

Massachusetts Institute of Technology
Engineering Systems Division

Working Paper Series

ESD-WP-2006-01

MANAGING SHARED ACCESS
TO A SPECTRUM COMMONS

William Lehr¹, Jon Crowcroft²

¹Massachusetts Institute of Technology

²Cambridge University

February 2006

Managing Shared Access to a Spectrum Commons

William Lehr
Massachusetts Institute of Technology

Jon Crowcroft
Cambridge University

** draft **

Abstract- The open access, unlicensed or spectrum commons approach to managing shared access to RF spectrum offers many attractive benefits, especially when implemented in conjunction with and as a complement to a regime of market-based, flexible use, tradable licensed spectrum ([Benkler02], [Lehr04], [Werbach03]). However, as a number of critics have pointed out, implementing the unlicensed model poses difficult challenges that have not been well-addressed yet by commons advocates ([Benjam03], [Faulhab05], [Goodman04], [Hazlett01]). A successful spectrum commons will not be unregulated, but it also need not be command & control by another name. This paper seeks to address some of the implementation challenges associated with managing a spectrum commons. We focus on the minimal set of features that we believe a suitable management protocol, etiquette, or framework for a spectrum commons will need to incorporate. This includes: (1) No transmit only devices; (2) Power restrictions; (3) Common channel signaling; (4) Mechanism for handling congestion and allocating resources among users/uses in times of congestion; (5) Mechanism to support enforcement (e.g., established procedures to verify protocol is in conformance); (6) Mechanism to support reversibility of policy; and (7) Protection for privacy and security. We explain why each is necessary, examine their implications for current policy, and suggest ways in which they might be implemented. We present a framework that suggests a set of design principles for the protocols that will govern a successful commons management regime. Our design rules lead us to conclude that the appropriate Protocols for a Commons will need to be more liquid ([Reed05]) than in the past: (1) Market-based instead of C&C; (2) Decentralized/distributed; and, (3) Adaptive and flexible (Anonymous, distributed, decentralized, and locally responsive).

1 INTRODUCTION¹

In recent years, policymakers, academics, and industry participants have been debating appropriate models for reforming spectrum management regimes to be more flexible and responsive to market forces. The debate has focused on whether spectrum should be managed via a *property rights/licensed/exclusive use* model or as *open*

access/unlicensed/spectrum commons.² Proponents of the latter model have not adequately explored the mechanisms that might be used to manage a spectrum commons. This paper attempts to address this gap in the literature by considering the range of multi-disciplinary issues (regulatory/political, business/economic, and technical) that would be involved in managing a spectrum commons (*i.e.*, non-exclusive, shared-use spectrum).

One approach to managing unlicensed spectrum is to rely on a technical protocol to allocate and manage shared access. Network engineers have designed an array of protocols that enable shared access in various wired and wireless contexts at different layers of the protocol stack.³ These offer a variety of trade-offs and reflect the perspectives of engineers seeking to solve a diverse set of problems. However, the evaluation of alternative protocols has taken place largely in technical forums wherein the relative merits are evaluated according to engineering performance criteria. While this perspective is obviously important, it is not sufficient within the context of the spectrum reform debate. An appropriate unlicensed management framework will need to address non-technical policy and business implementation issues as well.

This paper addresses this need by taking a more holistic view of the technical and institutional features of an effective management framework for a spectrum commons. We focus on shared access to non-exclusive use spectrum because we believe this will play an increasingly important role in future wireless services. While this may be implemented as a public commons, we do not expect the "commons" to be an unregulated free-for-all. The framework established to manage the "commons" will impose restrictions on what devices, uses, or users are permitted (*i.e.*, only those that are in compliance with the rules). This perspective encompasses such frameworks as the current Part 15 devices that operate in the ISM band (e.g., WiFi) and (semi-private or private) restricted

¹ Corresponding author: William Lehr (email: wlehr@mit.edu)

² See, for example, [Benkler02], [FaulFarb02], [Hazlett01], [Kolodzy02], [Lehr04], [Ofcom05], [Reed02], or [Werbach03].

³ See, for example, [Peha97], [RoyerToh99].

"commons" that may be established in flexible use licensed spectrum.

The goal of spectrum reform is to enhance the allocation and use of spectrum so as to foster innovation, competition, and the efficient use of spectrum. To be consistent with these goals, we argue that an appropriate framework for managing unlicensed spectrum ought to aspire to be as minimally constraining as possible and still be consistent with orderly management of the shared access spectrum. We believe better rules are needed for the operation of unlicensed devices that would co-exist with licensed devices via both secondary use easements (underlays or overlays) or as primary uses in dedicated unlicensed spectrum. The design of these rules/frameworks needs to go beyond the theoretically pure world of polar extremes (property rights v. commons) to address the real-world environment of mixed regimes (both models are in use but the bulk of spectrum is allocated via flexible use licenses).

Our "world view" is futuristic in that we speculate about an improved (but not ideal) framework for managing unlicensed that does not yet exist, and we largely ignore the all-important political/business issues associated with moving from today's messy status quo (e.g., resolving NIMBY issues). In this world, there will be more wireless services and products playing an even larger role in our lives and global economy. These will be based on a wider-array of heterogeneous technologies: legacy and new, centralized and decentralized, closed and open architectures, congestion tolerant and intolerant will all co-exist in the market. Radio systems will be smarter (i.e., able to share spectrum dynamically), but will fall short of the cognitive radio ideal. Importantly, an enduring feature of this world will be on-going fundamental uncertainty regarding what the high value services will be or the best business models and technologies to deliver those services. There will continue to be a need to accommodate substantial innovation and growth. In this world, there will be an increased role for market forces, but there will still be a centralized regulator. There will be reduced spectrum scarcity, but congestion will still occur.

In this worldview, rules will be needed, but these can be more market-responsive than the traditional command & control framework of legacy spectrum management. The goal of this paper is to discuss the key features of the rules that are likely to be needed to manage a spectrum commons effectively. An appropriate framework will promote innovation and minimize regulatory distortions. To accomplish the latter goal, the regulatory rules should strive for technical and market neutrality, while still promoting competition where possible. Finally, a good framework must be capable of evolution. This includes being able to support marginal adjustments between licensed and unlicensed, and within unlicensed, to support changing protocols without being biased in favor of incumbents or entrants.

Furthermore, we explain why an appropriate framework ought to favor distributed/decentralized management to the extent feasible in order to realize the maximal benefits of a "commons." Finally, we argue that the evaluation of management mechanisms for a spectrum commons ought to be consistent with the sort of environment and the requirements of those sorts of services that are most likely to find operation in a spectrum commons acceptable in a post-reform world (i.e., one which includes command and control, flexible-licensed, *and* commons spectrum). In this world, the commons spectrum will be most attractive to applications which are adaptive, reasonably tolerant of congestion, but intolerant of the higher spectrum opportunity costs associated with using licensed spectrum.

Having established the above, we posit that a suitably effective yet minimalist framework will require the following types of rules: (1) No transmit only devices; (2) Power restrictions; (3) Common channel signaling; (4) Mechanism for handling congestion and allocating resources among users/uses in times of congestion; (5) Mechanism to support enforcement (e.g., established procedures to verify protocol is in conformance); (6) Mechanism to support reversibility of policy; and (7) Will be consistent with security and privacy.

The balance of this paper is organized into three sections. In Section II, we define our terms and establish the context within which we expect a spectrum commons to operate. Section III then presents our taxonomy for evaluating unlicensed management frameworks. Section IV concludes with an analysis of certain key issues and suggestions for further research.

2 REGULATORY MODELS AND SHARED ACCESS SPECTRUM COMMONS

2.1 Defining terms

2.1.1 C&C, Licensed, and Commons

The debate over spectrum management regimes has been complicated by an inconsistent use of terms. Earlier discussion often focused on the relative merits of extreme theoretical versions of the various regimes, for example, comparing pure property rights (fee simple licenses) to pure unlicensed (no regulation, free and open access). Critiques of unlicensed often failed to recognize that many of the benefits ascribed to a property rights regime are also realizable in an unlicensed regime ([Lehr04]) and visa versa ([Faulhab05], [Goodman04]).

For this paper, there are three models for spectrum management that are relevant. The first, corresponding to today's *status quo* may be stylized as *Command & Control* ("C&C"). In this regime, market forces have only limited impact. A regulatory agency such as the FCC in the US or OfCom in the UK acts as a centralized administrator responsible for directly managing spectrum use. Under this top-down approach, the government controls the choice of

technology, spectrum uses, and users. Because government regulators may lack the expertise to make informed decisions, because regulation is often slow and expensive, and because it is vulnerable to influence costs, this approach is often criticized relative to market-based approaches.⁴ The name, C&C, is intended to conjure images of an inefficient Soviet-like bureaucracy.

In contrast, the two paradigmatic alternatives for spectrum management are the so-called *property rights, exclusive flexible use, or licensed* ("Licensed") model and the so-called *open access, unlicensed, or spectrum commons* ("Commons") model.⁵ Both of these approaches are stylized as market-based because decision-making power is decentralized to the market. Service providers, equipment makers, and end-users interact and compete in the marketplace to determine how spectrum is used, subject to the regulatory frameworks put in place to govern how the market operates.

Under the Licensed approach, an exclusive use license is assigned which may be traded in secondary markets. This is often referred to as the "property rights" regime because the license confers a property right on the licensee for exclusive use of the spectrum.⁶ The licensee is responsible for making all substantive choices as to how the spectrum is used. The PCS spectrum in the US provides a good example of this regime in practice. PCS licensees have flexibility in the choice of technology, services offered, and are allowed to trade the usage rights conferred by the license.⁷

In contrast, under the Commons approach, the right to access or use the spectrum is shared among the users subject to a *protocol, etiquette, or framework* (Protocol) that

⁴ Much of the delay and cost associated with regulatory processes is the result of open access/due process rules that are designed to enhance the quality of information shared and to protect the process from regulatory capture. Standards bodies face the same problems and for the same reasons (see, for example, [Lehr92]).

⁵ For purposes of paper, we will use Licensed and Commons with capital letters as shorthand for the concepts as discussed in this paper.

⁶ In truth, this nomenclature is unfortunate since all of the regimes convey different sets of property rights. In the licensed regime, the licensee has the sole right to determine who gets to use the spectrum. In the commons, there is no such exclusive right. In the licensed regime, the licensee has a property right for interference protection that is enforceable by the Courts or by the regulator. In the commons, there is no explicit right of protection from interference. Any protection is implicit, and a consequence of the protocol.

⁷ Even PCS spectrum is not pure "property rights" since there are rules which limit tradability and licenses are subject to term limits.

embodies the mechanism for managing the spectrum.⁸ As with the licensed model, the legacy terminology is misleading because it suggests a more extreme form of the regime than is likely to be either practical or desirable. For example, the spectrum will not be "open access" to all, but only those who conform with the unlicensed protocol. And, "unlicensed" does not mean unregulated. There will need to be mechanisms to allocate resources during periods of congestion, to recover the costs of management, to determine (and potentially modify) the rules that govern the Commons, and to enforce compliance with the rules and resolve conflicts among users. In contrast to the C&C regime where the decision-making authority resides with a central planner (the government) or in the licensed regime where decision-making authority is decentralized to a (presumptively)⁹ competitive price-based market, in a Commons, the decision-making authority is decentralized to those who share access to the Commons and is governed by the Protocol that is put in place.¹⁰

Furthermore, Commons spectrum cannot be "free" in the sense that costs that are incurred need to be borne by users (either directly through access payments or indirectly through taxes, protocol implementation costs, or congestion-related quality of service effects).¹¹ These costs include any

⁸ We will use Protocol with a capital "P" as shorthand to refer to the set of rules and mechanism that will instantiate the Commons regime. This will likely include "protocols" (i.e., running code for a software radio or a technical standard), "etiquettes" (i.e., rules for device behavior that guide protocol design), or "manners" (i.e. more general codes of conduct, as discussed by [Friedman05]). The more complete Protocol will also include the regulatory mechanisms needed to enforce compliance with the governance model.

⁹ We say "presumptively," because the efficiency benefits of transitioning to a general regime of flexible, tradable licenses that rely on market forces instead of government regulation presumes that there would be competitive secondary markets for spectrum licenses. If this is not the case, then at one extreme, the transition from direct government "command and control" is simply replaced with some form of public utility regulation (regulation of a privately-owned monopoly), or at least, remains governed by "antitrust." Thus, just as the C&C regime is not immune to market forces (e.g., through lobbying that reflects market interests), the "market" is never wholly free of regulation. Moreover, the feasibility of deregulation via market forces depends in part on the extent of competition and the robustness of secondary spectrum markets (liquidity, low transaction costs, etc.).

¹⁰ The choice of the Protocol may be made by the government or by the market via industry standardization.

