

The Framing Years: Policy Fundamentals in the Internet Design Process, 1969–1979

Sandra Braman

Department of Communication, University of Wisconsin, Milwaukee, Wisconsin, USA

Those responsible for the technical decision making that created the Internet found they had to think through a number of social policy issues, from privacy and intellectual property rights through the definition of common carriage and environmental problems, along the way. Such issues were framed by conceptualizations of the nature of the network, goals to be served by the network, users and uses of the network, and the design criteria that served as policy principles developed during the early years of the design process. This article examines such policy fundamentals as they developed through the technical document series that records the Internet design process, the Internet Requests for Comments (RFCs) during the first decade of that process, 1969–1979.

Keywords ARPANet, Internet, protocols, RFC, sociotechnical, technical standards, telecommunications policy

I think these meetings will turn out to be more important than we ever wanted them to be. (Steve Crocker, RFC 82 [1970], p. 18)

Within a few months of the first U.S. government grant to develop a network linking computers at different sites in 1969, a technical document series was launched to record the thinking and decision making of those who met face to face and conversed at a distance to design what we now know as the Internet. The Internet Requests for Comments (RFCs) series continues today; as of June 2011 there were over 6,200 documents. Over time, these documents became a venue through which technical standards—protocols—were developed for what was initially referred to as ARPANet because it was funded by the Advanced Research Projects Agency (ARPA)

(Abbate 1999). ARPANet was only one among the “network of networks” (RFC 1122)¹ that comprise the Internet. It was the U.S. invention of the 1973 Transmission Control Protocol (TCP) (RFC 604), though—a protocol that specifies just how to send information from one program on one computer to a program on another computer in a reliable and orderly stream—that came to define the Internet, first on its own, and then in combination with the Internet Protocol (IP) as TCP/IP (Mowery and Simcoe 2002; RFC 675; RFC 714).

The document series has uses that go beyond Internet design as well. It is used in the classroom to teach networked computing (Leiner et al. 2010). The RFC process is a model that has been taken up for use by decision makers working on other large-scale sociotechnical systems (RFC 2555). Specific documents serve as guides for those who design organizational networks (e.g., RFC 356) and for users (e.g., “netiquette” first appeared in RFC 1855). The documents are now considered scholarly publications (Carpenter and Partridge 2010), drawn upon not only by computer scientists but also by those in the social sciences (e.g., Franklin et al. 2007; Gupta et al. 2007) and humanities (e.g., Flack 2006).

Designers were well aware that “Network topology is a complicated political and economic question” (RFC 613, p. 1). In the course of solving technical problems, therefore, those involved in creating the Internet discussed legal and policy issues, engaged in explicit and implicit policy analysis and policymaking, and developed formalized decision-making processes and entities (Braman 2010). Analysis of the RFCs from a policy perspective thus should contribute to the building of a shared epistemic and discourse community that includes technical, legal, and social science experts. Communications within such a sociotechnical community would help reduce the extent to which legal and technical decisions work at cross-purposes even when values and goals are shared. Technical insights into the intricacies of policy issues as they work out on the ground can valuably inform policymaking, and vice versa. Understanding the treatment of social policy issues in the

Received 31 August 2010; accepted 1 July 2011.

Address correspondence to Sandra Braman, Department of Communication, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA. E-mail: braman@uwm.edu

course of designing the Internet in its current form should also be useful input into consideration of alternative scenarios for the future of the Internet (Levä et al. 2010).

Drawing on the first stage of a comprehensive inductive reading of the first 40 years of the documents,² this article uses discourse analysis to examine the policy frames developed during the decade of the Internet design process. The RFCs are worth studying as a discourse because they are considered by insiders to be the “documents of record” for the history of the Internet (Leiner et al. 2010). Discourse analysis has been increasingly popular in the study of law and policy since the mid-1990s rhetorical turn in this field (Fischer and Forester 1993; Schön and Rein 1994). The RFCs, however, are unusual as a policy discourse. Rather than issuing formally from established political entities, the RFC process was conceived in 1968 during a meeting of a few graduate students and staff members from the four institutions then involved with network design, and it was many years before the publication process became formalized. Instead of being mediated by other players before reaching the public, the RFCs were collected, archived, and aggressively edited by just a couple of individuals integral to the process for most of their existence. In addition to providing a record, RFC narratives range from jokes to the establishment of policymaking organizations and processes themselves. Most policy discourses that receive attention are comprised of a series of assertions, while the content of the RFCs is dialogic in nature.

Using discourse analysis to study the RFCs provides an opportunity to explore thinking about legal and policy issues as they irrupt within documents otherwise devoted to describing technical problems as well as in documents focused on policy matters. The great diversity of issues of concern to those actively involved in the network design process during its early years—and of opinions on those issues—become visible. Because the analysis is comprehensive and inductive, it recuperates details of Internet history that have otherwise been forgotten (even, upon occasion, by those who were involved at the time).

Most importantly, the method is an invaluable means of engaging in sociotechnical boundary work, revealing the “imaginary” (Bowker 1996; Latour 1993) that was as important as—indeed, drove the creation of—the hardware and software that constitute today’s network. The RFCs are the reflective dimension of the Internet, serving as the public sphere for those engaged in decision making for the network as it took place, and as the memory of decision-making processes once concluded (Rasmussen 2007). This study’s emphasis on the ways in which the computer scientists, electrical engineers, and others involved in the Internet design process thought about the nature of society, communication, politics, and the law replaces the more common topography of the relative tech-

nical importance of specific documents with an alternative map of ambitions and ethics, beliefs and ideas, dreams and fears, responsibilities and passions.

What this discourse analysis of one among the many conversations about the technical aspects of networked computing does not offer is an in-depth examination of how any individual decisions were made. The trajectory of this project is thus orthogonal to that of DeNardis (2009), who studied conversations across decision-making venues to unravel the intricacies of the thinking that led to specific technical outcomes.

The many histories of the Internet were already striking in the diversity of their perspectives by the late 1990s (Rosenzweig 1998), but there is little in existing Internet histories that attends to ways in which concerns about social policy were incorporated into the design process (Haigh 2004). Authors who make a point of emphasizing the democratic nature of RFCs (e.g., Hauben and Hauben 1997; Turner 2006) provide a story somewhat different from that offered by those who examine the specifics of actual Internet decision-making processes (e.g., Davidson et al. 2002; Mueller 1989). By the mid-1990s, however, the need to treat RFCs as policy documents was explicitly discussed (Kahin and Keller 1997) and even the notion of requiring “social impact statements” for system design features had been introduced (Schneiderman and Rose 1995). Within a few years the distinct features of Internet decision making as policy making processes were also receiving attention (Yu 2004). The research findings presented here contribute to a nascent but now growing literature that relies upon the Internet RFCs and related materials to examine such topics as the history of Internet governance (Mueller 1989), the struggle over the transition to the IPv6 protocol that will increase the volume of Internet space (Dell 2010; DeNardis 2009), battles over security issues involving the Internet “root” (Kuerbis and Mueller 2011), and the political and economic benefits that accrue to those who have influentially participated in the conversation (Tötterman 2009).

The network, its uses, its users, and policy principles presented as design criteria were conceptualized during the first decade in powerfully path-dependent ways that have been foundational for the treatment of legal and policy issues then and later. These are the “framing years” during which ways of thinking about the network—as well as the network itself—were being designed. Discussions of such things as the heterogeneity of the expected user population, the likelihood that the network would be commercialized, and a commitment to what can be described as “technological democracy” formed policy frames within which specific policy issues, such as privacy (Braman in press) were addressed.

Because the texts have such effects, specific choices regarding models, metaphors, and language were extremely

important. James E. White, in RFC 707, used a comparison of “command/response” versus “request/reply” as an example to illustrate his argument that the words used in protocols influence ways in which both applications and systems programmers understand the network. Irrespective of the substantive impact of a technical detail, he noted, each generates a perceptual model that affects whether or not a programmer identifies a particular problem as pertinent to what s/he does and therefore worthy of attention. Once engaged, the choice of model constrains and enables the types and range of protocol developments considered. One set of language decisions may suggest a rich set of possible extensions, while another may lead nowhere or set programmers off down the wrong path altogether.

This article begins with a deeper look at the RFCs as a discourse. It goes on to look at the policy frames developed by perceptions of uses and users, and at network design criteria as policy principles.

THE RFCS

The RFCs have been central to creation of the Internet and are an invaluable resource for research on its development. The technical decision making documented in the RFCs has been managed since 1986 by the Internet Engineering Task Force (IETF); others involved on the technical side include the Internet Advisory Board (IAB), root server operators, and other groups in a complex suite of formal and informal organizations that Nickerson and zur Muehlen (2006) describe as a standards process ecology. *Technical* decision making, however, is only one of the three types of decision making that constitute Internet governance as broadly conceived and can only be fully understood within the broader context. *Management* of the Internet has since 1998 been under the purview of the Internet Corporation for Assigned Names and Numbers (ICANN), an organization that grew out of discussions within the RFCs on the problem of how to find each other through the network that began in the first year of the RFC design discussion. *Uses* of the Internet are the subject of regulatory efforts by geopolitically recognized governments and international organizations.

