



University of Pennsylvania
ScholarlyCommons

Legal Studies and Business Ethics Papers

Wharton Faculty Research

2010

Castle in the Air: A Domain Name System for Spectrum

Kevin Werbach
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/lgst_papers



Part of the [Business Law, Public Responsibility, and Ethics Commons](#)

Recommended Citation

Werbach, K. (2010). Castle in the Air: A Domain Name System for Spectrum. *Northwestern University Law Review*, 104 (2), 613-640. Retrieved from https://repository.upenn.edu/lgst_papers/34

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/lgst_papers/34
For more information, please contact repository@pobox.upenn.edu.

Castle in the Air: A Domain Name System for Spectrum

Abstract

This article envisions the foundational infrastructure for a true wireless Internet. The domain name system (DNS) for addressing allowed the Internet to scale as a decentralized, loosely-coupled system. A similar system for the wireless communication would allow devices to negotiate frequently assignments and other attributes dynamically. The traditional, static approach to spectrum allocation creates massive inefficiencies, which will become increasingly problematic as wireless demand grows. A DNS for spectrum could be based on the database the Federal Communications Commission recently mandated for devices operating in the “White Spaces” around broadcast television channels. Such an infrastructure would enable rapid growth and innovation in next-generation mobile devices and applications.

Disciplines

Business Law, Public Responsibility, and Ethics

CASTLE IN THE AIR: A DOMAIN NAME SYSTEM FOR SPECTRUM

By Kevin Werbach*

This article envisions the foundational infrastructure for a true wireless Internet. The domain name system (DNS) for addressing allowed the Internet to scale as a decentralized, loosely-coupled system. A similar system for the wireless communication would allow devices to negotiate frequently assignments and other attributes dynamically. The traditional, static approach to spectrum allocation creates massive inefficiencies, which will become increasingly problematic as wireless demand grows. A DNS for spectrum could be based on the database the Federal Communications Commission recently mandated for devices operating in the “White Spaces” around broadcast television channels. Such an infrastructure would enable rapid growth and innovation in next-generation mobile devices and applications.

INTRODUCTION

In the early 1980s, the Internet faced a problem. The network of networks was growing rapidly... too rapidly. The increasing number of connected systems was overwhelming the simple system that tracked Internet host addresses. A single, static list couldn't keep up with the growing complexity and fluidity of the Internet. The solution, defined in 1983, was an elegant mechanism called the domain name system (DNS).¹ The DNS established a separate, distributed system for dynamic management of Internet addresses. Every time you type the uniform resource locator (URL) of a website or send an email message, you

* Assistant Professor of Legal Studies and Business Ethics, The Wharton School, University of Pennsylvania. Contact: werbach@wharton.upenn.edu. Thanks to the participants in the Maturing Internet Studies workshop at Northwestern Law School, and especially its organizer Jim Speta, for helpful comments.

¹ Paul Mockapetris, Domain names -- Concepts and Facilities, RFC 882 (November 1983); Paul Mockapetris, Domain names -- Implementation and Specification, RFC 883 (November 1983).

invoke the DNS. More than twenty-five years later, through massive growth and development of the Internet, the DNS continues to function effectively.

Today, wireless communication faces a similar problem. The rapidly increasing volume of traffic, as well as the growing number and sophistication of devices, threaten to overwhelm the simple, static spectrum allocation system managed by the Federal Communications Commission (FCC) and other government agencies. Command-and-control allocation of frequencies is wildly inefficient. As wireless data connections become increasingly common, the absence of a dynamic spectrum allocation mechanism becomes an increasingly severe restraint on the evolution of the Internet into a ubiquitous mobile platform.

The answer is a solution similar to the DNS: a hierarchical distributed database that brokers among those seeking and providing access to spectrum. The foundations for such a system are being laid today, without policy-makers realizing it. The FCC voted in late 2008 to authorize unlicensed wireless “White Space Devices” (WSD), subject to the creation of a database of available transmission slots throughout the country. This White Space Database could be the foundation for a Spectrum Networking Database (SND) that plays a DNS-like role in wireless communication. To achieve its potential, however, such a system would need to address both technical and policy requirements. Successful deployment of an SND would facilitate the continued development of a ubiquitous wireless Internet.

I. SCALING: WIRELESS AND THE INTERNET

A. *The Spectrum Challenge*

As billions of mobile phones gain increasingly sophisticated data capabilities, and trillions of wireless sensors are integrated with physical spaces and objects, conventional approaches to spectrum management will crumble. Current spectrum policy debates focus on initial allocations, such as the rules and auction in 2008 for licenses in the 700 MHz band reclaimed from UHF television stations.² The real challenge, however, comes in implementation. The only way to meet growing wireless capacity demands will be to view the spectrum as an ocean of potential capacity, theoretically available to any local device at any time. No centralized mechanisms, whether publicly or privately determined, can be

² In re Serv. Rules for the 698-746, 747-762 & 777-792 MHz Bands, 22 F.C.C.R. 15289 (Aug. 10, 2007) (second report and order)

sufficiently dynamic and localized for such a roiling ocean of wireless communication.

To a first approximation, the future of the Internet is wireless. Today, there are over 225 million mobile broadband users worldwide.³ That number nearly doubled from 2008 to 2009.⁴ For telephone service, there are already more than twice as many users of mobile phones as landlines worldwide. That means more than two billion people who today own a mobile phone have never owned a personal computer – and may never do so.

For most of its history, the Internet was primarily delivered to end users through wired network connections. Wireless systems simply did not offer the capacity and stability of their wired cousins, such as the landline telephone network and coaxial cable television connections. Until the past decade, high-speed wireless data systems and the mobile devices to make use of them simply were not available.⁵

In recent years, there has been an explosion of growth in wireless data. Unlicensed wireless hotspots using the WiFi and related protocols for short range connection have rapidly proliferated, as have laptops and other devices with the capability to make use of these networks.⁶ Wide-area wireless data services for the commercial market became feasible with latter version of digital mobile phone technology (so-called 2.5G systems), and have taken off with the growth of third-generation 3G mobile wireless technologies. In the US, all the major wireless carriers including Verizon, AT&T, Sprint and T-Mobile now offer nationwide wireless data coverage, with downstream speeds over one megabit per-second. New entrants such as Clearwire promise even faster data speeds.

All this mobile data growth is causing huge increases in demand for wireless capacity. The change is not merely in the absolute numbers, but in the quality of usage. Apple's phenomenally successful iPhone is the best example.⁷ The iPhone connects to the same networks as other mobile phones, but it encourages much more active usage because it provides a "real" Internet experience. Google found the average iPhone user searched the Web fifty times more frequently than users of less-capable phones.⁸ In the first few weeks after Apple introduced the iPhone

³ See Nick Wood, *Global Mobile Broadband Subscriptions Near Quarter of a Billion*, TOTAL TELECOM, July 22, 2009, at <http://www.totaltele.com/view.aspx?ID=447509>.

⁴ See *id.*

⁵ Wireless data networks began to be deployed in the 1980s for special-purpose applications, such as paging and fleet dispatch services. Metricom's Ricochet system provided end-user wireless data service beginning in 1994, but it was a proprietary offering with limited speed and coverage area, which never gained significant traction.

⁶ See *Wi-Fi: It's Fast, it's Here – and it Works*, BUS. WEEK, Apr. 3, 2002.

⁷ Leslie Cauley, *iPhone Gulps AT&T Network Capacity*, USA TODAY, June 17, 2009.

⁸ Maija Palmer and Paul Taylor, *Google Homes in on Revenues from Phones*, FT.COM, Feb. 13, 2008, at <http://www.ft.com/cms/s/0/667f13de-da60-11dc-9bb9-0000779fd2ac.html>.

3GS, featuring enhanced video recording capabilities, YouTube saw mobile video uploads increase 400%.⁹

The iPhone is not an outlier; it is a harbinger of things to come. AT&T Wireless predicts that broadband wireless data capacity demands will increase by a factor of 250 to 600 times between 2008 and 2018.¹⁰ More efficient technologies will provide some increase in capacity. For example, fourth-generation wireless systems such as Long Term Evolution (LTE) and WiMax promise to deliver greater capacity in the same spectrum as current 3G systems. These upgrades alone, however, will never match the scope of increased demand that is forecasted.

