

ICANN and Internet Governance: Leveraging Technical Coordination to Realize Global Public Policy

Hans Klein

School of Public Policy, Georgia Institute of Technology, Atlanta, Georgia, USA

The Internet Corporation for Assigned Names and Numbers (ICANN) was created in 1998 to perform technical coordination of the Internet. ICANN also lays the foundations for governance, creating capabilities for promulgating and enforcing global regulations on Internet use. ICANN leverages the capabilities in the Internet domain name system (DNS) to implement four mechanisms of governance: authority, law, sanctions, and jurisdictions. These governance-related features are embodied in seemingly technical features of ICANN's institutional design. Recognition of ICANN's governance mechanisms allows us to better understand the Internet's emerging regulatory regime.

Keywords domain name system (DNS), global public policy, governance, ICANN, Internet

INTERNET GOVERNANCE

The Internet has often been hailed as a domain of benevolent anarchy, a place where free communication is securely in place. It is a “modern Hydra” capable of circumventing regulation (Froomkin, 1999, p.129) and a “space of no control” (Lessig, 1999, p. 24). As Internet bard John Perry Barlow says, “Governments of the Industrial World, . . . You have no sovereignty where we gather” (cited in Lessig, 1999, p. 218).

Stated less colorfully, the Internet presents challenges to *governance*. By governance I mean the existence of some authority able to make globally applicable rules for In-

ternet usage backed up by sanctions. Internet governance exists in various partial forms (e.g., AOL chat rooms or government regulation of computers within national territory), but overall the Internet does not have a coherent and effective system of authoritative rule making and enforcement. The reasons for this lie both in characteristics of the technology, which make control difficult, and in the global reach of Internet communications, which creates jurisdictional conflict among government regulators.

This “ungovernability” of the Internet, however, is changing. In his book *Code*, Lawrence Lessig (1999) documents various strategies to eliminate user anonymity and thereby facilitate law enforcement. Likewise, in a recent court case over Internet auctions of Nazi memorabilia that was illegal in France, the Yahoo! Corporation was ordered to detect viewer's location and apply local content regulations to them (AFP, 1999). As with other technologies preceding it, as the Internet becomes important to the society around it, attempts are made to integrate it within existing regulatory structures (Hughes, 1983).

The most significant development in the trend to render the Internet governable is the creation of the Internet Corporation for Assigned Names and Numbers (ICANN). Created in 1998, ICANN is a private, nonprofit entity whose official mandate is to perform technical coordination of core Internet resources, most notably domain names (e.g., mycomputer.org). Its site of incorporation lies in California but its authority extends, directly or indirectly, over all users of the Internet.

ICANN has the potential to radically change the nature of the Internet. By putting in place all the mechanisms needed for the creation, promulgation, and enforcement of regulations, ICANN makes effective Internet governance possible for the first time. No longer is the Internet a Hydra that is impossible to control. ICANN's mechanisms for governance can serve any number of possible regulations. The frequency of past attempts to regulate the Internet suggests that the realization of these mechanisms could attract

Received 2 May 2000; accepted 13 August 2001.

Research for this study was performed while the author was a resident at the Center for the Sociology of Innovation at the Ecole des Mines, Paris, with support from a Chateaubriand Fellowship from the French Embassy in the United States.

Address correspondence to Hans Klein, Assistant Professor, School of Public Policy, MC: 0345, Georgia Institute of Technology, Atlanta, GA 30332-0345, USA. E-mail: hans.klein@pubpolicy.gatech.edu

widespread interest in their utilization (Froomkin, 1997). Depending on one's viewpoint, creation of a capability for governance holds the promise or the threat of taming the electronic frontier.

In what follows I offer a detailed analysis of ICANN. My purpose is to render comprehensible the interrelationship between technology, administration, and governance, explaining how a computer network addressing system makes possible a system of governance. To do this I explain what governance is, how the Internet's domain name system (DNS) works, and how the former is realized through the latter.

For policymakers and Internet users, this account can help inform policy. That the Internet has a point of central control and that ICANN makes global public policy possible are not widely admitted. Recognition of these governance capabilities justifies the application of normative criteria of legitimacy, accountability, and equity to the institution and its processes. This study contributes to a growing body of policy literature that analyzes ICANN from the perspective of its historical origins (Mueller, 1999; Klein, 2001a), its legal status (Froomkin, 2000; Klein, 2001c), and its institutional design (Post, 1998).

This account is also relevant to theoretical debates over the relationship between technological systems and their social context. Recent scholarship in the social construction of technology has emphasized the influence of social factors in shaping technological change (Bijker, 1995; Bijker et al., 1987; Klein & Kleinman, 2002). Writers on technology policy have likewise emphasized how designing systems is comparable to writing law, insofar as both activities create social structures that constrain human behavior (Lessig, 1999; Kapor, 1990; Klein, 2000). ICANN offers stark evidence of such social structuring, with the domain name system defining important parameters of governance.

In what follows, I first consider the technological and institutional factors that have inhibited Internet regulation and then continue with a more general discussion of what governance is. Following that I examine the domain name system, both as a technological and administrative hierarchy, and I identify characteristics that allow for governance. There then follows an analysis of the mechanisms of Internet governance as realized in ICANN. Then, as an illustration of how the governance mechanisms work in practice, I examine ICANN's global public policy defining property rights in domain names. I consider the problem of legitimacy in ICANN and speculate about possible future areas of regulation.

The Problem of Internet Governance

Regardless of whether one supports or opposes specific regulations, it is generally recognized that regulation of

the Internet has proven difficult. Copying music, software, and other forms of intellectual property has become simple, and the growth in legal actions against property violators likely represents only a small fraction of incidents of unauthorized copying (Fryer, 1995). Attempts by national governments to control content have foundered on the Internet's global nature and the interjurisdictional conflicts in regulation (Andrews, 1999).

Barriers to regulation arise in part from characteristics of the technology. Internet communications do not pass through a central channel but are instead passed between many independent networks, and even the messages themselves are broken into packets that may follow different itineraries from source to destination (Cerf & Kahn, 1974). With multiple independent parties sending multiple independent packets through multiple independent channels, there is no central communication channel that could serve as a control point for promulgating and enforcing regulations.

Regulation also founders on institutional factors. The Internet challenges established jurisdictions (Johnson & Post, 1997; Perrit, 1997). Public authority resides in the state, whose foundational characteristic is the exercise control over a geographically defined domain (Schroeder, 1998). Yet the "spaceless" nature of the Internet violates the geographical underpinning of public authority (Holitscher, 1999). The mismatch between a network that is global and regulations that are national undermines many attempts at regulation (Froomkin, 1997).

To make conceptual sense of this situation, it is useful to step back from the details and consider governance from a theoretical level. What is governance? What is needed in order to govern? What is needed for governance of the Internet?

In *Democracy and Its Critics* (1989), Robert Dahl defines what governance is and what is needed to achieve it. He identifies a set of "assumptions of a political order" (pp. 106–107) that specify the minimal conditions for a system of governance. I call these *mechanisms of governance*. Paraphrasing Dahl's definition, we can identify four such mechanisms. The first is an *authority*. Governance requires a governor or a sovereign. An entity, be it an individual or a group, must make policy decisions that apply to the members of the polity. A second governance mechanism is *law*. Laws implement policy decisions. They might take the form of a tax, a license, or simply a binding rule. Third, there must be some mechanism for imposing *sanctions*. This allows for punishment of those who violate laws. Finally, governance requires the definition of *jurisdiction*. Jurisdiction defines the space over which the authority makes decisions and within which the laws apply and are enforced by the threat of sanctions. These four mechanisms make governance possible: the governing *authority* can make a policy decision that applies within its

jurisdiction, embodying that decision in *law* and imposing *sanctions* on whomever disobeys. [A similar discussion can be found in the appendix of *Code* (Lessig, 1999).]

The Internet's vaunted ungovernability results from the absence of these four mechanisms. Regulation is difficult because authority, law, sanctions, and jurisdictions are not in place.