The RFC process was launched as an informal method of documenting and circulating ideas, information, and decisions. At the start, texts were distributed via the U.S. postal service, going online (RFC 580) and becoming machine-readable (RFC 582) in 1973. (The documents are now freely available, hosted by the Internet Engineering Task Force [IETF] at www.ietf.org.) Famously, it was the position of the network design group from the start that “Notes may be produced at any site by anybody and included in this series [the RFCs]” (RFC 3). Though ultimately the process of using the RFCs to reach agreement on an Internet protocol became formalized (RFC

2026; RFC 2028), participants in the conversation during its early years experienced the RFCs as creating

a positive feedback loop, with ideas or proposals presented in one RFC triggering another RFC with additional ideas, and so on. When some consensus (or at least a consistent set of ideas) had come together a specification document would be prepared. Such a specification would then be used as the base for implementations by the various research teams. (Leiner et al. 2010)

Documents published during the first decade of the series can be characterized in terms of dominant authorship, author uses of the RFC process, and the flavor of the discussion. About 75% of the documents published during the first decade appeared during 1971–1973.

Authorship

By the close of 2009, individuals from 44 countries had contributed to the RFCs, but during the first decade all documents were authored by individuals employed by institutions headquartered in the United States, with two exceptions: An author from Belgium submitted an announcement for an international conference on computer communications (RFC 631), and a Norwegian author commented on the proposed “host–host” protocol (RFC 65).

About 40% of the sole or first authors of documents were individuals employed by universities. Another 25% worked for consulting firm Bolt, Beranek and Newman (BBN), the ARPA grantee with responsibility for managing the network development process documented by the RFCs. Authors from research centers, dominated by the RAND Corporation with its U.S. Department of Defense contracts, also had a significant presence. Other institutional employers of RFC authors during the first decade of the document series included government agencies and corporations. MITRE Corporation, an extremely important contributor, is technically a nonprofit organization, but of a very particular type: It was created in 1958 to serve federal government agencies—including the U.S. Department of Defense—by supplying technologies, research, engineering, and project management on a contractual basis. For-profit corporations found the series useful as a competitive opportunity to promote their technologies, and by the fourth decade of the process had become the dominant players.

The initial desire to welcome all comers to the document series was tempered by practicalities. The quickly growing number of participants in the process became a design problem in and of itself (RFC 722). This introduced a tension. On the one hand, the need to set up committees to work on single issues needing focused attention was recognized fairly early (RFC 131). On the other hand, though, there were invitations that expanded the number

of participants in order to bring into the conversation individuals with pertinent expertise (RFC 632).

The numerical sequence of RFC documents does not always map onto the chronology of the discussions, decision making, and actions. In some cases, large clusters of documents all pertaining to interrelated decision making during a very short period of time inevitably had to be sequenced by number even though they were produced essentially simultaneously. In other cases, there is a lag between making a decision, acting upon that decision, and the ultimate production of a document for the series recording that decision. Some early documents disappeared altogether, and some were not made available to the public for intellectual property rights reasons. As the design process continued, documents not initially included in the RFC series were added and there were other idiosyncrasies produced by the process of collecting and indexing these materials; the 685 documents available from the launch of the series in 1969 through the close of 1979, numbered 1–758, thus actually include 1 item from 1968 and 1 from 1982.

Author Uses of the RFC Process

The inclusion of jokes, commentary on the design process itself, references to social interactions, and other nontechnical material in the document series provides evidence of the socialization functions of the conversation for those in the network community. Almost one-third of the documents during the first decade dealt with the development of protocols, while the rest of the texts provided such things as details of administrative support, history, communication theory, technical experiments, and use studies. (The percentage of documents devoted to protocol development went up after agreement was reached on TCP/IP in 1981, after the period discussed in this article.) Authors used the RFCs to seek reactions to possible approaches for resolving a problem (e.g., RFC 451), willingness to accept a new protocol (e.g., RFC 725), or responses to specific protocol details (e.g., RFC 573).

Foundational questions had to be addressed during the early years, such as just what those involved were making decisions about, and how consensual, usable, and effective decisions might be reached. It took some time before the need to begin by reaching agreement on the design goals themselves could be articulated (RFC 316), for “The tasks facing the ARPANet design teams were often unclear” (RFC 635, p. 1). Once that requirement was enunciated, there was then concern that it might not be accomplishable (RFC 647).

The fact that technical solutions continued to change once put in place was particularly vexing. The early concern that it would be impossible to make changes once the network was up and running (RFC 72) gave way to

questioning just how best to write protocols in such a way that change processes are facilitated (RFC 722). It was difficult to identify the moment in which change was dramatic enough that it justified treating a protocol as “new” rather than an incremental change to an existing standard (RFC 569). There was a sense that protocol change processes slowed down network development (RFC 72; RFC 569) and created dysfunctionalities (RFC 103; RFC 559)—but also that they stimulated further valuable innovations (RFC 569).

Some skeptics viewed the habit of tinkering with existing protocols as a matter of trendiness (RFC 451), and there were those within the network community who feared even incremental change (RFC 435). Efforts to quell such fear included providing examples of how to implement changes simply (RFC 435), and clarifying the difference between common practices and formally recognized protocols (RFC 724).

Flavor

Three factors combined to create a particularly informal, personal, and lively discussion during the first decade of this document series. Authors clearly perceived themselves to be members of a relatively small albeit fast-growing community involving individuals who, for many years, all knew each other. There appears to have been no expectation that anyone outside of that community would ever be looking at the documents. And, as expressly mandated in RFC 3, authors were encouraged to be informal, not to worry about format or length, and to present ideas at every stage of their development.

As a consequence, personalities ran rampant. Jon Postel presented a particularly vivid presence from the start, announcing himself the “czar” for names and numbers—what ultimately became the domain name and IP addressing systems—while still a graduate student (RFC 349). There are complaints about requests for information from each host (RFC 459), resentment of the fact that not everyone involved in the network community provides help to others (RFC 369), and outright insults (RFC 514).

The emergence of a sense of a network community within the first year of the design process (e.g., RFC 53) was of particular importance for everything that followed. This community was viewed as the arbiter of decision making for the establishment of what became known as “official” protocols (RFC 53), and had responsibility for finding flaws (“bugs”) in those standards (RFC 374). Those who deliberately fostered a sense of community in support of decision making were praised (e.g., RFC 101), while RFC authors who did not consult the community before offering a proposal were criticized (e.g., RFC 724). The desire to speak to the entire network community,

irrespective of technical specialization, affected how RFCs were written; those involved in developing the first draft graphics protocol, for example, explicitly took into account the possibility that the sheer unfamiliarity of what they were proposing could frighten those who were not themselves working with the problem (RFC 549).

Subcultures within the network community were recognized, including a distinction between the East Coast and West Coast cultures (RFC 557) (Hafner and Lyon's 1996 history of the period emphasizes the first of these, while Turner's 2006 history of the same period tells the story of the latter). Subcultural differences among site-specific communities of users with their own habits and practices (RFC 705) were also noted. The most common references, though, are to subcultures oriented around specific types of equipment, coding schemes, or languages. Insiders discuss the "revival meeting" air of get-togethers among those involved with a particular approach to coding, including piano sing-alongs and the development of an "emotional mystique" among users (RFC 468, p. 5). Such technical subcultures, of course, were linked at least as tightly to particular organizations as they are today (think Google), if not more so; sharing the details of a particular compression scheme, for example, was treated as a matter of "cultural enrichment" for those outside of IBM (RFC 468, p. 2).

USERS

Even in the early years, any site with a computer could become a network user—whether or not the site was supported by ARPA—simply by self-proclamation (RFC 597). RFCs reveal the attitudes of network designers towards users as well as distinctions among types of users considered pertinent to the design process. Some of the documents published have the explicit purpose of providing support to users.

Attitudes Toward Users

Although expanding the community of network users was put forward as a goal in the document that launched the series (RFC 1), the realities of the development process generated ambivalence and reluctance among designers regarding expansion. They became aware that each new type of user introduced demands that changed requirements for the network (RFC 468); even within the first decade, Internet designers had to rethink the fundamental Telnet protocol in order to support new types of facilities (RFC 731). Network designers gave lip service to heterogeneity among users and uses, and noted that cooperation among diverse types of users is key to achieving interoperability among systems (RFC 741). They found it difficult, however, to incorporate that awareness into their work,

typically assuming that all users will be working in technical environments like those of network designers (RFC 678).