Adding new capacity by allocating licenses to new spectrum bands is also not a viable long-term solution. There have been no significant usable “empty” frequency bands for some time. The “low-hanging fruit” for clearing small-scale incumbent uses to reallocate frequencies for more valuable broadband data uses has already been tapped.¹¹

The terrible irony is that spectrum utilization today is wildly inefficient.¹² Spectrum is not actually scarce, at least not to anywhere near the extent it seems.¹³ The frequency allocation table shows a dense checkerboard. However, measurements of actual wireless activity show a radically different picture: Most frequency bands are not actually used for transmission most of the time, in most places. In fact, a study of wireless utilization in the frequencies below 3 GHz showed that only about five percent of bands were in use, on average.¹⁴ Even in

⁹ YouTube Blog, *YouTube Mobile Uploads Up 1700% in 6 Months; 400% Since iPhone 3GS Release*, at <http://www.youtube.com/blog?entry=kbaLH7fmm-g>.

¹⁰ Rysavy Research, *Mobile Broadband Spectrum Demand*, Dec. 2008, at 12.

¹¹ See Susan P. Crawford, *The Radio and The Internet*, 23 BERKELEY TECH. L.J. 933, 934 (2008) (calling the January 2008 auction of 700 MHz frequencies, “probably the last competitive auction for a substantial amount of spectrum for the next few decades....”).

¹² See Kevin Werbach, *Supercommons: Toward a Unified Theory of Wireless Communication*, 82 TEXAS L. REV. 863 (2004); Phil Weiser, *The Untapped Potential of Wireless Spectrum*, Brookings Institute Report, available at http://www.brookings.edu/papers/2008/07_wireless_weiser.aspx; Victor Pickard and Sascha D. Meinrath, “Revitalizing the Public Airwaves: Opportunistic Reuse of Government Spectrum,” *Wireless Future Working Paper*, New America Foundation (June 2009), forthcoming in *International Journal of Communications* (2009).

¹³ See *Supercommons*, *supra* note 12.

¹⁴ See Mark McHenry, *Dupont Circle Spectrum Utilization During Peak Hours*, A Collaborative Effort of The New America Foundation and The Shared Spectrum Company, New America Foundation Issue Brief (2003), available at http://www.newamerica.net/files/archive/Doc_File_183_1.pdf; Mark McHenry, *NSF Spectrum Occupancy Measurements: Project Summary*, Shared Spectrum Company (August 2005), available at <http://www.sharedspectrum.com/measurements/>.

New York City, the most densely populated metropolitan area, the number was only thirteen percent.¹⁵

The reality is that most spectrum licensees do not, or cannot, use their frequencies in anywhere near the most efficient possible manner.¹⁶ In some cases this is due to the inflexibility of FCC licenses and the legacy characteristic of old systems. For example, broadcast television allocations were made in the 1940s. They required a substantial percentage of the frequencies to be dark as guard bands, because receivers were not sophisticated enough to distinguish signals otherwise.¹⁷

Both the demand and the supply of spectrum are inherently local, short-term phenomena.¹⁸ Whether one system can transmit without inhibiting the ability of other systems to do so depends on a multitude of factors, including the technical characteristics of the transmitters and receivers; the nature of the services involved; and the physical geography of the area. And the situation will change over time. The wireless capacity that is “wasted” is bounded by time, geography, and technical characteristics. It cannot be specified globally for all time. Only a system with up-to-date local information can identify the full opportunity space for wireless communication.

The existing spectrum allocation structure is so inefficient because it is inherently static and centralized. Whether a device in Omaha, Nebraska can transmit at ten watts of power at 488 MHz (UHF television channel 17) at 2:00pm this Thursday is a determined by a decades-old table of frequency allocations issued by the FCC.¹⁹ The frequency table answers the question of who is legally permitted to transmit in a designated frequency band across a large geographic area. It does not address two more important questions: who is actually transmitting in those frequencies, and who else could transmit without disturbing the licensee.

¹⁵ See *id.* And this was during the Republican National Convention in New York, a time of extremely high wireless use.

¹⁶ The FCC itself acknowledges that we face not a spectrum capacity problem, but a spectrum access problem. See Federal Communications Commission, *Spectrum Policy Task Force Report*, ET Docket No. 02-135 (November, 2002), available at <http://www.fcc.gov/sptf/reports.html>.

¹⁷ See Randy Hoffner, *White Space Devices: Threat to Broadcast TV?*, TV Tech., Dec. 5, 2007, at <http://www.tvtechnology.com/pages/s.0079/t.10086.html>; Sascha D. Meinrath & Michael Calabrese, *White Space Devices” & the Myths of Harmful Interference*, 11 N.Y.U. J. LEGIS. & PUB. POL’Y 495 (2008).

¹⁸ See Supercommons, *supra* note 12.

¹⁹ Moreover, spectrum allocations until recently were specific as to service. A television broadcaster cannot decide its license would be more efficiently or profitably used to offer mobile data services; it must operate a broadcast system or nothing at all.

The static spectrum allocation system is a historical artifact. Until recently, wireless systems generally lacked the ability to sense the environment around them. They could merely transmit or receive in a narrow range of frequencies, using specific waveforms. A television receiver, for example, could only look for television transmitters with sufficient signal strength to reach it. In the case of older services such as television, those receivers were assumed to be very poor by today's standards at identifying the relevant signals. As a result, huge swaths of spectrum were designated as buffer space between allocated frequencies. All those decisions were made centrally, by the FCC. It issued licenses that authorized certain transmitters and forbid all others.

This "command and control" spectrum allocation process was workable when wireless communications systems were limited, but breaks down today. When wireless networks were predominantly one-way, large-area, fixed radio and television systems or point-to-point relays, and overall wireless demand was small compared to the present, the inefficiencies were relatively small. The higher the level of demand, and the more dynamic the systems involved, the less a command-and-control approach can approach optimal efficiency. In a world of ubiquitous wireless broadband, the viability of the existing spectrum allocation process is increasingly called into question.²⁰

The only viable solution for the wireless capacity crunch is therefore a distributed one. Spectrum allocation decisions must be made locally and for limited durations. Cellular phone systems already do something like this in their licensed spectrum, dynamically allocating capacity to handsets within range of towers. The operator cannot anticipate all possible usage scenarios ahead of time; it allocates some decision-making to software in each base station. However, the cellular example is very limited. Cellular networks operate in defined licensed spectrum bands for the purpose of supporting a particular uniform service, with the ability to control handset technology directly.

The two main proposals for significant reform of spectrum allocation both employ dynamic, local decision-making.²¹ Under the "property" approach, spectrum licensees would be granted property rights to use or sell their spectrum as they pleased.²² This approach traces its intellectual history to a famous 1959

²⁰ See JONATHAN E. NUECHTERLEIN & PHILIP J. WEISER, *DIGITAL CROSSROADS: AMERICAN TELECOMMUNICATIONS POLICY IN THE INTERNET AGE* 239 (MIT Press 2005).

²¹ For a comparison of property and commons models for spectrum, See Supercommons, *supra* note 12; Yochai Benkler, *Overcoming Agoraphobia: Building the Commons of the Digitally Networked Environment*, 11 HARV. J. L. & TECH. 287 (1998); Gerald R. Faulhaber & David Farber, *Spectrum Management: Property Rights, Markets, and the Commons* (working paper), at http://rider.wharton.upenn.edu/~faulhabe/SPECTRUM_MANAGEMENTv51.pdf; Eli Noam, *The Fourth Way for Spectrum*, FT.com, May 29, 2003, at <http://news.ft.com/comment/columnists/neweconomy>.

²² See Supercommons, *supra* note 12.

article by Nobel Laureate Ronald Coase, urging the FCC to assign property rights rather than more-limited licenses.²³ In a property system, decisions about what to do with frequency blocks devolve from a central government regulator to a collection of private owners. These owners can make local decisions about use and access to their spectrum, including whether to sell that spectrum to another owner who values it more.