¹¹ Ignore case of subsidies (i.e., government pays to create open access spectrum and then provides free of use) since this is false distinction among models (e.g., licensed

costs of setting up and operating the management mechanisms (processing costs of implementing the sharing protocol, enforcement and congestion costs).¹²

Several additional distinctions between the Licensed and Commons models are worth noting. First, Licensed and Commons are both "shared" in the sense that multiple devices and end-users simultaneously access and use the spectrum. For example, mobile operators share spectrum over multiple users, and competition among operators offers competition across technologies and markets.¹³

Second, Licensed and Commons are both "market-based" in contrast to C&C. In the licensed regime, the spectrum is explicitly priced and traded on secondary markets. This allows ready observation of the opportunity cost of using spectrum (assuming secondary markets are efficient) which provides strong incentives to use spectrum efficiently since congestion externalities are internalized and helps ensure that scarce spectrum is allocated to higher value uses first. In the Commons regime, use of the spectrum is governed by the Protocol. However, use is still "market-based" in that users respond to private incentives (individual utility or profit maximization) to determine how they operate in the unlicensed spectrum, constrained by the spectrum etiquette. This raises the potential risk of a Tragedy of the Commons in which users lack adequate incentives to use spectrum efficiently (they are "greedy") and spectrum may fail to be allocated to the highest value uses ("garbage" uses may crowd out more valuable unlicensed users).¹⁴ To address this challenge, clever design

spectrum use could be subsidized by licensee because use produces other benefits for licensee as would be the case if Microsoft or Intel bought licensed spectrum and then created a "commons" for use by Intel or Microsoft devices).

¹² Note, these costs are also borne in licensed regime. Some analysts (Faulhaber05) ignore transition costs

¹³ In this sense, the Licensed model offers dynamic, spectrum access already and movement by end-users via roaming and via switching among operators already implements dynamic secondary markets. Viewed this way, the mobile customers are secondary licensees who get to use the spectrum on the basis of the rules established by the licensed operators.

¹⁴ That is, individual users do not bear the full cost of the negative congestion externality they impose on other users and so are inclined to use too much spectrum (e.g., by using too high power). Protocols that allow such self-interested behavior which results in a Tragedy of the Commons are termed "greedy" ([SataPeha00]). An additional problem is that when spectrum is scarce, a Commons may fail to allocate spectrum efficiently so that lower value uses crowd out high value. Such low value uses are termed "garbage" uses. However, because markets are often imperfect, one must be careful before equating a low ability to pay for access to "garbage" use.

of the Protocol may be used to substitute non-price (e.g., performance) incentives to use spectrum efficiently.

Which approach is better depends on the context and the relative scarcity of spectrum. If congestion is not a threat, then the opportunity cost will be low in both regimes (once again, assuming efficient secondary markets for licensed spectrum).¹⁵ The relative difference in opportunity costs under the two regimes should to induce applications/users to self-select so that users willing to pay more for strong interference protection will choose to locate in licensed spectrum, while users that are unwilling to pay for such protection will choose to locate in unlicensed spectrum. The latter will include both applications that are robust to interference (e.g., because the applications/devices are adaptive) and low value uses "garbage" uses.

Third, a key difference between Licensed and Commons spectrum is associated with the right to exclude. In Licensed, the right to exclude belongs to the licensee and so is centralized. Whereas, in the Commons, the right to exclude other uses is shared, and so decentralized.

This characterization of the regulatory models presented above is stylized in that it mischaracterizes the complexity of the status quo regime. The actual situation differs significantly from C&C. Different bands are subject to different regulatory regimes. Some like PCS spectrum, allow significant latitude for technology choice and trading of licenses, coming closest to the ideal of the Licensed model in the current environment. Other spectrum like the television broadcast or much government spectrum remains encumbered with legacy regulations and more nearly reflects the C&C model as described. A version of the Commons exists in the form of the "Part 15" rules which govern the operation of low power unlicensed devices in the ISM and U-NII bands.¹⁶ The Part 15 rules specify the operating parameters for low power unlicensed devices. The Part 15 rules seek to limit the interference caused by such devices for other users while seriously circumscribing the

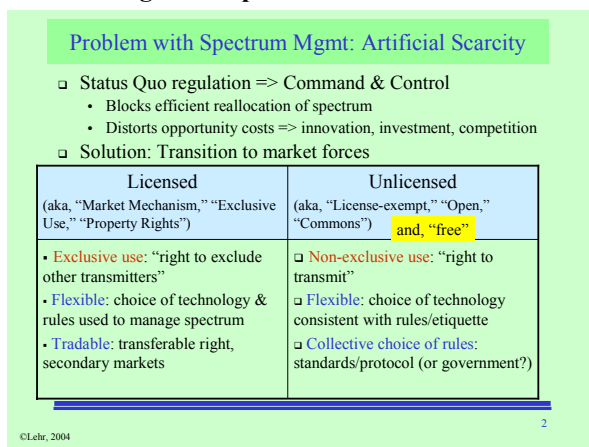
¹⁵ If secondary markets are not efficient either because of market power or because of high transaction costs (thin markets, information asymmetries, regulatory costs), then the price for licensed spectrum may be substantially above the economic value of the spectrum. The existence of unlicensed spectrum provides a safety valve and check to help keep the prices for licensed spectrum in line with the scarcity value of spectrum.

¹⁶ This refers to the relevant section of the U.S. Communications Law that governs the use of unlicensed devices (see http://www.access.gpo.gov/nara/cfr/waisidx_02/47cfr15_02.html for Title 47 Code of Federal Regulations Part 15, 47CFR15). The Industrial, Science and Medical (ISM) and the Unlicensed National Information Infrastructure (U-NII) bands include spectrum in the 900Mhz, 2.4GHz, and 5GHz bands in which cordless phones, garage openers, and WLAN devices (802.11a/b) devices operate.

interference protection granted to unlicensed devices. The FCC regulates compliance via the certification of devices (radios).

The characterization of the status quo as C&C is intended to highlight the benefits of further spectrum reform. These benefits will include increased reliance on market forces to allow spectrum to be more flexibly reallocated to new uses and users. This enhanced efficiency both in a static (spectrum goes to highest value uses) and dynamic (new technologies can be deployed) sense will reduce artificial scarcity (regulatory induced). This artificial scarcity distorts opportunity costs and hence adversely impacts innovation, investment, and competition. Figure 1 summarizes how the three regulatory models are related.

Figure 1: Spectrum Reform



Finally, it is worth noting that spectrum management involves several layers of activity or aspects that might be logically disaggregated. This might make it feasible to apply different management regimes for different aspects. For example, we might view spectrum management broadly as addressing three categories of issues: (1) structure of authority to determine the rules; (2) structure of rights ownership; and (3) the structure of communication (market mechanism) to transfer authority, ownership and usage (delegation). As in other areas of regulation, the extent of centralized, hierarchical control may differ. For example, the distributed control of a Commons may be more appropriate for enabling low overhead communications (the third aspect), while a degree of centralized control may simplify management of the rules (e.g., like the management of the address space and domain level identifiers in the Internet).

2.1.2 Flavors of Unlicensed

There are three basic approaches for allowing unlicensed devices to operate that are under discussion currently. These are: (1) underlays, (2) overlays, and (3) dedicated unlicensed.

A spectrum *underlay* is a secondary easement for low-power devices that allows devices to operate in the noise floor of the primary, licensed spectrum user. The Part 15 rules provide one model for how such an underlay might be

implemented. The current debate over underlays has focused on Ultrawideband (UWB) technologies that would use code-division-multiplexing to spread signals over a very wide frequency band and thereby allow much lower power operation per bit of information transmitted.

Similarly, a spectrum *overlay* is a secondary easement that allows opportunistic use of licensed spectrum that does not interfere with the primary operators use (e.g., in time slots or geographic locations where the primary spectrum is not in use). The "listen-before-talk" protocol mandated for use in U-PCS spectrum is an example of an "overlay" approach to managing unlicensed devices. The current debate over overlays focuses on the potential for smart or "Cognitive Radios" to dynamically adjust their behavior to efficiently share spectrum.

Finally, *dedicated unlicensed* spectrum would be spectrum that treated all unlicensed devices as primary users. From a practical sense, the focus of unlicensed use is on permitting decentralized coordination and so an effective Protocol for managing a Commons would entail the same sort of functionality as an overlay. Additionally, because (as we explain further below), the power for devices operating in the common is likely to be quite limited and because there is very little spectrum that might be available that would have no incumbency rights to address (e.g., satellite downlinks which might be regarded by policymakers as "primary" users), a Commons may look a lot like an underlay in the first instance (limited power), and an overlay in the second (ability to yield or be pre-empted by certain "primary" uses).

For the purposes of designing a suitable Protocol, it may not really matter whether one thinks of the unlicensed devices as operating in an underlay, overlay, or dedicated spectrum; however, we will focus on the last model because we believe that the politically-practical limits on implementing effective underlays (power limits) or overlays (cognitive radios) will be sufficiently severe as to hamper realization of the full benefits and applicability of the unlicensed models. For example, the power limits for UWB are quite stringent, which may limit the applicability of this promising technology to quite short range uses (e.g., a few tens of feet to serve as wireless "USB" replacement). Moreover, operating in an underlay or an overlay if used intensively, is likely to result in some interference for primary users unless guard bands are sufficiently stringent, in which case, the effective spectrum available will be even less. However, if it is feasible to allow sufficiently robust (but still limited) power and primary (incumbency/licensed) rights are adequately circumscribed in the spectrum assigned to the Commons, than the choice of unlicensed flavor will be less important.¹⁷

¹⁷ [FaulFarb02] recommended the use of underlays/overlays to reconcile the coexistence of licensed (property rights) and unlicensed (commons) models. [Lehr04] explains why this may be less attractive than using

In summary, therefore, a spectrum Commons enables:

- (1) *Shared use* (multiple devices with heterogeneous technologies and apps share spectrum access – multiplicity of power levels, bandwidth requirements, and duration); which is
- (3) *non-exclusive use* (control of access is decentralized to etiquette, there is no central-planner who manages access) that is managed as a
- (3) *Commons* (management control, including decision-making authority over etiquette is decentralized). It is
- (4) *unlicensed* because we presume that devices/users do not need to first pay for or otherwise acquire rights to access/use the spectrum. Instead, any device that adheres to the etiquette (and is therefore "conforming") is
- (5) *free* to use the spectrum.¹⁸
- (5) Limited *a priori* knowledge of who is participating in network (new devices may enter/leave at will). Makes it difficult to generate reputation effects.
- (6) Distributed/decentralized network management. There is a lack of or lack of a requirement for a centralized controller.

2.1.3 Relationship to Ad hoc networking

The significant progress in wireless technology and the growth of wireless services has provided the principal impetus for reforming spectrum management and the transition toward increased reliance on market forces. While many wireless technologies contribute to both the viability and desirability for managing spectrum via unlicensed (smart wireless system technology including software/cognitive radios, smart antennas, and MIMO), the benefits of unlicensed wireless are perhaps best anticipated in the context of "ad hoc" networks. Some of the key features that characterize "ad hoc" network environments include ([MackCor99]):

- (1) Mobile devices (moving autonomously). There is no central controller.
- (2) Energy conservation is a major concern since devices are likely battery powered. Devices are low powered so range-limited.
- (3) Multi-hop routing with most devices acting as both receivers and transmitters. This is a direct result of the fact that nodes are low powered and so wide-area communications requires multi-hop operation. This also turns out to be efficient ([XieKumar04]).
- (4) Random distribution of devices. Beyond density, very little a priori structure may be assumed for network topology. This results in difficult dynamic routing problem.

These features have a number of important technical/business implications for protocol design. First, in ad hoc networks there is a need to generate incentives to cooperate. The nodes need to cooperate to support multi-hop routing because of their limited range (power). This is challenging because acting as a repeater for other users traffic consumes power and so is privately costly.

Second, the randomness of the nodes participation and network topology provides limited opportunities for long-term contracting among devices to manage sharing. This may reduce the ability to rely on repeated game reputation effects to induce cooperative/efficient behavior.¹⁹ By analogy, the transactions among ad hoc nodes are more like what happens in anonymous arms-length trade than what happens in a firm that is under centralized management (which is like the Licensed regime). This suggests that protocols that use price-based mechanisms may offer more fruitful approaches for inducing cooperation. In order to avoid the transaction costs of using actual money (i.e., billing) which are unlikely to be warranted given the relative value of the services provided in each transaction (i.e., forwarding another node's bits), the protocol may be based on notional prices that reward repeating with sending rights.²⁰

Third, the need to keep track of distributed resources makes implementing traditional MAC approaches such as time (TDMA), frequency (FDMA), or code (CDMA) multiple access difficult. The mobility of the nodes, their limited resources, and the dynamism of the communications environment means that link-level and network-level reliability issues are an enduring challenge and necessitates multi-layered protocol design (see, [Kawad04]).

Fourth, the limited information that may be assumed about nodes location, behavior, and even identity (to protect privacy) imposes limits on the interference models that may be implemented. The nodes must respond to local information and information gleaned from adherence to the "protocol" to determine whether interference or communication problems they are experiencing are due to congestion, local factors, or other wider-area factors.

dedicated spectrum to address the needs of unlicensed devices.

¹⁸ The etiquette may require a payment (e.g., a license fee to use the technology) to allow recovery of spectrum management costs or other costs (transition costs associated with spectrum clearing).

¹⁹ In any case, relying on reputation effects to generate incentives to cooperate may raise privacy concerns. This is another potential advantage of the anonymity inherent in using a notional price-based approach.

²⁰ See [Crow03] for how such a price-based protocol might be implemented to generate cooperation incentives in an ad hoc network.

