded network called NSFNet, and it was the NSFNet backbone that was later opened up to commercial traffic, and then in 1995 replaced entirely by private service providers. The importance of this point is the explicit link at the time between institutional participation and connectivity, and the very idea that a network is conceived around an understanding of quite who and what will be permitted to connect.

The more interesting technical point for us here though concerns the relationship between ARPANET and CSNET. In a minor sense, CSNET was "in competition with" ARPANET; it was, after all, designed as a network for those who were institutionally denied access to ARPANET. However, in all the ways that matter it was entirely collaborative with ARPANET; key players in the ARPANET project, such as TCP/IP co-designer Vint Cerf at Stanford, participated in the CSNET effort, and the networks were bridged together. The basis for that bridging was CSNET's adoption of the TCP/IP protocols that had been designed within the ARPANET effort. Through this bridging, ARPANET became a subnet of CSNET. However, we normally think of arrangement of subnetting and internetworking as providing seamless interconnection, but the interconnection between CSNET and ARPANET was not quite so seamless, since they adopted different protocols for delivering services at the higher levels. For instance, the MMDF network messaging facility developed as part of CSNET [7] was needed to be able to bridge between the "phonenet" and TCP/IP components of CSNET, and that meant that messages destined for CSNET recipients would need to be routed explicitly to CSNET rather than simply dispatched using the usual protocols used on purely TCP/IP networks (such as SMTP, the standard Internet email transfer protocol). In other words – both ARPANET and CSNET implemented the Internet protocols (TCP/IP) but not all the other protocols of what we is sometimes called the "Internet protocol suite"; accordingly, even though they were connected together through a common TCP/IP infrastructure, they remained in some

other important, user-facing senses, distinct networks, suggesting intriguingly that there may be more to being "an Internet" than running TCP/IP on interconnected networks.

FACILITIES

Briefly examining some of these "othernets" lets us place our contemporary experience of the Internet – in its multiple capacities as a configuration of technologies, a constellation of services, and an object of cultural attention – in context. Some of that context is a "design space" – that is a space of possibilities that are available as outcomes of specific design decisions. Some lies in "historical circumstance" – that is, ways that different configurations arose reflecting their own historical trajectories. In order to reflect on these in a little more detail, we can take a different cut at the same question – of the nature of the Internet within a space of alternatives – by approaching the it in terms of the kinds of facilities that works can offer.

Naming

Consider even this simple question: does a computer "know" its own name? In the naming arrangement of traditional TCP/IP, the answer is no. A computer knows its address but not necessarily the name by which it might be addressed by others, and in fact can operate in the TCP/IP environment quite effectively without one, as signaled by the fact that both TCP nor IP use addresses, not names, internally. So, facilities are built into the Domain Name Service (DNS) protocols [26] that allow a computer, on booting, to ask a network server, "What's my name?" Naming is entirely delegated to the name service; that is, the network authority that answers the query "what is the address at which I can reach www.cnn.com?" is also the authority that tells a computer that it is www.cnn.com in the first place.

By contrast, other networking protocols – like AppleTalk, for example – delegate to individual computers the right to assign themselves names. This can be implemented in different ways – by having each computer register a name with a naming facility, for example, or by entirely distributing the name lookup process rather than electing particular computers to implement a name or directory service – and these different mechanisms have their own differences in terms of both technological capacities and organizational expectations. The abilities for a computer to assign itself a name and to advertise that name to others are facilities that the designers of the IETF "Zeroconf" protocol suite felt it important to add, in order to support the server-free "plug-and-play networking" approach of AppleTalk [38].

The issue of naming raises a series of questions that highlight the relationship between specific technical facilities and the social organization of technological practice. Who gets to name a computer? Who is responsible for the names that network nodes use or call themselves? How is naming authority distributed? How visible are names? How centralized is control, and what temporalities

shape it? Even such a simple issue as naming presents a microcosm of my larger argument that questions of online experience need to be exemplified in their technical and historical specificities.