The major alternative to the property approach, the so-called “commons” model, is even more distributed and dynamic.²⁴ Under a commons approach, anyone can transmit, subject to technical standards. Devices must negotiate themselves to avoid interference. The best example in practice of this model is the WiFi technology operating in the 2.5 GHz and 5 GHz unlicensed bands. Millions of WiFi base stations offer short-range connections in homes and businesses, coexisting despite the absence of exclusive rights.²⁵ The WiFi protocol allows individual devices to mitigate interference locally, and individual device owners can also coordinate transmission channels for more efficient operation. As with the property approach, these decisions are superior to those of command-and-control allocation because they are made “on the ground” in response to local conditions, rather than being specified, indefinitely and ahead of time, by a central regulator.

Technology is making such a dynamic spectrum allocation model feasible. Historically, wireless devices could only operate on a limited range of frequencies and transmit in very specific ways.²⁶ These characteristics had to be defined when the device was built, and built into the physical hardware. Moreover, devices had little or no ability to understand the local spectral environment. Numerous technical advances are changing radios into something entirely different: adaptable connected digital platforms.²⁷ Wireless devices today can be “tuned” to different frequencies and waveforms through software,²⁸ and can manage interference through power limits or technical protocols. These

²³ Ronald H. Coase, *The Federal Communications Commission*, 2 J. L. & ECON. 1, 1 (1959).

²⁴ See Supercommons, *supra* note 12.

²⁵ See *supra* note 6.

²⁶ See Supercommons, *supra* note 12; Kevin Werbach, Radio Revolution, New America Foundation Working Paper (2002).

²⁷ See Radio Revolution, *supra* note 26; Supercommons, *supra* note 12; Preston F. Marshall, A Potential Alliance for World-Wide Dynamic Spectrum Access, New America Foundation Wireless Future Program Issue Brief #25, June 2009; Crawford, *supra* note 11.

²⁸ See William Lehr, et al., *Software Radio: Implications for Wireless Services, Industry Structure, and Public Policy*, Massachusetts Institute of Technology Program on Internet and Telecoms Convergence (August 2002), at <http://itc.mit.edu>; Authorization and Use of Software Defined Radios First Report and Order, 16F.C.C.R. 17373 (2001); Spectrum Policy Task Force Report, ET Docket No. 02-135, Comments of Vanu, Inc. at 1-2.

developments allow for a grand shift of spectrum utilization from static to dynamic.²⁹

Dynamic Spectrum Access (DSA) means that wireless devices could transmit based on actual availability of capacity rather than pre-set rules of exclusive frequency allocation.³⁰ DSA techniques are the subject of extensive technical research today.³¹ DARPA, the research and development arm of the US Department of Defense, funded significant development work on adaptive radio technologies which could be employed for DSA through its XG program.³² Academics have suggested technical rules and pricing mechanisms for “secondary” users to access licensed spectrum on an as-need and as-available basis.³³ Even Google has entered the fray, proposing continuous real-time auctions for spectrum capacity, analogous to the way it sells advertising on its search engine.³⁴

The limitation of all these proposals is the absence of a coordination mechanism. Each envisions a particular DSA mechanism for a particular purpose, even though the DSA concept is potentially applicable to the entire usable spectrum.³⁵ What is needed is a way to “flip the switch” from the static to the dynamic paradigm. The missing piece for far-reaching spectrum reform is a meta-infrastructure with the capacity to incorporate virtually all forms of dynamic access, and both the property and commons allocation mechanisms. The criteria for this notional spectrum meta-infrastructure bear a striking parallel to an existing resource: the Internet’s domain name system.

B. Lessons from Internet Addressing

²⁹ See Radio Revolution, *supra* note 26.

³⁰ See John Chapin & William Lehr, *The Path to Market Success for Dynamic Spectrum Access Technologies*, IEEE COMMS, May 2007, at 96; Qing Zhao & Brian Sadler, *A Survey of Dynamic Spectrum Access*, IEEE SIGNAL PROCESSING, May 2007, at 79; Marshall, *supra* note 27.

³¹ See Douglas Sicker, *The Technology of Dynamic Spectrum Access and its Challenges*, IEEE COMMS., JUNE 2007, AT 48. *IEEE Communications*, a leading technical journal in the field, devoted a special issue to the topic in 2007, and research has only intensified since then. See *id.*

³² See Christian Bourge, *New Tech Feeds Spectrum Debate*, WIRELESS NEWSFACTOR, June 30, 2003, at <http://www.wirelessnewsfactor.com/perl/story/21828.html>.

³³ See Jon Peha and Sooksan Panichpapiboon, *Real-Time Secondary Markets for Spectrum*, TELECOMMS. POL’Y, Aug.-Sept. 2004, at 603.

³⁴ John Markoff, *Google Proposes Innovation in Radio Spectrum Auction*, N.Y. Times, May 22, 2007.

³⁵ One exception is spectrum dedicated for radio astronomy, where any radiated energy interferes with scientific research.

Every computer network must have an addressing system. Without unique addresses, data could not flow from one point to another. On a single network, addresses can be assigned and managed centrally, because one authority controls all aspects of the system. The Internet, however, is a network of networks.³⁶ There is no central regulator. Addresses must be assigned by agreement of the network operators or users. Remarkably, the Internet developed an addressing structure that is robust, scalable, and flexible, despite these limitations. That infrastructure is the domain name system (DNS).

The original Internet addressing mechanism was called the host name table. In those days, prior to the mid-1980s, machines were directly connected to the network, generally at universities or other research and government institutions. Addresses were maintained in a single flat text file, maintained by a graduate student named Jon Postel.³⁷ As the Internet grew, this simple mechanism became unworkable. The DNS was developed in 1983 by Paul Mockapetris to replace it.³⁸

The technical requirements and applications of the DNS were spelled out in a series of protocol documents that were revised over time.³⁹ There are many aspects to the DNS, but at its core, the DNS is a special kind of database. Paul Vixie, a long-time technical expert on Internet addressing, has called it “a distributed, coherent, reliable, autonomous, hierarchical database, the first and only one of its kind.”⁴⁰ The DNS establishes a “domain name space” of hierarchical zones, each served by an authoritative nameserver. For example, the domain name “en.wikipedia.org” means the English-language sub-domain of the Wikipedia domain, which is part of the .org (organization) generic top-level category within the DNS. Information is cached and replicated both horizontally – across multiple “root servers” or parallel local nameservers – and vertically – from a higher-level server down to sub-domains.

Functionally, the DNS can be thought of as providing two interfaces: a “front-end” resolution service for end-users, and a “back-end” registration service for network-connected resources. For end-users, the DNS seamlessly connects

³⁶ Kevin Werbach, *The Centripetal Network: How the Internet Holds Itself Together, and the Forces Tearing it Apart*, __ U.C. DAVIS L. REV. __ (forthcoming 2009).

³⁷ Jay P. Kesan & Rajiv C. Shah, *Fool Us Once Shame on You - Fool Us Twice Shame on Us: What We Can Learn From the Privatizations of the Internet Backbone Network and the Domain Name System*, 79 WASH. U. L.Q. 89, at 169 (2001); Brett M. Frischmann, *Privatization and Commercialization of the Internet Infrastructure: Rethinking Market Intervention into Government and Government Intervention into the Market*, COLUM. SCI. & TECH. L. REV. (2001).

³⁸ See Kesan and Shah, *supra* note 37.

³⁹ Jon Postel, *RFC 1591: Domain Name System Structure and Delegation* (Mar. 1994), at <ftp://ftp.isi.edu/in-notes/rfc1591.txt>; ELLEN RONY & PETER RONY, *THE DOMAIN NAME HANDBOOK* (1998).

⁴⁰ Paul Vixie, *DNS Complexity*, ACM QUEUE, April 2007.

human-readable domain names such as *icanhascheezburger.com* with the numerical Internet Protocol (IP) addresses that the network's routers understand.⁴¹ For those who wish to be reached, the DNS provides a mechanism to request a name, verify its availability, and associate it with an individual or organization.

As the DNS evolved, three distinct components developed: registries, registrars, and resolvers. A registry is the authoritative database for a top-level segment of the DNS, such as *.com* or *.cz*. It maintains the records of which names map to which IP addresses, and the official contact points for that registration. Registrars interface directly with those who wish to register names. The registrar function was originally combined with that of the registry, but at least for the generic top-level domains such as *.com* and *.org*, there is now competition among authorized registrars.⁴² Finally, individual devices and service providers can operate local resolvers that match a user's request for an address with the entry specified in the registry. The resolver usually operates on a local cached copy of the registry at the users Internet service provider.