ICANN realizes these four mechanisms through its control of the Internet's domain name system (DNS). Although Internet communication has no central control point, Internet addressing, as realized in the DNS, is centralized. DNS provides the control point from which to regulate users. Moreover, the DNS is also an essential resource, so it provides a means of sanctioning users: denial of access to domain names is the equivalent to banishment from the Internet. The DNS also defines jurisdictions on the Internet. The logical organization of the DNS allows authority to be mapped onto distinct zones. Finally, the contractual foundations of the DNS provide opportunities to promulgate regulations. Taken together, these features render ICANN capable of governance.

DNS AND GOVERNANCE—A SIMPLE ACCOUNT

In order to understand ICANN, one must first understand the domain name system. Here I analyze DNS in two passes. I first present the DNS in a simplified form, treating it as a single, nondistributed system. Seen this way, the governance features of the DNS are most easily recognized. In a later section I examine the distributed inner structure of the DNS and present the various mechanisms used to realize coherent administration and policymaking.

DNS: The Control Point of the Internet

I begin with a little-recognized fact: The Internet really consists of two "systems," one for communications (the "TCP/IP" protocols) and one for addressing (the DNS). The communication system is the Internet as we commonly know it. It is extremely decentralized—so much so that it is really not a "system" at all but rather just a set of protocols by which independent computer networks can send data packets to each other. It is this decentralized system that informs most public understanding of the Internet and that underlies claims about ungovernability.

In marked contrast to this, the addressing system—the domain name system (DNS)—is centralized (Albitz and Liu, 1998). Nearly all Internet communications rely on this single system. The DNS can be thought of as the Internet's telephone book and directory assistance service. Before one computer can communicate with another, it must do the equivalent of contacting directory assistance with a name of the party it wishes to call and receiving back the number to dial. This is a necessary prelude to communication.

Internally, the DNS consists of a database and a dynamic lookup service. The database includes pairs of domain names and IP numbers. Domain names are alphanumeric (and hence human-friendly) identifiers of computers on the Internet. IP (Internet Protocol) numbers (or addresses) are machine-friendly numeric identifiers. For example, a given computer's domain name might be *mycomputer.org*, and its corresponding IP number might be *12.34.56.78*. The DNS resolves domain names into IP numbers. In name resolution, the DNS accepts a domain name from a user and returns the corresponding number. The DNS computers performing name resolution are called *name servers*. Only after resolution has occurred can the user-to-user e-mail or web communication begin.¹

This two-step procedure is immediately visible on most web browsers (e.g., Netscape Navigator). Once a user enters a domain name, the browser will indicate that it is interacting with the DNS by posting a message like "Looking up host . . ." As long as a few seconds may pass before resolution is performed, and IP number is returned, and actual communication may begin. Sometimes name resolution fails, as when a misspelled name generates an error message like, "Unable to locate host . . ." and no number is returned. By watching the status messages on a browser's screen, a user can observe the name resolution process.

At the heart of the DNS is the Internet's *name space*. The name space lists (nearly) all computers on the Internet.² At the time of this writing the name space contains tens of millions of entries. When one reads statistics about the growth of the Internet, the numbers usually refer to the size of the name space. It provides a rough approximation of the number of individual users: Since most computers listed in the DNS are gateways into private networks with many users, the number of users is much greater than the number of entries in the name space.

In a manner of speaking, the name space *is* the Internet. In order to exist on the Internet, a computer must be listed in the name space. Without a listing (without a domain name and an IP number) a computer cannot be found by others. Removal of a computer's listing from the name space constitutes a kind of banishment, for a computer disappears from the list of addressable computers. Whatever entity controls the name space database effectively controls the Internet. These points are discussed in detail next.

As currently designed, the name space must obey certain design principles (IAB, 2000; ICANN, 2001). The system's designers claim that the name space must be *unique* and it must be managed by a single entity. There can only be one database that constitutes the definitive listing of computers on the Internet. Copies may exist, but independent name spaces cannot, because they could evolve to have different contents. Were multiple, independent name spaces to exist, a given domain name could resolve to different IP addresses depending on which name space

was used, which would render communication unreliable. This technological imperative of uniqueness underlies the centrality of the DNS, for all communications must use a single, authoritative name space. The Internet's use of a unique name space (with a single administrator) "is a technical necessity, not a policy choice" (IAB, 2000). (Were this design feature not necessary, then numerous possibilities for policy choices would be opened).

Administration

DNS is more than a technical system; it is also an administrative and policy system. Continuing with our simplified view of DNS as a single, nondistributed database, we can examine the DNS in terms of a single administrator and a single policy authority. The *policy authority* entity makes general rules for changes to the name space, such as allowable domain names, cost of registration in the name space, and restrictions on the addition or deletion of names. The *administrator* implements these decisions, adding, deleting, and modifying the database entries to reflect the entry, exit, and changed status of computers. The administrator also ensures the reliable operation of the name server.

The DNS's uniqueness requirement means that the policy authority and the administrator exercise monopoly power. There must be a unique name space, and it must be managed by a unique administrator, who, in turn, must be subject to a unique policy authority. "Both the design and the implementation of the DNS protocol are heavily based on the assumption that there is a single owner or maintainer" (IAB, 2000). Only in this way can the name space be guaranteed to function reliably. Directly or indirectly, this one DNS administrator contracts with every network connected to the Internet. Thus, paralleling the DNS's technical centralization is administrative and policy centralization.

The DNS administrator is also called a *registry*. For a computer to be available on the Internet, the user must approach the administrator and request to be *registered*. The registry registers the computer by adding the user's name-number pair to the name space.

The legal mechanism used to connect the central policy authority with users is the *contract*. The Internet is a network of networks; most computers registered in the name space are gateways to private networks managed by network administrators. Each listing in the DNS is accompanied by a contract between the central DNS administrator and a *network administrator*. The contract specifies rules and conditions for inclusion in the name space, such as the provision of contact information, the payment of an annual fee, acknowledgment of the role of the DNS administrator, and so on. Thus, every network in the Internet has a contract with the single entity overseeing the DNS. These contracts implement policy centralization.

DNS and Internet Governance

In this simplified version of DNS it is easy to recognize the feasibility of implementing governance. Only relatively minor modifications would be needed to realize the four mechanisms of authority, law, sanctions, and jurisdictions.

The DNS defines a central authority for the Internet. The uniqueness requirement of the name space requires a single central authority, and its decisions apply to all servers in the name space. To render the DNS policy authority a true regulatory entity, its domain of decision making would simply have to expand to public issues, such as intellectual property regulation or content control. Since there are few technical barriers to such an expansion, it would be a policy choice. Thus, realizing governance on the Internet would require simply broadening the range of topics regulated by the DNS policy authority.

The DNS also defines the second governance mechanism: law. The law of the Internet is contained in the domain name registration contracts. The provisions of the contracts with network administrators specify the detailed regulations for their actions. To regulate on broader topics, the language in the contracts would simply have to expand.

Third, DNS provided a powerful mechanism for sanctions: domain name denial (i.e., the deletion of a user's name-number listing from the name space). This is the power of banishment: Network administrators who refused to obey the regulations in their contracts could be delisted from the name space and made to disappear. Name registration could be treated as a privilege, revocable if a user violated the rules.

The DNS neatly solves the problem of jurisdiction as well. The jurisdiction of the DNS policy authority extends to every computer on the Internet but no farther. The registration contract is the manifestation of jurisdiction. Every network administrator is contractually bound to the DNS's policy authority.

Thus, the domain name system provided the means to realize mechanisms of governance. Relatively minor changes to DNS could put each mechanism in place. The DNS policy authority would need to merely broaden its regulatory scope and include those broader regulations in contracts with network administrators. Domain name denial provides an adequate mechanism to sanction rule breakers. The jurisdiction of the policy authority would exactly cover the Internet, no less and no more.