Fifth, the unpredictable entry/exit of nodes and of traffic patterns that makes ad hoc networks so dynamic provides a good test case for unlicensed use because it mirrors the process of innovation and entry/exit and the complexity of heterogeneous flexible uses that characterizes a competitive, dynamic wireless industry and is consistent with our view of the wireless future. Ad hoc networks present the most demanding challenge for supporting unlicensed use. Many of the uses of unlicensed that we actually expect to see will not fit the ad hoc model (e.g., fixed wireless broadband access or in-home broadband wireless). However, a regulatory framework that is robust to the deployment of ad hoc network technology will hopefully be robust to other innovations. In effect, ad hoc networking is the "Canary in the Mine" against which to evaluate unlicensed protocols.

Sixth, while Ad Hoc networks pose a challenge for dynamic routing and the design of management mechanisms, their decentralized/distributed nature may offer robustness benefits. With ad hoc networks there is no single point of failure. Contrast this with the vulnerability of a hierarchical telecom network which relies on a centralized signaling framework such as Signalling System 7 (SS7). This is not to say that this makes ad hoc networks more reliable than traditional hierarchical, centralized networks (which may benefit from redundant facilities and quick recovery mechanisms), but rather that they offer a different kind of robustness and reliability. A world with "Carrier grade" telecom networks offering "five 9's reliability"²¹ and ad hoc networks is likely to offer improved reliability over one with only one kind of network.

Seventh, and finally, the unpredictability inherent in ad hoc network management makes it difficult to know who will participate and what they will want to do. When these situations prevail in a macro market, it makes it difficult to develop a retail business model or to plan infrastructure. When you cannot do these things, it makes it difficult to raise financing. Such market situations may be best addressed by flexible and scalable business models, characterized by distributed experimentation and loose organizational structures. That is, "ad hoc" business/service models. These business models are at a disadvantage relative to the traditional service provider model when it comes to being able to assemble the resources to acquire access to licensed spectrum. For example, the traditional service provider model involves capital-intensive network infrastructure and retail operations. Such investments may require stronger interference protection because of the scale of investment that is at risk and may be co-specialized with a particular frequency (although with smart radio systems, this level of co-specialization is decreasing). Therefore, we might expect service providers to be biased in favor of licensed spectrum. While this is an important model for

²¹ The standard for modern telecommunication networks is to offer 99.999% availability.

deploying wireless services, it is certainly not the only business model or type of service environment that needs to be provided for to encourage a vigorous wireless future.

2.2 Characteristics of good mechanism

Before proposing a framework for evaluating alternative management protocols for shared, non-exclusive use spectrum, it is worthwhile considering the goals that a desirable system ought to promote, and to that, it is necessary to comment on the type of environment that is likely to prevail.

2.2.1 Features of future "unlicensed" environment

To provide additional context for evaluating potential protocols for managing unlicensed spectrum, there are additional features that are likely to characterize the unlicensed wireless environment. These include:

2.2.1.1 *Heterogeneous devices and uses.*

We should assume that the traffic may be bursty with heavy-tailed distributions. It will vary in duration, bandwidth, and power/range requirements for end-to-end communications. An ideal framework would allow as many types of traffic to co-exist as possible, but it may prove necessary to deviate from this ideal. One option is to think about designing separate commons for different classes of applications (e.g, a low power commons for mesh applications v. a higher power commons for multiple classes of applications). However, even if the Protocol does impose some limits on the classes of applications that are admitted, the class should be as broad as possible.

2.2.1.2 *Lots of potential operators and devices.*

Although the number of potential users should be large, it is not unlimited. An oft-heard critique of commons management regimes is that decision-making does not scale to large numbers.²² Fortunately, the relatively low power operation will likely bound the number of devices that need to be coordinated. Precisely what this bound might be is uncertain but it might be quite large.

2.2.1.3 *Relatively abundant spectrum.*

In the future we have described, radio systems will be more efficient and regulatory reform will have eliminated most of the artificial scarcity due to inefficient legacy regulations. This does not mean that the opportunity cost for using spectrum will be zero (no scarcity), but only that the marginal value of any particular frequency band in any particular locale will be much lower than is suggested by an analysis of historic spectrum auction data or wireless asset

²² For example, the social norms that coordinate cooperation in resources that are managed as a commons in other contexts (common pasture, fisheries) depend on group cohesiveness. As the group or scale of economic activity gets large, a market-based process based on property rights to manage the resources may become preferable (see [Demsetz02]).

transactions. Because such data are embedded in most business model forecasts used to estimate the marginal value of spectrum, these estimates likely are systematically biased upward. Moreover, because unlicensed applications are relatively new, such analyses are biased to value licensed spectrum significantly more than unlicensed. To the extent the social value of unlicensed is associated with innovation,²³ failure to incorporate this value into spectrum calculations when estimating what the appropriate allocation should be between licensed and unlicensed spectrum is biased in favor of licensed.

If we assume that spectrum will be relatively abundant, then the marginal value of improvements in spectral efficiency over other characteristics that might be important in a good Protocol will be less. When comparing two protocols it would be worthwhile considering which would be preferred if there was no scarcity (i.e., mutual use may be a positive sum game). Note that this is not the same as saying that the probability of interference is zero. Even if traffic is relatively sparse, unless it is coordinated, nodes may interfere with each other unnecessarily. However, the need to sort high value from low value uses would no longer arise and incentives to cooperate may be stronger. In contrast, in a licensed regime with zero scarcity the market price for spectrum might be sufficiently close to zero to be below the transaction costs that would be incurred by a seller to make it available to a buyer, so efficient use might be deterred in a licensed regime if scarcity is sufficiently low.

For example, it may be the case that a successful Commons protocol, as a first approximation, may assume that devices do not want to be "greedy" except out of ignorance and that there are no "garbage" applications. This does not avoid the need to incorporate the capability to address such circumstances when they arise, but it does increase the tolerance for relatively inefficient mechanisms when those are used infrequently.

Moreover, the assumption of relative spectrum abundance for unlicensed is also supported because we assume that there will be both licensed and unlicensed spectrum, which will encourage applications to self-sort on the basis of their relative demands for strong interference protection. Thus, devices that are relatively interference tolerant will be more likely to locate in unlicensed spectrum. And, even intolerant applications will find locating in unlicensed acceptable if the prospect of congestion is sufficiently low (e.g., in rural areas).

2.2.2 Future is shared spectrum

The above factors, and the nature of "ad hoc" networks suggest a future where the spectrum is heavily shared, not just among multiple users but by multiple uses and operators (see Figure 2). The increased capabilities of smart

radio systems have expanded the space of wireless architectures, making it possible for many very different architectures to support similar services (e.g., centralized and decentralized architectures are both options) and to decouple services from specific frequency bands. Customer expectations are also pushing us toward a world of heterogeneous, mixed networks. Customers want 24/7, ubiquitously available services, and do not care what the technical platform is that supports those services. They just want it to work. Moreover, with the transition to broadband, user traffic will become more bursty and heterogeneous. Regulatory reform and industry dynamics are increasing competition all across the value chain (e.g., wired and wireless broadband are nearer substitutes). This encourages operators and businesses to try and decouple infrastructure from spectrum and investment. An operator whose service is locked into a specific frequency band or can only provide service over a single platform will be at a disadvantage relative to more flexible operators. Just as mobile operators are now seeking ways to integrate 3G services with WiFi, so they will seek to be able to integrate other technologies.²⁴ As applications become more robust to operating across heterogeneous platforms, equipment makers and service providers will find it more advantageous to support such frequency/architecture flexibility.

Figure 2: Future is Shared spectrum

Future is shared spectrum: decoupling of spectrum frequencies from infrastructure investment & applications	
Technology (Capabilities) → frequency agility, expanded capacity for sharing	Smart radio systems, spread spectrum, transition to broadband platform architectures
Revenue (Customer experience) → 24/7 availability, simplicity of use, seamless mobility	Heterogeneous networks (3G/WiFi, wireless/wired, global roaming)
Costs (Network provisioning) → lower costs, take advantage intermodal competition	Bursty traffic, multimedia services, fat-tailed usage profiles
Policy (Spectrum reform) → reduced <i>artificial scarcity</i> due to legacy regulations	Transition to expanded flexible market-based licensing and unlicensed spectrum mgmt regimes

©Lehr, 2005

When evaluating Commons management regimes, it is worthwhile having some concrete ideas of the sorts of business models/applications that may need to be supported in unlicensed spectrum. A good Protocol ought to support as many of these as possible. Several possible scenarios for unlicensed use that should be considered include:

2.2.2.1 WiFi model of unlicensed device use

In this scenario, mass market consumers deploy new wireless devices in a viral manner, without central coordination of use. For example, the deployment of

²³ See [Odlyzko04] for discussion of role of unlicensed in promoting innovation in telecom services.

²⁴ See [LehrMcK03] for a discussion of integration of WiFi and 3G.

Bluetooth and potentially UWB in home appliances could follow such a model. Once such devices are widely deployed, it is reasonable to expect that someone might develop the applications or middleware to tie these together into larger networks (e.g., wireless grids).

2.2.2.2 *Mobile operators sharing 3G spectrum*

Today, spectrum licenses to provide mobile services provide an entry barrier that gives incumbent licensees a strategic advantage. However, with robust competition and the threat of increased allocations for competing wireless technologies on the one hand and the prospect of having to pay to acquire licenses for additional spectrum in order to support new 3G wireless broadband services on the other, mobile operators may be more inclined to share spectrum. The value of holding an exclusive license to a particular frequency band is diminished in the world we have described above. With bursty, wireless broadband traffic, it will be more efficient for mobile operators to dynamically share spectrum than for each of them to have to provision sufficient spectrum to support their peak bandwidth needs. The technologies to share such spectrum are becoming more readily available.

In such a world of agile infrastructure, providers may care less which spectrum they own and more about ensuring that no one else has a right to exclude them. Collective ownership of the 3G spectrum and management of it as a closed common (i.e., for qualified service operators) may offer an efficient management regime. Of course, were the operators to attempt such a structure, they might raise antitrust concerns. To forestall that, the operators might make the spectrum "open access" under the management of an industry-defined protocol. A similar approach is used to justify the openness of Internet peering arrangements (i.e., anyone can peer who meets the capacity requirements, but those are set so as to practically restrict the number of potential peering partners).

2.2.2.3 *Community last-mile networks*

In a future with lots of deep neighborhood fiber (FTTx where x may be the curb, block, building, etc.), wireless may play an important role in providing connectivity over the last few tens of feet to the home. If the FTTx is deployed by a large national service provider like a telephone or cable company, then they may find it practical and not difficult to acquire licensed spectrum. However, in other situations, communities may wish to deploy their own last-mile infrastructure. While not appropriate in all contexts, municipal networking is likely to be an increasingly common feature of the future broadband landscape, and wireless will play a big role in that.

Municipalities may prefer to locate in unlicensed spectrum because of the reduced transaction costs and their bias in favor of capital/equipment instead of service-provider models. That is, municipalities are typically relatively ill-equipped to operate retail services, but may have certain advantages in financing the deployment of

long-lived, "natural monopoly"-like infrastructure.²⁵ If used to support basic infrastructure, the municipalities may desire strong interference protection to protect the value of their local infrastructure and to minimize the need to address interference-related issues. The efforts of WISPs to gain special interference protection for their wireless broadband services deployed in unlicensed spectrum provides insight into how this might work. Alternatively, municipalities may be inclined to share their public safety spectrum if appropriate sharing mechanisms can be devised.

2.2.2.4 *Mesh networks*

A final example to consider is the case where there is a desire to deploy a large mesh network. This might arise in the context of an emergency (flood, hurricane) or special event (conference, festival). The proto-typical situation would be one in which it was desired to set up a mesh network very quickly to cover a campus-sized or larger area. Alternatively, the mesh might be a large ad hoc network. Possible examples might include a taxicab dispatch network or (vehicle) traffic management system. For example, such a system could be deployed in a city to better schedule public transport and alleviate congestion (e.g., in downtown London).²⁶

As noted earlier, ideally, an appropriate management regime should admit all of these possible scenarios. However, that does not preclude creating different commons for different applications. For example, there are a number of compelling reasons why mobile operators or municipalities might prefer a closed commons that is restricted to a relatively small number of participants. These include the desire to enforce market power (a cartel is easier to sustain with a few members); technical efficiency (having a small number of well-known parties simplifies the design of the sharing protocol and enforcement); economic efficiency (having a small number of well-known parties simplifies long-term contracting which can reduce coordination costs in the face of uncertainty and sunk cost investments). In contrast, the viral adoption and mesh models may prefer a more open model to ensure that the costs for new nodes to join is minimized.

2.2.3 **Goals for successful management protocol**

To summarize, relative to the type of environment specified above, a successful Commons management protocol should strive to realize the following goals:

2.2.3.1 *Promote/facilitate innovation in wireless devices, services, and business models.*

This means minimizing barriers to entry/adoption of new technologies and switching costs. It also means avoiding lock-in that may block further deployments of new

²⁵ See [LehSirGil04] for discussion of municipal wireless.

²⁶ These are real examples that are currently in operation.

technology. In many cases, the success of a technology and the strategic interests of its sponsors promote lock-in. For example, consumers stay with Microsoft Windows because it satisfies many of their needs, offers them a huge selection of compatible software and hardware options (positive network effects), and changing to another operating system would incur switching costs. Microsoft benefits from its position and has no incentive to encourage switching. While it is not the responsibility of the management framework to dictate user tastes or business models that often are the cause for lock-in, the framework does need to strive to avoid regulatory impediments to deploying new technologies. For example, the growth of an installed base gives rise to vested interests that can be expected to appropriate the Commons for their own, seeking special protection against new technologies. We have already seen a hint of this in the efforts by WISPs and campus networks that have deployed infrastructure in unlicensed spectrum seeking to block or limit competing uses that may interfere with their use. Such efforts, while predictable need to be resisted.

