

Malcic, Steven

Article

The problem of future users: how constructing the DNS shaped internet governance

Internet Policy Review

Provided in Cooperation with:

Alexander von Humboldt Institute for Internet and Society (HIIG), Berlin

Suggested Citation: Malcic, Steven (2016) : The problem of future users: how constructing the DNS shaped internet governance, Internet Policy Review, ISSN 2197-6775, Alexander von Humboldt Institute for Internet and Society, Berlin, Vol. 5, Iss. 3, pp. 1-13, <https://doi.org/10.14763/2016.3.434>

This Version is available at:

<https://hdl.handle.net/10419/214028>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/3.0/de/legalcode>



The problem of future users: how constructing the DNS shaped internet governance

Steven Malcic

University of California Santa Barbara, United States

Published on 30 Sep 2016 | DOI: 10.14763/2016.3.434

Abstract: Before the emergence of internet governance bodies like the Internet Corporation for Assigned Names and Numbers (ICANN), early network designers learned how to govern the internet in their work building the Domain Name System (DNS). Using original archival research, this article follows conversations among network designers in their daily struggle to keep the Advanced Research Project Agency Network (ARPANET) and early internet in working order. Drawing from social constructivism and path dependence theory, this history helps to conceive “internet governance” beyond its institutional focus, considering how the work of ordering the internet necessarily exceeds the parameters of governance authorities.

Keywords: Internet governance, Internet history, Network design, Domain Name System (DNS), Path dependence

Article information

Received: 27 Apr 2016 **Reviewed:** 17 Jun 2016 **Published:** 30 Sep 2016

Licence: Creative Commons Attribution 3.0 Germany

Competing interests: The author has declared that no competing interests exist that have influenced the text.

URL:

<http://policyreview.info/articles/analysis/problem-future-users-how-constructing-dns-shaped-internet-governance>

Citation: Malcic, S. (2016). The problem of future users: how constructing the DNS shaped internet governance. *Internet Policy Review*, 5(3). DOI: 10.14763/2016.3.434

This paper is part of Doing internet governance, a special issue of Internet Policy Review guest-edited by Dmitry Epstein, Christian Katzenbach, and Francesca Musiani.

INTRODUCTION

Like so many engineers building the Advanced Research Projects Agency Network (ARPANET), Elizabeth "Jake" Feinler struggled each day to get the network into some working order. Feinler was head of the Network Information Center (NIC) at Stanford Research Institute (SRI). Because the NIC functioned as the administrative clearinghouse for ARPANET and the early internet, Feinler had to keep track of everything, but without a standardised addressing system

to help. In the archives at the Computer History Museum (CHM) in Mt. View, California, I found a collection of printer paper that Feinler stapled together in June 1973. This hand-written directory, which Feinler titled “Changes or Reverifications”, lists what sites were online, the point of contact at each site, and even minutia such as the current phone numbers of the liaisons for the contacts (SRI ARC/NIC Records, Lot X3578-2006). One can imagine how unwieldy this task would become, as institutions connected sites to ARPANET at a rapid pace, often installing multiple computers, each of which required a unique identifier. Feinler’s desk reference, a historical precursor of the Domain Name System (DNS), is evidence of a basic quandary that all early network designers faced. As ARPANET’s ill-conceived addressing schema fueled frustration among the networking community, designers started doing the work of internet governance to solve a fundamental problem of design: the need to address future users.

Through a critical reading of documents circulated among ARPANET and early internet engineers, this article shows 1) how “the problem of future users” motivated the social construction of the DNS, and 2) how this historical process itself constitutes the preformation phase of internet governance. To do this, I draw from two theoretical approaches, showing how a social constructivist critique can inform path dependent theories of technological and organisational lock-in. On the one hand, social constructivists “claim that technological artifacts are open to sociological analysis, not just in their usage but especially with respect to their design and technical ‘content’.” (Bijker, Hughes, and Pinch, 1987, p. 4). On the other hand, path dependence theory “stresses the importance of past events for future action or, in a more focused way, of foregoing decisions for current and future decision making” (Sydow, Schreyögg, and Koch, 2009, p. 690). Whereas social constructivists are often concerned with issues of ideology, theorists of path dependence are concerned with self-reinforcing social and economic mechanisms that guide technologies and organisations toward “increasing stability and lock-in” (Dobusch and Schübler, 2012, p. 618). Despite their differences, both approaches regard historical evidence as “process data”, which “consist largely of stories about what happened and who did what when—that is, events, activities, and choices ordered over time” (Langley, 1999, p. 692). This conceptual dovetail opens a window, allowing one to consider how ideology—understood as values supported by material relations—can itself become a self-reinforcing mechanism of path dependence, setting constraints for the DNS, for ICANN, or for any other technological or organisational development.