In order to fully realize Internet governance, two additional considerations would have to be addressed. The first is practical: Some means would be needed to extend policy authority to the individual user. Since the domain name registration contract is between the central authority and a network administrator, individual users are not immediately subject to regulations. Regulation of individual users could be achieved using a *flow-down contract*. Network

administrators generally require users to sign an agreement when they obtain an account, and this user contract could repeat the provisions from the administrators' contract. In this way a single set of regulations could "flow down" from the central DNS administrator to network administrators and from there to all users. Indirectly, all Internet users could be regulated by the central DNS administrator. Violation of the user contract could lead to loss of Internet access for the user account. Network administrators who failed to enforce flow-down contracts on their users would find themselves subject to domain name denial, that is, banishment from the Internet. While such a user contract remains hypothetical, its feasibility is not. In a later section I survey the types of regulations that have been or could be implemented through such a system.

A second consideration about governance is more normative in nature. Were the DNS policy authority to become a general-purpose regulator, then careful thought would have to be given to its *legitimacy*. As the scope of its policymaking expanded, its authority would have to be grounded on some appropriate principle. This could be realized by placing ultimate policy authority in the hands of governments or in the hands of a newly constituted representational institution. As discussed later, when policy authority was located in ICANN, it adopted a representational mechanisms to ensure legitimacy.

The discussion so far has built on one simplifying assumption: that the Internet name space is a single, centralized database. In the early phases of the Internet's development this was true. In the 1970s the entire namespace was contained in one file called "hosts.txt" (Mockapetris, 1983). By 1983, however, the continued growth of the network had led researchers to redesign the name space, and break up the name space into multiple, interconnected pieces. The name space is less centralized than has been presented here. Decentralization renders governance of the Internet much more complex. I turn now to an analysis of that more complex architecture.

DNS AND GOVERNANCE—THE DISTRIBUTED SYSTEM

In fact, the name space is a *distributed* database. In theory, all name-number pairs could be held in one central database as described in the simplified account just given. However, since thousands of name resolution queries occur each second, a centralized DNS computer might be overwhelmed. Instead, the name space is distributed among multiple computers to share the workload.

The name space exists as a collection of partial, separate databases running on separate computers. Each partial database is called a *zone file* (or *zone*). A zone contains a subset of the total list of name-number pairs. To each zone is associated a *name server* (or *server*—a software pro-

gram for name resolution) and a *host computer* (or *host*—the hardware that hosts the zone file and name server). Thus, the entire name space is a distributed database-and-name-resolution system whose building block is the triad of a zone file, name server program, and host computer.

As in any distributed database, the relationship between the parts is carefully structured. The different zone are linked to each other to form a top-down pyramidal hierarchy or an inverted tree (with its root at the top). At the apex of the hierarchy is a single zone, the *root*. The root zone links to multiple zones just beneath it, and those zones in turn link to multiple zones beneath them, and so on. (This is the same structure as the files on a personal computer). The levels in the hierarchy are clearly identified: The root zone links to "top-level" zones, these link to "second-level" zones, then the "third-level" zones, and so on.

While a given zone may link downward to multiple zones, it can link upward to just one zone. Directly or indirectly, all zones link upwards to the single root zone. The existence of one root in the name space fulfills the uniqueness condition.

Subtrees in this distributed database are called *domains*. A domain consists of a zone and all zones beneath it in the hierarchy. Domain names are often referred to by their level in the tree. Domains beginning at top-level zones are *top-level domains* or TLDs; domains beginning at the next level are *second-level domains*, and so on. The domain of the root is the complete name space. The entire system constitutes the domain name system or DNS. The terms *zone* and *domain* are often used interchangeably, but the former refers to one single file and the latter refers to that single file and all lower files in its subtree.

A domain has a name—which, not surprisingly, is called a *domain name*. Well-known top-level domain names are .com, .org, and .net. The largest domain in the name space, .com links to millions of lower-level domains. An Internet address like mycomputer.com consists of a second-level domain (mycomputer) and a top-level domain (.com). A string of domain names, with the different levels separated by dots, uniquely identifies any computer in the name space.

This distributed hierarchy defines relationships of top-down control. Any zone file can be modified to link (include) or delink (exclude) the zones below it in the name space. This is the power of virtual life and death. When a name server is connected to the root via some series of links, then it exists in the name space. Should a zone file be modified to eliminate a link, the computer or computers below it in the hierarchy will be cut off from the name space. Each server in the hierarchy controls the path to the root for the servers below it.

An example may illuminate this. Suppose I want Internet e-mail services for my company. I already have an

internal company network, and now I want to connect my network to the Internet. To do that, I must link a host computer in my network into the name space, that is, register the host's domain name and IP address in a DNS zone file. Since the name space is a distributed database, there are many zone files to which I could link: A registry in Virginia maintains a zone file called .com, a registry in England maintains a zone file called .uk, my parent company hosts a zone file called .holdingcompany (which is itself linked to the .com zone file). By my registering my host into an available zone file, it becomes part of the name space and exists in the Internet. Likewise, should my entry in the zone file ever be delinked (the domain name canceled), my host would disappear from the Internet. The modification to the zone file that allows me to enter or exit from the Internet is not made by me but by the administrator of the next higher zone file, upon whom I continue to depend for my presence on the Internet.

Administration

In the simplified explanation earlier, DNS administration and policy authority were held by a single pair of organizations. In the distributed DNS, every domain has such an administration-policy pair (which I refer to simply as an "administrator"; in some cases it may be one and the same organization, anyway). These organizations are organized according to the distributed structure of the name space: Total DNS administration is a multiorganization hierarchy, with each administrator exercising control over lower level administrators. At the apex of the hierarchy is the root administrator.

Each administrator exercises monopoly control over its immediate zone file (to ensure uniqueness of its portion of the name space). Moreover, each administrator has authority over the entire domain beneath it. When it registers a lower level host, it delegates some authority to the lower level administrator, who exercises monopoly control over that lower zone file. Authority flows down the hierarchy from the root zone administrator, with responsibility for the entire name space, to the individual host computers at the lowest level zones. Each administrator is subject to the policies of higher level entities. In this way, the policies made at the root can be transmitted down through the levels of the hierarchy to apply, directly or indirectly, to all administrators in the DNS. As a group, the administrators serve as the gatekeepers to the name space and hence to the Internet.

Just as zone files are joined by links, administrators are joined by contracts. The root administrator formalizes its delegation of authority to top-level administrators in a contract. Some provisions of that contract may be required to be included in subsequent contracts. In this way, regulations may flow down the entire hierarchy to the individual

network administrator or possibly even the individual user, as discussed earlier.

As one would expect, the administration of the root zone is particularly important. All other hosts on the Internet access the name space only with a delegation of authority directly or indirectly from the root. Policy authority over the root—the power to add or delete top-level domains—confers direct control over all top-level domains and indirect power over all lower level domains. Authority over the root zone extends to the entire Internet.

In summary, this more accurate account of the DNS reveals another order of magnitude of technical and administrative complexity. Seen as a whole, the DNS is a centralized control point for the Internet. However, because it is a decentralized system, the DNS has an internal structure that relies on hierarchical control and contracts to achieve unified policy capabilities.

Historical Factors in DNS

At this point we can switch from technical analysis to historical analysis. The discussion so far outlined the functioning of the DNS and its structure as a distributed database. This historical analysis outlines evolution of the DNS namespace and its administrative hierarchy.

The Internet began as a research project in the 1970s, and the computer scientists developing it shaped the evolution of its administrative and policy institutions (Hafner & Lyon, 1998). This research community was centered in institutions like the Internet Engineering Task Force (IETF), the Internet Architecture Board (IAB), the University of Southern California's Information Sciences Institute (ISI), and the Internet Society (ISOC) (Leiner et al., 2000).

One person, in particular, played a key role in the DNS development: Dr. Jon Postel, a computer scientist at the University of Southern California. Working under a research contract from the US government, Postel exercised policy authority over the root, a function eventually called the Internet Assigned Number Authority (IANA). In his IANA role, Postel maintained the root zone file, authorized the addition of new top-level domains (TLDs), selected administrators to whom to delegate authority, and performed myriad other tasks. Postel had first assumed this task as a graduate student in the 1970s. As the Internet grew, the significance of Postel's decisions increased accordingly, and by the 1990s his decisions had global implications. Yet policy authority over the root continued to reside in him personally. Since he worked as a government contractor, final authority officially lay with the U.S. government—but for most of the 1980s and 1990s, Postel exercised personal authority over the DNS.