Some of the technical problems that had to be addressed during the first decade derived from particular user practices and desires. Designers had to find a way of preventing a single user from monopolizing the network (RFC 105). Expectations of operator errors, and the need to repair damage caused by such errors, influenced design decisions (RFC 122). Organizational issues, such as the possibility that a remote user might be denied the support of staff at a host, were also taken into account (RFC 364). In some cases, explicit efforts by user groups regarding a technical decision made a difference, as when intense user lobbying prevented abandonment of a particular compression scheme (RFC 468).

Early authors of the RFCs recognized the need to understand how users experience the network and what they will actually do with it (see, e.g., RFC 150). As early as 1971, therefore, an RFC titled "User's View of the Data-computer" appears (RFC 219). Computer scientist James Howard, of the University of California–Santa Barbara, operated as a naive social scientist when he conducted the first study of the experience of users outside of the design community in 1972 in research that looked at the ways in which computer science graduate students figured out how to log into the network, and what they did when they got there (RFC 302; RFC 369). In 1973 an ARPANet Users Interest Group, open to anyone interested in participating, was formed to represent the needs of those beyond the design community within the decision-making process (RFC 584). (Later, nontechnical users were separated from those with technical expertise who joined the "USING" interest group.) Internationalization of the RFC discussion (beginning with RFC 101 in 1971) and of the network itself (first mentioned in RFC 67 in 1970) added to the heterogeneity of research interests to be served by the network and highlighted the importance of attending to language and cultural differences in the development of the protocols themselves.

Distinctions Among Types of Users

Those in the network community during the first decade became aware of several distinctions among types of users that were pertinent to design decisions. Given the overweening presence of military concerns, it is not surprising that the computer scientists and electrical engineers involved are constantly thinking about the possibility of malicious, as well as benign, users. Differences between programmers and nonprogrammers, and between Internet insiders and other computer scientists, provide further evidence of the development of a sense of a network community. The distinction between "human users" (human

beings who are launching, using, or reacting to processes) and “daemon users” (software programs and processes), is of particular interest from a policy perspective (Braman 2002).

Benign vs. malicious. Most of the types of malware we experience in the early 21st century were foreseen as possibilities early on. These show up first in discussion of how to design a user–host interface in RFC 49, where descriptions of “illegal privileges” to be protected against include the ability of one process to masquerade as another, eavesdrop on communications not intended for it, force the host to send out spurious network commands or messages, and/or prevent legitimate users from accessing a particular host. Other malicious user activities discussed during the first decade included feeding bad information to output (RFC 221) or input (RFC 505) devices, inappropriate capture of passwords (e.g., RFC 426), and the junk mail (RFC 463) that came to be known as spam. Additional language used to describe malicious users who would engage in such destructive activity includes “pathologic” foreign hosts (e.g., RFC 215), “hostile” programs (e.g., RFC 61), and entities attempting “unauthorized” use (e.g., RFC 542).

Since experimentation with telephone “phreaking”—hacking of the telephone network—began in the 1950s, it was to be expected that hackers showed up early in the life of the Internet. In December 1973, a humorous piece listing types of security violations included as an item the fact that there is “lingering affection” for breaking into networks (RFC 602, p. 1). (Bob Metcalfe, then of Xerox PARC and author of that RFC, went on to make clear he didn’t find this particularly impressive because it was widely known that it was easy to break into, and to crash, networks.) Two high school students in Los Angeles had by that point already compromised a system password. A minimal security standard was defined as putting in place enough protection so that users would not be “victimized by every first year Computer Science student with access to the net” (RFC 725, p. 1).

Technical insider vs. technical outsider. In the first decade, two sets of distinctions between technical insiders and outsiders were discussed. The first was between those with programming expertise and those without such skills. The second was between those technical experts who had been a part of the network design process and those who had not.

The initial users of the network were either programmers or research scientists supported by project-specific programmers, but network designers did appreciate the need to ensure that those who are not programmers and/or would not have access to programming expertise could also be users. This goal was one among the arguments used

to justify standardization of as many network processes as possible (RFC 231). For those who study relationships between tacit and codified knowledge, it is interesting to note that designers treated informal procedures developed at local sites to deal with inadequacies in the formal procedures that had been put in place as a valuable source of expertise regarding design for nonprogrammers (RFC 364).

The second type of distinction between technical insiders and outsiders—between those who had been involved in design of the network and those who had not—became visible as the network expanded. When new sites expressed an interest in joining the network, there was pressure on those within the existing community to provide assistance. This led to tension over whether ARPANet was a service network (which would require those involved to help new sites, particularly those of government agencies, onto the network) or an experiment in network research (which would justify focusing exclusively on network development) (RFC 446). When individuals outside of the network community began to connect their sites to the network on their own, however, jealousy arose. RFC 647 includes an outright attack on subcontractors who claim to have the skills necessary to do this work, asserting that all such efforts by network outsiders have resulted in horror stories.

Human vs. daemon. One of the most difficult matters for Internet designers has been the need to deal simultaneously with nonhuman and human users. In the language of the RFCs, a user could be a human or a “daemon”³ (e.g., RFC 114), a computer program or process also variously referred to as a user program (e.g., RFC 310), automaton (e.g., RFC 414), or higher level process (e.g., RFC 172). Some network processes were designed in such a way that no human user would be involved (RFC 230). Others took both human and daemon users into account. In a few instances, this was done in a manner neutral towards the distinction, as when it is noted in RFC 539 that a receiver would do the same thing with information received about a service interruption, whether that receiver were a process or a person. When either a human or a process could accomplish something, the daemon is preferred because protocol design is “weighted towards automata” (RFC 542, p. 6). Amusingly, though a particular mail protocol was not intended for use by humans, it was recognized that “nothing can stop a human user from [using the protocol] if he has access to a TELNET user program and is determined to do so” (RFC 555, p. 4, emphasis in the original). Similarly, File Transfer Protocol (FTP) was considered to be directly usable by a human, although it was designed primarily for use by daemons (RFC 354).

The need to accommodate human users in a design environment oriented more typically toward daemons required additional programming that often appeared to be undertaken grudgingly, even though thinking about human users also could generate results that expanded the options for network use. A discussion of a protocol for remote usage of a computing facility, for example, included the comment that

in recognition of the fact that the official protocol is biased heavily toward use by programs, and is therefore rather cumbersome for human users, an alternate, optional, command syntax has been provided. An attempt was made to make this alternate syntax, called ‘local syntax’, as ‘natural’ as possible to a human user. It also provides some features not available with the standard syntax. (RFC 477, p. 1)

In another formulation of the same point, a reply to an FTP command is described as including both a code and a text string, with the former intended for use by programs and the latter “usually” intended for use by humans (RFC 542).

Some RFCs manifest outright disdain for the ways in which humans make meaning—what humans need from and perceive in the information and messages delivered through the network. RFC 555, for example, argues that there is no point in defining numeric shades of urgency for reading or response to particular mail messages unless there is corresponding code for alternative computing processes in addition to alerting the human user. Other RFCs discuss human foibles, as in the comment that because human users care deeply about the names associated with their sites, “human considerations should probably prevail” when establishing the naming scheme (RFC 273) (emphasis added). In another example of bowing to perceived human preferences, it was argued that closing a session with “bye” rather than “quit” was just too cute to be acceptable (RFC 368).

Support for Users

Concern about designing the network in such a way that it would serve all users—current and future, for uses known and unknown, across cultures, and requiring varying degrees of technical skill—runs throughout. This in turn led to appreciation of the need for various kinds of user support. Occasionally, unforeseen advantages to users would accrue from technical decisions, as when new protocols for the network introduced local process to process communication within sites where that had not previously been possible (RFC 712). More typically, however, the result was deliberate efforts to provide administrative and didactic user support.

Administrative support. The RFCs were not the only support system for the Internet design process. Over the

first decade of the effort to build the network, additional forms of administrative underpinnings were put in place to ensure that the network would succeed, with documentation of their existence appearing within the RFCs. There were schedules for the incorporation of new technologies into the network (e.g., RFC 111), documentation policies (e.g., RFC 115), procedures for collecting and distributing statistics about network functioning (e.g., RFC 388), and the establishment of user directories (e.g., RFC 621). On the behavioral side, these years saw the development of foundations for or models of practices still in widespread use, such as accounting procedures (e.g., RFC 136) and establishment of what we now call a help desk (RFC 356). Though at least one RFC author considered a request for information to be included in a software repository a joke (RFC 514), this type of service also was invaluable to network users of the era.

Guides. In subsequent decades, the Internet RFCs became the source of guides for Internet use by the general public such as the first publication of rules for netiquette and instructions for what school systems should do when they connect. During the first decade, instruction was targeted more narrowly towards those involved in the Internet design process itself. There are guides for how to write an RFC (RFC 154), how to contribute to and draw upon network documentation (RFC 543), and how to use specific protocols (RFC 365). There are also a number of guides specific to particular systems (e.g., RFC 274) and pieces of equipment (e.g., RFC 395). There was even a manual on how to deal with daemons (RFC 714).