These concrete functions only partially capture the significance of the DNS. The unified addressing system is, on some level, what makes the Internet the Internet.⁴³ Without the ability to know that *en.wikipedia.org* is *the* English-language Wikipedia, there would be a collection of loosely connected private networks, rather than a single Internet.⁴⁴ And the DNS serves as foundational infrastructure for new Internet applications and features, because it is so universal.⁴⁵

A review of the functions and history of the DNS reveals three key elements that allowed the system to scale effectively: resolution separated from transmission, distributed redundancy, and governance separated from technology.

First, the DNS is a specialized piece of Internet infrastructure. Its only function is to resolve addresses. The protocols for encoding and transmitting data packets are completely separate, as are those defining particular applications

⁴¹ This served three purposes: people remember names better than numbers, the names can stay stable when the numbers (tied to network topology) change, and allowed a single host to correspond to multiple network addresses. See Jon Klensin, *Role of the Domain Name System (DNS)*, RFC 3467, Feb. 2003, at 3.

⁴² This was a significant outcome of the transition of domain name governance to ICANN. See *infra*.

⁴³ “[T]he DNS provided critical uniqueness for names, and universal accessibility to them, as part of overall “single Internet” and “end to end” models....” See *id.* at 5.

⁴⁴ See Centripetal Network, *supra* note 36; Christopher Rhoads, *Endangered Domain*, WALL ST. JOURNAL, Jan. 19, 2006, at A1 .

⁴⁵ See Klensin, *supra* note 41, at 2 (“In recent years, the DNS has become a database of convenience for the Internet....”).

such as the World Wide Web and voice communication.⁴⁶ This separation helps the Internet to grow and develop. Scaling the addressing system is a challenge for the addressing system, not for the operators or users of other components of the Internet.

Second, as Paul Vixie noted, the DNS is a distributed, hierarchical, redundant system.⁴⁷ There is one authoritative DNS database for each top-level domain, but that database exists in many places. Billions of DNS queries take place every day, but the overwhelming majority do not go all the way to the central root servers. They are handled locally by caches and resolvers.

Third, the DNS is agnostic to how names are used. The DNS is a technical creation to manage an operational requirement of a network of networks. However, domain names are at the center of important policy and economic activities, which become increasingly significant as the Internet grows. Domain names potentially intersect with the law of trademarks and defamation, for example.⁴⁸ And the structure of the DNS shapes the degree to which the Internet is truly international, or inherently tilted toward the US and Western countries.⁴⁹ ICANN, the governance body established to oversee certain DNS management functions, has waded into all these disputes. However, the technical architecture of the DNS has remained the same. If the DNS does not solve all the non-technical problems around Internet addressing, at least it permits the battles around them to be fought in non-technical domains.

II. FROM WHITE SPACES TO SPECTRUM DNS

The rudiments of a DNS for the wireless Internet are being developed as an outgrowth of the FCC's decision to authorize unlicensed White Space Devices around former broadcast television bands. To implement the White Spaces decision, the FCC required devices to query a real-time database of utilization in the relevant frequency bands. Properly designed, this system could be the basis for a distributed dynamic routing database, analogous to the DNS on the wired Internet.

⁴⁶ Applications can be assigned to port numbers, which are part of the larger Internet addressing system, if not the DNS itself. And the DNS has a specialized component built in for email, known as the MX record.

⁴⁷ See Vixie, *supra* note 40.

⁴⁸ See *generally* MILTON MUELLER, RULING THE ROOT: INTERNET GOVERNANCE AND THE TAMING OF CYBERSPACE (2002) (detailing the tortured history of domain name management); Kesan and Shah, *supra* note 37.

⁴⁹ See Centripetal Network, *supra* note 36; Geoff Huston, *Addressing the Future of the Internet*, ISP COLUMN, Feb. 2007, <http://www.potaroo.net/ispcol/2007-02/address-paper.html>;

To achieve such a result, however, the database must not be limited to White Space devices alone. The FCC and industry must also take care to avoid the mistakes and failings of the current DNS infrastructure. These include the imposition of artificial scarcities, the creation of a private monopolist, and the bureaucratization of technical management functions. The Internet may have been a happy accident, but there is no excuse today for ignoring the infrastructural demands of its next instantiation.

A. *The White Space Database*

The White Space Database presents a singular opportunity to apply the DNS model to wireless communication. When the FCC allocated spectrum for broadcast television in the 1940s, it deliberately left many channels dark to avoid interference.⁵⁰ For example, channel 4 was licensed in New York but not in Philadelphia, and vice versa, so that receivers in each city would not become confused by overlapping signals. Many other frequencies in the TV broadcast bands are un-used as interference protection. Moreover, every available channel is not licensed in each city, and many licensees do not actually transmit, especially with most viewers receiving television content through cable or satellite connections rather than over the air. A 2005 study by the public interest group Free Press confirmed that most TV broadcast channels are simply not in use.⁵¹

These so-called TV White Spaces are perhaps the most egregious example of the inefficiency of the current spectrum allocation regime.⁵² They are considered “beachfront spectrum” because of their location on the frequency chart. Generally, the lower the frequency, the better the propagation of a wireless signal. A lower-frequency system can serve a larger geographic area and penetrate obstructions such as trees and building walls more easily than a higher-frequency system, all things being equal. Lower frequencies tend to be dedicated to older services, and thus are less likely to be made available for new systems. The broadcast frequencies have the capability to support valuable wireless broadband services, and may be especially useful for delivering broadband connectivity in rural areas.⁵³

⁵⁰ See Meinrath & Calabrese, *supra* note 17; NUECHTERLEIN & WEISER, *supra* note 20.

⁵¹ See Free Press study, at http://www.freepress.net/docs/whitespace_analysis.pdf. See also McHenry, *supra* note 14.

⁵² The designation is something of a misnomer, as there is nothing inherently linking the subject frequencies to television broadcasting beyond historical legacy.

⁵³ See Crawford, *supra* note 11; Meinrath & Calabrese, *supra* note 17; Pierre de Vries, Populating the Vacant Channels: The Case for Allocating Unused Spectrum in the Digital TV Bands to Unlicensed Use for Broadband and Wireless Innovation (New Am. Found., Working Paper No. 14, 2006), available at <http://www.newamerica.net/files/WorkingPaper14.DTVWhiteSpace.deVries.pdf>; J.H. Snider,

In 2004, the FCC sought comment on whether to allow unlicensed devices to transmit in the TV White Spaces.⁵⁴ Advocates and technical experts weighed in on both sides regarding the potential for interference. Broadcasters expressed concern that unlicensed white space devices would harm TV reception. Manufacturers and users of wireless microphones which operate in these bands expressed similar objections.⁵⁵ The FCC concluded in 2006 that unlicensed white space devices could operate without producing excessive interference.⁵⁶ It initiated a testing process for prototype devices to verify this finding.⁵⁷ Companies such as Microsoft, Philips, and Motorola developed equipment. The testing concluded in October 2008 with a finding that white space devices could detect and avoid nearby television systems and wireless microphones.⁵⁸

In November 2008, the FCC allocated the TV White Spaces for unlicensed use.⁵⁹ However, important issues remain unresolved before the White Space order can go into effect. The FCC left significant technical and implementation issues to a further stage of the proceeding.⁶⁰ In March 2009, the incumbent TV broadcasters sued the FCC to overturn the decision.⁶¹ And during the same period, President Obama's nominee as FCC Chairman, Julius Genachowski, took over for the Bush Administration's Kevin Martin.

There is at this moment, therefore, a significant opportunity to shape the White Space proceeding. The FCC has expressed its intent to move forward, but the actual structure can still be shaped. In parallel with the regulatory decisions,

Reclaiming the Vast Wasteland: The Economic Case for Re-Allocating the Unused Spectrum (White Space) Between TV Channels 2 and 51 to Unlicensed Service, New Am. Found., Feb. 2006, available at <http://www.newamerica.net/files/archive/Doc File 2898 1.pdf>.