Routing

Elsewhere [12], I discuss the case of Internet routing and explore the history of the development of routing protocols as being entwined with the history of the expansion of the Internet infrastructure. That paper contains a fuller argument, with a particular focus on the materiality of routing, but I want to summarise from it one relevant element here, which is the role of so-called “autonomous systems” in today’s Internet routing.

If there is one cornerstone of the argument that the Internet is a “decentralized” phenomenon, it is the case of network routing. Rather than requiring a central authority to understand the structure of the network and determine how traffic should flow, the design of the IP protocol, which gets Internet packets from one host to another, relies on each computer along a path deciding independently which way a piece of information should travel to take it closer to its destination, and creates the conditions under which coherent behavior results from this distributed process. Routing protocols are the protocols by which the information is distributed upon which these local decisions are to be made.

Early internets relied on fairly simple protocols for distributing routing information. As the Internet has grown from a small-scale research project to an international utility, the complexities of routing, and of the distribution of routing information, have also grown, and different protocols have arisen to solve emerging problems. The major routing information protocol operative in our current Internet is BGP, the Border Gateway Protocol. The particular relevance of BGP here is that, like its predecessor protocol (the Exterior Gateway Protocol or EGP), BGP is designed not to distribute the information necessary to route between networks but rather to route between so-called “autonomous systems.” The essential consideration here is that corporations, network infrastructure providers, universities, and so on, have networks that are themselves both quite complicated and autonomously managed. This in turn leads to the idea that the unit of network routing should be one of these systems. This is in many ways a recognition of a fundamental feature of our Internet, which is that it is not simply a network of networks, but a network of institutions and semi-corporate entities, each of which wish to maintain control over the organization of and access to their own networks. The protocols by which routing information are passed around are protocols that reflect this particular arrangement, encoding and producing “autonomy” as much as they encode “connection.” As our Internet grew, the “network” was no longer an effective unit at which routing information could be encoded, and “autonomous system” arose as an alternative that reflected

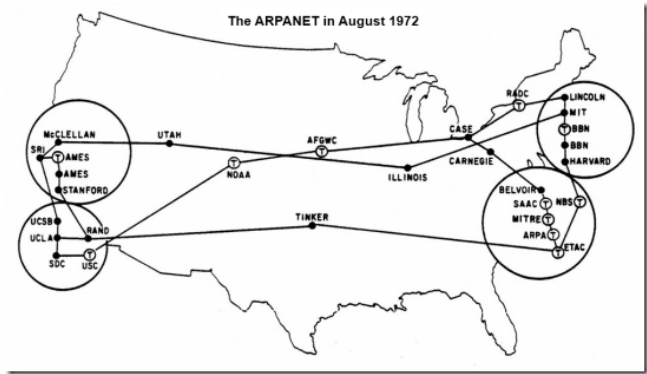


Figure 1. The Arpanet in 1972, exhibiting a net-like structure.

the practical realities of the technology as it had developed; but in the evolutionary cycle of network protocol development, this meant baking those arrangements right into the protocol.

Structure

The emergence of autonomous systems as a consideration for routing information has an analogy in the emergence of structure within the interconnection arrangements of the network itself.

The abstract model of interconnection that the Internet embodies is one of decentralized and amorphous structure. Indeed, the very term “network” was originally coined to convey not idea of computers connect together, but the topology of their connection as a loose and variable net-like mesh, in contrast to the ring, star, or hierarchical arrangements by which other connection fabrics had been designed. The basic idea of Internet-style networking is that computers can be connected together in whatever arrangement is convenient, and data units (“packets”) will nonetheless find their way from one point to another as long as some path (or route) exists from point A to point B. In a fixed topology, like a ring, star, or hierarchy, the path from one node to another is fixed and pre-defined. The key insight of the packet-switching design of the Internet is that there need be no prior definition of the route from one place to another; packets could find their own way, adaptively responding to the dynamics of network topology and traffic patterns. Amorphous topology is arguably one of the key characteristics of internet-style technologies.