USES

The first exposure of the network to the general public took place in 1972, at the International Computer Communications Conference in Washington, DC. That conference was planned because, as Robert Kahn of BBN noted, “The social implications of this field are a matter of widespread interest that reaches society in almost all walks of life” (RFC 371, p. 1). By that time, some computer sites offered access to anyone who wanted to take advantage of networked resources. The creepage familiar to 21st-century system administrators was immediately noted: People would get on for one purpose and quickly wind up doing other things as well, a matter of concern to those institutions supporting network experimentation with ARPA funds but billing for other, “casual,” uses (RFC 350, p. 1). Irrespective of the accounting and technical problems associated with network expansion, this continued to be a goal requiring support. Designers thus paid attention to the growth in the extent of usage following the development of new computational features (RFC 662), for example, and recommended that those

facilities wishing to expand their customer base should think in marketing terms (RFC 231).

When RFC 5 noted that there would in future be a great heterogeneity of network “uses,” it was referring to what we today call “applications.” Within the document series, such applications are explored as a means of identifying the network functions—and the contexts for and constraints upon them—that would ultimately be needed. Both governmental and nongovernmental uses were discussed during the first decade of the design process. Uses for the purposes of this article, then, are applications, and functions are network features that must be in place in order for applications to work.

Government Uses

It is commonly, and accurately, believed that military concerns and the desire to expand research capacity were behind U.S. government interest in funding the effort to network computers (Norberg and O’Neill 1996). There is relatively little reference to military and research needs within the *texts* of RFCs during the first decade, however. One can find discussion of the need to get all of the ARPA contractors onto the network and talking with each other as quickly as possible (RFC 46). Concerns of the Air Force and of the general research community are mentioned (RFC 164), and the word “military” first appears in RFC 316, but few documents address such needs directly.

It was not, of course, that such conversations were not going on, but they were largely not taking place within the very public and unclassified RFCs. Within the document series, national security concerns show up indirectly, most often evident in document *authors* or *assumptions*, rather than texts. MITRE, a nonprofit consulting firm that had and has contracts with numerous U.S. government agencies including significant ongoing relationships with the Department of Defense, is mentioned on the second distribution list for the documents (RFC 24). It is MITRE that calls the meeting discussed in RFC 99, and with RFC 100, MITRE takes over the task of organizing and indexing the document series. MITRE authors begin to appear on substantive documents with RFC 138. In addition, there are regular references to the needs of the Department of Defense as the contracting agency, awareness that most national security-sensitive discussions are taking place in other venues, an underlying appreciation of security needs that colors the discussion as a whole during its early years, and repeated references to the possibility that there would be those who would attempt to use the network for destructive purposes in documents authored by entities working under contract to the U.S. Department of Defense or Department of Energy.⁴

Other network needs of the U.S. government discussed included air traffic control (RFC 549), the criminal justice system (RFC 144), and education (e.g., RFC 313). The weather service was of particular importance because it relied upon mammoth amounts of data acquired from and used by entities all over the world and supported a wide range of research agendas (e.g., RFC 420). Uses of the network for seismic monitoring (RFC 542), libraries (e.g., RFC 286), and medical institutional and informational needs (RFC 144) were not always directly under government control, but often received funding from the U.S. or other governments and served the public interest. What we now refer to as “e-government,” meaning network delivery of government services, was already under discussion in 1972 (RFC 371).

Nongovernment Uses

Though the Internet was not commercialized until 1993, at least some participants in the design process assumed from the start that the network would ultimately be used for commercial purposes. As early as 1971, there was discussion of network uses by commercial entities such as banks, warehouses, and corporations in the health care industry (RFC 144). Recognizing that their current sources of government funding would not be extended indefinitely, network designers tried to plan not only for their current organizational circumstances but also for situations in which services might instead be offered by for-profit corporations (RFC 164). Network hosts, or service centers, were encouraged to take a market-oriented approach to users (RFC 231). Large firms such as General Motors regularly attended design meetings to provide input on the nature of their internal computing and networking needs (RFC 316). The interests of individuals in networked activities were acknowledged as well, particularly in the areas of personal correspondence (RFC 196) and the playing of games such as chess (RFC 369).

Functions in Support of Uses

Some functions that had to be designed into the network were identified by working backward from particular uses. This was the case with interactive graphics, for example—a network use envisioned in the very first RFC (RFC 1). Some of the government’s most salient needs required artificial intelligence (RFC 152) and networked simulations (RFC 316).

Other functional needs became visible only through experience. As soon as hosts began offering services to users at sites other than their own, the need for identity verification was clear (e.g., RFC 46). Synchronicity across sites, requiring an accurate and coordinated network clock, was necessary in order to support distributed databases (e.g.,

RFC 146) and collect measurements about network performance (e.g., RFC 177). Once data and files started moving around, the question of just where they would be stored arose (e.g., RFC 122), as well as the problem of how to ensure data consistency for databases held at multiple locations or distributed across the network (e.g., RFC 144). The complexities of genuine interactivity made it a target function in its own right (e.g., RFC 178). Designers became interested in using the network for participating in conferences with each other though geographically distant (e.g., RFC 441). Once people started sending files to each other they became interested in using the network itself for text and file manipulation (e.g., RFC 369). The problem of how to find information at remote sites was, of course, overweening (e.g., RFC 180).

A third set of functional needs was driven by the nature of the network itself. The flexibility of the network was increased, for example, through virtualization (RFC 468). For network designers of the period, the “virtual” was conceived of as the use of “pseudo-devices” created through logical rather than physical manipulations in order to alter the functions of elements such as operating systems.

DESIGN CRITERIA AS POLICY PRINCIPLES

For a sociotechnical system, design criteria are, in essence, policy principles. Some of those developed for the Internet during the first decade of the design process ran very close to the ground. These focused on logistical issues that, while technical in nature, have social, political, economic, and cultural ramifications. Other early but enduring design criteria for the Internet explicitly identify social goals to be met by the network.

These design criteria evolved in response to the need to make decisions about matters such as accounting and access control requirements (RFC 167). When new users and networks joined the network of networks, they exacerbated appreciation of how little could be known about local operating conditions (RFC 318), the fact that the system was become ever more complex (RFC 708), and intensification of the need for interfaces between network insiders and network outsiders (RFC 647). The constant expansion of the number, range, and types of design decisions for which standardization was considered necessary is a topic requiring detailed examination beyond the scope of this article.

Logistical Design Criteria

While it may be hyperbolic to refer to the RFCs as the equivalent of the Articles of Confederation that were a step on the way to the U.S. Constitution (Loshin 2000), identifying resonances between technical and legal documents that are foundational for decision making for large-scale

sociotechnical systems can be useful for those involved in policy analysis. Constitutions of geopolitical governments deal not only with aspirational goals, but also with details necessary for daily governance. The same is true of the Internet RFCs. Such logistical policy principles embedded into the design criteria put forward in this document series include the need to sustain the design process itself and to create protocols that will remain compatible with each other across network scale and despite technological innovation. The concept of reliability as applied to the network during its first decade had two different meanings: It was used to refer both to the integrity of content transferred, and to the functioning of the network itself.

Sustaining the process. The pragmatic nature of some of the design criteria of importance during the early years will be familiar to anyone who has conducted research on an external grant: Stimulate immediate use by a wide range of types of users (RFC 1), get all ARPA contractors on the network and talking with each other quickly (RFC 46), impress the U.S. Department of Defense (RFC 82), and provide a simple service that can be used quickly (RFC 164). The need to please funders generated considerable tension, for designers were aware of the path-dependent long-term consequences of decisions made as a result of prioritizing getting the network running quickly rather than doing the best job possible (RFC 139).

Achieving specific goals in a timely manner was a critical dimension of the sustainability problem. Deadlines were set but not met: The group failed to create the network as a “solid working entity” by 1972 (RFC 164, p. 28) in the manner intended, for example, and most sites did not move to a new Telnet protocol on the date when the entire network was to switch (RFC 595). In a few instances, blame is attached to particular groups considered responsible for slowing down the network development process (e.g., RFC 131). It was more common, though, to simply bemoan what was at the time considered to be unexpected slowness in achieving project goals (e.g., RFC 679).