⁵⁴ Unlicensed Operation in the TV Broadcast Bands, ET Docket No. 04-186, FCC 04-113 (2004).

⁵⁵ There are a handful of other authorized uses of the bands, including auxiliary transmission services for broadcasters and cable television head-end systems.

⁵⁶ First Report and Order and Further Notice of Proposed Rule Making in the Matter of Unlicensed Operation in the TV Broadcast Bands, ET Docket No. 04-186 and 02-380, October 18, 2006.

⁵⁷ See *id.*

⁵⁸ FCC Office of Engineering and Technology "Evaluation of the Performance of Prototype TV-Band White Space Devices Phase II," OET Report FCC/OET 08-TR-1005, Oc. 15, 2008, available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-08-2243A3.pdf.

⁵⁹ Federal Communications Commission, Unlicensed Operation in the TV Broadcast Bands, Second Report and Order and Memorandum Opinion and Order, ET Docket No. 04-186, ET Docket No. 02-380, FCC 08-2360 (released November 14, 2008) [White Spaces order].

⁶⁰ See *id.*

⁶¹ See Matthew Lasar, *Broadcasters Sue FCC Over White Space Broadband Decision*, ARS TECHNICA, March 3, 2009, at <http://arstechnica.com/tech-policy/news/2009/03/broadcasters-sue-fcc-over-white-space-broadband-decision.ars>.

the 802.22 subgroup of the IEEE is developing technical standards for devices that could operate in the White Spaces.

The FCC's November 2008 order introduced a new requirement of a database to protect against interference from white space devices.⁶² TV White Space devices would be required to incorporate geo-location capabilities, meaning they would have the ability to determine their physical location either directly or through communication with a fixed base station.⁶³ Before transmitting, each device would have to check the White Space database to identify broadcasters and others operating in that area.⁶⁴ This way, the device could be sure it was "in the clear" to use a vacant channel.

The White Space Database emerged from the FCC proceeding as a pragmatic solution to a narrow concern, but has the potential to become something much greater. The database was a response to broadcasters' objections that regulators could not rely on the spectrum-sensing capabilities of White Space devices themselves to protect pre-existing users. However, the few lines in the FCC order concerning the database leave open a great deal of room to determine how it should be structured, implemented, and governed. The database will be a piece of software and networking infrastructure, analogous to the DNS. And there is no fundamental reason it must be limited to broadcast white spaces.⁶⁵

A coalition called the White Space Database Group has taken the lead in recommending the architecture for the system.⁶⁶ The group proposed a structure broadly similar to the DNS, with the possibility for distributed data storage and a split between registries and registrars (which it calls repositories and service providers), as well as the potential for dedicated resolvers (called query services).⁶⁷ In short, the White Space Database could be the foundation for a DNS of wireless spectrum.

B. Comparing Interference to Addressing

Spectrum management and Internet addressing are both mechanisms for conflict resolution. In wireless communication, the potential conflict is between

⁶² See White Spaces Order, *supra* note 59. The order also encouraged other mechanisms such as spectrum sensing to prevent interference.

⁶³ See *id.*

⁶⁴ See *id.*

⁶⁵ See Michael Calabrese, The End of Spectrum 'Scarcity': Building on the TV Bands Database to Access Unused Public Airwaves, New America Foundation Wireless Future Program Working Paper No. 25 (June 2009).

⁶⁶ See Ex Parte Filing of the White Spaces Database Group, in ET Docket No. 04-186, April 10, 2009. Membership in the group includes Comsearch, Dell, Fox, Google, Microsoft, Motorola, NCTA, sMSTV, NetLogix, Neustar and the Public Interest Spectrum Coalition.

⁶⁷ See *id.*

two systems that cannot effectively communicate at the same time, a phenomenon known as interference.⁶⁸ In addressing, the conflict is between two systems that claim the same domain name.⁶⁹ The DNS solves this conflict by authoritatively linking an Internet protocol address (which is itself tied to the physical topology of the network, and unique) with a domain name.⁷⁰ The equivalent process for spectrum is usually handled by the FCC, through its command-and-control allocation of frequencies to designated licensees.⁷¹

Although they serve the same basic function, the DNS and current spectrum management techniques do so in very different ways. Spectrum management is prophylactic and fixed. The government decides ahead of time how a set of frequencies can be used, and writes those into its licenses. The DNS, on the other hand, does not associate a name to an IP address until the prospective user of that name registers it. Registration only requires a check that the name is unassigned. And the registration can be changed or transferred at any time. This user-controlled system is therefore dynamic where traditional spectrum allocation is static and inefficient.⁷²

At first blush, there are substantial differences between the resolution function of domain names and spectrum management.⁷³ “Interference” with a domain name means using the exact same string of alphanumeric characters. Two domain names that vary by a single character may co-exist without difficulty. With wireless systems, however, interference is a complex and contentious concept.⁷⁴ Whether a system interferes is not a stable physical property; it depends on the system allegedly interfered with.⁷⁵ The physics of wireless

⁶⁸ As I explain in *Supercommons*, interference is something of a misnomer. The wireless signals do not prevent each other from being received. They make it difficult for one or both sets of devices to distinguish the desired signal from unrelated noise. See *Supercommons*, *supra* note 12.

⁶⁹ Addresses do not necessarily need to be globally unique to avoid this problem. DHCP and NATs allow a hierarchical structure where only a portion of the network is subject to DNS. See text at note 81.

⁷⁰ The mapping need not be one-to-one. Several domain names can be aliased to one IP address, or multiple IP addresses can be pointed to the same domain name.

⁷¹ See *Supercommons*, *supra* note 12; Benkler, *supra* note 21; Noam, *supra* note 21; FCC Spectrum Policy Task Force, *supra* note 16. In unlicensed bands, the FCC establishes technical requirements for devices, which then manage interference locally. See Eli Noam, *Yesterday's Heresy, Today's Orthodoxy, Tomorrow's Anachronism: Taking the Next Step to Open Spectrum Access*, 41 J. L. & ECON., 765 (1998).

⁷² See *Supercommons*, *supra* note 12.

⁷³ Jim Speta first encouraged me to consider this point.

⁷⁴ See *Supercommons*, *supra* note 12; Yochai Benkler, *Some Economics of Wireless Communications*, 16 Harv. J.L. & Tech. 25 (2002); David Weinberger, *The Myth of Interference*, SALON, Mar. 12, 2003, at <http://dir.salon.com/story/tech/feature/2003/03/12/spectrum/index.html>.

⁷⁵ See *Supercommons*, *supra* note 12

communication produce many spillover effects between systems.⁷⁶ This is further complicated by the nature of exclusive licenses. Incumbent licensees have every incentive to make their devices cheap rather than robust to interference, and to use the regulatory process to claim interference is occurring.⁷⁷ The battles over low-power FM radio and White Spaces show how intense these conflicts can become.⁷⁸ A cursory duplication check and a simple first-in-time allocation rule, as the DNS uses, therefore do not suffice for wireless.

On further examination, however, the lines are not so clean in the DNS case either. Because domain names are human-readable strings, there are situations where a user may confuse them even if a computer does not. This becomes important when a particular domain name has economic value associated with it. This is widely the case with commercial websites and those overlapping trademarks.⁷⁹ Under both the Anti-Cybersquatting Consumer Protection Act in US law and the Uniform Dispute Resolution Process of ICANN, a registrant may be forbidden from using a domain name even if it is not identical to one already registered or trademarked.⁸⁰ So-called “cybersquatting” that is malicious or creates consumer confusion is prohibited. As the number of lawsuits in this area demonstrates, deciding when that has occurred is complicated and contentious. The analogy to spectrum interference debates is therefore not so far-fetched.

At a technical level, the closest analogy between dynamic spectrum brokering and Internet addressing is DHCP, the Dynamic Host Configuration Protocol.⁸¹ DHCP is one of many “hacks” to adapt the Internet architecture developed for small-scale research internetworking to the requirements of today’s commercial-scale global infrastructure.⁸² Because of the technical characteristics of the Internet Protocol (IP) and historical address-allocation policies, most Internet service providers cannot easily give each of their customers a unique IP address.⁸³ Mobile access and nomadic connections to WiFi hotspots complicate the problem even further. The solution is to assign users an address dynamically from a

⁷⁶ See *id.*

⁷⁷ See Supercommons, *supra* note 12.