In considering the social construction of the DNS as the preformation phase of internet governance, I show how Feinler’s mundane task of ordering the network by hand marks a catalyst in the development of design priorities and management functions that ICANN would inherit. Following an “identity crisis” that emerged during the shift from ARPANET protocol to the internet’s TCP/IP suite, designers needed to construct a standardised addressing schema. Initially, they did not solve the problem of future users by calling for the outright establishment of governmental institutions. First, they suppressed the visibility of numerical addresses, thereby hiding the historically contingent development of core infrastructure. Next, they harnessed the power of extensibility, choosing top-level domain names associated with generic social categories. These choices ushered new tasks of ordering the network into a preexisting discourse of social bureaucracy, which structured everyday work relations, emerging institutional affiliations, and the future ideology of internet governance.

ARPANET'S IDENTITY CRISIS: NAMES, NUMBERS, AND INITIAL CONSTRAINTS

The installation of ARPANET marks the triggering event in the development of universal digital addressing as embodied in the DNS, and as such constitutes the preformation phase of governance functions related to ICANN. Jörg Sydow, Georg Schreyögg, and Jochen Koch write that "history matters in the Preformation Phase", because in "organizations initial choices and actions are embedded in routines and practices" and "reflect the heritage [. . .] making up those institutions" (2009, p. 692). ARPANET became operational late in 1969, with sites at the University of California, Los Angeles (UCLA), the Stanford Research Institute (SRI), UC Santa Barbara (UCSB), and the University of Utah (UTAH). ARPANET was designed according to a two-layer architecture, allowing engineers to update, debug, or completely replace entire sections of the network without crashing the system. Even though it afforded designers much needed flexibility, this two-layer architecture established the material conditions for what I think of as "ARPANET's identity crisis", a debate among the engineering community about how best to order numerical identification in relation to site-specific names.

At first, designers assigned network addresses to sites according to the order in which machines were installed, setting a precedent that would become problematic as ARPANET grew. The first site, UCLA, had the address 1, while the fourth site, UTAH, had the address 4. The fact that UCLA was assigned address 1 had no overarching design rationale. Janet Abbate (1999) explains that ARPA chose it as the first site because Leonard Kleinrock and his students at UCLA were experimenting with "a mathematical tool called queuing theory to analyze network systems" (p. 58). This historical accident is exemplary of the fact that, as Sydow et al. write, "Since organizations are social systems and not markets or natural entities, triggering events in organizations are likely to prove to be not so innocent, random, or 'small'" (2009, p. 693). Even though it was random and erased from internet infrastructure, UCLA-1 set a precedent that would soon come to annoy many in the ARPANET community and motivate an ideological reorientation.

Designers did not find the numerical identification of the host layer satisfactory. In 1973, Feinler's colleague at the NIC, Jim White, wrote that "the fact that [. . .] Network software employs numbers to designate hosts, is purely an artifact of the existing implementation of the Network, and is something that the human user should NEVER see or even know about" (Gee Host Names and Numbers are Swell, May 11, 1973, SRI ARC/NIC Records, Lot X3578-2006, CHM). An addressing schema based solely upon the order in which machines were installed could not help but emphasise the historical contingency of ARPANET's initial design philosophy.

During these early years, the somewhat coincidental process by which ARPANET was assembled also drove heated debates involving site-specific naming conventions. Peggy Karp of the MITRE Corporation proposed a set of standardised names in 1971, citing problems related to the fact that each site "employs their own special list" of host mnemonics (Standardization of Host Mneumonics, Request for Comments 226), as evidenced by the hand-written desk reference introducing this article. Karp (1971) proposed a list of standardised host names, suggesting, for example, that host 1 remain "UCLA" and host 2 become "SRIARC" (ARC standing for the NIC's original and at that time still official department name, the Augmentation Research Center). Karp's proposal limited site names to their institutional affiliation, specifying neither the type of computer running at each address, nor each site's often more popular nickname.