One benefit of adopting a minimalist regulatory framework may be to limit regulatory hooks that might otherwise be used to lock-in particular technologies. For example, identifying a specific technical protocol (e.g., CDMA or TDMA) may make it easier for incumbents whose choice was dictated by the specific technology to argue for protection.²⁷ A minimalist regime helps precommit the regulator to letting the market dictate what technologies are successful.

2.2.3.2 *Minimize transaction costs for accessing spectrum.*

A principal benefit of unlicensed relative to licensed spectrum is that it minimizes the transaction costs associated with accessing spectrum. This chief virtue needs to be kept in mind when evaluating protocols for managing the Commons. Real-time usage fees (e.g., the per minute charges that characterize mobile telephony services) or the costs of accessing secondary markets to negotiate spectrum rights are obvious ways in which unlicensed usage may be deterred. However, transaction costs to accessing unlicensed may also arise in other ways. For example, an overly complex protocol for managing resource sharing will impose overhead costs on devices that can deter usage. Attempting to build in too strong interference avoidance protection into the Protocol may result in excessive requirements for device certification and enforcement that will further increase overhead costs associated with management of the Commons. Limiting devices expectations of protection may help control these overhead costs.

²⁷ [David86] describes the dilemma for policy-makers who promote technologies that are subsequently abandoned and must then confront the "angry orphans."

2.2.3.3 *Provide mechanism for congestion management*

While we expect spectrum to be less scarce in the future we are designing for, there still need to be robust mechanisms to address congestion when it occurs. As discussed earlier, however, we may be able to tolerate management mechanisms that are somewhat less allocatively efficient (i.e., emphasize coordination to allow all users to co-exist rather than seeking to selectively allocate better performance to higher value uses).

2.2.3.4 *Provide mechanism for modifying/establishing the etiquette*

An appropriate Commons management Protocol will need to include process for modifying and updating the Protocol. Such a mechanism is needed to allow the Commons protocol to evolve over time, and to facilitate the re-allocation of Commons spectrum as Licensed in the future should that be deemed desirable. A principal benefit of market-based spectrum reform is that it provides for a dynamic process for future adjustments.

In Licensed spectrum, the Invisible Hand of the competitive market (acting via secondary markets) provides the mechanism for modifying the management regime. This is a chief benefit of the property rights regime because it can take advantage of the general legal and regulatory infrastructure used to enforce and manage property rights in the general economy.²⁸

In contrast, the Commons regime, of necessity, will require specialized (spectrum-specific) attention to its mechanism. Because rights in the Commons are shared among all qualified users, it will be challenging to upgrade the mechanism since it will require consensus.²⁹ This suggests that reliance on industry standardization processes may offer the best option. These may be assisted by government action (e.g., by setting a time line for

²⁸ As noted earlier, since Licensed will not realize a pure property regime (fee simple rights), there will continue to be regulation. For example, the existence of term limits or secondary use easements will provide a sufficient (although not necessary) basis for ensuring a need for on-going regulation in the Licensed regime. Furthermore, to the extent the adjudication of disputes over spectrum license rights requires specialized technical knowledge, there may be a requirement for specialized spectrum courts that would, by their nature, look more like regulatory agencies than traditional general purpose Courts. [Goodman04] discusses the likely problems inherent in relying on either trespass or nuisance law to protect spectrum property rights, while noting that the prospects for managing a Commons are equally troublesome.

²⁹ [Heller98] explains how an "anti-commons" arises when the right to exclude is distributed too widely making communally-managed property unusable. That is, when too many individuals possess decision rights, reconciling their collective interests can be difficult, resulting in paralysis.

standardization), but the choice of the actual Protocol and its details may be best handled if delegated to an industry standards body. While de jure industry standardization processes are often cumbersome and vulnerable to many of the same capture risks as government regulatory bodies, they are the best that we have available.³⁰ When the stakes are high and interests diverge, it is only natural that there are costs associated with adopting and transitioning to new technologies. The de facto standardization that would operate in the Licensed regime does not avoid these adjustment costs. The Darwinian market process that encourages survival of the fittest has significant costs for participants.

Which standards organization might offer the best framework for managing the design of an appropriate Commons protocol requires further study. Whether it should be patterned on the IETF, NRIC, or the ITU is something that needs to be discussed, and will no doubt result in heated debate.

Another obviously important issue to address is the question of how to transition from today's status quo to the future anticipated in this paper. This requires overcoming NIMBY and clearing spectrum of incumbents. This is a problem that is faced by both Licensed and Commons models, and may best be addressed if both are confronted at the same time. If, as is hoped, reform will significantly lower spectrum opportunity costs, than the *ex post* value (after the transition) of the spectrum will be much lower than historical experience with auctions suggests. If we really believe this is the case, then perhaps a Big Bang auction along the lines suggested by [KwerWm02] would offer a practical way to effect the transition. In the auction they propose, a large chunk of spectrum is transitioned to flexible Licensed use. At the same time, spectrum could be "purchased" for unlicensed use; however, if we really expect the auction proceeds to be low (reflecting the low expected future value of spectrum in a spectrum abundant world) then it may be reasonable to approximate this purchase price as zero. That is, simply designate as unlicensed a portion of the spectrum that will be made available and use some of the aggregate auction proceeds for clearing.³¹

2.2.3.5 Promote fairness and non-discriminatory access

The Commons Protocol ought to promote non-discriminatory access. The Protocol should not attempt to sort uses of unlicensed spectrum into high and low value

³⁰ See [Lehr92] or [DavGreen90] for discussion of economics of industry standardization processes.

³¹ That is, do not require the government to purchase the unlicensed spectrum. Since everyone will be able to use the spectrum (including those that also or even preferentially may choose to operate in licensed spectrum) and since the unlicensed spectrum will lack strong interference protection that justifies paying more for licensed spectrum, such an approach seems reasonable.

users, but should leave that as much as possible to the market. The Protocol should not offer differential quality of service. The notion of non-discriminatory access for all qualified (compliant) users of the Commons is fundamental to the notion of a Commons.

2.2.3.6 Minimalist regulation

To realize the benefits of market-based reform and to protect against the inefficiencies inherent in the traditional C&C approach to spectrum management, the Commons regime ought to aspire to minimalist regulation. Achieving this goal is difficult because it requires a delicate balance between simplicity and flexibility. Often the simplest rules are bright line rules, but these can lack flexibility. On the other hand, simple rules that are highly flexible may not adequately constrain behavior to accomplish the goal of the rules. Vague rules can lead to higher ex post enforcement costs and higher ex ante uncertainty if they fail to provide clear guidance as to what should be expected.

3 TAXONOMY FOR EVALUATING SHARING PROTOCOLS FOR UNLICENSED SPECTRUM

In this section we offer our preliminary thoughts on the sorts of rules that we think are likely to be needed to meet the requirements of an appropriate Commons management Protocol as described above. Before examining the range of rules that we think are likely to be needed to effectively manage a spectrum commons, it is worth considering why defining a suitable resource sharing protocol for a Commons is difficult. Some useful lessons may be gleaned from examining how resource sharing is managed in wired telephone networks and the Internet.

3.1 Lessons from the past

The traditional Public Switched Telephone Network (PSTN) supported one basic service, the voice call. Because the statistical properties of voice calls were well understood and quite regular,³² it was relatively easy to standardize on an appropriate unit of capacity (e.g., Erlangs) and to map voice calls into standardized 4 KHz voice channels which could be encoded with 64Kbps PCM.³³ This regularity

³² That is the dynamic range of the human voice and ear for basic telephony calling is relatively narrow. Quality must be good enough (e.g., end-to-end delays should not exceed 250msec but there is a reasonably good tolerance to bit errors) and substantially better quality may not be very valuable. Of course, mobile telephony (which historically, has often been much lower quality) on the one hand and improved codecs (offering really high quality voice) on the other, we have learned that even voice calls can support a much wider range of quality of service than previously suspected.

³³ Of course, the monopoly/regulated industry structure helped solve the business problems that often plague standardization efforts in competitive markets. Note, that

simplified the design of resource sharing mechanisms in the PSTN.

In this environment, it was appropriate to use a centralized resource management policy that allocated units of capacity along each call route for the duration of the call.³⁴ A separate signaling channel was used to manage resource allocations (capacity requests, allocation, and set-up). When multiple carriers interconnected, this could be extended across carriers, with the allocation method offering strong isolation properties with hard guarantees of capacity and strict delay bounds (i.e., to limit interference between users and to ensure acceptable end-to-end quality).

With the growth of data traffic, which is inherently more diverse (unlike voice calls, data traffic differs significantly in its bandwidth requirements and in its delay/error rate tolerance), the resource allocation problem became more complex. Although more diverse, relative to voice (which has inelastic QoS requirements³⁵), data and data applications are typically more elastic (delay tolerant/flexible) and adaptive.³⁶ The movement to packet-based transport mechanisms offered a more flexible platform for dynamic resource allocation, but its implementation in the PSTN still heavily reflected the earlier "circuit-switched" mentality. For example, ATM allows variable bandwidth flows, but supports these in a hierarchical way that allows centralized resource allocation. Indeed, this has turned out to be a key virtue of ATM and explains why it is often used to support QoS-differentiated Internet services. A drawback to this approach, however, is that the capacity is still quantized based on the legacy

this is obviously a obviously an over-simplification because modern voice networks encode voice in variety of ways that require much less than 64Kbps, although 64Kbps channels are often still allocated because of the need to accommodate legacy technology.

³⁴ As we explain further below, this centralized approach becomes much more problematic once the traffic environment becomes more heterogeneous and complex which is precisely what we expect in a wireless environment.

³⁵ As note 32 *supra* explains, even voice is much more "elastic" than many might have you believe. Users of VoIP and mobile phones have demonstrated a willingness to use even very poor quality voice in return for other features (mobility, lower price, etc.).

³⁶ Because "data" implies using computers, the ability to implement software processing is fundamental and makes the adaptivity of data applications something we take for granted. Contrast this to the traditional telephone which lacked flexibility by design. With the computerization of the telephone network, voice telephony becomes more like data and a similar level of adaptiveness (e.g., voice compression, protocol conversion, etc.) can take place to respond to capacity constraints.

requirements of voice telephony (i.e., fixed cells instead of variable length datagrams) and the centralized management inherent in ATM imposes high overhead costs for short duration calls or data communications with relatively long datagrams.

In contrast, the Internet emerged as a packet-based transport overlay on top of the PSTN. It offered a common grade of service ("best effort") that treated all packets the same, and did not constrain packets to a fixed length. The transport of packets is segmented into IP (hop-by-hop) and TCP (end-to-end) control, with BGP providing interdomain routing support. In the Internet, resource sharing is decentralized. There is no "common channel" signaling to allow applications or users to coordinate resource allocation. There is no strong protection in the basic Internet infrastructure to protect different users and applications from interfering with each other as in the PSTN. The applications themselves are expected to address the problems of congestion or noise along the end-to-end communication path through various mechanisms like error concealment or recovery protocols such as the reliability mechanisms in TCP or the FEC capabilities of some VOIP systems.

Although lacking a centralized resource allocation mechanism,³⁷ the Internet does have decentralized resource sharing mechanisms. Today, the dominant mechanism is based on the congestion control algorithms of TCP, and indeed other transport protocols (such as DCCP and XCP) are being designed to support the notion of TCP friendliness, even for relatively inelastic applications such as VOIP.³⁸ In the Internet, end-nodes adjust their data rates to dynamically allocate bottleneck capacity according to the congestion avoidance algorithms in TCP or TCP-friendly protocols. These react to feedback from the network (in the form of packet loss, or in the form of Explicit Congestion Notification or ECN) and use the well known Additive Increase/Multiplicative Decrease control law.³⁹

Of course, as is increasingly apparent even in the wireline Internet, these decentralized sharing algorithms do a poor job when one seeks to integrate inelastic applications with the diverse elastic data applications that have

³⁷ Resources in the Internet are managed by network operators who can redirect traffic by manipulating BGP to implement policy-based routing and by dynamically altering the physical capacity available for Internet traffic (e.g., use ATM to support dynamic bandwidth allocation to virtual IP circuits).

³⁸ Many of these ideas are described at: http://www.icir.org/floyd/tcp_friendly.html

³⁹ This leads to a share for each source of the bottleneck capacity that is inversely proportional to the round trip time seen by each source to its respective destination, and otherwise proportional to the capacity at a shared bottleneck, amongst the n flows that co-exist there.

historically dominated Internet traffic (e.g., asynchronous file transfers). The need to accommodate such things as VoIP and streaming media have driven even wire-based Internet providers to seek improved mechanisms to dynamically allocate resources such as MPLS, DiffServ, or IntServ. However, to date, these have been implemented in a centralized way (i.e., within a single carrier domain) to support QoS differentiated services for enterprise VPNs. Industry is only now developing the frameworks to implement this across interprovider domains.⁴⁰

More typically, ISPs have found it easier to simply over-provision to address capacity constraints.⁴¹ This is due, in part, to the fact that the sort of contention-based mechanisms that are used in TCP and the Internet tend to work best in lightly-loaded networks. VoIP coexists fine with FTP, email, and other data traffic as long as the network is not congested. While centralized resource allocation mechanisms (e.g., ATM or token ring) might provide more assurance of bounded access delays, distributed protocols (e.g., TCP or Ethernet) provide similar delays when the networks are lightly loaded. However, the centralized approaches which may work well for *relatively slow changing source and traffic matrices*, are often less robust in very dynamic situations (as in Ad Hoc networks). In these latter situations, that may more reasonably characterize future wireless environments, a more decentralized/distributed approach may be the only feasible way to manage resources.⁴²

3.2 Rules for managing a Spectrum Commons

Now, let us consider what this suggests about the design of sharing protocols in wireless networks. First, note that much of the licensed spectrum is managed in a centralized way that is analogous to the traditional PSTN (since much of it is allocated to cellular telephone use or else to TV broadcast).⁴³ On the other hand, much of the unlicensed spectrum (e.g. in the ISM band for WiFi and other Wireless LAN and Bluetooth use) is managed in a decentralized way that is analogous to (or indeed is the same thing as) the Internet. In unlicensed, the applications are adaptive, and resource isolation can be less strict.