⁷⁸ See Stuart Minor Benjamin, THE LOGIC OF SCARCITY: IDLE SPECTRUM AS A FIRST AMENDMENT VIOLATION, 52 DUKE L.J. 1 (200) (lower power FM); *supra* note 53 (white spaces).

⁷⁹ See MUELLER, *supra* note 48.

⁸⁰ See Mueller, *supra* note 48; Anticybersquatting Consumer Protection Act, 15 U.S.C. 1125(d); Patrick Kelley, Note, *Emerging Patterns in Arbitration Under the Uniform Domain-Name Dispute-Resolution Policy*, 17 BERKELEY TECH. L.J. 181.

⁸¹ A. Michael Froomkin and Mark A. Lemley, *ICANN and Antitrust*, 2003 U. Ill. L. Rev. 1, n.19 (2003). TK -- analogy to DHCP in p. 4 of one of the technical papers below

⁸² Marjory S. Blumenthal & David D. Clark, *Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World*, in COMMUNICATIONS POLICY IN TRANSITION: THE INTERNET AND BEYOND (Benjamin M. Compaine & Shane Greenstein eds., 2001); *Upgrading the Internet*, ECONOMIST, Mar. 24, 2001, at 33, 34.

⁸³ See Centripetal Network, *supra* note 36.

common pool, each time they log on. Even if they have a broadband connection, most residential users do not remain connected all the time. DHCP makes dynamic IP address updates transparent to users.

With DHCP, a relatively small pool of IP addresses can be shared by a relatively large pool of users. Each user ties up an address only when he or she actually needs it. This efficiency gain parallels the greater efficiency that a more dynamic structure could bring to spectrum allocation. The problem with DHCP, and the related technique of Network Address Translation (NAT), is that they, as already noted, they were post-hoc hacks rather than core Internet infrastructure like the DNS. Because DHCP allocates addresses locally, there is no uniform mechanism to associate users and addresses across the network. This adds complexity to the Internet, and creates problems for some applications that would benefit from end-to-end visibility.⁸⁴ A DNS for wireless communication would marry the efficiency gains of DHCP with the canonical reliability of DNS.

As discussed below, conflict resolution is in truth only an intermediate function of both the DNS and spectrum allocation.⁸⁵ Both the Internet address space and the spectrum are notional constructs only.⁸⁶ Conflicts over names and frequencies are important only to the degree they have economic or social consequences. A system that minimized conflicts and controversy, but too heavily constrained productive use of the resource in question, would be a poor tradeoff. The ultimate goal of the DNS, as described above, was to scale the Internet.⁸⁷ The ultimate goal of spectrum management should likewise be to scale wireless capacity.⁸⁸

The DNS, for its many flaws, has allowed the Internet to grow from a purely research-oriented system linking a few thousand networks in 1983 to the global backbone of commerce, entertainment, communication, and government activity in 2009. To achieve something similar in spectrum requires an infrastructure for dynamic and distributed intermediation of spectrum allocation decisions, just as the DNS provides dynamic and distributed intermediation of address allocation decisions.

⁸⁴ See Blumenthal & Clark, *supra* note 82.

⁸⁵ John Klensin, Role of the Domain Name System (DNS), RFC 3467, February 2003.

⁸⁶ That the DNS is an arbitrary system created by Internet protocol definitions should be obvious. The illusory character of the spectrum is less intuitive, but equally accurate. See Supercommons, *supra* note 12.

⁸⁷ See *supra*.

⁸⁸ See Supercommons, *supra* note 12; Benkler, *supra* note 74.

III. THE SPECTRUM NETWORKING DATABASE

A. How the System Would Work

As has already been noted, a great deal of technical work is underway to facilitate real-time sharing of wireless capacity. The next step is to envision an intermediary that would broker requests for secondary access. A number of experts have incorporated some form of brokering into their technical protocols for dynamic spectrum access (DSA).⁸⁹ Viewing the problem from the other direction, though, the brokering mechanism should be the primary infrastructure for spectrum allocation, with particular technical schemes for sharing as subsidiary elements.⁹⁰ Putting the brokering engine at the center emphasizes the structural and policy elements to be addressed. I call the universal brokering engine the SND, or spectrum networking database.

The foundation for the SND is the reference architecture the White Spaces Database Group proposed to the FCC.⁹¹ This architecture involves three main elements: the repository, the registrar, and the query service.⁹² The repository would be the actual database of frequencies, locations, and authorized users. It is analogous to the DNS “root” database of domain names.⁹³ The repository would be fed from existing FCC spectrum allocation databases, such as the Universal Licensing Service (ULS), as well as from registrars. Registrars would take in information from spectrum users. One kind of users would be protected entities,

⁸⁹ See, e.g., M. Buddhikot, et al, DIMSUMnet: new directions in wireless networking using coordinated dynamic spectrum, World of Wireless Mobile and Multimedia Networks, 2005. WoWMoM 2005. Sixth IEEE International Symposium on a, 13-16 June 2005 Page(s):78 – 85; T. Maseng & T. Ulversoy, Dynamic Frequency Broker and Cognitive Radio, Cognitive Radio and Software Defined Radios: Technologies and Techniques, 2008 IET Seminar on, 18-18 Sept. 2008 Page(s):1 – 5; DSAP: a protocol for coordinated spectrum access, New Frontiers in Dynamic Spectrum Access Networks, 2005 First IEEE Dyspan International Symposium on 8-11 Nov. 2005 Page(s):611 – 614; Q. Peng et al, “A Distributed Spectrum Sensing Scheme Based on Credibility and Evidence Theory in Cognitive Radio Context”, 17th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, 2006, at 1-5.

⁹⁰ In a prior article, I envisioned a model for spectrum starting with a universal open entry privilege, cabined by tort-like and trademark-like liability rules. See Supercommons, *supra* note 12. The system described here could fit within that still-broader framework.

⁹¹ See White Spaces Database Group, *supra* note 66.

⁹² Some or all of these could be combined in a single company. The White Spaces Database Group also refers to device registrars and stand-alone query services as “service providers.”

⁹³ Like the DNS root, the repository would likely be replicated across many servers for redundancy and geographical dispersal. Unlike DNS, which has a single registry operator for each top-level domain, there might be more than one repository operator at the top level for the spectrum database. Parallel registry operators would have to mirror their contents regularly to assure consistency.

such as the incumbent broadcasters and wireless microphone operators. Another would be white space devices seeking secondary access to the spectrum. The registrar function parallels the competing registrars that users pay to register and maintain their domain names. Finally, the query service would check a device request against the available capacity in the database.⁹⁴ This is analogous to the resolver software that Internet service providers and personal computers use to determine the proper path to a domain name.

The SND would adapt architectural elements of the DNS to spectrum. At the core, the database would map blocks of spectrum to authorized users of that spectrum, just as the DNS maps domain names to servers on the Internet. The database would be granular as to frequencies and location, at a minimum. Before a device could operate, it would query the SND to identify available transmission opportunities. The database would return one of two results: the requested spectrum allows no authorized transmitters, or information about the incumbent networks and devices authorized to operate there. If the requested spectrum were empty, the device could register itself as operating in that band and location, and begin transmitting. If the requested spectrum were occupied, the device could not operate unless it either met criteria for low-power or ultra-wideband underlay operation consistent with the incumbent users, or it negotiated a license or lease from the incumbent.

Like the DNS, the SND structure could be hierarchical. This would avoid bottlenecks and facilitate greater local variation. At the top would be redundant root servers, parallel to the multiple root servers for the DNS. Local SND servers would query these root servers regularly and cache the results. Each SND server would therefore have an updated “map” of the spectrum in its geographical reach.

Both the DNS and the White Space Database Group proposal separate the functions of registries (which store canonical data) and registrars (which interface with end-users and insert data into the registry). This approach allows for competition and provides a check on the excessive power of a monopoly registry. However, it is not without its own dangers. For domain names, registrar competition is largely circumscribed to price. There is little incentive for new or better service, so registrars engage in a “race to the bottom” to offer the cheapest registrations. A high “wholesale” price for the .com domain registry exacerbated this error, giving registrars little room to maneuver. Spectrum registrars should have flexibility to innovate, so long as they meet the basic technical and business requirements to participate in the process.