Early network diagrams from the days of ARPAnet, the predecessor to the Internet, do indeed display this network character (figure 1). Over time, though, more hierarchical arrangements have arisen. These hierarchical patterns arise not least from two sources – geography and economics. The geographical considerations produce a distinction between metropolitan and long-haul networks, where metropolitan networks support relatively dense connectivity within small regions (for instance, the homes and businesses connected to the network in a town or neighborhood), while long-haul networks provide high-speed and high-capacity connections between urban regions [24]. The different demands upon

these different kinds of connection are best met with different technologies. This is where the economic considerations also come into play. Ever since the decommissioning of the NSFNet backbone in 1995 and the emergence of the commercial Internet, Internet service provision has been largely in the hands of commercial operators. To function as a competitive marketplace, network provision must involve multiple competing providers, each of whom in turn specialize in different aspects of service provision. One result of this specialization is a hierarchical segmentation of network service provision.

Broadly, we can distinguish four different commercial entities – and hence four different forms of service provision and four different technological arrangements – in contemporary Internet service. The four are content providers, content distribution networks (CDNs), internet service providers (ISPs), and transit carriers.

Content providers are corporations like Netflix, Amazon, Apple, or others whose business comprises (in whole or in part) delivering digital material to subscribers and consumers. Internet Service Providers, by and large, sell internet services to end-users, either individual subscribers or businesses. We think of our ISPs as delivering the content to us, but in truth they are generally responsible only for the last steps in the network, in our local cities or neighborhoods. Transit carriers are responsible for getting data from content providers to the ISPs. While the names of ISPs like Verizon and Comcast are well known, the names of transit carriers – such as Level 3, Cogent, or XO – are much less familiar, despite the central role that they play in bringing information to any of our devices. Content providers will typically contract directly with transit carriers for their content delivery needs. Transit carriers and ISPs connect their networks through a set of technical and economic arrangements collectively known as “peering” (as in the connection between two “peer” networks).

Peering arrangements are almost always bilateral, i.e. an agreement between two parties. They are generally bidirectional and often cost-free, although they may have traffic or bandwidth limits beyond which charges are levied. The term “peer” is a hold-over from the military and non-profit days preceding the development of the commercial Internet, and speaks directly to the notion of “inter-networking” (ie, connecting together different networks and transmitting traffic between them). The “peers” that are now linked by peering arrangements, though, are not peer academic or research networks but commercial entities engaging in economic relations. These arrangements are largely hidden from the view of end users, but become broadly visible when disputes arise around the adequacy of peering relationships, corporate responsiveness to changing conditions upon them, or the responsibilities of carriage associated with them. For example, recent (early 2014) debates around Internet streaming movie provider

Netflix paying ISP Comcast for access to its facilities and networks have largely ignored the fact that the problem to which Netflix was responding was a breakdown in relations between Comcast and Cogent, one of the transit carriers that Netflix pays to transmit its traffic to ISPs [32]. A dispute arose between Comcast and Cogent concerning whose responsibility it was to address bandwidth limitations when Cogent began to send more traffic onto Comcast’s network than their peering agreement allowed. In the face of this ongoing dispute, Netflix arranged to locate its servers with direct access to Comcast’s subscriber network, eliminating their dependence upon Comcast’s transit network. While this raised the specter in the popular press and in technical circles of an end-run around “net neutrality” arguments, the source of the problem in a dispute between carriers – and in particular at the boundary between an ISP and a transit carrier – is in some ways more interesting.

While the inspiration for the design of the internet protocols was to allow the creation of a network of networks, the emergent structure is one in which not all networks are equal. It’s not surprise that networks of course come in different sizes and different speeds, but the structure of contemporary networking relies upon a particular structural arrangement of networks and different levels of network providers, which in turn is embodied in a series of commercial and institutional arrangements (such as “settlement-free peering” [23]). As in other cases what we see at work here is an evolutionary convergence in which the network, as an entity that continues to grow and adapt (both through new technologies becoming available and new protocols being designed) does so in ways that incorporate and subsequently come to reinforce historical patterns of institutional arrangements.