With respect to licensed spectrum, the rationale for managing resources in a way that is analogous to the PSTN centralized approach is obvious in the case of mobile

telephony (i.e., similar services, similar ancestry, and need to interconnect seamlessly⁴⁴). In the case of broadcast television spectrum, the centralized approach is an artifact of legacy regulation. With respect to mobile telephony, the acute scarcity of spectrum, which itself is largely an artificial artifact of legacy regulation, provides powerful economic incentives to use spectrum efficiently.⁴⁵ Hence, mobile carriers deploy sophisticated centralized resource management techniques to achieve high spectrum reuse (dynamic power management, directional antennas, etc.). In contrast, the current regulatory regime provides little economic incentive for broadcasters to use spectrum efficiently.⁴⁶ Current regulations provide strong interference protection to broadcasters to protect the ability of the worst-case "dumb" receivers to receive the broadcast signal, while constraining the broadcast licensee's ability to benefit from using the spectrum more efficiently.⁴⁷ Consequently, the typical broadcaster uses a minimal set of transmitters to cover their licensed footprint. This architecture uses too much power for receivers that are close, and is vulnerable to even low levels of interference for receivers that are near the edge of the serving area. With smarter receivers, spectral efficiency could be enhanced significantly. Thus, broadcast networks, in contrast to communication networks, might benefit significantly if resource management responsibility were decentralized.⁴⁸

⁴⁴ That is, a key feature of mobile telephony service is the ability to make calls between mobile and fixed-line telephones.

⁴⁵ With a fixed amount of licensed spectrum, a carrier's only way to expand service is to use its spectrum more intensively. Indeed, the artificial constraints may even induce excessive investment in spectrum reuse.

⁴⁶ The opportunity to cram additional digital television channels into their licensed spectrum provides some incentive to increase spectrum use.

⁴⁷ The classic television (or radio) broadcaster is advertising supported and so its business model rewards it for reaching as many receivers as possible. Since the broadcaster does not control the choice of receivers (it is an open system in the sense that receivers and transmitters are designed independently), the broadcaster has an incentive to maximize the reach to the lowest quality receiver.

⁴⁸ That is, since interference occurs at the receiver and the receiver is in the best position to assess its local RF environment, the receiver is in the best position to adapt if it possesses the requisite intelligence. In an environment like the current TV environment in which the receivers do not have a transmitter capability, the current model of granting them interference protection is especially inefficient since it deters incentives for upgrading the receivers. Of course, to the extent that modern media broadcasters become more interactive (pay-per-view, interactive gaming), the receivers will need to become transmitters and the logic for decentralizing resource management are attenuated.

⁴⁰ For example, see [Briscoe05] or <http://cfp.mit.edu/qos> for current work on mechanisms to support interprovider QoS.

⁴¹ [Odlyzko98] has argued that it is easier and less expensive to simply over-provision than to try and implement complex mechanisms to support differentiated QoS in the wired Internet.

⁴² See [Kelly00] and [Reed05] for further discussion.

⁴³ We focus here on commercial spectrum, since the government spectrum is mostly C&C.

In contrast, resource sharing in unlicensed spectrum has been decentralized and more closely follows the model of the Internet. However, the communication environment that prevails in wireless Internet is substantially more complex than in the wired Internet. There are a number of reasons for this. First, the propagation characteristics of wireless signals in real-world environments (with buildings, rain, and lots of unintentional sources of interference) are much more complex than in the typical wired environment where the regularity of the transmission medium may be significantly controlled (both by nature and by design).⁴⁹

Second, wired communications constrain interference to the path followed by the wire,⁵⁰ whereas wireless signals propagate in all directions (in the absence of directional antennas).⁵¹ Thus, in wireless you not only have to deal with the interference to communications that intentionally share the communication path, but also communications that unintentionally share the communication path. That is, wireless devices cause "spectral pollution" that additively increases the noise floor for all users of the spectrum. In the context of wireless, the need to address such unintentional interference, poses an additional challenge that seldom arises in the context of wired communications. For example, wireless operators need regulatory mechanisms to adjudicate interference disputes, whereas wired operators do not.

Third, in part due to the nature of wireless communications and in part due to the more competitive industry structure, there is much more heterogeneity in how radio systems are designed compared to the wired communication systems in the traditional PSTN or Internet.⁵² Current wireless systems -- even for such well-defined services as 2G mobile telephony -- are quite diverse,

⁴⁹ That is, wired media is a controlled environment by its nature, and such media is purposefully designed to minimize many of the problems that arise in wireless (multipath, attenuation effects, etcera). For example, the refractive index of optical fiber and cable shielding are used to manage the propagation environment.

⁵⁰ Cross-talk problems aside. Moreover, electronic wired transmissions can result in additional wireless noise when the cables act as giant antennas.

⁵¹ Employing directional antennas is expensive and keeping them aimed appropriately entails additional cost, especially if mobility needs to be supported.

⁵² Although as noted, in the last decade, as a consequence of the growth of the IP traffic and telecom liberalization which has lead to increased wireline facilities-based competition, there has been a proliferation of next generation telecoms that are competing to offer QoS differentiated services. Thus, even wired infrastructure has become more heterogeneous than in the past (e.g., contrast Level 3 to Verizon network architectures today versus Verizon and SBC a decade ago).

embodying a number of modulation schemes (TDMA, CDMA), antenna designs (omni-directional and directional), and power modulation schemes. Each of these systems impose different constraints and demands on resource sharing. Additionally, the need to support mobility -- which is often a key value-added feature of wireless services -- adds further complexity (i.e., dynamic routing is hard, and especially so in ad hoc networks).⁵³

Fourth, the fact that wireless communications media are inherently lossy poses a special problem for the congestion algorithm used in TCP and the wired Internet. Without appropriate modifications, TCP too often interprets dropped packets to congestion when in fact it is due to poor link quality. Better protocols take advantage of good links when they can (high data rates) and don't try to send when the links are poor (regardless of the congestion condition, no point trying to send when you cannot establish link-level connectivity with the receiver node). Unlike in the wired world, the appropriate sharing protocol needs to be multi-layered in a way that may be avoided in the wired world (where link layer reliability is better assured).⁵⁴

Addressing this complexity in a decentralized, open access environment while allowing flexibility in the choice of technology poses a daunting problem. In the following subsections, we discuss the need/challenges associated with various types of regulatory constraints that we expect to be necessary. As we will discuss, these have both technical and institutional implications. Some may be more appropriately thought of as technical constraints (e.g., power limits) while others may more naturally be thought of as institutional constraints (e.g., enforcement mechanisms). The types of constraints we consider include:

- (1) Requirement that all transmitters also have receiver capability (to enable adaptation)
- (2) Power restrictions (to limit interference)
- (3) Signaling capability (to enable unlicensed devices to communicate resource needs)
- (4) Contention/allocation mechanism (allocate resources in congestion)
- (5) Enforcement mechanisms (to enforce compliance with other rules)
- (6) Reversibility (to allow resources to be de-allocated and deployed for other uses)
- (7) Security and Privacy protection (to ensure mechanism respects these needs)

⁵³ Supporting mobility in the wired Internet is also challenging.

⁵⁴ On the need for multi-layered protocols for ad hoc wireless networks, see [Kawad04].

3.2.1 No Transmit only devices, receive-only devices have no interference protection rights⁵⁵

Requiring every potential transmitter also to have a receiver capability⁵⁶ creates the potential for a feedback control loop that would allow the transmitter's behavior to be modified. Since the transmitters are the source of the signal that creates the possibility for interference, regulating their operation is an essential element in effective interference management.

Additionally, a key motivation for this rule is to allow the transmitters to be preemptible which would be a requirement for operating as an overlay (i.e., the primary licensed user gets to pre-empt unlicensed use).⁵⁷ Strong assurance of a reliable preemption capability would likely be a pre-requisite for sharing most government spectrum, especially, the spectrum allocated for primary use by public safety and emergency systems. Furthermore, the ability to be preempted would prove useful in dynamic reallocation of the spectrum (e.g., shutting off all the unlicensed transmitters if the Commons Protocol was to be changed, or if the spectrum was to be reallocated to another use).

Building the capability for real-time adaptability is essential in the context of a decentralized management framework. Lacking a centralized controller, devices will need to be able to sense their environment or communicate with each other to implement an optimal sharing mechanism. Otherwise, how would a transmitter learn that it was interfering with another receiver?

Building in a receiver capability is also essential to implement an ad hoc mesh. For nodes to act as repeaters or routers, they need to be able to both receive and transmit. Although acting in this way should not be a requirement of every node that wishes to operate in unlicensed spectrum,⁵⁸ the requirement that every transmitter have a receiver capability is consistent with system designs that encourage such cooperation. This is an added benefit, but not the basis for requiring the rule.

Complementing the above requirement, and recognizing that not every receiver will need to transmit, there should be no interference protection for receivers that do not have the ability to transmit. For example, there may

⁵⁵ We would like to credit David Reed and Andy Lippman of the MIT Media Lab for first suggesting the need for a rule like this. Actually, they argue that there should be no "receive only devices."

⁵⁶ The receive capability does not have to be implemented wirelessly. The upstream link could be provided via a wired network.

⁵⁷ Note, "listen-before-talking" offers one protocol strategy for pre-emption. However, this may not ensure a strong preemption capability (e.g., as might be enabled with a "kill code").

⁵⁸ That is, as we explain further below, the regulatory framework should not require that nodes cooperate.

be sensors that cannot transmit. Receivers that have no ability to transmit cannot let anyone know if they are being interfered with and so eliminate the possibility of dynamic adjustment. As long as they are willing to live with whatever interference they encounter, such devices may be allowed since they do not add to the interference experienced by others. Conversely, were such "receive-only" devices granted an entitlement to interference protection, this could not be readily implemented in a decentralized fashion. Potential transmitters would not know that the receive-only devices were there to be interfered with.⁵⁹

3.2.2 Power restrictions

In a wireless system, transmission power is probably the single most important design variable. The power of a wireless device determines the range of communication that can be achieved, and the geographic scope of any interference that transmissions may cause. Power management is also critical factor for mobile devices to maximize battery life. Power management also impacts other design characteristics of the radio system. For example, in a given part of the spectrum, different nodes can use some Media Access Protocol once they have agreed upon a common coding and multiplexing technique. However, the choice of coding and multiplexing itself (and the available antennae) are constrained by the available power budget. For example, Code Division Multiple Access (CDMA) requires much more constrained power to avoid distance effects with interference, whereas Time Division Multiple Access (TDMA) protocols can be sloppier about power allocation (unless a cell or region is adjacent to one that is using the same part of the spectrum for CDMA). Moreover, we should not expect to be able to limit unlicensed devices to just one of these choices. In the Software defined Radio Forum, even the lowest tier (tier 1) of radios (including some existing available cellular phones, e.g. quad band) can choose between CDMA and TDMA so the problem is endemic. Furthermore, if the radio is frequency agile, there is the relationship between antennae, frequency and power which must be addressed.⁶⁰ Consequently, it is hardly surprising that much of the research on models for managing shared access to unlicensed has focused on protocols for power management.⁶¹

⁵⁹ An exception that might be acceptable would be some spectrum cut-outs to address legacy users who may not be able to be efficiently relocated in spectrum that may be newly allocated for unlicensed use (e.g., satellite downlinks).

⁶⁰ That is, the higher the frequency, the more power that is required to cover an area with omni-directional antennae, but the less with dish/directional antennae (see, for example, [VanMatHaak98]).

⁶¹ See, for example, [SataPeha00], [Peha97], or [Kawad04].

The aggregate power from all the transmitters in a particular portion of the spectrum determine the level of interference (noise) realized at different points in space. In the worst case scenario of an ad hoc network, wherein we assume that nodes do not know *a priori* where other nodes are, the transmissions from all other nodes except the one that a receiving node is communicating with will be seen as noise. To maximize the capacity of such a system, it will generally be advisable to limit power and rely on multi-hop transmissions to communicate.⁶² Contrast this to the case for a closed system where devices may use knowledge of the spatial, temporal, and frequency distribution of signals to separate signals (e.g., using MIMO and MUD techniques), thereby rendering what might otherwise be considered as "noise," additional information. With suitable incentives for such "strong" cooperation, the devices may realize significant cooperation gain ([Reed02]), however, this assumes a level of cooperation that may be difficult to achieve and inappropriate to enforce in the context of the regulatory mechanism.⁶³

However, the optimal power limit for an individual device even in a mesh network depends on the node density. If nodes are sparse, you need higher power even to reach your nearest neighbor. Thus, extremely tight (per device) power limits may limit communications in unlicensed. As we argued earlier, since we expect unlicensed to be more likely to be favored (relative to licensed) in situations where congestion is less of a problem, restricting device power overly to maximize capacity and to limit interference seems inappropriate. This, however, is the approach that has been adopted for underlays (to protect interference to the primary users).