Karp's proposed site names generated a flurry of discussions throughout 1971, prompting Robert Braden of UCLA to declare, "Please, let's not perpetrate systems programmers' midnight decisions on all future Network users!" (Host Mnemonics Proposed in RFC 226, Request for Comments 239). He objected to "UCLA" because it does not specify the host computer, suggesting instead "UCLAS7 or UCLANM" because UCLA ran an NMC Sigma 7 computer. Braden also writes that "SRIARC" is "a poor choice[,] because "everybody calls it the NIC," and so suggests the name "SRINIC." Even though Braden's own proposals read as a programmer's midnight decisions, he is right to point out that mnemonics based upon the installation of ARPANET infrastructure were not "fully satisfactory", writing, "It is a set of historical accidents, and shows it." Braden ends up recommending that names be standardised according to codes at the NIC, based on its ability to function as an institutional reference.

Designers reached consensus around this idea, and the NIC accepted the task of standardising host names, which would fortify its institutional function of ordering the network into the foreseeable future. Writing on behalf of the NIC, Richard W. Watson (1971) emphasised the need to "recognize the expanding character of the Network, with the potential eventually of several hundred sites" (More on Standard Host Names, Request for Comments 273). The NIC standardised official site mnemonics based upon the rough structure of "institution name-computer" as initially proposed by Jon Postel (Standard Host Names, Request for Comments 236, 1971), then a graduate student at UCLA.

During the end of 1973 into early 1974, as the NIC secured centralised authority of the official host name list, a new project rumored to be underway stirred anxieties across the network. Up to this point, the work of ordering numerical addresses with site names functioned relative to ARPANET alone. Realising this might have been short-sighted, one concerned designer wrote, "There has been no general discussion of multi-network addressing—although there is apparently an unpublicized Internetworking Protocol experiment in progress—and some other convention may be more desirable" (L.P. Deutsch, Host Names On-Line, Request for Comments 606, 1973). This "unpublicized Internetworking Protocol experiment" brought the entire ARPANET identity schema under question. In dealing with APRANET's identity crisis, network designers realised that core infrastructural development must be suppressed from the user interface, an insight that would direct the early work of doing internet governance through its influence on the design philosophy of the DNS.

CONSTRUCTING THE DNS: THREE MECHANISMS OF POSITIVE GOVERNMENTAL FEEDBACK

The two-layer architecture of ARPANET could not accommodate network growth, offering designers a crash course in how to avoid negative governmental feedback. The idea of feedback or self-reinforcement is central to path dependence theory. Dobusch and Schüßler write, "Specifically, we argue that the mechanisms of positive feedback or self-reinforcement can be specified as a necessary condition for path dependence" (p. 618). In short, lock-in could not occur without such self-reinforcing mechanisms. Moreover, Dobush and Schüßler argue that as a conceptual construct, feedback "can act as an integrating factor—as a conceptual bridge to other theories that explain evolutionary processes characterized by increasing stability and lock-in" (p. 618). One such concept for which feedback can act as a bridge is ideology, itself a way social constructivists describe the self-reinforcing relationship between social values and work relations. Before the DNS could become locked-in as the internet's social interface, designers

had to renegotiate their values in order to foster positive governmental feedback.

Facing a future of exponential growth, designers adopted a specific orientation toward the past. To solve the problem of future users, designers not only needed to organise numbers in relation to names; they had to build a new layer of internet infrastructure, one that would not be deemed historically accidental in an ever-shifting internet network landscape. Working in response to the constraints of ARPANET infrastructure, designers constructed the DNS by negotiating three mechanisms of positive governmental reinforcement, implementing: 1) extensible field addressing; 2) domains of shared cognition; and 3) a hierarchical authority. The social construction of the DNS shows how the initial phase of path dependence is never open, and often restricted by a self-conscious goal to make a technology adoptable, when others had not yet been able to adopt it.