End-to-End

Gillespie [17] has discussed the importance of the so-called “end-to-end principle” [33] in not just the design of the Internet platform but also in the debates about the appropriate function of a network. The end-to-end principle essentially states that all knowledge about the specific needs of particular applications should be concentrated at the “end-points” of a network connection, that is, at the hosts that are communicating with each other; no other network components, such as routers, should need to know any of the details of application function or data structure in order to operate effectively. This implies too that correctly routing a packet should not require routers to inspect the packet contents or transform the packet in any way. So, for instance, if data is to be transmitted in an encrypted form, the intermediate nodes do not need to understand the nature of the encryption in order to transmit the packets. There is a trade-off in this of course; on the one hand, the network can be simpler, because every packet will be treated identically, but on the other, distinctions that we might want to make, such as between real-time and non-real-time traffic are unavailable to the network.

The end-to-end principle was originally formulated as a simple design principle, but it had, as Gillespie notes, many consequences that are organizational and political as well as technical. For instance, the separation between application semantics and packet transit also essentially creates the opportunity for a separation between application service providers (e.g. Netflix or Facebook) and infrastructure providers (e.g. ISPs like Time Warner or Verizon) by ensuring that applications and infrastructure can evolve independently. Arguably, these aspects of the Internet's design have been crucial to its successful evolution. However, the end-to-end principle is under assault as a practical rubric for design. Two examples of obstacles to the end-to-end principle are the rise of middleboxes and assaults on net neutrality.

Middleboxes is a generic term for devices that intervene in network connections. Unlike routers, which, designed according to the end-to-end principle, never examine packet internals or transform transmitted data, middleboxes may do both of these. Perhaps the most common middleboxes are gateways that perform Network Address Translation or NAT. Almost every residential network gateway or WiFi base station is actually a NAT device. NAT is a mechanism that partially resolves the problem of scarce network addresses. Most people buy a residential internet service that provides them with just a single network address, even though they have multiple devices using the service (e.g. a laptop and a smartphone). Network Address Translation allows both of these devices to use a single network address simultaneously. To the outside world, the laptop and the smartphone appear like a single device that is making multiple simultaneous connections; the NAT device keeps track of which connection are associated with which device and delivers responses appropriately. NAT violates the end-to-end principle, however, with various consequences. For exactly, because the individual devices are not technically end-points on the public internet, one cannot make a connection to them directly. Further, network flow control algorithms may be confused by the presence of multiple devices with essentially the same network address. Further, NAT gateways achieve their effects by rewriting packet addresses before the packets have been delivered to their end-point.

A second assault on the end-to-end principle is the emergence of threats to network neutrality. Net neutrality is a term to express the idea that network traffic should be treated identically no matter what it contains, and no matter where it originates or terminates. However, many organizational entities have reasons to want to violate this principle. For instance, a network provider which is owned by a media company might want to prioritize the delivery of its own media traffic to its customers, giving it a higher priority than other traffic to ensure a smoother service, or an ISP might want to limit the bandwidth available to high-traffic users in order to encourage more equitable access to the network. These kinds of situations often involve

deploying technologies like so-called “deep packet inspection” (mechanisms that examine not just the headers of a packet but its contents in order to decide how it should be treated) that violate the end-to-end principle.

As discussions such as those of Saltzer et al. [33] make clear, the end-to-end principle was a foundational and significant principle in the design of the Internet, representing not just a specific technical consideration but also a commitment to an evolutionary path for network service provision. There is no question then that the Internet was designed around this principle. However, as the rise of middleboxes and challenges to network neutrality make clear, that doesn't imply that the end-to-end principle is a design feature of our (contemporary) Internet. More generally, this case illustrates that technical design principles are themselves subject to revision, reinterpretation and revocation as the technology is implemented and evolves (c.f. [28]).