In 1984, in a document known as RFC920, Postel and colleague Joyce Reynolds defined the evolutionary trajectory of the name space (Postel & Reynolds, 1984). RFC920

defined the number of top-level zones and the names they would bear. While the name space would always have just one root zone file, Postel and Reynolds announced in 1984 that the top level would consist of some 250 zone files. RFC920 served as the blueprint for the structure of the namespace and its future growth. This number of zone files had no basis in technical necessity; it could have been smaller or larger.

RFC920 also specified the character strings to identify the zone files. The 250 TLDs would be divided into two naming classes: There would be six “generic” TLDs (gTLDs: .gov, .edu, .com, .org, .mil, and .net), and there would be some 244 “country code” TLDs (ccTLDs; based on the ISO 3166-1 standard list of two-character country codes like .uk for United Kingdom, .fr for France, .jp for Japan, and so on). Again, the particular character strings used in the domain names had no technical significance; they merely had to be unique. However, as chosen, the strings had major policy significance, for their meanings implied that different zones would have different uses. The 250 TLDs defined a name space partly based on function (.com for commercial, .mil for military, etc.) and partly based on geopolitical identifiers (country names). Decided long before the Internet’s global significance emerged, the number of TLDs and the meanings attached to them would have lasting consequences.

While RFC920 defined the DNS structure and naming conventions, its implementation proceeded incrementally over many years. Implementation of the TLDs required selecting an administrator to maintain the zone file and to operate the name server. Furthermore, the implementation of the generic TLDs proceeded very differently than the country code TLDs: The U.S. government selected the former, while IANA selected the latter.

While Postel/IANA possessed policy authority over the root, the administrator of the root was a private company: Network Solutions, Inc. (NSI). NSI took its orders from IANA, but ultimately operated under contract with the US government. NSI both administered and exercised policy authority over .com, .net, and .org.

Growth in .com made NSI both wealthy and powerful. After the U.S. government opened the Internet to commercial use in 1994, registrations in .com exploded. By the late 1990s, .com had grown to over 10 million registrations—more than half of the entire name space. This concentration of growth in the namespace was not an inherent feature of the DNS, but rather occurred as an unforeseen development—a combination of good marketing by NSI and widespread acceptance of the DNS naming convention, which identified .com as the commercial domain. Ultimately, the .com domain contained so much of the total name space that it rivaled the root for its importance for the overall network (Mueller, 1999). Charging an annual fee of \$35 per registered name, NSI collected hundreds of

millions of dollars of revenue from its monopoly of the Internet’s one commercially named domain.

In contrast to NSI, administrators of country code TLDs resembled IANA: They were usually nonprofit organizations, often affiliated with university research centers. Since IANA had defined zone files in terms of country codes and had created just one zone file per country, there was just one administrator in each country. Each of them constituted an implicit national monopoly: the .fr registry was France’s only registry, the .uk registry was the sole UK registry, and so on. Although there was no technical basis for national monopolies, the naming convention in RFC920 implied such a system. Organizationally, the system of national ccTLD monopolies was reminiscent of the system of national telephone companies (PTTs), which operated as national monopolies in many countries.

By October 2000 the full name space consisted of over 30 million name–number pairs (NetNames, 2000). Since IANA had not expanded the number of top-level domains since issuing RFC920, most growth in the name space occurred in second-level domains. Most of that was contained in a single TLD, .com, where NSI had registered over 18 million hosts. NSI’s .org and .net TLDs contained another 5 million hosts. The rest of the name space was mostly distributed in various country code TLDs. Above it all, Jon Postel at IANA oversaw delegations of authority to new administrators.

Thus the DNS as it existed toward the end of the 1990s was considerably more complex than the system described earlier. First, it was decentralized. The change from hosts.txt to the decentralized DNS occurred in 1983 (Mockapetris, 1983). Second, over the years numerous nontechnical developments had shaped the system. Most TLDs bore country code identifiers, which associated them with geopolitical entities (countries). One zone file in the hierarchy, .com, contained nearly the entire name space (thereby somewhat defeating decentralization). In the community of nonprofit DNS administrators, Network Solutions was emerging as a commercial giant. Most importantly, policy authority for the entire DNS lay with one person, Jon Postel. The DNS was complex—and rife with potential conflicts.

Governance

Nonetheless, the DNS could still be used to realize the mechanisms of governance. Although more complex arrangements would be needed than those discussed earlier, a decentralized DNS could still be used to realize authority, law, sanctions, and jurisdictions. In this section I consider in the abstract the manner in which the DNS technology would render governance possible.

Decentralization would not significantly affect the realization of two mechanisms: law and sanctions. Even

with the addition of multiple levels of hierarchy, flow-down contracts could still bring rules down to users. Decentralization would require a longer cascade of flow-down contracts but would not otherwise affect this mechanism. Likewise, domain name denial could still serve as a sanction. With each administrator in the name space, from the root on down, exercising monopoly control over its zone file, each one could delink lower level hosts from the name space. Name registration could still be treated as a privilege, revocable if a registrant violated the rules.

Decentralization would render the other two governance mechanisms considerably more difficult to realize, however. Decentralization fragmented authority and jurisdiction, especially in the country code domains.

Policy authority and jurisdiction would still be unified at the root. IANA could make rules for the entire name space and promulgate them down the hierarchy. However, at the top-level domains (TLDs) it would encounter another authority that might challenge it. The distinction between generic TLDs (gTLDs) and country code TLDs (ccTLDs) would inhibit unity of authority and jurisdiction.

ccTLDs are associated with countries, and most countries already have policy authorities: their national governments. National governments could claim jurisdiction in ccTLDs bearing their country code. Although national governments' domains were lower in the DNS hierarchy than IANA, it would be awkward for IANA to assert authority over them. An Internet governance institution would be in a poor position to challenge a national government's right to make public policy. Even if a national government was unaware of the DNS (as was often the case), the proactive exercise of policy authority by IANA might provoke a national government to act. Thus it would be difficult for the policy authority at the root to exercise policy authority in the ccTLDs.

Making matters more complex, among the ccTLDs there would be a plethora of authorities. Country code TLDs were independent from each other, and each could make its own policies in its domain, potentially resulting in a host of divergent and conflicting policies. Decentralization of the DNS had created hundreds of authorities, each with an implicit claim of jurisdiction over their ccTLD. Thus it would seem that integrated governance of the Internet would be impossible. The engineers' decision to organize the name space according to political lines, as codified in RFC920, had fragmented authority and jurisdiction.

In the gTLDs, in contrast, integrated governance would be possible. IANA could regulate domains like .com, .org, and .net, because these domains had no accountability to authorities outside of the DNS. Any authority they

possessed in the DNS was only by way of a delegation from IANA. Furthermore, although gTLDs were only a small number of the total set of domains, they contained the lion's share of all users. Effective authority in just these domains would still affect most Internet users.

Thus fully integrated governance would not be possible in the DNS—but a workable degree of governance would be. The realization of a single authority and a single jurisdiction in .com, .org, and .net would be straightforward. Moreover, such a jurisdiction would include most users, given the heavy concentration of registrations in the gTLDs. Policy authority at the root could regulate them by flow-down contracts backed by the threat of domain name denial. This would reduce but not overcome the problem of fragmented authority.

To further mitigate the fragmentation of authority, IANA and national governments could seek to coordinate policies. Although there were hundreds of ccTLDs, registrations were unevenly distributed among them. Domains like .uk or .jp contained many more registrations than others like .bg (Bulgaria). Coordination between IANA and just the largest ccTLDs would bring the Internet much closer to integrated governance by bringing most of the outstanding users under the same policies.

One could speculate that additional policy coherence could be achieved by pressuring any recalcitrant ccTLD authorities. Were a small ccTLD to resist enforcing some policy backed by IANA and by large governments, the larger parties could challenge the policy authority of the smaller national governments. IANA could exercise its ability to delink a top-level domain or reassign it to a more compliant administrator. In this way, smaller ccTLDs might be coaxed or bullied into adopting policies agreed on by larger players. Overall policy coherence would be improved.

