

Published in Proceedings of The Telecommunications Policy Research Conference (TPRC), 1996, Section 3, pp. 15-29. Also published in "Interconnection and the Internet : Selected Papers From The 1996 Telecommunications Research Conference," G. Rosston and D. Waterman (Eds.), Mahwah, NJ: Lawrence Erlbaum Associates, 1997, pp. 49-75.

Spectrum Sharing Without Licenses: Opportunities and Dangers¹

Durga P. Satapathy² and Jon M. Peha³

Carnegie Mellon University

Department of Electrical and Computer Engineering

Abstract

The current spectrum management policy typically gives exclusive and unlimited access to license-holders within their domain, and offers meager transmission opportunities for non-license-holders. This paper addresses spectrum management techniques in which no licensing is required and individual devices have real-time access to shared spectrum. An example is the 30 MHz of unlicensed spectrum allocated by the Federal Communications Commission in the new Personal Communications Services band, and industry is already requesting hundreds of MHz more of the same. Such spectrum has several advantages. It eliminates the delays of the licensing process. It facilitates mobility, as a license is not required wherever a system may operate. It also promotes spectrum sharing, as one device may transmit while others in the area are idle. This paper discusses some of the challenges to be overcome. Foremost among these is an inherent Tragedy of the Commons resulting from the fact that device designers lack an incentive to conserve the shared spectrum resource. This phenomenon is quantitatively demonstrated in a practical scenario. Some options for this problem are also discussed.

1 Introduction

The primary method of spectrum allocation today is based on licenses granted by the Federal Communications Commission (FCC) which give recipients *exclusive rights* to spectrum for a limited duration. It has also been suggested that these temporary licenses be replaced with permanent deeds, like property (Pressler, 1996). This too is a spectrum management policy in which users are given exclusive rights to spectrum. Generally, those with exclusive rights to a given block of spectrum have *exclusive access* at all times.

There are wireless applications that cannot be supported efficiently under a system based on permanent exclusive access to spectrum, but would be well served with real-time access to spectrum, even if that spectrum is shared (Peha, 1994-95). Such applications include mobile wireless applications that would desire the ability to access spectrum anywhere within a wide area, but require much smaller coverage at any given time, such as a mobile wireless Local Area Network (LAN), or a wireless Private Branch Exchange (PBX). Granting permanent exclusive access to such a device for any location at which it might ever be operated would be grossly inefficient. Other applications that are poorly served with exclusive access are those that need only sporadic access to spectrum, and can tolerate widely varying access delays, like a wireless electronic mail service. These applications can share spectrum with minimal penalty. For any application, sharing spectrum is inherently more efficient (Peha, 1997; Salgado-Galicia, 1995,1997). For example, consider the case where 8 PBXs are sharing enough spectrum to support 32

¹We gratefully acknowledge the support of Bellcore through Carnegie Mellon University's Information Networking Institute. However, the views expressed in this paper are those of the authors, and do not necessarily reflect the views of Bellcore.

²Doctoral Candidate; Email: d.p.satapathy@ieee.org; <http://www.contrib.andrew.cmu.edu/usr/dsaq/home.html>

³Associate Professor; Address: Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213-3890; Email: peha@ece.cmu.edu; <http://www.ece.cmu.edu/afs/ece/usr/peha/peha.html>.

simultaneous calls. Calls arrive according to a Poisson process, and their duration is exponentially distributed. If they share the spectrum, they can sustain a traffic load of 68.9 % while blocking only 1% of the calls (Peha, 1997), where a call is blocked when no channels are free at the time the call is attempted. In contrast, if each of the 8 PBXs is given exclusive access to 4 channels, the 1% blocking probability is not achievable with a load over 21.7 % (Peha, 1997). Thus, sharing makes it possible to carry over three times as many calls. Finally, there are also novel and rapidly evolving applications and technologies for which experimentation is important. Such applications are poorly served in a system where access to spectrum involves long administrative delays.

Consequently, there is a need for shared spectrum that allows real-time access. Real-time sharing is made difficult by three problems. The first is that, as devices do not have exclusive access, they may interfere with each other's transmissions. To deal with this mutual interference, a set of rules are required which dictate when, where, and how devices may transmit. The second problem is that, since shared spectrum with real-time access would be valuable for a wide variety of applications and devices, there is motivation to create bands supporting diverse applications rather than create many different bands of shared spectrum. These applications may vary greatly in terms of average data rate, transmission duration, or even the technology used. Such variations make it difficult to enforce efficient utilization for all applications. The third problem is that, since spectrum is shared, there is no inherent incentive to use the spectrum efficiently, which may result in a tragedy of the commons (Hardin, 1968; Hundt, 1995; Peha, 1997). This problem made the Citizen Band radio service highly inefficient and undependable in crowded regions, where users wasted spectrum with high-power transmitters. This problem can occur if too many devices are deployed in shared spectrum, or if individual devices waste spectrum. In the former case, the problem is relatively easy to avoid by requiring a fee to be paid for each device deployed, and for a number of reasons, this is a good policy (Peha, 1996). However, in the latter case, the problem is more difficult.

There are three possible approaches to offer real-time access for shared spectrum. One is *unlicensed spectrum*. This spectrum is under FCC control, and any device is allowed to use it - a public park for wireless devices. The other possible approaches are appropriate for either licensed spectrum or spectrum for which there are property rights. In these cases, the profit-driven license-holder or spectrum-owner would demand compensation for the use of spectrum, and there are two ways to do this. The first is usage-based pricing, where each device is charged a fee that depends on how much spectrum it uses (Noam, 1995). The second option is that a one-time fee be charged for each device deployed, independent of how much spectrum the device actually uses (Peha, 1997).

In a system with usage-based pricing, a centralized authority must monitor, control, and regulate usage. Devices must explicitly obtain permission from this centralized authority before transmitting, and where there is conflict, those that are willing to pay more will gain access. As in wired networks like the Internet (Wang, 1996, 1997), this approach has two advantages. First, since fees depend on the amount of spectrum resources consumed, there is an incentive to conserve spectrum. Thus usage-based pricing can prevent a tragedy of the commons. Second, when there is conflict, resources go to applications whose value exceeds the price. However, there are many complications in implementing usage-based pricing for real-time access to shared spectrum. One problem is that it is difficult to avoid mutual interference in a system with centralized control, unless each device has the ability to convey signal measurements to the central controller. Another problem is that supporting a diversity of devices increases the system complexity, since the central controller would need to support a variety of wireless communication interfaces to enable communication with different devices. Most importantly, in order to have mobile applications communicate with the access provider over a wireless link from anywhere in the access provider's region of coverage, a highly complex and massive infrastructure would be needed. This would raise the transaction costs so high that usage-based pricing for spectrum access would be impractical.