1. EXTENSIBLE FIELD ADDRESSING

In 1977, Jon Postel, who had become a researcher at the University of Southern California (USC), proposed a solution to the addressing problem through finding a way to order numbers and names hierarchically, unlike the schema associated with ARPANET. Postel (1977) wrote, "The addressing scheme should be expandable to increase in scope when interconnections are made between complex systems[.]" and concluded that the best solution to the problem "is to always represent the address by fields" (Extensible Field Addressing, Request for Comments 730). Fields are discrete categories that structure a database. The organisation of fields in a database indicates how categories of data relate to each other. Databases can accommodate a specific instantiation of data by positioning it in its proper field. For example, Postel (1977) proposed this addressing schema: Network / Switching Machine / Host / Message-Identifier. The original address for UCSB, the third node on the ARPANET, would read: ARPANET / 3 / 3 / [message-id]. Postel (1977) considered this hierarchical structure "a natural way" of addressing, because "the most general field should come first with additional fields specifying more and more details" (Extensible Field Addressing, Request for Comments 730), it seeming more natural, I suspect, in comparison to the prior ARPANET addressing, a non-hierarchical schema that nobody liked.

Extensibility refers to the ability to accommodate future infrastructural development seamlessly at the user interface. An extensible field model of address afforded designers the opportunity to layer over the physical history of network design and all the accidents that came with it. Designers could choose what categories of data each field would embody. They could label fields according to named concepts such as "network", "host", or even "domain", and put unique numerical identification into its place. Extensible field addressing facilitated the creation of a database that allowed designers the choice of how network entities could be represented at the user interface.

By embodying the value of extensibility, the DNS was designed to interpolate all future users, ushering new sites into their respective social categories like so many *Matryoshka* nesting dolls. In the year leading up to the installation of the master table, Paul Mockapetris (1983), who outlined the technical specifications of the DNS, wrote that while the DNS "database will initially be proportional to the number of hosts using the system", it "will eventually grow to be proportional to the number of users on those hosts as mailboxes and other information are added to the domain system" (Domain Names—Implementation and Specification, Request for Comments 883). This database would itself become the discursive body of network infrastructure, structured by domains of shared cognition.

2. DOMAINS OF SHARED COGNITION

Designers developed a new addressing schema based upon the extensible field model, reaching a consensus around the concept of "Internet Name Domains". Deciding what a domain actually was, however, required much discussion. D.L. Mills first proposed this system. Mills (1981) writes that since "every internet host is uniquely identified by one or more 32-bit internet addresses and that the entire system is fully connected[,] a "hierarchical name-space partitioning can easily be devised to deal with this problem" (Internet Name Domains, Request for Comments 799). Mills discussed this schema in relation to email. He suggests the structure "<user> . <host> @ <domain>", with specific network mnemonics, such as ARPA or COMSAT, placed in the domain field. Like Postel's proposal, this domain model also positions networks at the top of the address hierarchy. Even though this model suppresses the visibility of core infrastructure, it still maintains a site-specific, historical reference through the "host" field.

Considering the rapid growth of internetworking, David D. Clark of the MIT offered a rather contemplative response. Clark (1982) begins, "It has been said that the principal function of an operating system is to define a number of different names for the same object, so that it can busy itself keeping track of the relationship between all of the different names" (Names, Addresses, Ports, and Routes, Request for Comments 814). He goes on to argue that network protocols such as TCP/IP are no different. He suggests that the "scope of the problem" had not yet been accurately judged, writing,

One of the first questions one can ask about a naming mechanism is how many names one can expect to encounter. In order to answer this, it is necessary to know something about the expected maximum size of the internet. Currently, the internet is fairly small. It contains no more than 25 active networks, and no more than a few hundred hosts. This makes it possible to install tables which exhaustively list all of these elements. However, any implementation undertaken now should be based on an assumption of a much larger internet. The guidelines currently recommended are an upper limit of about 1,000 networks. If we imagine an average number of 25 hosts per net, this would suggest a maximum number of 25,000 hosts.

Even with what we now know to have been low estimates, Clark argues that the potential breadth of the internet requires the complete suppression of infrastructural fields, such as "<host>", at the directory interface in order to implement an acceptable management strategy.