In closing this section, a final issue must be addressed: the role of the U.S. government. The U.S. government employed Jon Postel and Network Solutions, and it claimed final authority over the root zone file. Although IANA was the highest policy authority in the DNS, IANA itself operated under policy authority of the United States. Were Internet governance mechanisms to be implemented, and were the status of the United States not to change, then the United States would be the final authority over the Internet. This again might cause tensions with other national governments, which would find themselves subordinate to the United States.

In summary, a decentralized DNS would not allow for full realization of mechanisms of governance. Law and sanctions could be easily realized, but authority and jurisdiction would be fragmented. Located in the gTLDs, most users could be regulated by IANA. Bringing the entire name space under IANA's authority, however, would

require negotiation with many autonomous national authorities.

ICANN

Having reviewed the DNS, we can now turn to the Internet Corporation for Assigned Names and Numbers (ICANN). Created in 1998 and still evolving at the time of this writing, ICANN realizes the governance potential in DNS, leveraging Internet addressing to achieve global governance. Not only has it created the capabilities for regulation, it has even employed them: In 1999 ICANN promulgated global public policy that defined intellectual property rights in domain names. In what follows I identify the specific features of ICANN by which it realizes authority, jurisdiction, law, and sanction.

I begin by setting the historical scene. By the late 1990s the DNS had come under severe stress from a variety of sources. The Internet had rapidly outgrown its original institutions, most notably the very personal nature of IANA, whose legitimacy was based on the reputation of one man. Should something have happened to Jon Postel, IANA could have become unstable. Another source of stress arose from entrepreneurs wishing to compete with NSI's monopoly: They began proposing alternate name spaces, new TLDs (e.g., .web), and independent registries (Mueller, 1998). This threatened to fragment the name space. The global nature of IANA was another issue. The United Nations' International Telecommunications Union (ITU) became involved and sought to assume authority over the name space. National governments and the European Commission became interested, too; they perceived a threat to their sovereignty from U.S. control of this new global information infrastructure. Disputes over sovereignty and jurisdiction were heating up. Intense conflicts also began to emerge over domain names that matched trademarks (e.g., coca-cola.com). The United Nations' World Intellectual Property Organization (WIPO) and U.S. interest groups applied intense political pressure to install trademark regulations in domain names (Shaw, 1997). Making this political mix all the more volatile was that these conflicts developed in "Internet time"; every passing month witnessed exponential growth in the size of both the network and the political stakes.

The process by which the research community, trademark interests, communication businesses, and national governments came together to create a new institution to replace IANA is documented elsewhere (Mueller, 1999; Klein, 2001a). Here we are interested in the product of that long and contentious process: ICANN. In the following institutional analysis I dissect a snapshot of ICANN, as it existed around year 2000.

ICANN is best understood as a set of semi-autonomous institutions. That set includes not only ICANN as a corporation but also some external entities like a committee of national governments and the TLD administrators. To distinguish ICANN-the-set-of-institutions and ICANN-the-corporation, I refer to the former as the "ICANN system" and the latter simply as "ICANN."

The four mechanisms of governance are mixed deep in ICANN's administrative system and so can be difficult to identify. In what follows, I analyze ICANN's features in terms of their governance-related functions. First I focus on how ICANN realized mechanisms for authority and jurisdictions, and in the following section I focus on the mechanisms for policy and sanctions.

Authority and Jurisdiction

The new ICANN corporation replaced Jon Postel as the policy authority over the root. ICANN solved the problem of stability: A person was replaced by an institution, so that the IANA could function independently of any one individual. ICANN also partially solved the problem of inter-governmental conflict: ICANN was private, and its bylaws explicitly forbade government officials from serving on the board. Thus although its authority would extend globally, that authority was ostensibly nongovernmental and would not conflict with national governments' sovereignty. Furthermore, with a mission to engage in simple technical coordination of the Internet, ICANN claimed to have no public policy role.

The problem of legitimacy was addressed by the composition of the board of directors. A person, Postel, was replaced by a collection of representatives; legitimacy through expertise and personal reputation was replaced by legitimacy through accountability to stakeholders. ICANN's board represented different functional and geographic constituencies. Of 19 directors, 9 represented technical expert groups, another 9 represented users, and the final director was the organization's top staff person.

ICANN's board, however, was itself subject to a higher authority: the U.S. government. The U.S. Department of Commerce (DOC) retained ultimate control of the root, leaving ICANN policy decisions subject to a potential veto. Despite the much-publicized privatization, the United States never completely ceded its hold over the Internet. As an official "fact sheet" of the DOC stated, "The Department of Commerce has no plans to transfer to any entity its policy authority to direct the authoritative root server" (DOC, 1999). Thus the Internet was internationalized and privatized but only under the watchful oversight of the U.S. government. (Whether the United States will eventually cede full authority to ICANN is not currently known.)

Beneath the root, contracts extended the authority of ICANN and the United States down to the administrators of the gTLDs and the ccTLDs. The generic TLDs proved more willing to sign on, since NSI administered nearly all of them and was under pressure from the United States to participate in ICANN. Following some bargaining over conditions, NSI and ICANN reached agreement in 1999. ICANN thereby achieved policy authority in the most populous domains. The ccTLDs proved more circumspect, and as late as 2001 ICANN was still reporting small progress in this area (ICANN, 2001). Top-down policy authority in these domains failed to be established and remained one of the most difficult issues in the system.

The implicit conflict of authority between ICANN and national governments manifested itself in the Governmental Advisory Committee (GAC). The GAC was an official ICANN advisory committee in which national governments could meet, discuss, and coordinate their actions. Individually, each national government could assert policy authority over the zone file bearing its country code. Together, in GAC, the national governments could coordinate policy.

GAC's first acts were to establish the legitimacy of members' claims to policy authority. First, GAC declared, "The Internet naming system is a public resource in the sense that its functions must be administered in the public or common interest" (GAC, 2000). By defining the DNS as a public good, similar to electromagnetic spectrum, GAC prepared the way for governmental oversight. Then GAC linked that public interest to national governments' authority: "ultimate public policy authority over the relevant [country code domain] rests with the relevant government" (GAC, 2000). This justified the claim by national authorities that ccTLD domains were under their jurisdiction.

Thus ccTLD administrators found themselves under two authorities—and they asserted a third of their own. ICANN claimed that ccTLDs' authority derived from its higher authority over the root; if administrators did not follow ICANN's policies, ICANN could redelegate authority to another party. For their part, national governments claimed that their zone file was a public resource under their authority. A third approach was backed by ccTLD administrators, who cited policy documents that located authority in the "local Internet community" rather than in ICANN or in governments (Postel, 1994). This prescription would render the administrators accountable to Internet users in their home country rather than to their government or ICANN.

GAC members sought to resolve this ambiguity in their favor by requesting from ICANN a veto power over ccTLDs similar to the U.S. veto power over the root. GAC proposed that ICANN's power of redelegation be given to national governments: "when ICANN is notified by the relevant government or public authority that the [administra-

tor] has contravened the terms . . . ICANN should act with the utmost promptness to reassign the delegation" (GAC, 2000). Country code managers would have access to the root only as long as their national government allowed it. ICANN resisted this arrangement, which would have subordinated it to national governments. At the time of this writing the fragmentation of authority over the ccTLDs remained unresolved.

The GAC forum allowed national governments to coordinate on a variety of other policies. GAC began developing a "best practice" document for country managers, so that national authorities could standardize their operations (GAC, 2000). Once common policies were defined, each national government could promulgate and enforce those practices in its own jurisdiction.

Thus the multiplicity of authorities caused jurisdictional fragmentation. ICANN claimed jurisdiction over the entire name space and, hence, over all users. Likewise, the U.S. jurisdiction extended over ICANN and so over the entire name space. At the top level of the name space, however, jurisdictional conflicts merged. The jurisdictions defined by the generic TLDs posed little problem; there ICANN prevailed. In the country code TLDs, however, national governments claimed jurisdiction. This prevented ICANN and the United States from realizing one unified jurisdiction in the name space. Still, the vast majority of Internet users found themselves in ICANN's sole jurisdiction.