Content reliability. The computer scientists and electrical engineers who dominated the Internet RFCs document series during its first decade took a firm content-neutral approach, following the admonition to “assume nothing about the information and treat it as a bit stream whose interpretation is left to a higher level process, or a user” (RFC 172, p. 6)—a bit is a bit is a bit. The expectations were when content characteristics had concomitant technical requirements, as with graphics. Maintaining the integrity of all content transmitted, though, was considered extremely important. Possible transmission errors were divided into two types: failure to receive an expected message, and picking up an unexpected messages (RFC

203). There are, therefore, documents focused on trying to prevent data pollution (RFC 82) and responding to the reception of erroneous messages (RFC 46). Issues raised by distributed databases led designers to the problem of ensuring that updates in one location did not conflict with data held in another (RFC 144). There were efforts to prevent messages from being garbled (RFC 203) or lost (RFC 492). Even the sequence in which data managed by Telnet processes is presented on a terminal could affect its integrity (RFC 728). In the days in which there were few nodes on the network, individual sites were rated for their reliability on these matters (RFC 369). Expectations were realistic:

It is not our intent to propose the design of a perfect communication channel, rather it is our contention that in the real world there can be no perfect channels and that no amount of protocol can insure [sic] the error free transfer of information. Our goal is to explicate the various types of errors that are possible and to provide for each techniques of detection and recovery that, at a cost, can be made arbitrarily good. In this way the mean time between undetected errors can be made as large as necessary. (RFC 203, p. 1)

Network reliability. Concerns about network reliability were serious enough that they were used as a justification for not attempting to develop graphics protocols right away (RFC 282). Noting that “Our idea of the Network has evolved as the Network itself has grown” (RFC 528, p. 1), by 1973 designers realized that their initial expectation that only the communication circuits could fail was simply not correct. Not only was it possible for any element of the system to fail, but often those operating the network had to remotely diagnose problems that were neither categorized nor understood.

Early experiences with network reliability were dismal by contemporary standards. In 1971 it was considered a success when the network crashed only every couple of days (RFC 153). (A “crash” was defined as the total failure of an “IMP,” the equipment linking a computer to the network, possibly affecting equipment with which it was linked [RFC 528].) No respondent in a 1972 survey reported local experience with “mean-time-between-failure” at more than 2 hours, and the average time for “trouble-free” operation was only 35% (RFC 369).

Events that could bring the network down of which designers were aware during the early years included sending too much traffic to a single site on the network (RFC 334), operator inexperience (RFC 369), corrupted (RFC 467) or incorrect (RFC 528) control messages, and problems with the electronic equipment (RFC 528). Some types of crashes resulted from what communication theorists would describe as the “metadiscourse” of network communication (Craig 2008), meaning communications about the communication taking place. These occur when there

are miscommunication about network status (RFC 386), or users do not know that commands have been aborted (RFC 122).

Compatibility. Several types of compatibility are discussed, in addition to the familiar concepts of that which runs backward (RFC 687) and forward (RFC 663). Commands for a protocol at one level of the network must be interpretable at higher levels of the network (“upward” compatibility) (RFC 292). Development of “front-ending” techniques is recommended as a way of making legacy equipment compatible with networking requirements (RFC 666). There was also concern about compatibility across networks, given the conceptualization of what was being built as a multinet environment (RFC 730).

Forward compatibility, also referred to as “extensibility,” was considered necessary because it was assumed the network would continue to evolve under conditions of path dependency. The still-familiar habit of 6-month targets for upgrades was recommended as a means of bundling a number of incremental changes and implementing them in such a way as to minimize disruptions due to network improvements in RFC 82. This evolution was expected to take place along two dimensions: (1) *Growth in network size*: Protocols had to be designed in such a way that they would continue to work with an increase in the number of network nodes (RFC 192),⁵ network volume (RFC 251), and the spatial expanse of the network whether wired (RFC 355) or wireless (RFC 346). (2) *Multiplication of network processes*: The network had to develop in such a way that it could handle the addition of diverse network functions and processes (RFC 310), service to users not associated with a host site (RFC 462), and irrespective of complex and perhaps idiosyncratic uses of computers linked to the network (RFC 647).

Social Design Criteria

A number of design criteria oriented toward society rather than technology were also formulated by Internet designers during the period 1969–1979. The explicit intention to design the network for all imaginable and not-yet-imagined users and uses was discussed earlier. Additional social design criteria of particular importance for policy analysts include embedding a proprioceptive capacity into the network, technological democracy, a consistent focus on facilitating telepresent distributed and distant computing, and the assumption that it is a social good to continue to promote innovation *qua* innovation.

Network proprioception. Proprioception is the ability of an organism to sense its own organs and movements. A network becomes proprioceptive when collecting

information about the network and network activity and using that information as an input into networking processes are required design features. Doing so became an Internet design criterion because it was believed necessary in order to further the design process, support users, and maximize use of the network (RFC 573). Experience had shown that without such information, the network would crash (RFC 626).

Complaints about the inadequacy of such information appeared early on (RFC 369). The structure of the network at the time, however, created methodological difficulties. Initially, there were so many steps between where events occur and the site at which such knowledge was compiled and analyzed that gathering and using the information placed great demands on machinic and human resources and introduced delays and errors. The design response was to embed information collection throughout the network as well as ongoing processes transmitting data collected to a central site for analysis (RFC 462). This in itself, however, slowed down the network by adding so many messages to the ongoing flow (RFC 619). Those involved in designing the Internet during its first decade were also aware that it was impossible to know if network conditions were actually the same from one test to the next (RFC 508). Successfully balancing the various requirements for network proprioception—the network’s knowledge of itself—was not achieved during the first decade, though the collection, distribution, and use of information about the network was mandated as a key design criterion/policy principle.

Technological democracy. One of the most critical design criteria *cum* policy principles for the Internet was to insist upon technological democracy—creation of a network that can be used by those with any type of equipment, from the most humble and limited to the most sophisticated and capacity-rich. While many factors influence the differential rate of diffusion of the Internet across and within societies, it was this insistence upon the principle of technological democracy that has made it possible for the network to become genuinely global and access and use to be nearing universal.

This principle grew out of early assumptions that the number of computers on the network would keep expanding, and that users would continue to vary in the level of sophistication of their equipment, computing capacity, and network bandwidth/speed. It was thus believed necessary to design for a highly heterogeneous hardware (e.g., RFC 55), software (e.g., RFC 80), and content (e.g., RFC 82) environment. Designers were encouraged to think in terms of situations more complex and general than those raised by the details of the problems to which they were immediately responding (RFC 94) and to design for worst-case scenarios (RFC 60). The best way of doing so, it was suggested, was to “stay aloof from the eccentricities of

present day machine organization” (RFC 150, p. 2) while simultaneously creating protocols that could be easily implemented under the size and complexity restrictions then in place for the smallest of the hosts involved in the network (RFC 167). One manifestation of this approach in practice was the computer industry design philosophy of beginning with the most primitive elements of a program and then building ever-more-powerful tools upon that base (RFC 192).

Achieving the goal of technological democracy was made difficult by the heterogeneity of elements of the network themselves (RFC 285), the fact that not every node was on the network at all times during the early years (RFC 54), and differences between physical and logical structures (e.g., RFC 58). Specific types of hardware were more difficult than others to accommodate (e.g., RFC 282), and sometimes it was felt necessary to in essence “trick” the network in order to make something work (RFC 60). Market pressures for diversity were also acknowledged (RFC 192).

Telepresent distant and distributed computing. In a technological equivalent of the concept of the “glocal”—manifestations of globalized phenomena and processes in the local, and of the most local phenomena and processes in the global—another design criterion *cum* policy principle for the Internet was to achieve a computing environment in which it would be possible to experience both processing at a distance and distributed processing as if they were taking place at the physical site of the user. This technologically mediated sense that the distant is locally present has been termed “telepresence” by social scientists (Steuer 1992). Achieving telepresence was a design criterion, whether that computing was done in batch mode at a single machine or involved processing distributed across several computers. Diverse activities all pointed to the importance of this principle, including data sharing (RFC 144), cooperation among software processes at several sites (RFC 302), and the particular problems raised by the processing of large databases that may themselves be spread across multiple hosts (RFC 299). In one early formulation, success in making the user’s console appear to be directly connected to a machine in another location was described as having achieved network transparency (RFC 339). It was understood that transmission speed was essential if real-time conferencing were to become possible (RFC 508), a practice soon referred to as teleconferencing (RFC 647).

Stimulate innovation qua innovation. Recognizing that technological innovation would be ongoing (RFC 184), Internet designers were constantly moving back and forth between the ideal and the real, and between ultimate goals and current needs, in their thinking. This required

attention to the sequencing of decision making about matters that were new to all involved, and appreciation of the inevitable shaping of future possible paths of innovation through early design decisions. Interactions between networking and the need to promote innovation to increase computing capacity—an issue treatable by policymakers through the use of such tools as research funding, tax breaks, adaptation of antitrust (competition) laws, and government procurement—were acknowledged (RFC 420).

Design approaches to supporting and encouraging ongoing technological innovation included incorporating the ability to send equipment-specific information through the network as part of a protocol (RFC 184), designing the naming system to be independent of specific computers that will become obsolesced in favor of orienting around longer lived organizations (RFC 237), and using device buffers to accommodate demands for new features (RFC 365). There were mandates to design protocols in such a way that they would be usable by new equipment as it came on the market (RFC 192), and to provide a foundation for the creation of protocols operating at higher levels of the network (RFC 435) under conditions of greater network complexity (RFC 524). The need to continually update users' guides to incorporate innovations was discussed (RFC 386), as was the fact that things that worked one day might not the next (RFC 420).