Having come to understand core infrastructure as historically accidental to the user interface, designers recuperated domains by redefining them according to abstract concepts of network governance rather than according to site installation. Postel and Zaw-Sing Su (1982) of SRI defined a domain as "a region of jurisdiction for name assignment and of responsibility for name-to-address translation" (The Domain Naming Convention for Internet User Applications, Request for Comments 819). The intent of a domain-based addressing schema, they wrote, "is that the Internet names be used to form a tree-structured administrative dependent, rather than a strictly topology dependent, hierarchy." In defining domains as spaces of administration and jurisdiction, engineers opened a way to organise the directory interface according to categories of bureaucratic discourse. In other words, defining domains as spaces of network governance allowed a new set of names to restructure existing sites, by occupying the top of the extensible field hierarchy.

While TCP/IP became the universally adopted protocol suite in January 1983, the DNS became

operational on 15 December 1984, when the NIC acquired the master table of top-level domain names and their associated servers (Postel, Domain Name System Implementation Schedule—Revised, Request for Comments 921). Designers reached consensus around five top-level domains: GOV, EDU, COM, MIL, and ORG. (Initially, ARPA itself was a sixth top-level domain, although it was restricted for use of network experimentation). While this decision has no direct technical rationale, Postel and Reynolds wrote that the intention of the system was "to provide an organization name [. . .] free of undesirable semantics" (Domain Requirements, Request for Comments 920, 1984). Names indicating what might become historical accidents of network design were avoided: UCLA, for example, could not reside in a top-level field. A specific institution would reside within its conceptual category, 'Education' or 'Government' or 'Commerce,' and so on. The philosophy of extensible field addressing allowed designers to position institutional modes of social identification at the top of the hierarchy, ushering the DNS into a preexisting discourse of governmental functions that exist independently of the internet itself.

Because the DNS has social concepts at the top of its extensible field hierarchy, the system could both reflect the society into which it was implemented, while simultaneously making room for future users outside the ARPANET community. With networks addressed according to governmental concepts rather than institutions themselves, designers made what they perceived to be a pragmatic decision: build a flexible, layered network organised by a hierarchy of institutional signifiers that already exist in the world. The DNS provided common terms through which the general public could understand how the internet is organised in relation to society, proving to be a solution to ARPANET's identity crisis. With the DNS, designers had a stable addressing system in place. Now all they had to do was to find a way to make it work on an everyday basis, and into the future.

3. HIERARCHICAL AUTHORITY

In order to govern future users, Paul Mockapetris introduced the concept of authority. He writes, "Although we want to have the potential of delegating the privileges of name space management at every node, we don't want such delegation to be required" (Mockapetris, Domain Names—Concepts and Facilities, RFC 882, 1983). If such delegation were required, each network or specific institution would have final authority over its users, leading the system back into the realm of "historical accident" that designers needed to avoid. Instead, Mockapetris recommended investing authority with a name server, which would have "authority over all of its domain until it delegates authority for a subdomain to some other name server". He proposed principles of authority that require a name server administrator to "register with the parent administrator of domains" and also to "identify a responsible person[,] someone "associated with each domain to be a contact point for questions about the domain, to verify and update the domain related information, and to resolve any problems (e.g. protocol violations) with hosts in the domain" (Mockapetris, Domain Names—Concepts and Facilities, RFC 882, 1983).

Mockapetris borrowed the term "responsible person" from Jon Postel. In order to establish a domain, Postel (1981) wrote, "There must be a responsible person to serve as a coordinator for domain related questions" (The Domain Names Plan and Schedule, Request for Comments 881). He goes so far as to cordon off a special section in order to define "responsible person" precisely:

An individual must be identified who has authority for the administration of the

names within the domain, and who takes responsibility for the behaviour of the hosts in the domain in their interactions with hosts outside the domain.

[. . .]

If some host in a domain somehow misbehaves in interactions with hosts outside the domain (e.g. consistently violates protocols), the responsible person for the domain must be able to take action to eliminate the problem.