Policy and Sanction

Although ICANN regulated users, it did not have direct contact with users. Instead, a four-tiered system was implemented, with ICANN at the top, users at the bottom, and two kinds of organization—registries and *registrars*—in between. At the top ICANN used its authority to make regulations. Beneath it, registries maintained the zone files and operated the servers (as described earlier). Beneath the registries came the registrars, who served as the retail interface to users. They performed customer-oriented tasks of leasing and servicing domain names to users, often bundling these with additional services like Internet service provision. Finally, at the bottom tier were users (or network administrators, who, in turn, contracted with individual users).

Flow-down contracts spanned these levels. ICANN's regulations were embodied in contracts with registries, which included the regulations in their contracts with registrars, who included them in their contracts with network administrators. Policies "flowed down" from ICANN to registries to registrars and ultimately to private networks. The terms of the contract defined the laws of the Internet.

At each level, ICANN backed up contracts with the threat of domain name denial. Registries who disobeyed could have their domain redelegated. Registrars who

disobeyed could lose their access to registries, so they would no longer be able to offer domain names to users. Users who disobeyed could have their domain names removed from the name space or assigned to someone else.

ICANN's Registrar Accreditation Contract (ICANN, 1999a) was the primary mechanism for promulgating law. Any organization that wished to serve as a registrar had to obey the terms of this contract. It included an open-ended requirement: "Registrar shall comply . . . with all ICANN-adopted Policies" (Section II.D.1.b.i).³ As ICANN policies changed and the accreditation agreement evolved (as foreseen in Section II.O, "Right to Substitute Updated Agreement"), so could the conditions imposed on domain name usage. It was this contractual blank check that most clearly gave ICANN the right to exercise broad governance activities. The contract's provisions had to be repeated in lower level contracts between registrars and users, guaranteeing that regulations would flow from ICANN to the registrars and eventually to users. These regulations were enforceable with clear sanctions: "The [domain name] holder shall agree that its registration of the [domain] name shall be subject to suspension, cancellation, or transfer pursuant to any ICANN-adopted policy . . . for the resolution of disputes" (Section II.J.7.i).

Thus the basic governance mechanisms were flow-down contracts backed up by domain-name denial. The accreditation contract stipulated the regulations for the Internet, and the power to revoke domain names provided the sanction for enforcement. Internet users could only enjoy access to the name space if they obeyed ICANN's rules; if they broke the rules, they could see their domain name suspended, canceled, or transferred.

GLOBAL PUBLIC POLICY

The preceding discussion of DNS, governance, and ICANN's institutional design has been largely descriptive and analytical. What is reviewed here is mostly established fact. If one accepts the four-part definition of governance, and if one admits that those mechanisms are indeed present in ICANN, then it is not controversial to state that ICANN has the capability to engage in Internet governance. So far, little has been said about whether ICANN exercises that capability or whether the system of governance is legitimate. In this section I engage some of these more contentious topics.

Global Public Policy

ICANN not only has the capacity to govern, it has also done so. ICANN has made global public policy. In this section I explain what it means to make public policy and how ICANN has done it.

Shortly after ICANN's governance capabilities were implemented, they were put into practice. In August 1999

ICANN promulgated its first major policy: the Uniform Dispute Resolution Policy (UDRP), which mandated a procedure for deciding who has intellectual property rights in a domain name (ICANN, 1999b). The UDRP illustrates how ICANN's governance mechanisms work in practice. The UDRP also constitutes ICANN's first global public policy.

In the late 1990s domain names became valuable, with names like yahoo.com and amazon.com becoming important business assets. As the value of domain names rose, disputes arose over name rights. Some disputes arose when individuals allegedly registered trademarks in anticipation of selling them to their owner; other arose when owners allegedly attempted to wrest control over desirable character strings from other users. Sometimes ownership rights conflicted with rights of fair use or free speech (Kleiman, 1999). The trouble with such conflicts was that existing trademark law was inadequate for many disputes: Trademark law was national, whereas many disputes were international. Legal mechanisms to settle international disputes in domain names were expensive and awkward to employ.

ICANN's UDRP defined procedures for resolving domain name disputes, thereby effectively setting rules of ownership and property. Disputes would be settled through an alternative dispute resolution procedure, in which certified private arbitrators would decide the question of rights based on criteria defined by ICANN. Arbitration decisions could be enforced by removal or transfer of the disputed name. It was a "voluntary" system in that parties who were unsatisfied by the outcome of arbitration could still resort to existing judicial forums for judgment. However, since existing forums were extremely costly, the UDRP would in most cases provide the final decision on property rights. The UDRP had de facto the force of law.

The implementation of the Uniform Dispute Resolution Policy (UDRP) illustrates ICANN's use of all four governance mechanisms. First, the UDRP was developed with input from staff and various parties and ultimately approved by the ICANN Board in an exercise of its authority. Second, the policy was codified into law through the Registrar Accreditation Agreement. ICANN made the UDRP a condition for registrar access to the name space, and registrars had to include the UDRP in their retail contracts with users (it had to "flow down"). Third, the UDRP included sanctions: Users who refused to agree to the policy in advance were denied access to the name space, and users who were found to be in violation of the UDRP could have their names deleted or reassigned (banishment). Finally, the UDRP applied in ICANN's jurisdiction. The policy regulated domain name usage in the .com, .net, and .org domains. In country code domains, where ICANN's authority did not immediately apply, the decision to adopt

the UDRP lay with country code administrators. At the time of this writing, the policy had been adopted in some domains but not in all.

In making the UDRP, ICANN made global public policy. The UDRP regulates something of public value: rules of property. Rules of property, such as trademark, copyright, and intellectual property, are traditionally made by governments. The U.S. Constitution, for instance, specifies rules of intellectual property protection in patents. At the global level, the reason for a lack of regulation is not that property rules are somehow less public, but that there has been no recognized public institution to make such rules. ICANN stepped into this void. By making global rules on property, ICANN made a decisions on public values. Although the UDRP may not have been a policy of enormous import—property rights in domain names is a relatively small area of regulation—it was a significant first step into policymaking.⁴

Legitimacy

If ICANN makes global public policy, then it is appropriately evaluated by such policy criteria as legitimacy, accountability, and equity. Indeed, it is around issues like these that most controversies have erupted (Weinberg, 2000; Froomkin et al., 1999; Klein, 2001c). Here I briefly review some of the issues that have arisen around the ICANN board's legitimacy.

U.S. policy for the creation of ICANN was laid out in the Department of Commerce "White Paper," which defined principles for ICANN. Two principles there were particularly relevant to legitimacy: ICANN should be committed to "private, bottom-up coordination," and it should be committed to "representation...[providing] input from the broad and growing community of Internet users" (DOC, 1998b). Some of these principles became embodied in ICANN's bylaws, especially in the mechanisms for representation on the board (Klein, 2001a).

In a number of instances these principles of legitimacy were not convincingly upheld. I mention three here. ICANN's first board of directors was a nine-person interim board. In a move that generated considerable public outcry (and hearings in the U.S. Congress), the first set of interim directors was appointed with no public participation or consultation. Instead, the selection process was conducted behind closed doors in a process later described even by Jon Postel as "undemocratic and closed" (Daley, 1998). However, it was this board that promulgated the Uniform Dispute Resolution Policy.

The ICANN board was also implemented unevenly. The directors for the nine expert representatives were seated within approximately 1 year of ICANN's creation, and those directors quickly moved to weaken the seats reserved

for Internet user representatives. In a series of board meetings in 1999 and 2000, the appointed and expert directors sought to eliminate, reduce, or delay the implementation of the elected directors (Klein, 2000a). In so doing, they repeatedly revised the corporate bylaws that constrained board actions. As one top government official declared to them at their meeting in July 2000, "the Board is increasingly giving the impression of being extremely cavalier in changes to the by-laws" (Wilkinson, 2000). Shortly thereafter, the board decided to modify the bylaws again to defer the final round of director elections until 2002—fully 4 years after the creation of ICANN.