THE FEUDAL INTERNET

The emergent structure of our Internet – the market niches of transit carrier and ISP, the practical solution of CDNs, the fragmentation produced by middleboxes, the relationship between mobile carriers, telecommunications firms, and media content producers, and so on – draws attention to a simple but important truth of internetworking: the Internet comprises a lot of wires, and every one of them is owned by someone. To the extent that those owners are capitalist enterprises competing in a marketplace, then the obvious corollary is that, since all the wires can do is carry traffic from one point to another, the carriage of traffic must become profit-generating. The mechanisms by which traffic carriage over network connections becomes profitable is basically either through a volume-based or per-byte mechanism – a toll for traffic – or through a contractual arrangement that places the facilities of one entity's network at the temporary disposal of another – a rental arrangement. This system of rents and tolls provides the basic mechanism by which different autonomous systems, each of which provisions its own services, manages its own infrastructures, and then engages in a series of agreements of mutual aid.

If this system seems familiar, it is not so much that it encapsulates contemporary market capitalism but more that it is essentially feudal in its configuration. Marx argued for feudalism and capitalism as distinct historical periods with their links to material means of production – “The hand-mill gives you society with the feudal lord; the steam-mill society with the industrial capitalist” [25]. Recent writers, though, have used the term “neofeudal” to describe the situation in late capitalism in which aspects of public life increasingly become private, gated domains – everything from toll lanes on the freeway and executive lounges at the airport, on the small end, to gated communities and tradeable rights to pollute the environment issued to large corporations at the other (e.g. [34]). The essential

consideration here is the erasure of public infrastructure and the erection of a system of tariffs, tolls, and rents that govern the way we navigate a world made up of privatized but encompassing domains, within which market relations do not dominate.

Beyond the general use of the term “neofeudal” to refer to the privatization of public goods, let me take the metaphor of the feudal Internet more seriously for a moment to point to a couple of significant considerations.

The first is that operating mechanism of feudalism is not the market transaction but rather long-standing commitments of fealty, vassalage, and protection. These are not the instantaneous mutual engagements of market capitalism but temporally extended (indeed, indefinite) arrangements with little or nothing by way of choice or options. Indeed, the constraints upon feudal relations are geographical as much as anything else: infrastructural, if you will. One can see, arguably, some similarities to the way that geographical and infrastructural constraints lead to a pattern of relations between internet providers that also relies upon long-term, “residence”-based, partnerships. The ties that bind individuals to their service providers in semi-monopolistic conditions of the US broadband market, or perhaps even more pointedly, the links that connect large-scale data service providers such as Netflix with transit carriers like Level 3 are not simply conveniently structured as long-term arrangements, but rather can only operate that way because of the infrastructure commitments involved (such as the physical siting of data stores and server farms.) Similarly, the need for physical interconnection between different networks makes high-provider-density interconnection nodes like One Wilshire in downtown Los Angeles (see, e.g., [12, 40]) into “obligatory passage points” in Callon’s language [5] – that is, to get interconnection between networks, you need to be where all the other networks are, and they all need to be there too. For all that we typically talk about the digital domain as fast-paced and ever-changing, these kinds of arrangements – not simply commitments to infrastructure but commitments to the institutions relationships that infrastructure conditions – are not ones that can change quickly, easily, or cheaply. These relations are more feudal than mercantile.