Rather than limiting the power of individual devices, it would be better to limit the aggregate power of all the transmitters together. That is, one use may be two devices that want to communicate over a relatively large area and need higher power to communicate; whereas a group of other devices may be organized into a mesh that uses multihops to provide connectivity. Each of these will appear as noise to the other, so a limit on the aggregate system power would be more appropriate. This, in effect, is the idea behind defining an interference temperature. In this approach, the interference temperature is measured as the

⁶² [XieKumar04] use information theoretic analysis to show that "decode and forward" is order optimal for network capacity.

⁶³ That is, any closed system in the unlicensed network may choose to adopt a cooperative approach. However the regulatory framework should eschew mandating that devices cooperate except in some minimalist fashion. For example, it seems likely that any rules that specified that devices implement MIMO or MUD would need to enforce a level of cooperation that might make it more difficult (costly) to introduce new technologies or may conflict with promoting competition.

local noise level arising from all of the transmissions that are received at that locale.

Interference temperature was proposed as an approach for measuring the local noise floor, and hence, as a better way to regulate signal power over a transmitter's footprint.⁶⁴ The interference temperature could be used to set limits on the maximum level of interference a receiver could expect to tolerate. With limits on the interference temperature allowed in a region, a transmitter could determine whether it could transmit or not based on how its transmissions would incrementally change this temperature *if* it could observe the interference temperature dynamically and has a good model of the interference implications of its transmissions. Simply observing the local interference conditions at the transmitter would likely not be sufficient.⁶⁵ While there are obviously numerous problems associated with implementing an interference temperature approach, the opposition to this proposal was also motivated by the recognition that such an approach, if adopted successfully, would permit further constraints on the use of exclusive use licensed spectrum (e.g., by facilitating the implementation of secondary use underlay easements).

Another problem arises in the context of mixed power uses of unlicensed spectrum. As suggested above, there may be multiple, incompatible systems seeking to share the unlicensed spectrum. Indeed, enabling such a scenario is precisely the intent of promoting the unlicensed model (as discussed earlier). When devices with mixed power share spectrum, the hidden node problem arises (i.e., high power devices interfere with low power devices that the high power devices fail to see). In the absence of a mechanism to address this problem (e.g., a signaling mechanism), it may be preferable to limit the interference rights of low power devices. That is, if they cannot operate because admissible higher power devices do not see them, that is their problem.

Finally, in part due to NIMBY concerns, the opposition to UWB and attempts to develop an interference temperature approach to enable underlays has merit. Primary uses, especially those that are willing to pay for an exclusive license and the added interference protection it implies, have a valid fear that spectral pollution will raise the noise floor. Moreover, the current use of primary licensed spectrum may not provide an accurate measure of where the floor should be since primary uses may extend their ability to operate in the noise floor (introduce more sensitive receivers) and may themselves raise system noise as a consequence of their own usage growth. On the other hand, if the power limits are too strict, that will severely

⁶⁴ See [FCC03a].

⁶⁵ This may not be true. Rapid estimates of the number of other nodes or the level of interference might be possible using techniques from large deviation theory as were used to address the problem of measurement based admission control for Internet traffic which is self similar or otherwise heavy-tailed. See, for example,

limit the flexibility of how the unlicensed spectrum is used (range of devices, applications that can be supported). While unlicensed devices are likely to be power limited even in dedicated spectrum, the per device limits that could be used and the aggregate flux limits could be higher than is likely to be practical in an underlay context. This suggests that the unlicensed Commons may operate best if implemented in dedicated unlicensed spectrum.

3.2.3 Signaling capability

A key attribute of ad hoc networks is that they will be decentralized and devices will have access to only local information. A signaling capability would enable devices to acquire and share information on wider-area concerns. Devices would signal their intent and aspects of their mode of operation in a particular region, thereby letting other devices know their whereabouts and operating parameters.⁶⁶

This could play an important role in a number of contexts. First, it could help address the mixed power problem. Second, a signaling capability also may help address the mixed power problem (see above). Third, it could be used to simplify coordination and to share global information that could prove useful in interference management (e.g., regarding the interference temperature at different points in space). Fourth, it could be useful in supporting enforcement of the management Protocol, including implementing a payment mechanism. For example, it could be useful in determining the identity of devices. Fifth, it could be used to provide location-based services and context-dependent resource allocation adjustments.

It may be best to implement the signaling capability in an out-of-band common signaling channel.⁶⁷ This might simplify coordination. Additionally, it would allow faster responses and adjustments (e.g., a control channel to induce collective responses such as preemption). Moreover, since wireless is inherently a broadcast medium (in contrast to wired communications), multi-destination protocols are inherently very efficient. Informing other users of ones' intent and even getting consensus, while difficult to achieve atomically, is far more efficient than in a multi-party wireline collaboration.

Mandating a signaling capability is problematic, however, since it poses a threat to privacy. Although we might presume intra-group trust might exist among a mesh of users, it seems ill-advised to presume that inter-group trust will exist among all devices sharing the unlicensed

spectrum. The privacy issues raised will depend critically on what is "signaled." If devices signal their identity, then the privacy concerns are likely to be great. If the devices only signal their power, then privacy might be protected and there may still be sufficient global information available to manage interference. Lack of an ability to identify devices, however, may limit opportunities for enforcement.⁶⁸

Because of the privacy concerns, it seems desirable to limit any regulatory requirement for signaling to power.⁶⁹ Precisely what about power needs to be signaled still needs to be determined.

3.2.4 Contention/allocation mechanism

When congestion occurs, there will need to be rules to address how resources should be allocated. To allow maximum technical flexibility, it seems desirable that the regulations be cautious in specifying any particular mechanism. For example, it would clearly be wrong to mandate use of the TCP congestion management approach for the reasons cited above, however, it may be desirable to mandate that a "TCP-like" or "TCP-friendly" protocol be implemented to enable contention management. (This has the added advantage of providing easier integration with the wired Internet.)

[Friedman05] suggests that an appropriate approach may be to specify "manners" instead of a specific protocol, wherein the appropriate rule would allocate power using a Wireless Fair Share (WFS) rule that is analogous to the fair-share algorithm used in TCP in the Internet. This appears consistent with fair queuing and the ECN-based feedback mechanism proposed by [Briscoe05]. The "manners"-approach offers a way to generalize the rule (allowing flexibility), while still capturing the benefits of structure and the efficiency from fair queuing. This contrasts with a rule that specified a particular sharing algorithm.

Furthermore, for many applications that have elastic requirements (are delay tolerant), the slow adaptation process of in-band management mechanisms is likely to be fine. However, if the spectrum is also to be shared with delay-intolerant (non-rate-adaptive) applications, then there will also be a need for an out-of-band control mechanism. The out-of-band signaling channel would help in implementing this.

3.2.5 Enforcement mechanisms

A critical component of an appropriate Protocol for a Commons is the enforcement mechanism. There are two levels at which these need to operate. First, the absence of

⁶⁶ Of course, the low power devices would have to be sufficiently high power to use the signaling channel protocol, and receivers would have to be sufficiently

⁶⁷ That is, "out-of-band" in the sense that the signaling takes place in a way that is reserved for signaling and is common to all devices sharing the spectrum. It could be "in-band" in the sense that it shares the spectrum used by the unlicensed devices.

⁶⁸ Enforcement mechanisms that are embedded in the protocol (e.g., non-complying devices are collectively spammed by the other devices) may still be used. But, enforcement based on reputation effects may not be applicable.

⁶⁹ Note this does not limit the amount of information that devices may voluntarily share in the interests of supporting enhanced services.

strong interference protection in the Commons signals to devices that it is a "Jungle environment" and that the best way to survive in such an environment is for devices to be robust (adaptive) to congestion in their local environment. To the extent unlicensed devices are encouraged to become more robust (i.e., not rely on the strong interference protection offered to licensed users), the need for enforcement is reduced. Thus, self-reliance is a substitute for external enforcement.

Second, the technical protocol should implement decentralized enforcement mechanisms. That is, the protocol should incorporate mechanisms to make appropriate behavior (spectrum efficient, cooperative behavior) incentive compatible and self-enforcing.

Third, in the event that users deviate from the protocol (e.g., by using non-complying devices or by altering the code for their software radios), there needs to be an appropriate legal/regulatory institutional framework to resolve conflicts and discourage such "illegal" behavior. The existence of both such mechanisms reduces the burden on either.

Thus, for example, the protocol can specify a game (i.e., consisting of a well-defined set of admissible strategies, sequence of moves, and payoffs for different strategies) that has many desirable equilibria (so it is not constraining) but is so structured as to preclude undesirable equilibria.⁷⁰ There is already a large literature of protocol games that induce cooperative equilibria.⁷¹ Earlier, we discussed the desirability of requiring any admissible protocol to demonstrate that it is "TCP-friendly," while eschewing specifying that the protocol be required to implement any specific game.

Of course, even if you have restricted the class of games to those which enforce desirable cooperative behavior, you have to make sure that devices play the designated game. If devices are free to choose which game they play, they may be tempted and able to exploit the expectation of cooperation opportunistically. While this is true, it is not up to the protocol to ensure that everyone complies. Extra-regulatory mechanisms can help resolve that issue. For example, Part 15 behavior is enforced via a mandatory certification process. While it is possible for end-users to modify equipment after the fact, the threat of legal sanctions (fines, prison) is sufficient to induce general compliance. Nevertheless, the decentralized commons does raise important enforcement issues that need to be discussed

⁷⁰ For example, [Neel05] recommend using the theory of supermodular games to constrain admissible sharing protocols for distributed management of wireless networks. These are attractive because they have a number of desirable properties, including a stable steady-state equilibrium which allows one to ensure convergence.

⁷¹ See, for example, [Friedman05], [JohTs03], [SataPeha00].

further. These are addressed in the following two sub-sections.

3.2.5.1 *Technical protocol enforcement mechanisms*

One of the challenges posed for management of the Commons is that enforcement is a public good like national defense. If it is costly to detect and punish deviant behavior, then nodes have an incentive to free-ride on the enforcement efforts of others. Therefore, protocols that minimize the individual costs of enforcement can make it more likely that nodes will comply. One way to do this might be to extend the idea of notional pricing protocols to compensate nodes for enforcement (e.g., by granting "enforcing" nodes preferential forwarding treatment).⁷² Additionally, if spectrum is relatively abundant so that cooperation may not be a zero or negative sum game, relatively light-weight mechanisms may be sufficient to enforce cooperation (e.g., akin to what is used in BitTorrent and other "incentive-compatible" peer-to-peer systems). This is a rich area for further exploration.

3.2.5.2 *Institutional enforcement mechanisms*

In addition to the intrinsic enforcement mechanisms, there will need to be external, institutional enforcement mechanisms. The most obvious approach here is to rely on a "Part 15"-like certification process. To minimize the regulatory intrusion, the certification authority ought to be delegated to an industry association or some other quasi-private entity, rather than a government agency like the FCC.

The likelihood that software radios will play an important role in future unlicensed systems raises additional problems for how to enforce the certification rules. Hardware certification works relatively well because it is relatively difficult for the typical user to modify hardware and because there are a relatively few number of hardware manufacturers (so that government oversight of their actions and the assignment of liability is relatively easy). In contrast, certifying software is much more difficult because there are likely to be many more versions of software to certify and because it can be changed much more easily after the fact.⁷³

In the noisy, stochastic wireless environment, it will be hard to detect and enforce violations ([Goodman04]). While true, this problem is unavoidable if we wish to promote environments that support ad hoc wireless operation (i.e., anonymous, random entry/exit of nodes, and hence, mirror the ideal of competitive markets).

⁷² See [Crow03].

⁷³ These features lead [Faulhaber05] to be skeptical as to the viability of cognitive radios in practical consequences. Indeed, at a conference in May 2005, Professor Faulhaber referred to cognitive radios as the "Kalashnikov rifle" of the wireless future because of the threat they pose for "hit-and-run" violations.

Furthermore, although determined criminals may not be deterred by legal sanctions that are hard to enforce, the average user, device or application designer/manufacturer is likely to be deterred and hence legal sanctions (fines for deviations) are likely to be broadly effective. This is analogous to the issue of compliance with pollution control in automobiles: although it is relatively easy to modify a car's catalytic converter to enhance mileage performance, most car drivers comply with the pollution abatement standards. Similarly, while it is possible to modify implementations of TCP to behave "greedily" and so enhance individual node performance in the wireline Internet, most users do not do this.

Several strategies may prove useful for supporting the certification of software radios. First, bright-line rules to define that which must be certified and that which is open to user modification will be useful. Thus, the operating system for the software radio may be certified, but the applications that are enabled by that operating system may not need to be. Limiting the functionality of the operating system would offer a mechanism for controlling system behavior. For example, the operating system may be allowed to support modulation selection and dynamic power adjustment, but not frequency agility, which may still require hardware certification.⁷⁴

Second, the protocols that implement the Commons ought to be designed to support compliance testing. For example, self-certifying code⁷⁵ or restrictions on the structure of any games implemented by the protocol (as suggested by [Neel05]) can help make it easier to certify whether a device or implementation of the Protocol is compliant.