The other two schemes for spectrum sharing, unlicensed spectrum and one-time deployment fees to a license-holder, are quite similar. Both employ decentralized control, i.e., both require the devices to follow an *etiquette* to determine when, where and how they may transmit. An example is the etiquette set forth by the FCC for unlicensed operation in the 2 GHz Personal Communication Services (PCS) band

(FCC, 1994; Steer, 1994). An etiquette solves the problem of mutual interference, and it's possible that a well-designed etiquette can help alleviate a tragedy of the commons. The difference between the two schemes lies in the motivation of the entity in control; the FCC is motivated to serve the public good, and a license-holder is motivated by profit. Profit-seeking entities tend to be more efficient, but there are two dangers to be addressed. One is that, since revenues come from devices at deployment time, a profit-seeking entity has less incentive to protect the performance of devices that have already been deployed. It may therefore change the etiquette to favor new devices over those already deployed. The other danger is that a profit-seeking entity that controls a critical resource may engage in anti-competitive behavior. For example, it may exclude devices from some manufacturers in return for suitable compensation from their competitors, or it may overcharge. This danger may be mitigated through competition, but only if multiple identical bands of this type can be created and controlled by different profit-seeking entities. Besides these differences, the two schemes are essentially the same, so we will discuss the problems of real-time sharing with distributed control in the context of unlicensed spectrum for the rest of the paper.

Given the significant advantages and potential problems of real-time sharing in unlicensed spectrum, it is important to evaluate whether allocating additional unlicensed spectrum is justified. The opportunity costs of allocating spectrum for unlicensed use are high, as shown by the billions of dollars fetched by the recent PCS auctions (Pressler, 1996). The FCC has already allocated 30 MHz of unlicensed spectrum in the 2 GHz PCS band (1910-1930 MHz, 2390-2400 MHz), and is now considering additional allocations in the 5 GHz range (FCC, 1996) and at 59-64 GHz (Marcus, 1994, 1996). These allocations would not be justified if the utilization of unlicensed spectrum is likely to be highly inefficient. Although many spread spectrum devices have been developed for unlicensed use in the Industry, Science and Medicine (ISM) bands (902-908 MHz, 2.4-2.48 GHz and 5.725-5.85 GHz), it remains to be seen whether or not a diverse group of unlicensed systems can coexist efficiently, even with an etiquette. It is therefore important to evaluate whether the provisions present in the etiquette are sufficient to prevent the tragedy of the commons. It is equally important to evaluate whether these provisions are necessary, since an overly restrictive etiquette can both reduce the spectral efficiency and increase the cost of unlicensed devices. In this paper, we will demonstrate that there is a potential risk of the tragedy of the commons occurring in unlicensed spectrum, and suggest possible techniques to avoid the problem through modifications to the spectrum etiquette if required.

In Section 2, we discuss proposals for real-time access to shared spectrum with distributed control that have been previously suggested. Section 3 discusses the strategies that designers may adopt for unlicensed devices that might result in a tragedy of the commons. In Section 4, we demonstrate the potential risk of a tragedy of the commons through analysis and simulation. In Section 5, we suggest etiquette modifications that may be used to deal with a tragedy of the commons if required. Finally, we present our conclusions in Section 6.

2 Previous Proposals for Real-time Sharing with Distributed Control

In this section, we will look at four previous proposals for real-time sharing with distributed control, and examine their potential to avoid a tragedy of the commons. The most far-reaching plan is based on the premise that emerging technology will eliminate spectrum scarcity, thereby eliminating both the need for licenses for exclusive access, and the potential for a tragedy of the commons. Instead of getting licenses, users will access spectrum through frequency-nimble devices that allow them to find spectrum as needed without having a band exclusively reserved for them (Gilder, 1994). This approach is inappropriate, at least at this time, for two reasons. One is that the technology required for this plan to work is still too expensive (McGarty, 1994), and the other is that even though technical innovations are constantly increasing the effective availability of spectrum, demand for spectrum is also increasing rapidly. There are no signs that supply will surpass demand.

Spectrum sharing with real-time access is possible without such radical change. There are existing unlicensed bands within the current licensed system. The Industrial, Scientific and Medical (ISM) bands are the simplest, where unlicensed devices are allowed secondary access, and they have to share the bands

with licensed devices that have primary access. Unlicensed devices must not cause harmful interference to licensed devices, and must accept any interference from licensed devices as well as other unlicensed devices. Unlicensed operation in this band is restricted to devices using spread spectrum modulation only, and the maximum transmission power of these devices is restricted to 1 Watt. The low signal energy density of spread spectrum modulation and the limit on power reduce interference to some extent, but devices are always at risk of unmitigated interference from licensed devices and unavoidable interference from unlicensed devices. There is obviously no scope for diversity, since only spread spectrum devices are allowed access. The only provision to deal with the risk of the tragedy of the commons is an upper limit on power, which is a solution with limitations. Although an upper limit prevents excessive use beyond the limit, it provides no incentive to use only what is necessary. Devices may therefore choose to always transmit at the maximum power to achieve a high signal-to-noise ratio (SNR), even if lesser power would be adequate. In this case the power limit is low. This solves some problems, but creates inefficiency for other reasons, since devices cannot transmit at greater power even when no interference or spectrum contention would result.

The other form of unlicensed band involves the use of an etiquette, as used in the 2GHz PCS band. This etiquette (FCC, 1994) uses a "Listen Before Talk" (LBT) approach, which requires devices to first sense the channel for a specified time and determine whether there is a transmission underway. If the received power is sufficiently low that they are unlikely to experience or cause interference, they can transmit. The LBT feature inherently provides much better protection from interference than a system of power limits where there is no sensing. The etiquette has no restriction on the technologies that devices may use except that they must follow the etiquette, and therefore it supports more diversity. (However, the etiquette has an inherent limitation that it cannot distinguish between applications of low value and of high value, and restricts all applications equally.) The etiquette also includes many provisions to improve spectrum efficiency. For example, under one such provision, devices requiring bandwidth less than 625 kHz in the isochronous band must search the band from left to right, which reduces inefficiencies due to spectrum fragmentation. A few features of the FCC etiquette were also designed to inhibit excessive use of spectrum, thereby reducing the potential for a tragedy of the commons. For instance, the etiquette encourages operation at low power by raising the threshold used to determine whether a channel is free or not for low power transmissions, thereby reducing the access delay. Features that inhibit excessive use of spectrum often complicate access strategy design, like the requirement that transmission in the isochronous band must cease if an acknowledgment has not been received within the last 30 seconds. Whether the restrictions in the etiquette are necessary for efficient operation, or are sufficient for efficient operation, remains to be seen.