The second interesting point that a feudal approach draws attention to is the persistence of pre-existing institutional structures – perhaps most obviously, the nation-state. Although John Perry Barlow’s [2] classic “Declaration of the Independence of Cyberspace” famously argues that the “governments of the industrial world... [have] no sovereignty” in the realm of the digital, and notwithstanding the IETF’s famous motto that “rejects kings [and] presidents” in favor of “rough consensus and running code” [20], the truth is that governments and presidents continue to manifest themselves quite significantly in not just the governance but the fabric of our Internet. National and regional concerns arise in a variety of

ways – in the provision of specific linguistic content [36], in the regional caching of digital content, in the question of international distribution rights for digital content (e.g. which movies can be viewed online in which countries), in assertions of national sovereignty over information about citizens (e.g. Vladimir Putin’s public musings that data about Russian citizens should be stored only on servers in Russia [22]), in the different regimes that govern information access (e.g. the 2014 EU directive known popularly as the “right to be forgotten”), and in local debates about Internet censorship (from China’s “Great Firewall” and Singapore’s self-censorship regime to discussions of nationwide internet filters in Australia and the UK). The very fact that a thriving business opportunity exists for commercial Virtual Private Network (VPN) services that allow users to get online “as if” they were located in a different country signals the persistent significance of nation-states and national boundaries in the experience of our Internet. Similarly, significant debate has surrounded the role of national interests in Internet governance [e.g. 27], and the International Telecommunications Union or ITU – a United Nations organization whose “members” are not technical experts or corporations but nation-states – remains a significant body in the development of network technologies and policy. Just as feudalism reinforced and made all aspects of everyday life subject to the boundaries of the manor, the shire, and the nation, so too does our Internet – not necessarily any Internet, but certainly our Internet – retain a significant commitment to the relevance of similar geographical, national, and institutional boundaries.

The third point to be emphasized here is the way that these are simultaneously social and technical arrangements – not social arrangements that give rise to technological designs, nor technological designs that provoke social responses. Middleboxes, deep packet inspection, and the autonomous systems of BGP should be regarded as, and analyzed as, both at once. This requires, then, a view that both takes specific technological configurations (rather than principles, imaginaries, and metaphors) seriously as objects of analytic attention, and that identifies and examines the social, political, and economic contexts within which these come to operate. To the extent that a feudal regime of hierarchical relations based on long-term structures of mutual commitment can be invoked as an account of our Internet, it can be done only within with context of a sociotechnical analysis that is both historically specific and symmetric.

CONCLUSIONS

The larger project of which the current exploration forms a part is an attempt to take seriously the materialities of information and their consequences. It is critical to this project that we move beyond accounts simply of information infrastructure, but also recognize the relevant materialities of representation, and their consequences [11, 13]. As Powell [31] has argued in her study of open hardware projects, patterns of historical evolution torque

the design principles that often provide not only technical arrangements but also the social imaginaries that are mobilized in our discussions of technology. To bring this perspective to contemporary networking, then, means not simply that we need to think about the pragmatics of undersea cables [37], satellite downlinks [29], and server farms [40], but also the way that specific representations of action and encodings of data are designed to be manipulated, transmitted, and moved in particular sorts of ways, with consequences for the resultant experience of the network as offering particular opportunities for individual and collective action.

In other words, the primary concern is to see networked arrangements as historically particular crystallizations of not just technical but also institutional, economic, and political potentialities. To do this, particularly with respect to the technology that we rather glibly call “the Internet,” I have suggested that two moves are needed.

The first move is from the idea of “the Internet” to that of “an Internet” – that is, to re-encounter our contemporary network as not the only possible Internet that could have been built, but as one of a range of possible networks. When we consider a number of possible networks, we start to pay attention to the range of expectations, institutional arrangements, policies, technical configurations, and other dimensions that might characterize the space of potentials. The second move is from “an Internet” to “this Internet” – that is, to narrow down once more and so grapple with the historical, geographical, political and social specificities that constrain and condition the particular network with which we are engaged at any particular moment in time. This Internet is not the one that was conceived of by those like Paul Baran or Donald Davies, designed by those like Vint Cerf or Bob Taylor, or opened to commercial operation by the NSF – it has elements of each of those, but it is a historically specific construction which has encompassed, transformed, extended, and redirected any of those particular networks. This Internet is something that we can grapple with empirically. We are not in much of a position to make general statements about “the Internet”; but when we ask questions about “this Internet,” we may have a starting point for investigation.