The nature (RFC 72) and complexity (RFC 167) of the design process itself were seen as limits on achieving perfect solutions to problems at any given time. Designers challenged each other on the correctness of the assumptions underlying protocol proposals (RFC 195). Trade-offs between multiple goals had to be made in the specifications for any given protocol (RFC 60). Finally, but far from trivially, Internet designers admitted that they simply did not fully understand the network issues with which they were trying to grapple (e.g., RFC 313)—as did governmental policymakers of the period (Braman 1995).

CONCLUSIONS

Those responsible for technical design of the Internet from 1969 to 1979 were aware that their preliminary, often tentative and experimental, decisions were important, but did not always foresee the long-term consequences of their decisions from the perspective of their functions as policy frames. This should not be a surprise: These individuals and their employer organizations had to meet the immediate demands of funders, feed ambitious imaginations of possible futures, and serve users with both operational and developmental needs. And they had to do these things under operating conditions marked by constant shifts in the subject of their efforts, in the technologies and theories available, in the terminology used to discuss what they

were trying to build, and in the level of formalization of decision-making processes and the attendant discourse.

The development of foundational frames for policy-making and analysis was an unintended side effect of the network design process. All social policy is the result of compromises among very different views about extremely complex systems. Inevitably, then, the suite of frames put in place by the close of the first decade was neither comprehensive in nature nor internally coherent. The ideal of an open decision-making process that welcomes all comers is familiar from the world of geopolitical democracy—as are the consequent dismay over the problems that come with such openness and a desire to restrict the group of those involved in number and/or by subject expertise. Logistical design criteria/policy principles are relatively concrete, just as are rules about voting and establishment of parliamentary bodies in politically constitutive documents. Social design criteria/policy principles, on the other hand, are more aspirational in nature, as is again the case with political principles.

One of the most interesting questions raised by the study of large-scale sociotechnical systems is the extent to which thinking about the design of social and of technical systems converges and the ways in which it diverges. This study of policy frames developed by those responsible for technical design of the Internet during the first decade of that design process found three types of answers to that question. (1) In some instances, we can identify parallels between the policy principles established during the first decade of the Internet design process and those that are important in the geopolitical world, as with the goal of technological democracy. In another example, network proprioception has parallels with constitutional supports for geopolitical state acquisition of information about itself. (2) In a second grouping are concepts that are relatively unformed and likely to remain contested over time, as with the multiple layers of network insiders and outsiders and consequent issues of compliance that these distinctions suggest. (3) A third set of the policy frames described here departs from those familiar to those who study geopolitical decision making. Most important among these is the simultaneous focus on serving machinic as well as human users and the development of the habit of preferencing the former over the latter should needs come into conflict.

Work in progress on analyses of specific policy issues within the RFC document series, on the evolution of formal policymaking processes out of this informal conversation, on the transformation of technical problems into social policy issues, and on techniques for decision making for large-scale sociotechnical infrastructure under conditions of profound instability will examine how the policy frames identified here become translated into design decisions and practices. Perceptions of problematics and

policy issues by those responsible for technical design of what we now refer to as the Internet during its first decade have had path-dependent effects for conversations, ideas, and policy proposals within both the legal and technical communities. Though the languages of computer science, social science, and the law are very different from each other, analysis of the Internet RFCs as a discourse makes clear that there are so many agreements over identification of the critical issues and resonances in the ways in which they are thought about—as well as differences in approaches to thinking about solutions to problems—that it would be useful to all to enter a common conversation in order to learn from each other and think together about this shared sociotechnical space. Doing so is as likely to shed light on the nature of our geopolitical systems as on the Internet.

NOTES

1. When RFC documents are occasional entries in standard reference lists, the Internet Engineering Task Force (IETF) recommends a standard scholarly reference style beginning with authors' last names. However, in texts devoted to the RFCs, as in oral conversation and long-standing habit, the documents are referred to by number, and that is the citation practice used here.
 2. This material is based on work supported by the National Science Foundation under Grant No. 0823265. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author. Thanks to networking experts David Stack and David Crass of the University of Wisconsin-Milwaukee's University Information Technology Services for their assistance in translating the technical language of the RFCs into terms understandable by the layperson, to project manager Alyse Below, and to research assistants Liza Barry-Kessler, Andrew Cole, and Jolie LeGate for all of their contributions.
 3. In ancient Greek mythology, the concept of the "daemon" referred to agents capable of action that are neither human nor divine. The concept entered Western philosophy via Plato's *Symposium*.
 4. It is the Department of Energy that handles nuclear matters in the United States.
 5. According to Jon Postel, when the process of designing the Internet began they had had no idea that the network would ever include as many as 256 hosts. When the size of the address field was expanded in 1975, the change made it possible to have as many as 16,777,216 hosts on the network (RFC 690).
- Braman, S. 2002. Posthuman law: Information policy for the machinic world. *First Monday* 7(12). http://firstmonday.org/issues/issue7_12/braman/index.html (accessed July 1, 2011).
- Braman, S. 1995. Policy for the Net and the Internet. *Annual Review of Information Science and Technology* 30:5–75.
- Carpenter, B. E., and C. Partridge. 2010. Internet Requests for Comments (RFCs) as scholarly publications. *SIGCOMM Computer Communication Review* 40(1):31–33.
- Craig, R. T. 1999. Metadiscourse, theory, and practice. *Research on Language and Social Interaction* 32:21–29.
- Davidson, A., J. Morris, and R. Courtney. 2002. Strangers in a strange land: Public interest advocacy and Internet standards. Paper presented to the Telecommunications Policy Research Conference, Arlington, VA, September.
- Dell, P. 2010. Two economic perspectives on the IPv6 transition. *Info* 12(4):3–14.
- Denardis, L. 2009. *Protocol politics: The globalization of Internet governance*. Cambridge, MA: MIT Press.
- Fischer, F., and J. Forester, eds. 1993. *The argumentative turn in policy analysis*. Durham, NC: Duke University Press.
- Flack, C. 2006. The subtle violence of nonviolent language. *Cross Currents* 56(3):312–329.
- Franklin, J., V. Paxson, A. Perrig, and S. Savage. 2007. An inquiry into the nature and causes of the wealth of Internet miscreants. Presented to CCS '07, October/November, Alexandria, VA.
- Gupta, A., and S. Seshasai. 2007. 24-Hour knowledge factory: Using Internet technology to leverage spatial and temporal separations. *ACM Transactions on Internet Technology* 7(3):article 14. doi: acm.org/10.1145/1275505.1275507.
- Hafner, K., and M. Lyon. 1996. *Where wizards stay up late: The origins of the Internet*. New York: Simon & Schuster.
- Haigh, T. 2004. The history of computing: An introduction for the computer scientist. In A. Aker and W. Aspray, eds., *Using history to teach computer science and related disciplines*, pp. 5–26. Washington, DC: Computing Research Association.
- Hauben, M., and R. Hauben. 1997. *Netizens: On the history and impact of Usenet and the Internet*. Los Alamitos, CA: IEEE Computer Society Press.
- Kahin, B., and J. Keller, eds. 1997. *Public access to the Internet*. Cambridge, MA: MIT Press.
- Kuerbis, B., and M. L. Mueller. 2011. Securing the root. In L. Denardis (Ed.), *Opening standards: The global politics of interoperability*, pp. 45–62. Cambridge, MA: MIT Press.
- Latour, B. 1993. *We have never been modern*. Cambridge, MA: Harvard University Press.
- Leiner, B. M., V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. Wolff. 2010. *A brief history of the Internet*. www.isoc.org/internet/history/brief.shtml (accessed Oct. 24, 2010).
- Levä, T., H. Hämmäinen, and K. Kilkki. 2009. *Scenario analysis on future Internet*. Presented to the International Telecommunications Society, Tokyo, June.
- Loshin, P. 2000. *Big book of Lightweight Directory Access Protocol (LDAP) RFCs*. San Diego: Academic Press.
- Mowery, D. C., and T. Simcoe. 2002. Is the Internet a US invention? An economic and technological history of computer networking. *Research Policy* 31(8–9):1369–1387.
- Mueller, M. L. 1989. *Ruling the root: Internet governance and the taming of cyberspace*. Cambridge, MA: MIT Press.

REFERENCES

- Abbate, J. 1999. *Inventing the Internet*. Cambridge, MA: MIT Press.
- Bowker, G. C. 1996. How things change: The history of sociotechnical structures. *Social Studies of Science* 26(1): 173–182.
- Braman, S. In press. Privacy for networked computing, 1969–1979. *New Media & Society*.
- Braman, S. 2010. The interpenetration of technical and legal decision-making for the Internet. *Information, Communication & Society* 13(3): 309–324.