The issue of promoting efficient spectrum sharing with real-time access has been addressed using an approach similar to that of the unlicensed PCS band, but it does so in the more limited case of common carriers offering the same service using cellular infrastructure (Salgado-Galicia, 1995, 1997). As with the FCC etiquette for unlicensed PCS, all devices are required to follow a specific protocol in acquiring access to spectrum. In this case, the etiquette is based on Dynamic Channel Allocation (DCA), through which channels are dynamically shared amongst the competing operators on a call by call basis. All devices search channels in the same order, and the first channel that meets the minimum Carrier-to-Interference (C/I) ratio is selected. This approach was found to offer far more efficient spectrum utilization than a traditional licensing approach based on exclusive access to spectrum, provided that all of the firms were willing to invest in equipment that would enhance spectral efficiency. Unfortunately, this sharing also led to a tragedy of the commons. The reason is that as channels saturate due to increased usage, cells should be made smaller to increase system capacity through frequency reuse. However, as the spectrum is shared, the operator investing in a new cell site bears all the costs, but the benefits are shared by all. Consequently, there is little incentive to increase effective capacity of a given spectrum block. The authors of this approach addressed this problem by limiting the number of transceivers deployed per broadcast tower. This helps to alleviate the risk of a tragedy of the commons in this particular scenario, but is not applicable for all types of devices, so another approach is needed in unlicensed spectrum supporting greater diversity. Section 5 will focus on suitable provisions for the unlicensed spectrum etiquette.

3 What Causes the Tragedy of the Commons?

In all wireless systems, design decisions are exclusively based on the self-interest of the users of the device being designed. The design of the access strategy involves a trade-off between competing goals and interests. One goal is to conserve spectrum; others might be to reduce equipment and operating costs, or to optimize some measure of performance like access delay or reception quality. In licensed spectrum, where the spectrum consumed is typically the exclusive domain of the users of the device, the goal of conserving spectrum is important to device designers. What sets unlicensed spectrum apart is that, although conservation of spectrum is no less important from a system perspective, there is considerably less incentive to design individual devices to conserve the shared spectrum, as mentioned in Section 1. Thus, in unlicensed spectrum, it is more likely that the best design decision from the selfish perspective of the designer of a given device is also a *greedy* approach, where the more a device is designed to waste shared spectrum unnecessarily in favor of its own goals, the more we consider it to be greedy. The amount of resources a device consumes with a transmission depends on three factors: the transmission duration, bandwidth, and coverage area, the latter of which is a function of transmission power. Thus, the transmissions of a greedy device have greater duration, bandwidth, or power, than is necessary. We will refer to these three factors as the three dimensions along which devices may manifest greedy behavior. We will now present examples where designers of unlicensed devices have motivation to be greedy along each of these three dimensions and the tradeoffs involved, beginning with transmission duration.

If a device could always access spectrum within an acceptable delay, there would be no reason for it to transmit longer than necessary. However, in unlicensed spectrum, there can be no such guarantees and the access delay may vary considerably. Devices may therefore be designed to transmit longer just to avoid the access delay whenever they have a message to transmit again. For example, consider a wireless bridge operating in the PCS unlicensed band that connects two wired Local Area Networks (LANs). Whenever a packet must be forwarded from one LAN to the other, the bridge has to wait for a given monitoring time before it may begin transmission as per the Listen Before Talk (LBT) protocol. Instead of releasing the channel at the end of packet transmission, the bridge may be greedy by continuing to transmit even if it has nothing to send. This way all packets that arrive after the first one are spared the access delay imposed by the LBT rule. Essentially, the device is hoarding spectrum that it may or may not need later. The bridge may continue to hold the channel for as long as the etiquette permits. However, there is a cost in doing so. The greedy bridge prevents other unlicensed devices from using the channel, and the queue of packets awaiting transmission at the other devices would grow. Thus, when the bridge finally releases the channel, it may take much longer to reclaim it. Consequently, this form of greed can be beneficial to the user, but isn't always.

We now consider greed in the bandwidth dimension. Application designers may be attracted to higher bandwidth choices because of the cost advantage of inefficient modulation schemes, or for improved performance such as better video quality. However, the disadvantage of using more bandwidth is that there would be a low probability of finding the needed bandwidth free since spectrum is shared. Consequently, as with the duration dimension, greed in the bandwidth dimension may or may not benefit the user.

Devices may likewise be designed to transmit at a higher transmission power than necessary in order to improve the signal-to-noise ratio (SNR), which in a digital transmission, decreases the bit error rate. Transmitting at a higher power also reduces the frequency reuse in the system. However, for a given device for which power consumption is not an issue, there would be no disadvantages to high power unless those disadvantages are imposed by the etiquette. As mentioned in Section 2, under the current FCC etiquette, the LBT noise threshold by which a device determines whether a channel is busy or free is a function of the transmit power. Consequently, if a device increases its transmit power, its noise threshold would be lower, resulting in greater access delay for the device. Of course, battery powered devices would have the additional disadvantage of a reduced battery life if they transmit at high power.

To evaluate the potential for a tragedy of the commons, we need to determine whether unlicensed applications will be designed to be greedy, and if so, to what extent and how greed would be manifested by different applications. If the potential greedy behavior of devices significantly degrades system performance, there is cause for concern and etiquette modifications to discourage greedy behavior would be necessary.

4 An Example of a Tragedy of the Commons

In this section we will consider a practical scenario where devices may be designed to be greedy in the duration dimension. In Section 4.1 we provide a description of this scenario and specify how devices may be greedy. Unfortunately, this is a difficult system to analyze. Section 4.2 gives a closed-form analysis that approximates the greedy behavior model well in many scenarios, particularly when the devices are greedy. Simulations are used in Section 4.3 to demonstrate the accuracy of this approximation. Section 4.4 discusses the implications of our findings.

4.1 Our Scenario

In this scenario, two devices compete for access to a wireless channel in the 2 GHz unlicensed PCS band. Devices are sufficiently close together that each device receives the other's transmissions, and no frequency reuse is possible. The devices follow the isochronous band etiquette, and their transmissions require the same bandwidth. Transmission power and bandwidth would therefore have no impact on device performance. We assume that messages arrive at the devices for transmission according to a Poisson process, and that the message lengths are exponentially distributed. Messages are queued in a buffer of infinite length until they can be transmitted. Devices may be greedy in the duration dimension by holding on to the channel longer than necessary, as discussed in Section 3. Nongreedy devices release the channel as soon as they transmit the last message in their queue, whereas a greedy device with a greed of duration T holds the channel for duration of at least T even if it has no messages to send. At the end of duration T , the greedy device releases the channel once its queue becomes empty.

If greed can improve performance, then equipment designers will select their greedy strategy based on their projections of the extent to which they will share spectrum with competing devices, and the extent to which those competing greedy devices will also be greedy. Since the strategies employed in one device can influence the optimal design for another device, a useful measure of resulting behavior is the *reaction function* $r_i(T_j)$ for each device i , which gives the optimal greed for a device in response to that of the other device. If Device 1 has a greed of T_1 , then $r_2(T_1)$ is the greed that minimizes Device 2's delay. Such dynamic reactions to another device's greed may occur in one of two ways. It is possible that greed on some devices can be adjusted by a system administrator, the way one might change the maximum packet length in a LAN. In other systems, the extent of greed will be fixed when the device is manufactured. However, greed can change over time when equipment is replaced. For example, if most CB radios in use have a power of P , then someone buying a new radio will buy one with power $r(P)$, and generation after generation, the power levels will change.