Third, alternatively, it may be possible to have a registry of pre-defined and tested behaviors and have software selection (rather than definition) of the appropriate algorithms. Such an approach would shift conformance tests off-line and is similar to the use of signed code. A more sophisticated intermediate approach might be to employ proof-carrying behavioral codes.

Fourth, to ease detection, the enforcement mechanism may benefit from employing witnesses. In wireless, triangulation can be used to locate misbehaving devices and kick them off the network or shut them down (e.g., by invoking the capability that makes the Protocol interruptible). Access to a signaling capability can facilitate enforcement by providing a mechanism for distributed reporting.

To be consistent with enforcement while respecting privacy needs, there must be secure identity mechanisms,

⁷⁴ We are not recommending such a limitation but rather suggesting the type of things that might be done. Indeed, this is close to the approach that the FCC has adopted.

⁷⁵ [Greaves03] suggests how application interfaces may be designed using "pebbles" to make them self-certifying.

and a variety of algorithms to assess, distribute and combine reports on node trust and reputation. This is essential for decentralized enforcement. The alternative is a centralized approach. The centralized approach will work for centralized services, but many of the unlicensed spectrum uses will have inherently decentralized applications, and users may not wish to submit to centralized enforcement. There are a number of unsolved problems in decentralized enforcement. At the very least, fully decentralized currency (or tokens of trust) and identity mechanisms are not solved problems, nor are all of the actual incentive alignment problems ([FeigShen02]).

An important aspect of mechanism design is the ability to be able to negotiate between agents (in our case cognitive radios, or their agent, the provider) without revealing unnecessary information. Luckily, progress has been made on how this might be accomplished (which also has promise for resolving some of the negotiation issues arising around settlements and QoS for wired ISPs and telecoms). For example, this includes solutions to the so-called Millionaires' paradox (two agents find out which is richer without finding out actually how much either of them have). In general, then, a set of nodes could negotiate a spectrum share at a given power level, with another set of nodes, without necessarily revealing their location or identity. Although designing the actual protocols to do this would be non-trivial, the progress in work on zero-knowledge protocols suggests this should be doable ([CanGold96]).

Finally, it is worth spending a little time thinking about the legal/regulatory enforcement apparatus that might be employed. A number of legal scholars have addressed this question, focusing on the types of property rights that might be assigned to unlicensed devices and the implications of this for enforcement costs.⁷⁶ Since incentives to enforce compliance are generally weaker in a Commons (i.e., because there is no licensee who internalizes both the costs and benefits of enforcement), it is better if the need to enforce is less. Fortunately, as we have discussed, this is likely to be the case and techniques exist to reduce the costs of enforcement in distributed environments. However, when external enforcement is required, it does not seem advisable to expect individual users to seek Court or even regulatory-agency remedies. Rather, this is likely to be the role for either the central regulatory agency, or its delegated authority which may be able to seek liability damages from manufacturers of infringing equipment. For this reason, the liability rules used to assign responsibility for controlling deviations will be critical.

In most cases, it may be unreasonable to believe that the node operators or infrastructure owners will be able to adequately control compliance and so assigning liability to the users of unlicensed devices may be inappropriate in

⁷⁶ For example, see [Goodman04], [Benjam03], [Werbach02].

many contexts.⁷⁷ Indeed, the current certification approach assigns principal liability to device manufacturers. Of course this might differ depending on the context (e.g., consider the case of a municipal wireless network as opposed to an anonymous ad hoc network).

3.2.6 Reversibility

A critical aspect of any mechanism used to manage spectrum is that it must be possible to reclaim the spectrum if it would be better used for another purpose. This has both a short-term and long-term aspect. In the short-term it must be possible to both allocate and de-allocate resources. In the longer-term, we need to be able to modify the protocol used to manage the Commons or potentially reallocate the spectrum as Licensed if that is desired. We need to build into the Protocol lease mechanisms to allow devices to be de-allocated spectrum (turned off).

In the short-term, there are already mechanisms for managing such access at a higher level. For example, WiFi operators already do this by DHCP address lease expiry, as do many (e.g. pay-as-you-go) cellular telephone operators. The same mechanisms that allow the spectrum to be preemptible may also be used to de-allocate spectrum. For rapid dynamic reallocation, a key challenge for unlicensed arises because many of the decentralized protocols adapt rather slowly. Thus it will be important to provide for a control plane for such allocation/re-allocation mechanisms.

With licensed spectrum, active secondary markets provide a valid mechanism to recycle spectrum when its value in its current use is less than the opportunity cost of the spectrum. In principle, these secondary markets for licensed spectrum could operate both in the short (real-time markets) and long term, although today only the latter exist for wholesale trade.⁷⁸ In contrast, unlicensed spectrum lacks such an obvious mechanism. In the longer term, we need a way to ensure that unlicensed spectrum could be cleared. Such a mechanism is necessary in case the spectrum is polluted by too many garbage applications. Although a key point of unlicensed spectrum is to permit experimentation, there needs to be a way to clear out bad experiments and low value uses if it turns out that spectrum is, in fact, relatively scarce. Applications that produce less surplus than the opportunity cost associated with the spectrum may be considered garbage applications (although, as noted earlier, we need to be sure not to classify as garbage valuable

applications that fail to generate appropriable or easily measured benefits).

Another option may be to impose clear term limits which specify when the spectrum will be reviewed. While superficially attractive, such term limits are problematic for the efficient working of markets. Term limits may deter investment in infrastructure that has an economic life that does not match the remaining term of the spectrum. This can give rise to a potential hold-up problem when the spectrum is up for review, or equivalently, a regulatory commitment problem if the policymaker is induced to protect incumbent users even against more efficient new uses for the spectrum. The problems with de-allocating broadcast television spectrum following the conversion to digital broadcasting demonstrates that this is a real risk.

To the extent the radio infrastructure is frequency agile, it will be less co-specialized with any particular spectrum and so may be able to relocate if the spectrum it had been using is de-allocated at the end of the term. While this might work in an ideal world, there are likely to be non-negligible costs associated with migrating to another spectrum band. Consequently, most property rights advocates oppose term limits for licensed spectrum, believing that such limits are unnecessary when there is a market.⁷⁹ For a Commons, however, the market-mechanism of simply transferring the exclusive use rights to a new licensee for a monetary payment would not work. Therefore, a regulatory commitment to review Commons usage periodically (i.e., term limits) is likely to be important. A clear ex ante commitment to such a policy will help reduce any distortionary effects of the term limits for unlicensed use. Moreover, at least with respect to ad hoc applications, we may be comforted that most of the infrastructure is relatively short-lived which means that term limits may be less disruptive than they might be in another context (e.g., radio systems with lots of capital-intensive, long-lived antenna sites).

3.2.7 Security and Privacy protection

The final class of rules we expect to be needed are ones to address privacy and security concerns. Within a Licensed model, it may be possible to centralize responsibility protecting security and privacy by assigning liability to the licensee. The licensee can, in turn, contract with users to implement appropriate approaches.

In an ad hoc, distributed/decentralized environment, protecting security and privacy raises additional challenges ([MackCor99]). Ad hoc nodes may be especially susceptible to denial of service attacks, spoofing, and sniffing that will challenge efforts to generate distributed trust models. Also, any use of device identity signaling which would simplify enforcement of the

⁷⁷ This does not mean that willful or intentional misuse of an unlicensed device should be free of legal sanctions, only that users who use the device in its original configuration (as certified) should be able to expect safe harbor relief from interference claims.

⁷⁸ At the retail level, mobile networks are continuously allocating and de-allocating spectrum to support different mobile callers. This is relatively easy in a centralized network management context.

⁷⁹ For example, Faulhaber ([Faulhab05]) and Hazlett ([Hazlett01]) oppose term limits, however Baumol ([Baumol05]) supports the use of term limits.

protocol, may pose an unacceptable risk for privacy. To protect privacy, it will be desirable to limit any signaling to the minimal information necessary to implement the sharing mechanism (e.g., limit signaling to power, but not device identity). In the end, both a secure out-of-band control channel (analogous to the signaling channel in Q931 or SS7) will be needed as well as the in-band, decentralized controls, since the capability and trust in alternative distributed privacy mechanisms may take many years to develop and deploy.

Thus, we advocate power control as the primal resource concern. We propose a hybrid of a control channel (e.g. licensed, protected and secure) as well as decentralized control for power, location identity to be expressed, and spectrum rights/grant responses. In such a situation, security and privacy concerns will be reflected in the users' choice of spectrum access. In a more trusting community, a set of users may select an ad hoc system dependent on mutual interest for good behavior, while in a more competitive context, users may call for a more hardened approach to security. Thus, privacy considerations also support a diversification of spectrum control mechanisms and policy.

4 CONCLUSIONS AND FUTURE RESEARCH

The legacy of command & control (C&C) spectrum management regulation has hampered the growth of wireless services and the deployment of new technologies. A consensus has emerged regarding the desirability of reforming spectrum management to rely more on market forces. There are two principal models for implementing such reform: property rights/licensed/exclusive use ("Licensed") and open access/unlicensed/commons ("Commons"). This paper examines how a Spectrum Commons might be managed effectively.

As wireless traffic become more like Internet traffic (heterogeneous, bursty or fat-tailed, long hold time for connectivity but variable link status because of ad hoc networking), we will need new ways to manage wireless resources. The traditional and time-tested hierarchical, centralized approaches to network management will need to be supplemented with, and in some cases, replaced by more distributed/decentralized network management techniques. This will require new Protocols both at the level of running code (protocols and standards) and at the level of institutional frameworks (spectrum management regimes). There are two paradigmatic approaches to decentralizing spectrum management to rely less on centralized regulatory authority (so-called, command and control) and more on markets: licensed and unlicensed spectrum.

The Protocols that are needed need to be more dynamically adaptive and responsive to local information. Following David Reed,⁸⁰ there needs to be a *phase shift* in

how portions of the RF spectrum are managed: from the *solid* approach of C&C which presumes a lot of long-range static order in the wireless ecosystem (predictable link quality, demand loads, and routing) to the more *liquid* approach of a commons (ad hoc networks requiring locally very dynamic demand and routing adaptation).

Policy-makers are committing to a dual regime of flexible licensed and commons/unlicensed spectrum to provide for the evolution from C&C to these more decentralized management regimes. While economists well understand how an ideal licensed regime based on decentralization of residual control over spectrum use (and technology choice) to markets (via tradable property rights to exclusive use) ought to work, there is less clear understanding of how a spectrum commons might work. This latter is closer to a barter or social market, than to the commodity goods market that best approximates the competitive ideal. This paper addresses this gap by focusing on the elements for a Protocol for managing a Commons.

Our analysis suggests a set of design rules, if you will, for evaluating alternative technical and institutional proposals for creating an appropriate management Protocol for a Spectrum Commons. This framework involves a characterization of the environment in which we expect the Protocol to operate (Table 1A), the characteristics that will distinguish a good mechanism (Table 1B), and our preliminary thoughts on the kinds of rules that will be needed (Table 1C). While much remains to be done to flesh out the Commons Protocol, this framework should prove useful in evaluating alternative proposals and rules.

There are multiple pressures on the design space for spectrum management. And, there are a number of different tools we can bring to bear. From the technical side, we have engineering control theory, information theory, computer science (complexity, computability, proof systems), and mathematics (modeling). From economics, political science and legal studies we have the theory of the firm and markets, asymmetric information games, mechanism design, and institutional theory (common resource management, institutional enforcement mechanisms, and pricing theory). Since the evaluation of Commons regimes will be inherently multidisciplinary, we will need to use all of these tools.

Proposed rules will need to be examined for their expected performance with respect to technical efficiency (ability to support ad hoc communications with low overhead but high QoS), industry efficiency (promote competition and innovation), and regulatory efficiency (minimize regulatory costs and distortions). Although we lack a science of design in such a complex space, we have a number of metaphors and design principles (e.g., the end-to-end principal in Internet architecture, efficiency of competitive markets in economics) that allow us to tell when one design is better than another for some reasonable definition of better (e.g., resilient, futureproof, tussel-space-compliant, etc.).

⁸⁰ See [Reed05].

Our design rules lead us to conclude that the appropriate Protocols for a Commons will need to be more *liquid* than in the past: (1) Market-based instead of C&C; (2) Decentralized/distributed (i.e., including Commons'

managed spectrum along with Licensed); and, (3) Adaptive and flexible (Anonymous, distributed, decentralized, and locally responsive).