ACKNOWLEDGMENTS

Many key ideas here arose in conversation with Jon Crowcroft, Alison Powell, and Irina Shklovski. I am indebted to Morgan Ames, Panayotis Antoniadis, Sandra Braman, Marisa Cohn, Courtney Loder, Lilly Nguyen, Katie Pine, and Chris Wolf, each of whom offered valuable critique at key junctures. This work is part of a larger project on the materialities of information in which Melissa Mazmanian has been a key partner, and it has been supported in part by the Intel Science and Technology Center for Social Computing and by the National Science Foundation under awards 0917401, 0968616, and 1025761.

REFERENCES

1. Antoniadis, P., Le Grand, B., Satsiou, A., Tassioulas, L., Aguiar, R., Barraca, J.P., and Sargento, S. 2008. Community building over Neighborhood Wireless Mesh Networks. *IEEE Society & Technology*, 27 (1): 48-56.
2. Barlow, J. 1996. A Declaration of the Independence of Cyberspace. Davos, Switzerland. Downloaded from <https://projects.eff.org/~barlow/Declaration-Final.html> on June 4, 2014.
3. Birrell, A., Levin, R., Needham, R., and Schroeder, M. 1982. Grapevine: An Exercise in Distributed Computing. *Communications of the ACM*, 25(4), 260-274.
4. Boggs, D., Shoch, J., Taft, E., and Metcalfe, R. 1980. *IEEE Transactions on Communication*, 28(4), 612-624.
5. Callon, M. 1986. Elements of a sociology of translation: Domestication of the Scallops and the Fishermen of St Brieuc Bay. In Law (ed.), *Power, Action and Belief: A New Sociology of Knowledge?* 196-233. London: Routledge.
6. Comer, D. 1983. The computer science research network CSNET: a history and status report. *Communications of the ACM*, 26(10), 747-753.
7. Crocker, D., Szurkowski, E., and Farber, D. 1979. An Internetwork Memo Distribution Facility – MMDF. *Proc. ACM Conf. Computer Communication SIGCOMM*, 18-25.
8. Dahlberg, L. 2001. Democracy via cyberspace: examining the rhetorics and practices of three prominent camps. *New Media & Society*, 3, 187-207.
9. Denning, P., Hearn, A., and Kern, W. 1983. History and overview of CSNET. *Proc. Symp. Communications Architectures & Protocols (SIGCOMM '83)*, 138-145.
10. Digital Equipment Corporation. 1982. DECnet DIGITAL Network Architecture (Phase IV): General Description. Order number AA-N149A-TC. Maynard, MA: Digital Equipment Corporation.
11. Dourish, P. 2014. NoSQL: The Shifting Materialities of Databases. *Computational Culture*.
12. Dourish, P. 2015. Packets, Protocols, and Proximity: The Materiality of Internet Routing. In Parks, L. and Starosielski, N. (eds), *Signal Traffic: Critical Studies of Media Infrastructures*, 183-204. University of Illinois Press.
13. Dourish, P. and Mazmanian, M. 2013. Media as Material: Information Representations as Material Foundations for Organizational Practice. In Carlile, Nicolini, Langley, and Tsoukas (eds), *How Matter Matters: Objects, Artifacts, and Materiality in Organization Studies*, 92-118. Oxford, UK: Oxford University Press.