4.2 How Greed Affects Performance

We will use analysis and then simulation to demonstrate the potential tragedy of the commons in this scenario. We will show that if one device is designed to be greedy, it increases the average queuing delay for the messages transmitted by the other device. Furthermore, this other device can always reduce its delay by increasing its own greed. This degrades performance for the first device, and it can reduce its own delay by increasing its greed again, thereby continuing the process whereby each device always responds with more greed than the other.

To make analysis tractable, we use a *fluid-flow* model (Anick, 1982; Tucker, 1988) for message arrivals where messages arrive at a constant rate. More precisely, the amount of data received in any period of duration τ is exactly $\rho\tau$. In practice, the arrival rate fluctuates, causing the delay calculated by the fluid-flow model to have an error in the range of a few message transmission times. Therefore, whenever the delay caused by waiting for the other device to release the channel is much greater than the delay caused by the arrival rate fluctuations, the error percentage would be small, and the fluid flow model would therefore be a good approximation. As a result this model should be more accurate when devices are greedy.

We will now define the variables we use for our analysis. Let ρ_i be the load of Device i . We assume $\rho_1 > 0$, $\rho_2 > 0$, and $\rho_1 + \rho_2 < 1$, i.e., the total message arrival rate at both devices does not exceed the total capacity. Device $i : i \in \{1,2\}$ has a greed of duration T_i , and it holds the channel for a period of duration H_i , where $H_i \geq T_i$. For a period of $X_i : X_i \leq H_i$, Device i has messages queued and is transmitting at the maximum rate possible. Devices are required by the FCC etiquette to start transmission only if they find the channel to be free for a duration M . We assume devices sense the channel with persistence, i.e., devices continuously monitor the channel until they detect it to be idle (Vukovic, 1996). Figure 1 shows Device 1's *unfinished work* as it varies over time, where unfinished work is the amount of time it would take to transmit all currently queued messages. We define time 0 to be the time when Device 2 starts monitoring the channel. At that time, Device 1 has just relinquished the channel, and presumably has emptied its queue of messages. Device 2 finds the channel free and begins transmission at time M . It then releases the channel at time $H_2 + M$. Device 1's queue increases from time 0 until $H_2 + 2M$, when Device 1 begins transmission after monitoring the channel for duration M . Its unfinished work built up at this time is given by $\rho_1(2M + H_2)$. Device 1's queue length then starts decreasing and reaches zero at time $2M + H_2 + X_1$. The device continues to hold the channel, transmitting messages as quickly as they arrive until time $2M + H_2 + H_1$. The process is then repeated.

The following theorems characterize the potential greedy behavior in this system. (See Appendix for proofs.) Theorem 1 shows the holding times H_1^* and H_2^* for devices 1 and 2 respectively, when neither device is greedy. Note that Device i is never greedy if $T_i \leq H_i^*$, because the device always holds the channel for at least H_i^* anyway. Therefore, delays are identical for any $T_i \leq H_i^*$. Theorem 2 shows the general reaction functions. Theorem 3 shows that even if device i is nongreedy, if $H_i^* < 2M$, then the other device is better off being greedy. To determine the ultimate results of these reaction functions, let Device 2 select an initial greed $T_2^{(0)}$. Devices 1 and 2 will then take turns responding to each other's greed, i.e., for $i > 0$, $T_1^{(i)} = r_1(T_2^{(i-1)})$ and $T_2^{(i)} = r_2(T_1^{(i)})$. Theorem 4 shows that if one device is greedy, it will cause both devices to escalate their greed to infinity. Theorem 5 shows that whenever a device's greed is increased, delay increases for the other device.

Theorem 1: If $T_1 = T_2 = 0$, then $H_1^* = \frac{2M\rho_1}{1-\rho_1-\rho_2}$ and $H_2^* = \frac{2M\rho_2}{1-\rho_1-\rho_2}$

Theorem 2: $r_1(T_2) = \max\{T_2, 2M\} \frac{1-\rho_2}{\rho_2} - 2M$
 $r_2(T_1) = \max\{T_1, 2M\} \frac{1-\rho_1}{\rho_1} - 2M$

Theorem 3: If $T_1 < H_1^*$ and $H_1^* < 2M$ then $r_2(T_1) > H_2^*$
If $T_2 < H_2^*$ and $H_2^* < 2M$ then $r_1(T_2) > H_1^*$

Theorem 4: If $T_2^{(0)} > H_2^*$ then $\lim_{i \rightarrow \infty} T_2^{(i)} = \infty$ and $\lim_{i \rightarrow \infty} T_1^{(i)} = \infty$

Theorem 5: $\frac{\partial D_1}{\partial H_2} > 0$ and $\frac{\partial D_2}{\partial H_1} > 0$

All together, these theorems show the potential for a tragedy of the commons. If $H_1^* < 2M$ or $H_2^* < 2M$, which occurs when $\rho_1 + \rho_2 + \min\{\rho_1, \rho_2\} < 1$, then it will inevitably lead to an escalation of greed until both devices hold the channel as long as possible, and neither has adequate performance.

4.3 Results from Simulations

Of course, as described in Section 4.2, the results of a fluid-flow approximation are accurate only when the delay caused by waiting for the other device to release the channel significantly exceeds the delay caused by fluctuations in the arrival rate of the messages. This is certainly the case when the greed of both devices is significant, but may not be when devices are not greedy. We will test these results through simulation.

We first present the case when each device has a load of 40 % and an average message transmission time of 0.5 ms. Device 1 then varies its greed from zero to eight hours while Device 2 remains nongreedy. Figure 2 shows that the greed of Device 1 results in increased delay for both devices. As predicted by the analytic approximation, Device 1's own delay is minimized at a greed of 0, so there is no incentive to make it greedy. If this were always true, there would be no risk of a tragedy of the commons. However, when the devices are less heavily loaded at 10 % load, each having an average message transmission time of 0.5 ms, we see in Figure 3 that Device 1's performance does benefit from being greedy. Again, this result was predicted through analysis. The delay of Device 1 is minimized at 160 ms, whereas the delay of Device 2 increases monotonically with the greed of Device 1. Thus, there is incentive to make Device 1 greedy in this case, and Device 2's performance will suffer as a result. Indeed, as Figure 4 shows, in some cases a device can decrease its delay by an order of magnitude through greed.