Table 1A : Past v. Future Environment for Design of Wireless Protocols

Past	Future
Circuit-switched, voice	Packet-switched, data
Fixed QoS, predictable traffic	Variable QoS, unpredictable traffic
Centralized/hierarchical network management of resource sharing	Decentralized/distributed network management of resource sharing
Spectrum scarce, expensive, narrow frequency bands	Spectrum (relatively) abundant, cheap, wider frequency bands
Fixed radio system design (omnidirectional antennas, frequency limited, dumb receivers)	Smart (software) radio systems (steerable antennas, frequency agile, adaptive radios)
Mobile service provider networks	Ad hoc mobile wireless networks
Licensed spectrum (strong interference protection)	Unlicensed Commons (weak interference protection)

Table 2B : Goals for a Successful Commons Management Protocol

Characteristic	Rationales and Features
Promote Innovation	<ul style="list-style-type: none"> Promote innovation in wireless devices, services, and business models. Complement C&C and flexible licensed
Minimize spectrum access costs	<ul style="list-style-type: none"> Minimize entry barriers for new nodes to access spectrum, new technologies to be deployed No real-time usage charges for access.
Provide mechanism for congestion management	<ul style="list-style-type: none"> Distributed/decentralized mechanism for managing congestion, coordinating usage. "TCP" friendly techniques
Provide mechanism for modifying/establishing etiquette	<ul style="list-style-type: none"> No tradable licenses to allow market mediation of changing of technical protocol Need structured process (industry standardization) to mediate change.
Promote fairness and non-discriminatory access	<ul style="list-style-type: none"> Promote open access. Distributed/decentralized is inherently fair.
Minimalist regulation	<ul style="list-style-type: none"> Decentralize to market forces. As minimal regulatory constraints as necessary to provide structure. Technology, business model neutral

Table 1C : Taxonomy of Rules for Management of Spectrum Commons

Rule	Rationales and Features
No transmit only devices, no protection for receive only	<ul style="list-style-type: none"> • Transmitters need to be able to receive to provide feedback loop. • Receive-only devices should have no interference protection rights since in decentralized management regime, they have no way to signal their presence.
Power	<ul style="list-style-type: none"> • Single most important design constraint for radio systems • Power rules should anticipate multihop use. • Limits on individual device power and some sort of interference temperature limits (aggregate flux) likely to be needed.
Signalling	<ul style="list-style-type: none"> • Need signaling capability to share global information • Should be provision to support common channel signaling • Will aid in supporting cooperation and enforcement • Will aid in making spectrum preemptible.
Contention	<ul style="list-style-type: none"> • Specify "contention" manners instead of specific protocol. • "TCP-friendly" fair queuing approach seems desirable. • Consistent with ECN approach recommended for Internet.
Enforcement	<ul style="list-style-type: none"> • Protocol should be designed to support incentive compatible cooperation, which includes mechanisms to support enforcement • Legal sanctions and other external institutional frameworks will complement and help support enforcement. • Software radio certification rules will be key element • Liability rules will be central feature.
Reversibility	<ul style="list-style-type: none"> • Mechanisms will be needed to support de-allocation of commons spectrum both in short term and long term. • Preemptibility function will aid short-term reallocation • Long-term reallocation may require term limits for regulatory review of Commons spectrum.
Security/Privacy	<ul style="list-style-type: none"> • Protecting security and privacy will be critical, and ad hoc networks raise additional challenges. • Decentralized privacy mechanisms are needed. • Also, will likely need secure out-of-band control channel. • Signaling of device information should be minimal, perhaps limited to power usage aspects to protect privacy

5 REFERENCES

[BarhamCros98] Barham, Paul, Simon Crosby, Tim Granger, Neil Stratford, Meriel Huggard and Fergal Toomey (1998), "Measurement Based Resource Allocation for Multimedia Applications," *Multimedia Computing and Networking (MMCN)*, 1998.

[Baumol05] Baumol, William (2005), "Towards Reasonable Rules for Spectrum Usage: Avoiding Ossification for an Unforseeable Future," paper presented to AEI-Brookings Joint Center for Regulatory Studies conference, Washington DC, April 19, 2005.

[Benkler02] Benkler, Yochai (2002), "Some Economics of Wireless Communications," 16 *Harvard Journal of Law and Technology* 25 (2002).

[Benjam03] Benjamin, Stuart Minor (2003), "Spectrum Abundance and the Choice Between Private and Public Control," *New York University Law Review*, vol. 78 (2003) 2007.

[Briscoe05] Briscoe, Bob, Arnaud Jacquet, Carla Di Cairano-Gilfedder, Andrea Soppera and Martin Koyabe (2005), "Policing Congestion Response in an Internetwork using Re-feedback," *Proceedings of ACM SIGCOMM'05*, *Computer Communications Review* 35(4) (Sep 2005).

- [CanGold96] Canetti, Ran Canetti, Uri Feige, Oded Goldreich, and Moni Naor (1996) "Adaptively secure multi-party computation," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 639-648. ACM Press, 1996.
- [Crosby97] Crosby, Simon, Ian Leslie, John Lewis, Raymond Russell, Fergal Toomey and Brian McGurk (1997), "Practical Connection Admission Control for ATM Networks Based on On-line Measurements," *Proceedings IEEE ATM '97*, June 1997, Lisbon.
- [Crow03] Crowcroft, Jon, Richard Gibbens, Frank Kelly and Sven Ostring (2003), "Modelling incentives for collaboration in Mobile Ad Hoc Networks," *Proceedings of WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, March 2003.
- [David86] David, Paul (1986) "Narrow Windows, Blind Giants and Angry Orphans: the Dynamics of Systems Rivalries and the Dilemmas of Technology Policy", Technology Innovation Project Working Paper Number 10, Center for Economic Policy Research, Stanford University, March 1986.
- [DavGreen90] David, Paul and Shane Greenstein (1990), "The Economics of Compatibility Standards: an Introduction to Recent Research", *Economics of Innovation and New Technology*, vol 1 (1990) p3-41.
- [Demsetz02] Demsetz, Harold (2002), "Toward a Theory of Property Rights II: the Competition between Private and Collective Ownership," *Journal of Legal Studies*, Vol. XXXI (2), June 2002, S653-672.
- [FankPlatt99] Fankhauser, George and Bernhard Plattner (1999), "Diffserv Bandwidth Brokers as Mini-Markets," *International Workshop on Internet Service Quality Economics (ISQE)*, December 1999
- [FaulFarb02] Faulhaber, Gerald R., and David Farber (2002), "Spectrum Management: Property Rights, Markets, and the Commons," AEI-Brookings Joint Center, Working Paper 02-12 (December 2002).
- [Faulhab05] Faulhaber, Gerald (2005), "The Question of Spectrum: Technology, Management, and Regime Change," paper presented to "Economics, Technology, and Policy of Unlicensed Spectrum," Quello Center, Michigan State University, May 16-17, 2005.
- [FCC03a] FCC (2003a), *Notice of Inquiry and Notice of Proposed Rule-making*, In the Matter of Establishment of an Interference Temperature Metric to Quantify and Manage Interference and to Expand Available Unlicensed Operation in Certain Fixed, Mobile and Satellite Frequency Bands, Federal Communications Commission, ET Docket No. 03-237, November 13, 2003 (hereafter, *FCC Interference NPRM (2003)*)
- [FeigShen02] Feigenbaum, Joan and Scott Shenker (2002), "Distributed Algorithmic Mechanism Design: Recent Results and Future Directions," in *Dial-M '02*, Atlanta, Georgia, USA 2002.
- [Friedman05] Friedman, Eric (2005), "Fair and Robust Power Allocation Rules for Multiple Access Channels," draft mimeo, School of Operation Research and Industrial Engineering, Cornell University, 2005.
- [Goodman04] Goodman, Ellen (2004), "Spectrum Rights in the Telecosm to Come," *San Diego Law Review*, vol 41 (2004) 269-404.
- [Greaves03] Greaves, David and Omitola, Tope (2003), "Towards Autonomous Dependable Systems," paper presented at Cabernet Workshop, Porto Santo, Portugal, November 2003 (available at: <http://www.cl.cam.ac.uk/Research/SRG/HAN/pebbles/towards.pdf>).
- [Hazlett01] Hazlett, Thomas (2001) "The Wireless Craze, The Unlimited Bandwidth Myth, The Spectrum Auction Faux Pas, and the Punchline to Ronald Coase's 'Big Joke': An Essay on Airwave Allocation Policy" *Harvard Journal of Law and Technology* (Spring 2001).
- [Heller98] Heller, Michael (1998), "The Tragedy of the Anticommons: Property in the Transition from Marx to Markets," 111 *Harvard Law Review* 621 (1998).
- [JohTs03] Johari, Ramesh and John Tsitsiklis (2003b), "Network Resource Allocation and the Congestion Game: The Single Link Case," *Massachusetts Institute of Technology*, April 8, 2003.
- [Kawad04] Kawadia, Vikas (2004), "Protocols and Architecture for Wireless Ad Hoc Networks," PhD Dissertation, Department of Electrical Engineering, University of Illinois, Urbana-Champaign, 2004 (adviser: P. R. Kumar).
- [Kelly00] Kelly, Frank (2000), "Models for a self-managed Internet," *Philosophical Transactions of the Royal Society*, A358 (2000) 2335-2348.
- [Kolodzy02] Kolodzy, Paul (2002), *Spectrum Policy Task Force*, Office of Engineering & Technology, Federal Communications Commission, November 2002.
- [KwerWm02] Kwerel, Evan and John Williams (2002), "A Proposal for a Rapid Transition to Market Allocation of Spectrum," OPP Working Paper #38, Federal Communications Commission, November 2002.
- [Lehr04] Lehr, William (2004) "Economic Case for Dedicated Unlicensed Spectrum Below 3GHz," paper prepared for New America Foundation Conference, April 2004.
- [Lehr92] Lehr, William (1992), "Standardization: Understanding the Process", *Journal of the American Society for Information Science*, vol 43, no 8 (September 1992) 550-555.

- [LehrMcK03] Lehr, William and Lee McKnight (2003), "Wireless Internet Access: 3G vs. WiFi?," *Telecommunications Policy*, 27 (2003) 351-370.
- [LehSirGil04] Lehr, William, Marvin Sirbu and Sharon Gillett (2004), "Municipal Wireless Broadband: Policy and Business Implications of Emerging Access Technologies," with, paper prepared for "Competition in Networking: Wireless and Wireline," London Business School, April 13-14, 2004.
- [MackCor99] Macker, J.P. and M.S. Corson (1999), "Mobile Ad Hoc Networking: Routing Protocol Performance Issues and Evaluation Considerations", Internet RFC 2501, Jan 1999
- [Neel05] Neel, James, Rekha Menon, Jeffrey Reed, and Allen MacKenzie (2005), "Using Game Theory to Analyze Physical Layer Cognitive Radio Algorithms," paper presented to "Economics, Technology, and Policy of Unlicensed Spectrum," Quello Center, Michigan State University, May 16-17, 2005.
- [Odlyzko04] Odlyzko, Andrew (2004), "Telecom dogmas and spectrum allocations," mimeo, written for the Wireless Unleashed blog at <http://www.dtc.umn.edu/~odlyzko/doc/networks.html>
- [Odlyzko98] Odlyzko, Andrew (1998), "The economics of the Internet: Utility, utilization, pricing and Quality of Service", mimeo, AT&T Labs, July 7, 1998 (available at: <http://www.dtc.umn.edu/~odlyzko/doc/internet.economics.pdf>).
- [Ofcom05] Ofcom (2005), Spectrum Policy Framework Review, June 28, 2005.
- [Peha97] Peha, Jon and Durga Satapathy (1997), "Spectrum Sharing without Licenses: Opportunities and Dangers," in *Interconnection and the Internet: Selected Papers from the 1996 Telecommunications Research Conference*, G. Rosston and D. Waterman (ed), Lawrence Erlbaum Associates Publishers: Mahwah, NJ, 1997.
- [Reed02] Reed, David (2002), "How Wireless Networks Scale: The Illusion of Spectrum Scarcity," presentation slides to FCC Technology Advisory Council, Washington, DC, April 26, 2002 (available at: <http://www.jacksons.net/tac/Spectrum%20capacity%20myth%20FCC%20TAC.pdf>)
- [Reed05] Reed, David (2005), "Liquid Protocols," presentation slides to MIT-Cambridge Communications Futures Program Bi-annual Meeting, June 29-July 1, 2005, Cambridge, (available at: http://cfp.mit.edu/events/slides/jun05/David-Reed_jun05.pdf)
- [RoyerToh99] Royer, Elizabeth and H. K. Toh (1999), "A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks," mimeo, 1999.
- [SataPeha00] Satapathy, D. P. and J. M. Peha (2000), "A Novel Co-existence Algorithm for Unlicensed Fixed Power Devices," *IEEE Wireless Communications and Networking Conference (WCNC)*, September 2000.
- [WalshDuff97] Walsh, Cormac and Nick Duffield (1997), Predicting QoS Parameters for ATM Traffic using Shape Function Estimation," Proceedings 14th IEE UK Teletraffic Symposium (UKTS 1997), February, 1997.
- [Werbach02] Werbach, Kevin (2002), "SUPERCOMMONS: Toward a Unified Theory of Wireless Communication," paper presented to 31st Annual Telecommunications Policy Research Conference, George Mason University, Arlington, VA, September 19-21.
- [Werbach03] Werbach, Kevin (2003), "Radio Revolution: The Coming Age of Unlicensed Wireless," white paper prepared for New America Foundation, Washington, DC, December 2003.
- [VanMatHaak98] Vanderau, J.M., R.J. Matheson, and E.J. Haakinson, "A Technological Rationale to Use Higher Wireless Frequencies," NTIA Report 98-349 National Telecommunications Information Agency U.S. Department of Commerce, February 1998 (available at: <http://www.its.bldrdoc.gov/pub/ntia-rpt/98-349/>).
- [XieKumar04] Xie, L. and P. Kumar (2004), "A network information theory for wireless communication: Scaling laws and optimal operation," *IEEE Transactions on Information Theory*, vol. 50, no. 5, pp. 748-767, 2004.