Finally, ICANN's early board gave evidence of preferring industry professionals to represent Internet users. In the first, partial round of director elections, the board filled most of the nominee slots for user representatives with candidates of its own choosing, selecting individuals from telecommunications giants like France Telecom, Fujitsu, Deutsche Telekom, and Verizon (Klein, 2000b). This tendency to favor the telecommunications supply industry over Internet users attracted the notice at ICANN's July 2000 meeting, of an Australian government official, who stated, "[ICANN] runs the risk of potentially becoming a de facto industry association" (Twomey, 2000). The legitimacy of the board to make decisions affecting all Internet users was again weakened by this tendency for some interest groups to seek (and arguably to achieve) disproportionate influence on the board.

With ICANN making global public policy, its lack of legitimacy was striking. Although the elections of 2000 brought some user representation to ICANN, they fell short of implementing the degree of representation called for in ICANN's original bylaws (Klein, 2001b, 2001c).

Future Policy

Institutions are not static entities; they grow over time and often expand their areas of activity. Such seems likely for ICANN. As an Internet governance entity, what policies might ICANN promulgate in the future? Here I briefly speculate.

Perhaps the most likely area of policy expansion is in intellectual property protection. Such expansion of rights has been actively championed since the beginning of the ICANN process and would be consistent with the initial direction of ICANN's activities (Froomkin, 1999). The UDRP could be expanded to give special registration rights to owners of celebrity names, famous marks, geographical names, and so on. ICANN could become a global regulator in the service of property and e-commerce.

Control over the name space could also be leveraged to promote social justice. ICANN and ccTLD monopolies could raise funds for a universal service fund to

overcome the global digital divide, allowing poorer countries to pay less for Internet access than wealthy countries. In private conversations with this author, some ICANN directors from developing countries have supported such policies.

ICANN's capabilities could also be used for content regulation. Sites violating content regulations could be censored by having their domain names revoked or redirected. Domain name denial was used in this way in the case of *voteauction.com*, which operated a site containing illegal content (an online mechanism for the buying and selling of votes). The registrar for the domain canceled the registration to suppress its content (Perritt, 2001). In theory, ICANN could enforce regulations broadly using similar mechanisms.

ICANN could become a vehicle for taxation, perhaps serving as a means whereby governments collect e-commerce taxes or whereby ICANN funds its own initiatives. With domain names available from one sole source, users would have to either pay the fee or suffer denial of access. Indeed, both U.S. legislators and disgruntled ccTLDs have accused ICANN of levying taxes (McCullagh, 1999; Ward, 2000).

Finally, ICANN could become a vehicle for U.S. national policy. In times of war or terrorism, countries opposing the United States could see their domains removed from the Internet. Individual registrations could be canceled or redirected to reduce the effectiveness of hostile entities. The relationship between ICANN policy and U.S. national policy had already come up when the U.S. Department of Commerce had approved the addition of the .ps domain—for Palestine—to the root zone. Although the United States did not pursue its narrow national interest, the case attracted attention because of the potential for conflict (Cisneros, 2001).

With governance mechanisms in place, the possibility of mission creep—the steady expansion of ICANN's regulatory scope—seems possible. The combination of effective mechanisms for governance and weak mechanisms for legitimacy could allow some parties to make global public policies that favor their interests.

CONCLUSIONS

The simple recognition that ICANN engages in Internet governance is significant. It contradicts established beliefs and it raises concerns about what kind of governance is being established. It forces us to ask what should be done.

ICANN contradicts the popular myth of a benevolent Internet anarchy. As it turns out, the Internet *can* be controlled. The DNS provides a basis for top-down control, and ICANN leverages that to make policy. The implications of this are far-reaching and will only be seen over

time. For this reason, ICANN's future should concern all Internet users.

I close with some observations about the relationship of technology to society. In ICANN we see three ways in which technology shapes society.

First, objective features of the technology shaped the administrative and regulatory system. In particular, the technical characteristics of a distributed database set important policy parameters. The need for a single name space with a unique root created a central control point. Likewise, the need for unique identifiers (so that a name identifies just one host computer) created problems of control and monopoly. With just one .com zone, within which there could be registered just one .coca-cola, a system of monopoly registries was created and the basis for trademark fights was laid.

It may be that these technology design features were not absolutely necessary (although convincing alternatives were not advanced even by ICANN's critics). At minimum, however, history rendered those technical features sufficiently embedded that they became equivalent to "necessary." Any attempt to change ICANN's status as regulator may have to begin by redesigning the underlying technology (particularly the requirement for a unique root under the control of a unique administrator).

A second way that technology shaped society was the role of engineers in making policy. The selection of country codes as domain names was a historical decision with profound policy consequences. This decision was made so early in the Internet development process that the only participants were research engineers. These engineers decided that Internet domains should be associated with geopolitical entities. Had they selected different alphanumeric identifiers—one could imagine colors, sequential numbers, the table of elements—then there would have been no basis for the one-per-country distribution of registries nor the subsequent assertion of national authority over registries. The engineers decided to organize the Internet like national PTTs. Engineers could make such decisions because they controlled the process early in its developmental history.

A third way that technology influences society is that it provides legitimacy for secretive decision-making. When policy decisions are categorized as "technical," then it becomes legitimate for them to be made behind closed doors by elite groups. Policy disappears from the public sight (Lessig, 1999). The groups that gained control of ICANN invoked this veil of technological legitimacy to discount their critics. Despite having no technical training, ICANN's lawyers justified their actions by claiming that they were making neutral choices on the basis of technical expertise (McLaughlin, 2000).

ICANN leverages control of Internet addressing to realize global public policy. In ICANN technology has shaped society, technologists have made profound policy decisions, and interested social groups have exploited technological legitimacy. Most importantly, the regulatory framework for the global information infrastructure of the next century has been created.

NOTES

1. This is a slight oversimplification. In fact, Internet communication is possible by using IP numbers directly, which avoids the need for interaction with the DNS. However, very few communications use IP numbers directly. A note to the reader: In this author's experience, for nearly every statement made about the DNS an exception can be found. The reader should bear this in mind when reading this section, which may contain some oversimplifications in the pursuit of clarity.

2. This sentence overstates the case a bit. Some computers may be listed more than once. Other may be listed not at all and may be reached by using their IP address directly. In the majority of cases, however, a computer on the Internet has one entry in the DNS name space. Furthermore, as discussed later, most computers in the DNS are not user computers but gateways to private networks within which individual user accounts exist.

3. This part of the Registrar Accreditation Agreement is important but not particularly succinct. The full text of Section D.1.b.i is: "D. General Obligations of Registrar. 1. During the Term of this Agreement: b. Registrar shall comply, in such operations, with all ICANN-adopted Policies insofar as they: i. relate to one or more of the following: (A) issues for which uniform or coordinated resolution is reasonably necessary to facilitate interoperability, technical reliability and/or stable operation of the Internet or domain-name system, (B) registrar policies reasonably necessary to implement Consensus Policies relating to the Registry, or (C) resolution of disputes regarding the registration of domain names (as opposed to the use of such domain names)."

4. One question that I do not address here is whether the UDRP is "good" or "bad" public policy. The substance of the regulation is not the issue. UDRP has been praised by some (Cohen, 2000) and condemned by others (Flynn, 2000; Mueller, 2001). What is important here is that fact that ICANN regulates at all.