We saw in Section 4.2 that in a fluid-flow model, if one device is made greedy, the other device will as well, and greed will escalate. Figure 5 shows that this phenomenon also occurs as predicted. This figure shows the reaction functions when both devices have a load of 10 % and the average message transmission length is 0.5 ms. When the greed chosen for Device 1 significantly exceeds monitoring time, Device 2's delay is reduced by making it more greedy, by about $\frac{1-\rho_1}{\rho_1} = 9$ times as much, as predicted in Theorem 2. For example, if the devices start with greed indicated by point A on Figure 5, they will progress to points B, C, D, etc. Eventually they reach point H, where both hold the channel for 8 hours at a time, which is the maximum transmission duration allowed by the FCC etiquette. Note that this is the only point where the reaction functions intersect, so it is the only equilibrium. If there were no such upper limit imposed by the etiquette, our analysis predicts that the greed of devices would tend to infinity. Figure 6 shows that as the greed of devices escalates from zero to 8 hours, the average delay increases monotonically. This is a tragedy of the commons. We have similarly observed this phenomenon in other reaction functions that were derived through simulation.

4.4 Implications

As described in Section 4.1, the reaction functions employed in Sections 4.2 and 4.3 are applicable when there are two devices competing for spectrum, and greed for each device is chosen independently. Initially, there will be few unlicensed devices competing for spectrum. Consequently, the first set of unlicensed devices produced by the industry are less likely to be designed greedy. As usage of the unlicensed PCS bands increases, and more devices compete with each other for access to spectrum, their performance would decrease. The reaction of equipment designers would depend to a large extent on the diversity of devices using unlicensed spectrum, and the corresponding industry structure. At one extreme, consider the case where a single manufacturer produces all of these unlicensed devices. This

manufacturer would only produce nongreedy devices, as there is nothing to be gained from a tragedy of the commons between two or more of its own customers. As we see in Figure 7, if there are two devices that have the same greed, their delay increases monotonically with their greed. However, the situation is very different if all manufacturers have so little market share that two devices from the same manufacturer rarely compete for spectrum. In this case, greedy devices will emerge and spectral efficiency will be reduced. Of course, reality is somewhere in between, and it will be explored in future work, as will the impact of greed when there are more than two competing devices.

The way to prevent manufacturers from marketing greedy devices would be to discourage greedy behavior through modifications to the FCC etiquette. We will provide examples of modifications that may be used to discourage greedy behavior in the following section.

5 Ways to Avoid the Tragedy of the Commons

One way to control greed is to set upper limits on each of the dimensions of greedy behavior, i.e., one might impose a maximum duration, power, or bandwidth. However, there are problems with this approach. For example, setting a time limit on a voice conversation would not be desirable as phone calls may then be prematurely terminated. Also, an upper limit may result in spectrum inefficiency unless it is chosen appropriately. As discussed in Section 2, a device may be designed to use as much of a resource as is allowed by the upper limit even if it is not necessary for adequate performance. Hence a higher than optimal limit would be inefficient. A lower limit may restrain the use of resources to the extent that spectrum is unnecessarily made unusable for some applications. Another option is to have a slightly more flexible form of upper limit. As mentioned in Section 3, the maximum transmitted power allowed by the FCC etiquette is a function of bandwidth. Narrow bandwidth applications are permitted higher power spectral density, which provides an incentive for devices to not use excessive bandwidth. Although there is some flexibility in the dimension of resource consumption a device chooses, this is still an upper limit, and there is no incentive to do better than the limit requires. Upper limits are indeed a blunt instrument to deal with such a delicate problem.

To produce further incentives, we might give devices that consume less spectrum some form of priority in accessing spectrum. For incentives to be effective, they must be based on parameters that strongly affect the device performance. We will now provide examples of such parameters. In the current etiquette, a device must monitor the channel before transmitting to ensure that detected power remains below a threshold throughout the monitoring period. Altering either *monitoring time* or *power threshold* would affect a device's chances of accessing spectrum. For example, if two devices begin monitoring at the same time, the one with the smaller monitoring period will get access. (Currently, monitoring period is fixed in both bands, but threshold is a function of transmit power.) Thus, factors like monitoring time and power threshold can be used to provide incentives; if devices seeking a large bandwidth were assigned a large monitoring time, there would be more incentive to use efficient modulation. Another factor that influences access to spectrum is *inter-burst gap*, which is the minimum amount of time a device must wait to transmit after completing a transmission. For example, if a need is discovered to induce a device to end its transmission early, the inter-burst gap following a long transmission could be made large compared to the monitoring times of other devices, whereas a short transmission can be rewarded with short inter-burst gaps. In the current etiquette, inter-burst gap also does not depend on resources consumed. It is 10 ms in the isochronous band, and it is selected randomly in the asynchronous band, where the distribution depends only on the number of previously unsuccessful attempts to access the spectrum. Another parameter which affects a device's performance is the *back-off period*, which is the minimum time a device has to wait before attempting to access a channel again once the channel is detected busy. Making the back-off period depend on the spectrum resources consumed by the devices might provide an incentive to design devices that conserve spectrum. Our future work will address and quantify the extent to which such parameters affecting access to spectrum can be used to induce efficient utilization of the spectrum resource.

6 Conclusion

Some wireless applications are not well-served under a system of exclusive rights to spectrum, and are better off in shared spectrum supporting real-time access. Real-time sharing has several advantages. It supports mobility of wireless applications, allows spectrum sharing, and facilitates experimentation and innovation. Unlicensed spectrum offers the potential to realize these benefits. However, as there is little inherent incentive for individual devices to be designed to use unlicensed spectrum efficiently, they may engage in greedy behavior, i.e., device designers may sacrifice the goal of spectrum efficiency to meet other design goals. In this paper we have demonstrated, in a simple scenario, that greedy behavior is sometimes rewarded, and it can lead to a tragedy of the commons. The severity of this problem will be determined in future work, which will consider more complex models. For example, we will consider more complicated wireless applications, other forms of greed, the result of competition for spectrum among more than two devices, and the impact of industry and market structure. If the tragedy of the commons proves to be significant, the most practical way to reduce the risk of a tragedy of the commons is offered by the etiquette for unlicensed spectrum. We have suggested examples of etiquette modifications that may prove effective in discouraging greedy behavior. These will be addressed in more detail in future work.

Meanwhile, demand for more unlicensed spectrum is high, and the FCC is moving forward. At this point, there is still little evidence that provisions in the current FCC etiquette are both necessary and sufficient. It is possible that a tragedy of the commons will lead to poor performance in this band, and it is also possible that many existing restrictions in the etiquette could be relaxed without penalty, thereby simplifying designs. Little has been published on this issue, and the industry groups backing unlicensed spectrum have kept their work mainly proprietary. Consequently, before industry invests significantly in the manufacture of unlicensed devices, and before the FCC releases more large blocks of unlicensed spectrum, caution is advised until more is known.