Postel conceives the "responsible person" not simply as a steward, but as a potential disciplinary authority, someone who has the power to decide what constitutes "misbehavior" and then to "eliminate the problem" accordingly. That Mockapetris adopted this term, too, suggests that in terms of internetwork administration, responsibility meant something very specific. A responsible person seems to be someone who shares values akin to those of internet designers. An irresponsible person, someone who "consistently violates protocols", for instance, is someone who does not share values akin to designers, someone who might very well warrant an administrative elimination.

Organisational representatives that used the internet for specific projects or business activities initially filled the role of "Responsible Persons", although this ended up contributing to administrative difficulties it was intended to resolve. Archival documents from the NIC show that the "Responsible Person" (RP) model was not effective in managing network access. This was largely due to the definition of RPs as organisational figureheads. Even though she no longer had to order the network by hand, Feinler still had to order RPs, which proved just as difficult. In a hand written memo, Feinler (1985) wrote,

The Responsible Persons are the wrong people to track who has permission to use the network. They are people such as very important PIs or Vice Presidents of companies and the like—people who deal in concepts and macro mgt; not administrative minutia. They either forget or outright refuse to do the job and yet they are listed as contacts. (Memo on Responsible Persons, SRI ARC/NIC Records, Lot X3578-2006, Computer History Museum)

When someone sought a password to access network services, "Responsible Persons" were the ones charged with managing this. However, due to the fact that many RPs were "very important PIs or Vice Presidents", this sort of "administrative minutia" fell through the cracks. This problem was compounded because, in this early system, "passwords [were] invalidated after 6 months[.]" at which point users had "to get permission again from the RP". Feinler (1985) wrote, "Unfortunately the RPs usually let their own passwords expire and can't reactivate their users" (Memo on Responsible Persons, SRI ARC/NIC Records, Lot X3578-2006, CHM). In defining "Responsible Persons" according to their institutional status in disparate organisations, the NIC was often unable to help users effectively get network access.

Another administrative problem emerged regarding the fact that multiple "Responsible Persons" were often affiliated with a single host computer. Users tended to think that host sites facilitated network access, not realising that "Responsible Persons" of specific organisations or projects served that function. In an email to Feinler, Bob Baker (1985) wrote,

There has been a lot of confusion caused by people failing to understand that the registration [. . .] is organization oriented and not host or site oriented. Thus some people have the mistaken idea that to get registered they should contact someone connected with the host they use, instead of the "responsible person" for the organization they belong to. (How to Announce TACACS, January 2, 1985, SRI ARC/NIC Records, Lot X3578-2006, CHM.)

In another memo to Feinler pointing out the main problems of the "Responsible Person" model, Johanna Landsbergen (1986) wrote, "When the password for the Responsible Person of an Org expires and he/she does not remember their old password", no one at the organisation has the ability to "get a new password" (ARPANET TACACS TOOL PROBS, January 15, 1986, SRI ARC/NIC Records, Lot X3578-2006, CHM). Feinler (1986) raised these issues at multiple NIC meetings, her notes indicating that "Responsible Persons" are "administratively confusing", for it was difficult to find the correct "RP if there are 4 at one org" (Notes, SRI ARC/NIC Records, Lot X3578-2006, CHM). With the DNS having signified network topology according to institutional concepts, it was difficult to clearly identify "Responsible Persons" as stable points of contact.

The association of "Responsible Persons" with the organisations they represent also led to ambiguities in the construction of network domain databases. In the draft of a proposed "User Database Host Schema" from Bolt Berenek and Newman (BBN), John V. DelSignore includes a glossary that attempts to distinguish the terms "user", "person", and "organization." He writes,

The word "user" is generally used to indicated [sic] a person that is or has logged into the database tool and is performing or performed a certain act or command. Also 'user' indicates the real-life person associated with a person record.

[. . .]

occasionally the words "person" or "organization" appear in sentences such as "The user created the person". We realize the users create "person records" and not "persons" per se. The terms "person" and "organization" are often used interchangeably with "person record" and "organization record" for the sake of brevity. (John V. DelSignore, Jr., ARPANET TAC Access Control User Database Host Schema, September 16, 1986, SRI ARC/NIC Records, Lot X3578-2006, CHM)

Feinler (1986) was right to conclude "that as a registration scheme it is an administrative nightmare" (Notes, SRI ARC/NIC Records, Lot X3578-2006, CHM), and that the RP model of internet governance could not hold.