For the system to operate, all dynamic devices would need unique identifiers, so they could be tracked and managed through the database.⁹⁵ A certification process would need to be established to ensure that devices operate as specified,

⁹⁴ This service might also be integrated into the registrar function.

⁹⁵ This would be one point of distinction between the new adaptive wireless devices and the legacy static devices entitled to protection.

including responding to shut-off commands from the database. That process could be run directly by the FCC or delegated to private bodies, as with most communications equipment certification today.

Devices would need the ability to communicate with the database, either through an existing Internet connection or via a dedicated wireless control channel. A wireless “control plane” for administrative communications would parallel the “out of band” architecture of the SS7 signaling network on the public switched telephone network (PSTN).⁹⁶ On the Internet, signaling information uses the same physical channel as the information payload, but utilizes specialized protocols such as Border Gateway Protocol (BGP) for connecting networks and the DNS protocol for domain name resolution queries.⁹⁷ The SND would need a separate wireless channel for administrative communication, because the entire purpose of the system is to establish a physical wireless connection in a particular spectrum block. This channel could utilize existing unlicensed frequencies, which are already available throughout the United States, and in some cases globally harmonized.⁹⁸

Protected incumbent allocations in the SND could be classified in several ways. Some existing uses would preclude any potentially interfering transmissions. Military systems, radio astronomy bands, avionics, medical devices, and heavily utilized commercial systems such as cellular telephone service could be designated for maximum protection.⁹⁹ Other bands might be subject to “easements” for low-power use, without permission from existing users.¹⁰⁰ Others could be available on an unlicensed basis when not in use, as with the TV White Spaces under the FCC order.¹⁰¹ Still others might be available subject to negotiation with the incumbent. The SND could manage interference because devices would need to query it prior to initiating transmissions. If a device were determined to be operating improperly, the system would issue a “no channel available” message, preventing it from transmitting.

⁹⁶ See Kevin Werbach, *The Internet's SS7 Network*, RELEASE 1.0, Dec. 1999, available at cdn.oreilly.com/radar/r1/12-99.pdf.

⁹⁷ Kirk Lougheed, A Border Gateway Protocol (BGP) (1981), <http://www.ietf.org/rfc/rfc1105.txt>.

⁹⁸ Revision of Parts 2 and 15 of the Commission's Rules to Permit Unlicensed National Information Infrastructure (U-NII) Devices in the 5 GHz Band, ET Docket No. 03-122 (rel. June 4, 2003).

⁹⁹ It is worth pointing out, however, that even these systems, with the exception of radio astronomy, could at some point become more flexible and dynamic. Public sector spectrum users could participate in the SND infrastructure on a case-by-case basis. For radio astronomy, though, there is no way to negotiate with a quasar, so a strict zone of inclusion is needed.

¹⁰⁰ See Farber and Faulhaber, *supra* note 21.

¹⁰¹ See White Spaces Order, *supra* note 59.

Technical standards for devices would need to be established. As noted previously, the IEEE 802.22 group is developing standards, but they envision a more limited White Space database compared to the broader SND vision sketched out here. The standards adopted would significantly influence the scope and flexibility of the SND. At an extreme, devices would have full flexibility in selecting spectrum bands, access models, and transmission modalities. The reality is that such flexible radios are not yet feasible to build, let alone affordable to sell, despite rapidly advancing research. The functionality of devices will involve cost/performance tradeoffs. The architects of the SND will have to determine how and where to restrict those tradeoffs. For example, how often must devices query the database, and should baseline etiquettes be mandated for DSA devices to co-exist with each other?

Furthermore, devices might operate on an ad hoc basis, handling all the spectrum access negotiation themselves, or they might be tied to fixed base stations. The value of a fixed base station is that a mobile device need only find and establish contact with the base station, which can have a dedicated link to the SND. The base station can ensure that the mobile device complies with the constraints of its authorization. The mobile device would not be able to operate without confirmation from a base station, so if it were moved outside the specified transmission area, it would cease to transmit without authorization from another base station. The architecture in this case would parallel that of cellular telephone networks. Mobile phone handsets operate by locating a transmission tower, which authorizes them to communicate, assigns them temporary frequency slots, and handles “backhaul” of the communication onto the larger network backbone.

B. Benefits

The power of this system lies in its flexibility. The same basic architecture could be applied in one set of frequencies (such as the former broadcast television bands) in a limited geographic area, or it could encompass the entire usable spectrum nationwide. Moreover, the structure is agnostic between the two competing allocation mechanisms for more efficient spectrum use: property and commons. If a frequency band is designated for unlicensed use, the SND would provide the technical restrictions such as power limits on that band, and authorize transmission. If the frequency band is licensed, the SND would allow the licensee to specify terms for secondary access. These could include anything from open entry so long as the licensee has priority and override capability, requirement of a reservation or “lease” for the system seeking access,¹⁰² or real-time access payments using either a fixed or auctioned price.¹⁰³

¹⁰² See Peha and Panichpapiboon, *supra* note 33; Maseng & Ulversoy at p. 3. Note that this could be the “unlicensed park” model suggested by spectrum property advocates, where an

The SND would allow for a new category of spectrum allocation that is neither property nor commons. In this scenario, rather than designating frequencies as either licensed or unlicensed, as it does today, the FCC could establish a set of initial access algorithms for the SND. These algorithms could vary depending on the nature of the spectrum involved and the use cases. For some very high “millimeter wave” frequencies, for example, the FCC has already established a reservation system, through which users can sign up to operate point-to-point links on a first-come, first-served basis.¹⁰⁴ In other cases, the FCC might designate a primary user (such as a wide-area high-power system) and a category of secondary or unlicensed users with rights to operate below certain power thresholds or when the primary user was not. The approach could be used for so-called ATC (ancillary terrestrial component), in which frequencies designated for satellite uplink are re-used for terrestrial service through devices able to determine the angle of signal arrival.¹⁰⁵ And it could be used for sharing spectrum with radar systems that rotate or operate only during circumscribed time periods.¹⁰⁶

As one concrete example, the FCC recently attempted to promote the deployment of a nationwide interoperable public safety wireless network using the so-called D Block of frequencies in the 700 MHz bands being vacated by analog UHF television channels.¹⁰⁷ The Commission sought to auction the block subject to the condition that the winner construct a network capable of both commercial service and public safety use, with the public safety users having override capability in times of need.¹⁰⁸ The funds from the auction were to be used to purchase equipment for the public safety users.¹⁰⁹ The auction in 2008 failed when no bidder offered the reserve price.¹¹⁰ The FCC is now evaluating

equipment vendor buys the spectrum and charges for devices. The SND establishes the foundational infrastructure to make this approach more feasible.

¹⁰³ See *ex parte* letter from Richard S. Whitt, Google, to Marlene H. Dortch, FCC, May 21, 2007.

¹⁰⁴ See Allocations and Service Rules for the 71-76 GHz, 81-86 GHz and 92-95 GHz Bands, 18 F.C.C.R. 23,318, 23,339-341 (2003); Phil Weiser and Dale Hatfield, *Policing the Spectrum Commons*, 74 *FORDHAM L. REV.* 101 (2005).

¹⁰⁵ Report and Order and Notice of Proposed Rulemaking, FCC 03-15, “Flexibility for Delivery of Communications by Mobile Satellite Service Providers in the 2 GHz Band, the L-Band, and the 1.6/2.4 Bands, IB,” adopted Jan. 29, 2003; released Feb. 10, 2003; Supercommons, *supra* note 12.

¹⁰⁶ Michael J. Marcus, *New Approaches to Private Sector Sharing of Federal Government Spectrum*, New America Foundation Wireless Futures Project Issue Brief #26 (June 2009).

¹⁰⁷ See Crawford, *supra* note 11.

¹⁰⁸ See *id.*; Auction of 700 MHz Band Licenses Scheduled for January 16, 2008, Public Notice, 22 F.C.C.R. 15,004 P 2 (Aug. 17, 2007).

¹⁰⁹ See Crawford, *supra* note 11.