7 APPENDIX

Theorem 1:

If $T_1 = T_2 = 0$, then $H_1^* = \frac{2M\rho_1}{1-\rho_1-\rho_2}$ and $H_2^* = \frac{2M\rho_2}{1-\rho_1-\rho_2}$

Proof for Theorem 1:

When $T_1 = T_2 = 0$, both devices are nongreedy, so $H_1^* = X_1$ and $H_2^* = X_2$.

Device 1 has $\rho_1(2M + H_1^* + H_2^*)$ unfinished work built up during the period from

0 to $2M + H_1^* + H_2^*$, which it transmits in time H_1^* . Therefore we have

$H_1^* = \rho_1(2M + H_1^* + H_2^*)$ and $H_2^* = \rho_2(2M + H_1^* + H_2^*)$ from symmetry.

The solutions to the above equations are given by

$$H_1^* = \frac{2M\rho_1}{1-\rho_1-\rho_2} \text{ and } H_2^* = \frac{2M\rho_2}{1-\rho_1-\rho_2} \quad (1)$$

Theorem 2:

$$r_1(T_2) = \max\{T_2, 2M\} \frac{1-\rho_2}{\rho_2} - 2M, \text{ and } r_2(T_1) = \max\{T_1, 2M\} \frac{1-\rho_1}{\rho_1} - 2M$$

Proof for Theorem 2:

Device 1 responds with the greed $T_1 = r_1(T_2)$ that minimizes Device 1's delay. The total unfinished work

$\rho_1(2M + H_2 + X_1)$ built up by Device 1 during the period from 0 to $2M + H_2 + X_1$ is transmitted in

duration X_1 , so $X_1 = \rho_1(2M + H_2 + X_1)$, i.e.,

$$X_1 = \rho_1 \frac{2M + H_2}{1 - \rho_1}. \text{ Similarly, } X_2 = \rho_2 \frac{2M + H_1}{1 - \rho_2} \quad (2)$$

The average delay D_1 for Device 1's messages is its average unfinished work divided by ρ_1 . As Figure 1 shows, Device 1's average unfinished work during the period from 0 to $2M + H_2 + X_1$ when it has queued messages is $0.5\rho_1(2M + H_2)$. The fraction of time Device 1 has queued messages is

$$\frac{2M + H_2 + X_1}{2M + H_2 + H_1}, \text{ so } D_1 = 0.5(2M + H_2) \frac{2M + H_2 + X_1}{2M + H_2 + H_1} \quad (3)$$

We will consider two cases: $T_2 > X_2$ and $T_2 \leq X_2$.

Case 1: $T_2 > X_2$, so Device 2's holding time is $H_2 = T_2$, independent of T_1 . Consequently,

$$\frac{\partial D_1}{\partial H_1} = -0.5(2M + T_2) \frac{2M + T_2 + X_1}{(2M + T_2 + H_1)^2} \text{ which is always negative. Thus, as long as}$$

$$T_2 > X_2 = \rho_2 \frac{2M + H_1}{1 - \rho_2}, \text{ Device 1 would choose to increase } H_1, \text{ which it can do by increasing } T_1.$$

Case 2: $T_2 \leq X_2$, so $H_2 = X_2$ and Device 1's average delay is

$$D_1 = 0.5(2M + X_2) \frac{2M + X_2 + X_1}{2M + X_2 + H_1} \text{ where}$$

$$X_1 = \rho_1 \frac{2M + X_2}{1 - \rho_1} \text{ and } X_2 = \rho_2 \frac{2M + H_1}{1 - \rho_2}.$$

$$\text{Differentiating, we get } \frac{\partial D_1}{\partial H_1} = \frac{0.5(2M + H_1\rho_2)[\rho_2(H_1 + 2M) - 2M(1 - \rho_2)]}{(1 - \rho_1)(1 - \rho_2)(2M + H_1)^2}.$$

$$\frac{\partial D_1}{\partial H_1} > 0 \text{ if and only if } \rho_2(H_1 + 2M) > 2M(1 - \rho_2), \text{ or equivalently,}$$

$$\text{if } X_2 = \frac{\rho_2(2M + H_1)}{1 - \rho_2} > 2M. \text{ Thus, in the case where } T_2 < X_2, \text{ Device 1 would choose to increase } H_1 \text{ (and thus } T_1) \text{ if } X_2 < 2M, \text{ decrease } H_1 \text{ if } X_2 > 2M, \text{ and leave } H_1 \text{ unchanged if } X_2 = 2M.$$

Considering both cases, we see that Device 1 would choose an H_1 at which $X_2 = \max\{T_2, 2M\}$ if possible. (It may be that H_1 will be larger than this even if Device 1 is nongreedy.) Consequently, we set T_1 equal to this value of H_1 , and by Equation 2, we have

$$T_1 = r_1(T_2) = \max\{T_2, 2M\} \frac{1 - \rho_2}{\rho_2} - 2M \quad (4)$$

$$\text{By symmetry, } T_2 = r_2(T_1) = \max\{T_1, 2M\} \frac{1 - \rho_1}{\rho_1} - 2M \quad (5)$$

Theorem 3:

If $T_1 < H_1^*$ and $H_1^* < 2M$ then $r_2(T_1) > H_2^*$

If $T_2 < H_2^*$ and $H_2^* < 2M$ then $r_1(T_2) > H_1^*$

Proof for Theorem 3:

Let $T_1 < H_1^*$ and $H_1^* < 2M$. From Theorem 2, we have

$$r_2(T_1) = \max\{T_1, 2M\} \frac{1-\rho_1}{\rho_1} - 2M = 2M \frac{1-2\rho_1}{\rho_1} \text{ as } T_1 < 2M.$$

$$\text{From Theorem 1, } H_1^* = \frac{2M\rho_1}{1-\rho_1-\rho_2}$$

$$\text{As } H_1^* < 2M, \frac{2M\rho_1}{1-\rho_1-\rho_2} < 2M.$$

Multiplying both sides by $\frac{1-\rho_1}{\rho_1}$ and subtracting $2M$, we get $\frac{2M\rho_2}{1-\rho_1-\rho_2} < 2M \frac{1-2\rho_1}{\rho_1}$, i.e.,

$$H_2^* < r_2(T_1). \text{ Hence if } T_1 < H_1^* \text{ and } H_1^* < 2M \text{ then } r_2(T_1) > H_2^*$$

By symmetry, if $T_2 < H_2^*$ and $H_2^* < 2M$ then $r_1(T_2) > H_1^*$

Theorem 4: If $T_2^{(0)} > H_2^*$ then $\lim_{i \rightarrow \infty} T_2^{(i)} = \infty$ and $\lim_{i \rightarrow \infty} T_1^{(i)} = \infty$

Proof for Theorem 4:

$$\text{Let } S_2^{(0)} = T_2^{(0)}. \text{ We define } f_1(S_2) = S_2 \frac{1-\rho_2}{\rho_2} - 2M \text{ and } f_2(S_1) = S_1 \frac{1-\rho_1}{\rho_1} - 2M.$$

Note that $\frac{\partial f_1}{\partial S_2} > 0$ and $\frac{\partial f_2}{\partial S_1} > 0$, i.e., f_1 and f_2 are monotonically increasing functions.