14. Elmer-Dewitt, P. 1993. First Nation in Cyberspace. *Time*, 49, Dec 6.
15. Fernaeus, Y. and Sundström, P. 2012. The Material Move: How Materials Matter in Interaction Design Research. *Proc. ACM Conf. Designing Interactive Systems*, 486-495.
16. Gaved, M., & Mulholland, P. 2008. Pioneers, subcultures, and cooperatives: the grassroots augmentation of urban places. In Aurigi, A. and De Cindio, F. (eds.), *Augmented urban spaces: articulating the physical and electronic city*, 171-184. England, Ashgate.
17. Gillespie, T. 2006. Engineering a Principle: 'End-to-End' in the Design of the Internet. *Social Studies of Science*, 36(3), 427-457.
18. Hauben, M. and Hauben, R. 1997. *Netizens: On the History and Impact of Usenet and the Internet*. Wiley/IEEE Computer Press.
19. Hindman, M. 2008. *The Myth of Digital Democracy*. Princeton University Press.
20. Hoffman, P. 2012. The Tao of IETF. Downloaded from <http://www.ietf.org/tao.html> on June 4, 2014.
21. Jungnickel, K. 2014. *DIY WIFI: Re-imagining Connectivity*. Palgrave Pivot.
22. Khrennikov, I. and Ustinova, A. 2014. Putin's Next Invasion: The Russian Web. Downloaded from <http://www.businessweek.com/articles/2014-05-01/russia-moves-toward-china-style-internet-censorship> on June 4, 2014.
23. Laffont, J.-L., Marcus, S., Rey, P., and Tirole, J. 2001. Internet Peering. *Annual Economic Review*, 91(2), 287-291.
24. Malecki, E. 2002. The Economic Geography of the Internet's Infrastructure. *Economic Geography*, 78(4), 399-424.
25. Marx, K. 1847 (1910). *The Poverty of Philosophy*. Tr: Quelch, H. Chicago, IL: Charles H. Kerr & Co.
26. Mockapetris, P. 1987. Domain Names – Concepts and Facilities. RFC 1034, Network Working Group.
27. Mueller, M. 2010. *Networks and States: The Global Politics of Internet Governance*. MIT Press.
28. Oudshoorn, N. and Pinch, T. 2003. *How Users Matter: The Co-construction of Users and Technology*. Cambridge, MA: MIT Press.
29. Parks, L. 2012. Satellites, Oil, and Footprints: Eutelsat, Kazsat, and Post-Communist Territories in Central Asia. In Parks and Schwoch (eds), *Down to Earth: Satellite Technologies, Industries, and Cultures*. New Brunswick: Rutgers University Press.
30. Powell, A. 2011. *Metaphors, Models and Communicative Spaces: Designing local wireless infrastructure*. *Canadian Journal of Communication*.
31. Powell, A. 2014. The History and Future of Internet Openness: From 'Wired' to 'Mobile'. In Swiss and Herman (eds), *Materialities and Imaginaries of the Open Internet*. Routledge.
32. Rayburn, D. 2014. Here's How the Comcast & Netflix Deal is Structured, With Data and Numbers. *Streaming Media Blog*, February 27 2014. Downloaded from <http://blog.streamingmedia.com/2014/02/heres-comcast-netflix-deal-structured-numbers.html>
33. Saltzer, J., Reed, D. and Clark, D. 1984. End-to-End Arguments in System Design. *ACM Transactions on Computer Systems*, 2(4), 277-88.
34. Shearing, C. 2001. Punishment and the Changing Face of Governance. *Punishment and Society*, 3(2), 203-220.
35. Shirky, C. 2008. *Here Comes Everyone: The Power of Organizing Without Organizations*. Penguin.
36. Shklovski, I. and Struthers, D. 2010. Of States and Borders on the Internet: The Role of Domain Name Extensions in Expressions of Nationalism Online in Kazakhstan. *Policy and Internet*, 2(4), 107-129.
37. Starosielski, N. 2015. *The Undersea Network*. Durham, NC: Duke University Press.
38. Steinberg, D. and Cheshire, S. 2005. *Zero Configuration Networking: The Definitive Guide*. Sebastopol, CA: O'Reilly Media.
39. Turner, F. 2006. *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. Chicago, IL: University of Chicago Press.
40. Varnelis, K. 2008. *The Infrastructural City: Networked Ecologies in Los Angeles*. Barcelona: Actar.