¹¹⁰ See *id.*; Alejandro Valencia, *The FCC's Regulatory Mulligan: Exploring the Options in the Wake of a Failed D Block Auction*, 10 *N.C. J.L. & TECH.* 313 (2009).

other options, both for the 700 MHz D Block and for a nationwide interoperable public safety system.

Under an SND model, the public safety goal could be achieved more readily. The FCC would designate frequencies in which public safety systems had a guaranteed override capability. The simplest option would be then to allocate the system for auction to a private operator, who would get the spectrum subject to that limitation. Or the spectrum could be made available as unlicensed, subject to the same condition. The public safety override would then be incorporated into the technical rules for unlicensed devices approved by the FCC to operate in the band, and to the relevant technical standards. As a variant, a royalty on each compliant device could be used to fund the purchase of public safety equipment, although that would require additional questions about the best funding mechanism for that goal. In other words, the SND itself does not solve the public safety problem, but it creates a much larger toolkit for policy-makers to employ.

A collateral benefit of this database-driven approach would be future-proofing. One of the great challenges of wireless communication is that as technology evolves and usage patterns change, bands may become more or less heavily used. Incumbent systems, however, can be expensive to clear out of bands, even when they are not very active. Clearing systems from the bands given to 2G and 3G cellular systems was extremely expensive, and in some cases the installed base of devices makes reallocation effectively infeasible. With an SND architecture, the FCC could change the allocation algorithm when conditions changed significantly. Because all devices would be required to operate under the constraints of the SND, migration and adaptation would be automatic.¹¹¹

Compared to a pure property rights regime, the SND would be more effective because the SND functioned as a market-making intermediary. Relying on individual spectrum owners to negotiate with potential entrants falls victim to the tragedy of the anticommons.¹¹² There are many externalities and complications of interference and wireless utilization that property owners are ill-suited to take into account.¹¹³ A pure property system for spectrum is like the old host table for Internet addresses. It becomes a bottleneck because it has no real-time global visibility.

The SND would serve the same purpose as the New York Stock Exchange (NYSE): creating the trust and liquidity necessary for a well-functioning market. The NYSE assigns unique ticker symbols to each stock, but this is a relatively trivial aspect of its activity. The more important functions of the NYSE include

¹¹¹ There would still be costs, and some old devices would stop working and have to be replaced. But the process would be far smoother and more efficient than it is today.

¹¹² See MICHAEL HELLER, *THE GRIDLOCK ECONOMY* (Basic Books 2008).

¹¹³ See *Superccommons*, *supra* note 12.

servicing as the locus for a set of rules, some private and some government-imposed, that create well-defined rights as the basis for transactions.

Collective mechanisms such as exchanges and government regulation become necessary when the efficient distribution of property rights is too difficult to achieve through bilateral transactions. As Thomas Merrill and Henry Smith have noted, property functions both *in rem* (as a thing against the world) and *in personam* (as a bundle of rights attaching to an owner).¹¹⁴ Coase advocated property rights in spectrum by emphasizing the *in personam* rights in wireless equipment, rather than the notional *in rem* rights in spectrum frequencies.¹¹⁵ The problem, as Merrill and Smith observe, is that the complexity and uncertainty of spectrum use rights makes private transactional regimes unworkable or far too limiting.¹¹⁶ The traditional, inefficient solution is command-and-control regulation. The SND provides a multi-lateral mechanism to incorporate new spectrum allocation techniques based on commons and property rights.

C. A Note About Governance

The final and perhaps greatest challenge for implementation of an SND is the governance structure. Domain names are the technical hinge around which the Internet rotates. It has long been clear that control over domain name allocation was perhaps the closest proxy to a central governance function for the resolutely distributed Internet. The US government, originally through the National Science Foundation, and later through the Department of Commerce and White House staff, exerts limited but explicit oversight over the DNS. Operational decisions, generally reside with ICANN, a quasi-public, quasi-international body that has been a locus of controversy since its inception.¹¹⁷

¹¹⁴ See Thomas W. Merrill & Henry E. Smith, *What Happened to Property in Law and Economics?*, 111 YALE L.J. 357 (2001). Merrill and Smith assert that the *in rem* character of property has been excessively submerged in contemporary law and economic scholarship. See *id.*

¹¹⁵ See *id.* at 371; Coase, *supra* note 23. See also Supercommons, *supra* note 12 (expanding on Coase's point).

¹¹⁶ See Merrill and Smith, *supra* note 114, at 374. See also Supercommons, *supra* note 12 (elaborating on the difficulty of applying a property regime to spectrum, due to the transactional complexities). Merrill and Smith would prefer an *in rem* emphasis on rights in spectrum itself. This, however, is even more problematic, because spectrum is simply not a "thing" in the sense of physical property. See *id.*

¹¹⁷ A. Michael Froomkin, *Wrong Turn in Cyberspace: Using ICANN to Route Around the APA and the Constitution*, 50 DUKE L.J. 17 (2000); Kesan and Shah, *supra* note 37; Jonathan Weinberg, *ICANN and the Problem of Legitimacy*, 50 DUKE L.J. 187, 209-12 (2000).

This is not the place to catalog ICANN's structural flaws and missteps.¹¹⁸ Suffice it to say that governance for the DNS has not been a smooth process. It is unlikely to be a simple one for spectrum. That being said, many of ICANN's challenges do not apply to the SND. ICANN attempts the impossible challenges of being both public and private, as well as being both global and subject to national laws. The SND would be squarely subject to the national regulatory authority of the FCC, but it would delegate many real-time decisions to private actors. ICANN is nominally a technical coordination body, which finds itself engaged in significant policy debates, without sufficient procedural and substantive protections.¹¹⁹ The SND is an implementation mechanism for policy decisions made and enforced elsewhere, just like the DNS.

CONCLUSION

The SND vision adapts the three architectural attributes of the DNS to spectrum.

First, the SND would allow a wireless device to ask the question, "Am I permitted to transmit here, and if so, how?" Today, the answer is hardwired into the device. Even an unlicensed device such as a WiFi node is strictly limited in the frequencies, power levels, and protocols it can employ, regardless of its actual local environment. Resolution – the process of matching the physical device with the virtual communications space – is paired with the function of transmission. The DNS breaks this linkage for Internet traffic. Special-purpose infrastructure handles the process of resolving the location of a network address. Sending and receiving information is a separate function, which can employ a wide variety of technical mechanisms for efficient transmission. The SND is similar because it establishes a special-purpose element of wireless communications infrastructure that is distinct from the transmission process itself.

Second, the SND would be a fundamentally distributed system. Wireless devices would query local copies of the database, just as Internet service providers query local domain name caches. Consistency of data would need to be assured, but as the DNS demonstrates, that is not incompatible with multiple redundant copies or competing service providers. For the White Space database, the FCC has not yet specified whether there can be multiple registrars or repositories. There is little disagreement on the value of competition between registrars, and allowing for multiple repositories could also have salutary effects.

¹¹⁸ See Centripetal Network, *supra* note 36.

¹¹⁹ See Froomkin, *supra* note 117.

The DNS has suffered because a single entity, NSI, was given control over the registry function for the .com top-level domain.¹²⁰

Third, the SND can be agnostic as to uses. As described above, virtually any policy regime and type of application can be mapped into the database, so long as it is designed with sufficient flexibility and scalability to handle them. Disputes over the boundaries of spectrum rights will not disappear, but they will be separated from disputes over the basic allocation of spectrum.

Creating a DNS for spectrum would generate tremendous benefits. Both wireless communication and Internet addressing are activities requiring conflict resolution. The conflicts in wireless communication are between two or more users who wish to communicate, and the conflicts in Internet addressing are between two or more users who want the same identifier. The ultimate goal, however, is not resolve the conflicts, but to promote communication and productive activity. What mitigates interference going forward is deployment of smarter devices that can use spectrum in more efficient ways – through both property rights transactions and commons.

The DNS removed the bottleneck of the host table system, and allowed for massive growth in the still-decentralized Internet. If the SND could have a small fraction of that success in the spectrum domain, it would be a tremendous boon for economic activity, innovation, and open communication.

¹²⁰ See MUELLER, *supra* note 48; Weinberg, *supra* note 117.