Let $S_1^{(i)} = f_1(S_2^{(i-1)})$ and $S_2^{(i)} = f_2(S_1^{(i)}) = f_2(f_1(S_2^{(i-1)}))$. Likewise, let $T_1^{(i)} = r_1(T_2^{(i-1)})$ and $T_2^{(i)} = r_2(T_1^{(i)}) = r_2(r_1(T_2^{(i-1)}))$.

$$\text{By Theorem 2, } r_1(T_2) = \max\{f_1(T_2), 2M \frac{1-\rho_2}{\rho_2} - 2M\} \geq f_1(T_2) \quad (6)$$

$$\text{and } r_2(T_1) = \max\{f_2(T_1), 2M \frac{1-\rho_1}{\rho_1} - 2M\} \geq f_2(T_1) \quad (7)$$

We will first prove three propositions, which we use to prove Theorem 4.

Proposition 1: $T_2^{(i)} \geq S_2^{(i)} \forall i \geq 0$

Proof: We will prove proposition 1 by induction.

Basis: Let $i = 0$. $T_2^{(i)} \geq S_2^{(i)}$ becomes $T_2^{(0)} \geq S_2^{(0)}$ which is true as $T_2^{(0)} = S_2^{(0)}$.

Inductive step: If $T_2^{(i)} \geq S_2^{(i)}$ for a given i , we show $T_2^{(i+1)} \geq S_2^{(i+1)}$.

From Equation 7 we have $T_2^{(i+1)} = r_2(r_1(T_2^{(i)})) \geq f_2(r_1(T_2^{(i)}))$. From Equation 6 we have

$r_1(T_2^{(i)}) \geq f_1(T_2^{(i)})$. Since f_2 is monotonically increasing, $f_2(r_1(T_2^{(i)})) \geq f_2(f_1(T_2^{(i)}))$. Therefore

$T_2^{(i+1)} \geq f_2(f_1(T_2^{(i)}))$. Since f_1 is also monotonically increasing, and $T_2^{(i)} \geq S_2^{(i)}$, we have

$f_1(T_2^{(i)}) \geq f_1(S_2^{(i)})$ and therefore $f_2(f_1(T_2^{(i)})) \geq f_2(f_1(S_2^{(i)}))$. We thus have

$T_2^{(i+1)} \geq f_2(f_1(T_2^{(i)})) \geq f_2(f_1(S_2^{(i)})) = S_2^{(i+1)}$, i.e., $T_2^{(i+1)} \geq S_2^{(i+1)}$.

Hence $T_2^{(i)} \geq S_2^{(i)} \forall i \geq 0$.

Proposition 2: $S_2^{(i)} = (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2})^i (S_2^{(0)} - H_2^*) + H_2^* \forall i \geq 0$

Proof: We will prove this by induction.

Basis: When $i = 0$, $S_2^{(i)} = (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2})^i (S_2^{(0)} - H_2^*) + H_2^*$ becomes

$$S_2^{(0)} = S_2^{(0)} - H_2^* + H_2^* = S_2^{(0)} \text{ which is true.}$$

Inductive step: If proposition 2 is true for a given i , then we show it is true for $i+1$. Assume that

$$S_2^{(i)} = (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2})^i (S_2^{(0)} - H_2^*) + H_2^* \text{ for some } i.$$

By definition, $f_1(S_2) = S_2 \frac{1-\rho_2}{\rho_2} - 2M$ and $f_2(S_1) = S_1 \frac{1-\rho_1}{\rho_1} - 2M$. Therefore, we have

$$S_2^{(i+1)} = f_2(f_1(S_2^{(i)})) = f_1(S_2^{(i)}) \frac{1-\rho_1}{\rho_1} - 2M = \frac{1-\rho_1}{\rho_1} \frac{1-\rho_2}{\rho_2} S_2^{(i)} - \frac{2M}{\rho_1}.$$

From Equation 1 we have

$$\begin{aligned} S_2^{(i+1)} &= (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2}) S_2^{(i)} - H_2^* \frac{1-\rho_1-\rho_2}{\rho_1\rho_2} \\ &= (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2}) (S_2^{(i)} - H_2^*) + H_2^* \\ &= (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2}) [(1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2})^i (S_2^{(0)} - H_2^*) + H_2^* - H_2^*] + H_2^* \\ &= (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2})^{i+1} (S_2^{(0)} - H_2^*) + H_2^* \end{aligned}$$

Hence $S_2^{(i)} = (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2})^i (S_2^{(0)} - H_2^*) + H_2^* \forall i \geq 0$.

Proposition 3: If $S_2^{(0)} > H_2^*$ then $\lim_{i \rightarrow \infty} S_2^{(i)} = \infty$ and $\lim_{i \rightarrow \infty} S_1^{(i)} = \infty$

Proof: From Proposition 2, we have $S_2^{(i)} = (1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2})^i (S_2^{(0)} - H_2^*) + H_2^* \forall i \geq 0$

Since $\rho_1 + \rho_2 < 1$, $1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2} > 1$. Also, $S_2^{(0)} - H_2^* > 0$.

Consequently, $(1 + \frac{1-\rho_1-\rho_2}{\rho_1\rho_2})^i (S_2^{(0)} - H_2^*)$ approaches ∞ as i approaches ∞

and so does $S_2^{(i)}$. We have $S_1^{(i)} = f_1(S_2^{(i-1)}) = S_2^{(i-1)} \frac{1-\rho_2}{\rho_2} - 2M$.

Thus if $S_2^{(i-1)}$ goes to ∞ as i goes to ∞ , $S_1^{(i)}$ must do the same.

Hence if $S_2^{(0)} > H_2^*$ then $\lim_{i \rightarrow \infty} S_2^{(i)} = \infty$ and $\lim_{i \rightarrow \infty} S_1^{(i)} = \infty$.

We now prove Theorem 4. From Proposition 3 we see that

if $S_2^{(0)} > H_2^*$ then $\lim_{i \rightarrow \infty} S_2^{(i)} = \infty$. From Proposition 1 we have $T_2^{(i)} \geq S_2^{(i)} \forall i \geq 0$. Therefore

$$\lim_{i \rightarrow \infty} T_2^{(i)} = \infty. \text{ We have } T_1^{(i)} = r_1(T_2^{(i-1)}) = \max\{T_2^{(i-1)}, 2M\} \frac{1-\rho_2}{\rho_2} - 2M.$$

Thus if $T_2^{(i-1)}$ goes to ∞ as i goes to ∞ , so does $T_1^{(i)}$. As $T_2^{(0)} = S_2^{(0)}$, we get if $T_2^{(0)} \geq H_2^*$ then $\lim_{i \rightarrow \infty} T_2^{(i)} = \infty$ and $\lim_{i \rightarrow \infty} T_1^{(i)} = \infty$.