Much like numbered host identification that developed in an historically contingent manner, the "Responsible Persons" model of network administration could not efficiently accommodate future users by virtue of the fact that RPs were associated with specific projects at specific organisations at specific moments in time. By the end of 1986, designers abandoned the RP model, instead situating host administrators as points of contact for network users to register in the DNS and acquire network access. The RPs still functioned as gatekeepers; however, they no longer had to manage passwords, directly correspond with users seeking access, or maintain records of host activity. The host administrator took on this task, working as a liaison between users, organisational representatives, and the NIC. The introduction of the host administrator

role brought the division of labour in line with the DNS addressing schema. To replace the historical accident of RPs, as well as to manage existing and future sites as if they had always been expected, designers created a new responsibility—one of maintaining the internet's order—abstracted from specific projects.

As early network designers introduced the concept of authority through the social construction of the DNS, they catalysed the development of bureaucratically independent internet governance functions like IANA, which paved the way for later institutions like ICANN. More than a material base, the DNS provided the conceptual structure for a hierarchical regime of internet governance that centralised administrative power within discursive categories inherited from historically naturalised social categories. Through the social construction of the DNS, early network designers better ensured that future users would themselves maintain the DNS as a foundation for governing the global internet.

CONCLUSION

By using social constructivist historical analysis in tandem with path dependence theory, this article shows how early network designers built the DNS through harnessing three modes of positive governmental feedback: 1) extensibility, which afforded ways to hide the contingent development of network infrastructure; 2) domains of shared cognition, which allowed non-expert users to navigate the internet in a socially legible manner; and 3) hierarchical authority, which established the initial structure for "internet governance" as an institutionalised function, and ensured that future users could themselves maintain the system and extend it further. After the installation of the DNS, Feinler continued as head of the NIC until 1989, the NIC transformed into InterNIC in 1993, and ICANN finally assumed all responsibilities of InterNIC with its foundation in 1998.

The lock-in of ICANN as the internet's primary governing body is indeed related to macro-level forces in the global political economy of the 1990s. In his article "ICANN between technical mandate and political challenges" (2000), Wolfgang Kleinwächter argues that ICANN became locked-in with its incorporation due to four macro-level problems: 1) "the need to demonopolise" (p. 556) registrars during the dot-com boom; 2) the need to settle "disputes between trademark holders and domain name holders" (p. 558), which led to ICANN's Uniform Dispute Resolution Policy (UDRP); 3) the need to recognise country-code top-level domains (cc-TLDs), codified in ICANN's Governmental Advisory Committee (GAC); and 4) the need to create new generic top-level domains (gTLDs), which motivated ICANN to work in tandem with the World Intellectual Property Organization (WIPO). Kleinwächter persuasively articulates the immediate historical context of ICANN's global lock-in; but such macro-level forces are themselves historically related to the micro-level decisions of people who built the DNS toward governmental ends. In a way, Kleinwächter himself reveals how ICANN's incorporation solved the problem of future users as it emerged again in the 1990s, but at a macro-level policy analysis.

As more people consider "internet governance" beyond its institutional focus, an idea promoted by scholars including Laura DeNardis (2012) and Francesca Musiani (2014), finding new ways to conceptualise the term will become more important. Understanding the prehistory of institutional bodies is one way to consider how people "do" governance in response to everyday working pressures. This article shows how internet governance has never been given, and has always been done. Internet governance is not a product or result of organisations like ICANN.

Even though it serves an important governance function, ICANN is itself based on the historical contingencies of ordering the early internet. ICANN must also constantly respond to how people use the internet in increasingly varied ways, and as more users understand the political significance of “doing” internet governance, complexities related to the problem of future users will only compound, as users themselves attempt to create—or alter—control structures of the internet.