- Nickerson, J. V., and M. zur Muehlen. 2006. The ecology of standards processes: Insights from Internet standard making. *MIS Quarterly* 30:467–488.
- Norberg, A. L., and J. E. O’Neill. 1996. *Transforming computer technology: Information processing for the Pentagon, 1962–1986*. Baltimore, MD: Johns Hopkins University Press.
- Rasmussen, T. 2007. *Techno-politics, Internet governance and some challenges facing the Internet*. Oxford: Oxford Internet Institute Research Report 15.
- Rosenzweig, R. 1998. Wizards, bureaucrats, warriors, and hackers: Writing the history of the Internet. *American Historical Review* 103(5):1530–1552.
- Schön, D., and M. Rein. 1994. *Frame reflection: Toward the resolution of intractable policy controversies*. New York: Basic Books.
- Schneiderman, B., and A. Rose. 1995. Social impact statements: Engaging public participation in information technology design. In *Proceedings, Symposium on Computers and the Quality of Life*, pp. 90–96. Philadelphia, PA: Association of Computing Machinery.
- Steuer, J. 1992. Defining virtual reality: Dimensions determining telepresence. *Journal of Communication* 42(4):73–93.
- Tötterman, J. 2009. *Internet standardization: A participant analysis*. Master’s thesis, Faculty of Electronics, Communications and Automation, presented to the Helsinki University of Technology, Helsinki, Finland.
- Turner, F. 2006. *From counterculture to cyberculture: Stewart Brand, the Whole Earth Network, and the rise of digital utopianism*. Chicago: University of Chicago Press.
- Yu, P. K. 2004. The origins of ccTLD policymaking. *Cardozo Journal of International and Comparative Law* 12:387–408.
- RFC 100, *Categorization and Guide to the NWG/RFCs*, P. M. Karp, February 1971.
- RFC 101, *Notes on the Network Working Group Meeting*, R. W. Watson, February 1971.
- RFC 103, *Implementation of Interrupt Keys*, R. B. Kalin, February 1971.
- RFC 105, *Network Specifications for Remote Job Entry and Remote Job Output Retrieval at UCSB*, J. E. White, March 1971.
- RFC 111, *Pressure from the Chairman*, S. D. Crocker, March 1971.
- RFC 114, *File Transfer Protocol*, A. K. Bhushan, April 1971.
- RFC 115, *Some Network Information Center Policies on Handling Documents*, R. W. Watson, J. B. North, April 1971.
- RFC 122, *Network Specifications for UCSB’s Simple-Minded File System*, J. E. White, April 1971.
- RFC 131, *Response to RFC 116: May NWG Meeting*, E. Harslem, J. F. Hafner, April 1971.
- RFC 136, *Host Accounting and Administrative Procedures*, R. E. Kahn, April 1971.
- RFC 138, *Status Report on Proposed Data Reconfiguration Service*, R. H. Anderson, V. G. Cerf, E. Harslem, J. F. Hafner, J. Madden, R. M. Metcalfe, A. Shoshani, J. E. White, D. C. M. Wood, April 1971.
- RFC 139, *Discussion of Telnet Protocol*, T. C. O’Sullivan, May 1971.
- RFC 144, *Data Sharing on Computer Networks*, A. Shoshani, April 1971.
- RFC 146, *Views on Issues Relevant to Data Sharing on Computer Networks*, P. M. Karp, D. B. McKay, D. C. M. Wood, May 1971.
- RFC 150, *Use of IPC Facilities: A Working Paper*, R. B. Kalin, May 1971.
- RFC 152, *SRI Artificial Intelligence Status Report*, M. Wilber, May 1971.
- RFC 153, *SRI ARC-NIC Status*, J. T. Melvin, R. W. Watson, May 1971.
- RFC 154, *Exposition Style*, S. D. Crocker, May 1971.
- RFC 164, *Minutes of Network Working Group Meeting*, J. F. Hafner, May 1971.
- RFC 167, *Socket Conventions Reconsidered*, A. K. Bhushan, R. M. Metcalfe, J. M. Winett, May 1971.
- RFC 172, *The File Transfer Protocol*, A. Bhushan, B. Braden, W. Crowther, E. Harslem, J. Hafner, A. McKenzie, J. Melvin, B. Sundberg, D. Watson, J. White, June 1971.
- RFC 177, *Device Independent Graphical Display Description*, J. McConnell, June 1971.
- RFC 178, *Network Graphic Attention Handling*, I. W. Cotton, June 1971.
- RFC 180, *File System Questionnaire*, A. M. McKenzie, June 1971.
- RFC 184, *Proposed Graphic Display Modes*, K. C. Kelley, July 1971.
- RFC 192, *Some Factors which a Network Graphics Protocol Must Consider*, R. W. Watson, July 1971.
- RFC 195, *Data Computers: Data Description and Access Language*, G. H. Mealy, July 1971.
- RFC 196, *Mail Box Protocol*, R. W. Watson, July 1971.
- RFC 203, *Achieving Reliable Communication*, R. B. Kalin, August 1971.
- RFC 215, *NCP, ICP, and Telnet: The Terminal IMP Implementation*, A. M. McKenzie, August 1971.
- RFC 219, *User’s View of the Datacomputer*, R. Winter, September 1971.
- RFC 221, *Mail Box Protocol: Version 2*, R. W. Watson, August 1971.
- RFC 230, *Toward Reliable Operation of Minicomputer-Based Terminals on a TIP*, T. Pyke, September 1971.

RFCS CITED

- RFC 1, *Host Software*, S. Crocker, April 1969.
- RFC 3, *Documentation Conventions*, S. Crocker, April 1969.
- RFC 24, *Documentation Conventions*, S. D. Crocker, November 1969.
- RFC 46, *ARPA Network Protocol Notes*, E. Meyer, April 1970.
- RFC 49, *Conversations with S. Crocker (UCLA)*, E. Meyer, April 1970.
- RFC 53, *Official Protocol Mechanism*, S. D. Crocker, June 1970.
- RFC 54, *Official Protocol Proffering*, S. D. Crocker, J. Postel, J. Newkirk, M. Kralej, June 1970.
- RFC 55, *Prototypical Implementation of the NCP*, J. Newkirk, M. Kralej, J. Postel, S. D. Crocker, June 1970.
- RFC 58, *Logical Message Synchronization*, T. P. Skinner, June 1970.
- RFC 60, *Simplified NCP Protocol*, R. B. Kalin, July 1970.
- RFC 61, *Note on Interprocess Communication in a Resource Sharing Computer Network*, D. C. Walden, July 1970.
- RFC 65, *Comments on Host-Host Protocol Document #1*, D. C. Walden, August 1970.
- RFC 67, *Proposed Change to Host/IMP Spec to Eliminate Marking*, W. R. Crowther, January 1970.
- RFC 72, *Proposed Moratorium on Changes to Network Protocol*, R. D. Bressler, June 1970.
- RFC 80, *Protocols and Data Formats*, E. Harslem, J. F. Hafner, December 1970.
- RFC 82, *Network Meeting Notes*, E. Meyer, December 1970.
- RFC 94, *Some Thoughts on Network Graphics*, E. Harslem, J. F. Hafner, February 1971.
- RFC 99, *Network Meeting*, P. M. Karp, February 1971.