Theorem 5:

$$\frac{\partial D_1}{\partial H_2} > 0 \text{ and } \frac{\partial D_2}{\partial H_1} > 0$$

Proof for Theorem 5:

As shown in the proof of Theorem 2, Device 1's delay is given by

$$D_1 = 0.5(2M + H_2) \frac{2M + H_2 + X_1}{2M + H_2 + H_1} \text{ where } X_1 = \frac{\rho_1}{1-\rho_1}(2M + H_2).$$

$$\text{Therefore } D_1 = 0.5(2M + H_2) \frac{2M + H_2 + \frac{\rho_1}{1-\rho_1}(2M + H_2)}{2M + H_2 + \max\{T_1, \frac{\rho_1}{1-\rho_1}(2M + H_2)\}}$$

$$\text{If } T_1 < X_1 \text{ then } D_1 = 0.5(2M + H_2) \text{ and } \frac{\partial D_1}{\partial H_2} = 0.5$$

$$\text{If } T_1 > X_1 \text{ then } D_1 = 0.5(2M + H_2) \frac{2M + H_2 + \frac{\rho_1}{1-\rho_1}(2M + H_2)}{2M + H_2 + T_1}$$

$$\text{and } \frac{\partial D_1}{\partial H_2} = \frac{0.5}{1-\rho_1} \frac{(2M + H_2)(2M + H_2 + 2T_1)}{(2M + H_2 + T_1)^2}$$

Since $M > 0$, $H_2 > 0$, and $T_1 > 0$, we have $\frac{\partial D_1}{\partial H_2} > 0$. By symmetry, the same is true for $\frac{\partial D_2}{\partial H_1}$.

This shows that greed always hurts the other device.

References

Anick, D., Mitra, D., & Sondhi, M. M. (1982). Stochastic Theory of a Data-Handling System with Multiple Sources. Bell Sys. Tech. Journal, 61, 8, 1871-1894.

FCC. (1994). Amendment of the Commission's Rules to Establish New Personal Communications Services. Memorandum Opinion and Order, FCC 94-144, Gen. Docket No. 90-314.

FCC. (1996). NII/SUPERnet at 5 GHz. Notice of Proposed Rulemaking, FCC 96-193, ET Docket No. 96-102.

Gilder, G. (1994). Auctioning the Airways. Forbes ASAP Supplement, 153, 8, 99-112.

Hardin, G. (1968). The Tragedy of the Commons. Science, 162, 1243-1248.

- Hundt, R. E & Rosston, G. L. (1995). Spectrum Flexibility Will Promote Competition and the Public Interest. IEEE Communications, 33, 12, 40-43.
- Marcus, M. J. (1994). Millimeter Wave Spectrum Management. Applied Microwave & Wireless, 6, 98-106.
- Marcus, M. J. (1996). Recent Progress in U.S. Millimeter Wave Spectrum Management Policy. IEEE MTT-S International Microwave Symposium Digest, 2, 505-507.
- McGarty, T. P., & Medard, M. (1994). Wireless Architectural Alternatives: Current Economic Valuations versus Broadband Options, The Gilder Conjectures Proc. 22nd Telecommunications Policy Research Conference.
- Noam, E. M. (1995). Taking the Next Step Beyond Spectrum Auctions: Open Spectrum Access. IEEE Communications, 33, 12, 66-73.
- Peha, J. M. (1994-95). Seamless Wireless Networks: Motivation and Challenges. International Engineering Consortium Annual Review of Communications, 48, 553-556.
- Peha, J. M. (1997). Developing Equipment and Services for Shared Spectrum: Is It A Good Gamble? To appear in International Engineering Consortium Annual Review of Communications, 50.
- Peha, J. M. (1996). A Proposed New Spectrum Management Policy. <http://www.ece.cmu.edu/afs/ece/usr/peha/peha.html>.
- Pressler, L. (1996). Spectrum Reform Discussion Draft. Congressional Record (Senate). S4928-S4936.
- Salgado-Galicia, H., Sirbu, M., & Peha, J. M. (1995). Spectrum Sharing Through Dynamic Channel Assignment For Open Access To Personal Communications Services. Proc. IEEE International Conference on Communications (ICC), June, 417-422.
- Salgado-Galicia, H., Sirbu, M., & Peha, J. M. (1997). A Narrow Band Approach to Efficient PCS Spectrum Sharing Through Decentralized DCA Access Policies. IEEE Personal Communications, 4, 1, 24-34.
- Steer, D. G. (1994). Wireless Operation in the Unlicensed Band. IEEE Personal Communications, 1, 4, 36-43.
- Tucker, R. C. F. (1988). Accurate Method for Analysis of a Packet-Speech Multiplexer with Limited Delay. IEEE Trans. Commun., 36, 4, 479-483.
- Vukovic, I., & McKown, J. (1996). Spectrum Sharing Under The Asynchronous UPCS Etiquette: The Performance of Collocated Systems Under Heavy Load. Proceedings of ACM Mobicom, Nov, 67-72.
- Wang, Q., Peha, J. M., & Sirbu, M. (1997). Optimal Pricing for Integrated-Services networks with Guaranteed Quality of Service. In J. Bailey and L. McKnight (Eds.) Internet Economics, (pp. 353-376). MIT Press.
- Wang, Q., Sirbu, M., & Peha, J. M. (1996). Dynamic Pricing for Integrated Services Networks. In G. W. Brock & G. L. Rosston (Eds.), The Internet and Telecommunications Policy: Selected Papers From The 1995 Telecommunications Policy Research Conference, (pp. 65-78). Mahwah, NJ: Lawrence Erlbaum Associates.




Figure 1: Device 1's unfinished work as a function of time under a fluid-flow model for message transmissions.




Figure 2: Delay of each device as a function of Device 1's greed. Device 2 is nongreedy. Each device has a load of 40% and an average message transmission time of 0.5 ms. Monitoring time = 10 ms.




Figure 3: Delay of each device as a function of Device 1's greed. Device 2 is nongreedy. Each device has a load of 10 %, and an average message transmission time of 0.5 ms. Monitoring time = 10 ms.




Figure 4: Delay of each device as a function of Device 1's greed. Device 2's Greed = 10,000 ms. Each device has 10 % load and an average message transmission time of 0.5 ms. Monitoring time = 10 ms.




Figure 5: Reaction Function for two devices at 10 % load each. Average message transmission time is 0.5 ms. Monitoring time = 10 ms.




Figure 6: Device 1's delay at its optimal greed vs. Device 2's greed. Devices are each at 10 % load and an average message transmission time of 0.5 ms. Monitoring time = 10 ms.




Figure 7: Delay of each device as a function of greed. Devices have equal greed. Each has a load of over 10 % and an average message transmission time of 0.5 ms. Monitoring time = 10 ms.