In designing and implementing the DNS, early designers paved the way for new technocratic functions that ICANN would inherit, although they did not initially call for the institution of governance bodies *per se*. Rather, they worked through technical issues of the finest complexity, and in so doing developed perspectives on issues such as the nature of historical contingency, the jurisdiction of virtual space, and the concept of authority itself. In learning how to recognise the historical contingency of network design, embracing extensibility, and reifying a new division of labour, early network designers ensured that future users could not only navigate the internet, but could also keep the system in working order on an everyday basis. The technocratic relations that rigidified around the DNS fostered values related to a particular brand of universality, one supported through the potentially infinite extension of social genera, as evidenced today in ICANN’s slogan: "One world. One Internet."

REFERENCES

- Abbate, J. (1999). *Inventing the Internet*. Cambridge: MIT Press.
- Baker, B. (1985). How to Announce TACACS (SRI ARC/NIC Records, Lot X3578-2006, Computer History Museum).
- Bijker, Wiebe E., Thomas P. Hughes, & Trevor Pinch. (1987) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. Cambridge: MIT Press.
- Braden, R. (1971). Host Mnemonics Proposed in RFC 226 (Request for Comments 239).
- Clark, D.D. (1982). Names, Addresses, Ports, and Routes (Request for Comments 814).
- DelSignore, J.V. Jr. (1986). ARPANET TAC Access Control User Database Host Schema (SRI ARC/NIC Records, Lot X3578-2006, Computer History Museum).
- DeNardis, L. (2012). Hidden Levers of Internet Control: An Infrastructure-based Theory of Internet Governance. *Information, Communication and Society* 15(5): 720-38.
- Deutsch, L.P. (1973). Host Names On-Line (Request for Comments 606).
- Dobusch, L. & E. Schüßler. (2012). Theorizing path dependence: a review of positive feedback mechanisms in technology markets, regional clusters, and organizations. *Industrial and Corporate Change* 22(2): 617-647.
- Feinler, E. (1973). Changes and Reverifications (SRI ARC/NIC Records, Lot X3578-2006, Computer History Museum).
- Feinler E. (1985). Memo on Responsible Persons (SRI ARC/NIC Records, Lot X3578-2006, Computer History Museum).
- Feinler, E. (1986). Notes (SRI ARC/NIC Records, Lot X3578-2006, Computer History Museum).
- Karp, P. (1971). Standardization of Host Mneumonics (Request for Comments 226).
- Kleinwächter, W. (2000). ICANN between technical mandate and political challenges. *Telecommunications Policy* 24: 553-563.
- Landsbergen, J. (1986). ARPANET TACACS TOOL PROBS (SRI ARC/NIC Records, Lot X3578-2006, Computer History Museum).
- Langley, A. (1999). Strategies for Theorizing from Process Data. *The Academy of Management Review* 24(4): 691-710.
- Mills, D.L. (1981). Internet Name Domains (Request for Comments 799).
- Mockapetris, P. 1983). Domain Names—Implementation and Specification (Request for Comments 883).
- Musiani, F. (2014). Practice, Plurality, Performativity, and Plumbing: Internet Governance Research Meets Science and Technology Studies. *Science, Technology & Human Values* 40(2):

272-286.

Postel, J. (1984). Domain Name System Implementation Schedule—Revised (Request for Comments 921).

Postel, J. (1981). The Domain Names Plan and Schedule (Request for Comments 881).

Postel, J. (1977). Extensible Field Addressing (Request for Comments 730).

Postel, J. (1971). Standard Host Names (Request for Comments 236).

Postel, J. & J. Reynolds. (1984). Domain Requirements (Request for Comments 920).

Postel, J. & Z.S. Su. (1982). The Domain Naming Convention for Internet User Applications (Request for Comments 819).

Sydow, J., G. Schreyögg & J. Koch. (2009). Organizational Path Dependence: Opening the Black Box. *The Academy of Management Review* 34(4): 689-709.

Watson, R.W. (1971). More on Standard Host Names (Request for Comments 273).

Wentheimer, E. (1971). Network Host Status. *Request for Comments* 252.

White, J. (1973). Gee Host Names and Numbers are Swell (SRI ARC/NIC Records, Lot X3578-2006, Computer History Museum).

ACKNOWLEDGMENTS

The author wishes to thank Andrea Hackl, Leonard Dobusch, and Francesca Musiani for their generous guidance in the revision process. The author also wishes to thank Sara Lott, Senior Archives Manager at the Computer History Museum, who helped in the early selection of materials.