- RFC 231, *Service Center Standards for Remote Usage: A User's View*, J. F. Heafner, E. Harslem, September 1971.
- RFC 237, *NIC View of Standard Host Names*, R. W. Watson, October 1971.
- RFC 251, *Weather Data*, D. Stern, October 1971.
- RFC 273, *More on Standard Host Names*, R. W. Watson, October 1971.
- RFC 274, *Establishing a Local Guide for Network Usage*, E. Forman, November 1971.
- RFC 282, *Graphics Meeting Report*, M. A. Padlipsky, December 1971.
- RFC 285, *Network Graphics*, D. Huff, December 1971.
- RFC 286, *Network Library Information System*, E. Forman, December 1971.
- RFC 292, *Graphics Protocol: Level 0 Only*, J. C. Michener, I. W. Cotton, K. C. Kelley, D. E. Liddle, E. Meyer, January 1972.
- RFC 299, *Information Management System*, D. Hopkin, February 1972.
- RFC 302, *Exercising the ARPANET*, R. F. Bryan, February 1972.
- RFC 310, *Another Look at Data and File Transfer Protocols*, A. K. Bhushan, March 1972.
- RFC 313, *Computer Based Instruction*, T. C. O'Sullivan, March 1972.
- RFC 316, *ARPA Network Data Management Working Group*, D. B. McKay, A. P. Mullery, February 1972.
- RFC 318, *Telnet Protocols*, J. Postel, April 1972.
- RFC 334, *Network Use on May 8*, A. M. McKenzie, May 1972.
- RFC 339, *MLTNET: A "Multi Telnet" Subsystem for Tenex*, R. Thomas, May 1972.
- RFC 346, *Satellite Considerations*, J. Postel, May 1972.
- RFC 349, *Proposed Standard Socket Numbers*, J. Postel, May 1972.
- RFC 350, *User Accounts for UCSB On-Line System*, R. Stoughton, May 1972.
- RFC 354, *File Transfer Protocol*, A. K. Bhushan, July 1972.
- RFC 355, *Response to NWG/RFC 346*, J. Davidson, June 1972.
- RFC 356, *ARPA Network Control Center*, R. Alter, June 1972.
- RFC 364, *Serving Remote Users on the ARPANET*, M. D. Abrams, July 1972.
- RFC 365, *Letter to All TIP Users*, D. C. Walden, July 1972.
- RFC 368, *Comments on "Proposed Remote Job Entry Protocol"*, R. T. Braden, July 1972.
- RFC 369, *Evaluation of ARPANET Services January-March, 1972*, J. R. Pickens, July 1972.
- RFC 371, *Demonstration at International Computer Communications Conference*, R. E. Kahn, July 1972.
- RFC 374, *IMP System Announcement*, A. M. McKenzie, July 1972.
- RFC 386, *Letter to TIP Users-2*, B. Cosell, D. C. Walden, August 1972.
- RFC 388, *NCP Statistics*, V. Cerf, August 1972.
- RFC 395, *Switch Settings on IMPs and TIPs*, J. M. McQuillan, October 1972.
- RFC 414, *File Transfer Protocol (FTP) Status and Further Comments*, A. K. Bhushan, December 1972.
- RFC 420, *CCA ICCC Weather Demo*, H. Murray, January 1973.
- RFC 426, *Reconnection Protocol*, R. Thomas, January 1973.
- RFC 435, *Telnet Issues*, B. Cosell, D. C. Walden, January 1973.
- RFC 441, *Inter-Entity Communication: An Experiment*, R. D. Bressler, R. Thomas, January 1973.
- RFC 446, *Proposal to Consider a Network Program Resource Notebook*, L. P. Deutsch, January 1973.
- RFC 447, *IMP/TIP Memory Retrofit Schedule*, A. M. McKenzie, January 1973.
- RFC 451, *Tentative Proposal for a Unified User Level Protocol*, M. A. Padlipsky, February 1973.
- RFC 459, *Network Questionnaires*, W. Kantrowitz, February 1973.
- RFC 462, *Responding to User Needs*, J. Iseli, D. Crocker, February 1973.
- RFC 463, *FTP Commands and Response to RFC 430*, A. K. Bhushan, February 1973.
- RFC 467, *Proposed Change to Host-Host Protocol: Resynchronization of Connection Status*, J. D. Burchfiel, R. S. Tomlinson, February 1973.
- RFC 468, *FTP Data Compression*, R. T. Braden, March 1973.
- RFC 477, *Remote Job Service at UCSB*, M. Krilanovich, May 1973.
- RFC 492, *Response to RFC 467*, E. Meyer, April 1973.
- RFC 505, *Two Solutions to a File Transfer Access Problem*, M. A. Padlipsky, June 1973.
- RFC 508, *Real-Time Data Transmission on the ARPANET*, L. Pfeifer, J. McAfee, May 1973.
- RFC 514, *Network Make-Work*, W. Kantrowitz, June 1973.
- RFC 524, *Proposed Mail Protocol*, J. E. White, June 1973.
- RFC 528, *Software Checksumming in the IMP and Network Reliability*, J. M. McQuillan, June 1973.
- RFC 528, *Software Checksumming in the IMP and Network Reliability*, J. M. McQuillan, June 1973.
- RFC 539, *Thoughts on the Mail Protocol Proposed in RFC 524*, S. D. Crocker, J. Postel, July 1973.
- RFC 542, *File Transfer Protocol*, N. Neigus, August 1973.
- RFC 543, *Network Journal Submission and Delivery*, N. D. Meyer, July 1973.
- RFC 549, *Minutes of Network Graphics Group Meeting*, J. C. Michener, July 1973.
- RFC 555, *Responses to Critiques of the Proposed Mail Protocol*, J. E. White, July 1973.
- RFC 557, *Revelations in Network Host Measurements*, B. D. Wessler, August 1973.
- RFC 559, *Comments on the New Telnet Protocol and its Implementation*, A. K. Bhushan, August 1973.
- RFC 569, *NETED: A Common Editor for the ARPA Network*, M. A. Padlipsky, October 1973.
- RFC 573, *Data and File Transfer—Some Measurement Results*, A. Bhushan, September 1973.
- RFC 580, *Note to Protocol Designers and Implementers*, J. Postel, October 1973.
- RFC 582, *Comments on RFC 580: Machine Readable Protocols*, R. Clements, November 1973.
- RFC 584, *Charter for ARPANET Users Interest Working Group*, J. Iseli, D. Crocker, R. Neigus, November 1973.
- RFC 595, *Second Thoughts in Defense of the Telnet Go-Ahead*, W. Hathaway, December 1973.
- RFC 597, *Host Status*, N. Neigus, E. J. Feinler, December 1973.
- RFC 602, *"The Stockings were Hung by the Chimney with Care"*, R. M. Metcalfe, December 1973.
- RFC 604, *Assigned Link Numbers*, J. Postel, December 1973.
- RFC 613, *Network Connectivity: A Response to RFC 604*, A. M. McKenzie, January 1974.
- RFC 619, *Mean Round-Trip Times in the ARPANET*, W. Naylor, H. Operderbeck, March 1974.
- RFC 621, *NIC User Directories at SRI ARC*, M. D. Kudlick, March 1974.

- RFC 626, *On a Possible Lockup Condition in IMP Subnet Due to Message Sequencing*, L. Kleinrock, H. Operdeck, March 1974.
- RFC 631, *International Meeting on Minicomputers and Data Communication: Call for Papers*, A. Danthine, April 1974.
- RFC 632, *Throughput Degradations for Single Packet Messages*, H. Operdeck, May 1974.
- RFC 635, *Assessment of ARPANET Protocols*, V. Cerf, April 1974.
- RFC 647, *Proposed Protocol for Connecting Host Computers to ARPA-Like Networks via Front End Processors*, M. A. Padlipsky, November 1974.
- RFC 662, *Performance Improvement in ARPANET File Transfers from Multics*, R. Kanodia, November 1974.
- RFC 663, *Lost Message Detection and Recovery Protocol*, R. Kanodia, November 1974.
- RFC 666, *Specification of the Unified User-Level Protocol*, M. A. Padlipsky, November 1974.
- RFC 675, *Specification of Internet Transmission Control Program*, V. Cerf, Y. Dalal, C. Sunshine, December 1974.
- RFC 678, *Standard File Formats*, J. Postel, December 1974.
- RFC 679, *February, 1975, Survey of New-Protocol Telnet Servers*, D. W. Dodds, February 1975.
- RFC 687, *IMP/Host and Host/IMP Protocol Changes*, D. C. Walden, June 1975.
- RFC 690, *Comments on the Proposed Host/IMP Protocol Changes*, J. Postel, June 1975.
- RFC 705, *Front-End Protocol B6700 Version*, R. F. Bryan, November 1975.
- RFC 707, *High-Level Framework for Network-Based Resource Sharing*, J. E. White, December 1975.
- RFC 708, *Elements of a Distributed Programming System*, J. E. White, January 1976.
- RFC 712, *Distributed Capability Computing System (DCCS)*, J. E. Donnelley, February 1976.
- RFC 714, *Host-Host Protocol for an ARPANET-Type Network*, A. M. McKenzie, April 1976.
- RFC 722, *Thoughts on Interactions in Distributed Services*, J. Haverty, September 1972.
- RFC 724, *Proposed Official Standard for the Format of ARPA Network Messages*, D. Crocker, K. T. Pogran, J. Vittal, D. A. Henderson, May 1977.
- RFC 725, *RJE Protocol for a Resource Sharing Network*, J. D. Day, G. R. Grossman, March 1977.
- RFC 728, *Minor Pitfall in the Telnet Protocol*, J. D. Day, April 1977.
- RFC 730, *Extensible Field Addressing*, J. Postel, May 1977.
- RFC 731, *Telnet Data Entry Terminal Option*, J. D. Day, June 1977.
- RFC 741, *Specifications for the Network Voice Protocol (NVP)*, D. Cohen, November 1977.
- RFC 1122, *Requirements for Internet Hosts—Communication Layers*, R. Braden ed., October 1989.
- RFC 1855, *Netiquette Guidelines*, S. Hambridge, October 1995.
- RFC 2026, *The Internet Standards Process—Revision 3*, S. Bradner, October 1996.
- RFC 2028, *The Organizations Involved in the IETF Standards Process*, R. Hovey, S. Bradner, October 1996.
- RFC 3978, *IETF Rights in Contributions*, S. Bradner, ed., March 2005.