REFERENCES

- Agence France-Presse. 1999. Internet pioneer Vinton Cerf pans French ruling against Yahoo! November 25.
- Albitz, Paul, and Liu, Cricket. 1998. *DNS and BIND*. Cambridge, MA: O'Reilly.
- Andrews, Edmund. 1999. German court overturns pornography ruling against Compuserve. *New York Times*, November 18.
- Bijker, Wiebe. 1995. *Of bicycles, bakelites, and bulbs: Toward a theory of sociotechnical change*. Cambridge, MA: MIT Press.
- Bijker, Wiebe, Hughes, Thomas, and Pinch, Trevor, eds. 1987. *The social construction of Technological systems: New directions in the sociology and history of technology*. Cambridge, MA: MIT Press.
- Cerf, Vint, and Kahn, Robert. 1974. A protocol for packet network interconnection. *IEEE Transactions on Communications* COM-22 (5):637-648.
- Cisneros, Oscar. 2001. Dot-PS: Domain without a country. *Wired News*, January 12. (<http://www.wired.com/news/politics/0,1283,41135,00.html>)
- Cohen, Jonathan. 2000. *Presentation on UDRP*. Korea Internet Forum, Seoul, Korea.
- Dahl, Robert. 1989. *Democracy and its critics*. New Haven, CT: Yale University Press.
- Daley, William. 1998. Letter to the Internet Corporation for Assigned Names and Numbers. 15 October.
- Department of Commerce. 1998b. Management of Internet names and addresses (white paper). *Federal Register* 63(111).
- DOC (U.S. Department of Commerce). 1999. Domain Name Agreements between the US Department of Commerce, Network Solutions, Inc., and the Internet Corporation for Assigned Names and Numbers (ICANN), Fact sheet, 28 September 1999.
- Flynn, Laurie. 2000. Trademarks winning domain fights. *New York Times on the Web*, 4 September.
- Froomkin, Michael. 1997. The Internet as a source of regulatory arbitrage. In *Borders in cyberspace*, eds. Brian, Kahin, and Charles Nesson, pp. 129-163. Cambridge, MA: MIT Press.
- Froomkin, Michael. 1999. *A commentary on WIPO's "The management of Internet names and addresses: Intellectual property issues."* Version 1.0, 17 May. (<http://www.law.miami.edu/~amf/commentary.htm>)
- Froomkin, Michael. 2000. Wrong turn in cyberspace: Using ICANN to route around the APA and the Constitution. *Duke Law Journal* 50(17):17-184.
- Froomkin, Michael, Post, David, and Farber, David. 1999. *ICANN-watch*. (<http://www.ICANNwatch.org>) Viewed 30 November.
- Fryer, Bronwyn. 1995. The software police: They hear from the snitch you copied that disk, they send in the marshals to bust your ass. No joke. *Wired*, May.
- Governmental Advisory Committee. 2000. *Principles for delegation and administration of ccTLDs*. Presented at ICANN Board meeting, 23 February. (<http://www.icann.org/gac/gac-ccldprinciples-23feb00.htm>)
- Hafner, Katie, and Lyon, Matthew. 1998. *Where wizards stay up late: The origins of the Internet*. New York: Touchstone.
- Holitscher, Marc. 1999. Debate: Internet governance. *Swiss Political Science Review* 5(1):115-116.
- Hughes, and Thomas P. 1983. *Networks of power: Electrification in Western society, 1880-1930*. Baltimore, MD: Johns Hopkins University Press.
- ICANN. 1999a. *Registrar accreditation agreement*. 12 May. (<http://www.icann.org/ra-agreement-051299.html>)
- ICANN. 1999b. *Uniform dispute resolution policy (UDRP)*. 24 October. (<http://www.icann.org/udrp/udrp-policy-24oct99.htm>)
- ICANN. 2001. *Third status report under ICANN/US government memorandum of understanding*. 3 July. (<http://www.icann.org/general/statusreport-03jul01.htm>)
- Intellectual Property Constituency. 2000. *IPC position paper on ccTLD issues*. Presented at ICANN Board meeting, 1 March. (<http://www.icann.org/cairo2000/ipc-position-01mar00.htm>)
- Internet Architecture Board. 1999. *Request for comments 2826: IAB technical comment on the unique DNS root*. (<http://www.rfc-editor.org/rfc.html>)
- Johnson, David, and Post, David. 1997. The rise of law on the global network. In *Borders in cyberspace*, eds. Brian, Kahin, and Charles Nesson, pp. 3-47. Cambridge, MA: MIT Press.

- Kapor, Mitch. 1990. *The software design manifesto*. (http://www.kapor.com/homepages/mkapor/Software_Design_Manifesto.html) Visited October 20, 2000.
- Kleiman, Kathryn. 1999. *Brief of amicus curiae Association for the Creation and Propagation of Internet Policies. Worldspport Networks Limited v. Artinternet S.A. and Cedric Loison*. U.S. District Court for the Eastern District of Pennsylvania, No. 99-Cv-616. (<http://www.domain-name.org/worldsport.html>)
- Klein, Hans. 2000a. Cyber-Federalist No. 4: An analysis of the ICANN-Named Board Nominees. August 8 [online]. (<http://www.cyber-federalist.org>)
- Klein, Hans. 2000b. System development in the federal government: How technology influences outcomes. *Policy Studies Journal* 28(2):313–328.
- Klein, Hans. 2001a. Online social movements and Internet governance. *Peace Review* 13(3):403–410.
- Klein, Hans. 2001b. The feasibility of global democracy: Understanding ICANN's at large election. *info* 3(4):333–348.
- Klein, Hans, ed. 2001c. *Global democracy and the ICANN elections* [special issue]. *info* 3(4).
- Klein, Hans, and Kleinman, Daniel. 2002. The social construction of technology: Structural considerations. *Science Technology & Human Values* January: 28–52.
- Leiner, Barry, Cerf, Vinton, Clark, David, Kahn, Robert, Kleinrock, Leonard, Lynch, Daniel, Postel, Jon, Roberts, Lawrence, and Wolff, Stephen. 2000. *A brief history of the Internet by those who made the history, including Barry Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Lawrence G. Roberts, Stephen Wolff*. August, (<http://www.isoc.org/internet/history/>)
- Lessig, Lawrence. 1999. *Code and others laws of cyberspace*. New York: Basic Books.
- McCullagh, Declan. 1999. ICANN too tax you. *Wired News* 18 June. (<http://www.wired.com/news/print/0,1294,20293,00.htm>)
- McLaughlin, Andrew. 2000. *ICANN: Myths and reality*. TIES Conference, Paris, April.
- Mockapetris, Paul. 1983. *Domain names—Concepts and facilities. Request for comments 882*. Published by Internet Architecture Board. (<http://www.rfc-editor.org/rfc.html>)
- Mueller, Milton. 1998. The battle over Internet domain names: Global or national TLDs. *Telecommunication Policy* 22(2):89–107.
- Mueller, Milton. 1999. ICANN and Internet governance: Sorting through the debris of “Self-regulation.” *info* 1(6):497–520.
- Mueller, Milton. 2000. Technology and institutional innovation: Internet domain names. *International Journal for Communication Law and Policy* 5:1–32.
- Mueller, Milton. 2001. Rough justice. *The Information Society* 17(3):151–163.
- NetNames. 2000. *Internet comes of age with 30 millionth domain name*. Press release. Boston: Netnames.com, 4 October.
- Perritt, Henry. 1997. Jurisdiction in cyberspace: The role of intermediaries. In *Borders in cyberspace*, eds. Brian Kahin, and Charles, Nesson, pp. 164–204. Cambridge, MA: MIT Press.
- Perritt, Henry. 2001. *Electronic commerce: Issues in private international law and the role of alternative dispute resolution*. WIPO Forum on Private International Law and Intellectual Property, Geneva, January.
- Post, David. 1998. Cyberspace's constitutional moment. *American Lawyer* November.
- Postel, Jon, and Reynolds, Joyce. 1984. *RFC:920: Domain requirements*. USC Information Sciences Institute. (<ftp://ftp.isi.edu/in-notes/rfc920.txt>)
- Postel, J. 1994. *RFC:1591: Domain name system structure and delegation*. (<ftp://ftp.isi.edu/in-notes/rfc1591.txt>)
- Schroeder, Ralph, ed. 1998. *Max Weber, democracy and modernization*, pp. 107–134. New York: St. Martin's Press.
- Shaw, Robert. 1997. Internet domain names: Whose domain is this? In *Coordination of the Internet 1997*. Cambridge, MA, MIT Press.
- Twomey, Paul. 2000. Spoken comments before the ICANN Board of Directors. (<http://cyber.law.harvard.edu/icann/yokohama/>)
- Ward, Mark. 2000. *Net groups in World Wide Wrangle*. BBC News, 4 July. (http://news.bbc.co.uk/hi/english/sci/tech/newsid_817000/817657.stm)
- Weinberg, J. 2000. ICANN and the problem of legitimacy. *Duke Law Journal* 50(1):187–260.
- Wilkinson, Christopher. 2000. Spoken comments before the ICANN Board of Directors. (<http://cyber.law.harvard.edu/icann/yokohama/>)