Chapter 7: Dynamic Frequency Selection[1]

"So Abram said to Lot, "Let's not have any quarreling between you and me, or between your herdsmen and mine, for we are brothers.
Is not the whole land before you? Let's part company.
If you go to the left, I'll go to the right; if you go to the right, I'll go to the left."
- Genesis 13:8-9

Over the past years the number of WiFi hotspots has exploded. We all know that if we go downtown or to a large apartment building, we can find dozens of open access points to log on to. In fact Philadelphia, San Francisco, and our own Roanoke have rolled out city-wide WiFi services. So even before the wide-scale deployments of 3G and WiMAX systems high speed data services are already ubiquitous.

While WiFi coverage has become less of a problem, external network interference has emerged as a significant problem as the networks fight for access to a limited number of channels (and frequently, the same channel!). In theory, this interference problem could be ameliorated by applying a frequency reuse pattern to the networks – a seemingly easily implemented approach as 802.11b has three nonoverlapping channels (1, 6, and 11) and 802.11a has eight minimally interfering channels in the US (nineteen in Europe) which are explicitly intended to facilitate frequency reuse in a minimally interfering manner. However, most people never modify their access points from the factory settings so many access points operate on the same pre-set channel.[2]

A few years ago this reliance on the factory settings became a problem for a friend of mine he set up a WiFi network in his house. A few months after setting up his network, his neighbors set up their own WiFi networks in their houses. As everyone had left the access points with their factory settings and had bought the same model of access point,

[1] This chapter is mostly taken from [Neel_06].

[2] An online acquaintance of mine has noted this same phenomenon with compasses in cars. Because magnetic north is not true north, the direction a compass points varies geographically. To combat this, compasses in cars intended for the US come equipped with fifteen different settings which calibrate the difference between magnetic north and true north. Invariably when he has bought a car, whether new or used, whether in California or in Virginia, the car's compass has been set for region 8 – Detroit, the factory site. So right now in the US, there's thousands of people who think they're driving north when they're really driving northwest or northeast because they never changed their compass from its default setting.

all of the networks were operating on channel 6 and the performance of my friend's network suffered noticeably from the interference. Being a PhD wireless engineer at Virginia Tech which confers a certain degree of mischievousness, my friend was not content to simply reconfigure his access point to operate on channel 1 or 11. At the time leaving an access point on its factory settings meant that the access point had no security and a common password, so he logged into his neighbors' access points and changed their operating channels so that all three access points were operating on non-interfering channels – a process he extended to his entire street as new access points were added. In this way, an optimal frequency reuse scheme was found for everyone's networks even if most parties were unwitting participants in the process.

Unfortunately when designing multi-channel networks, we can not count on the networks being deployed on streets with wireless engineers willing to tune the networks so another solution is required. Instead, we would prefer to construct self-tuning networks wherein the networks autonomously choose their parameters post-deployment indicating an opportunity for one of the killer applications of cognitive radio identified in Chapter 1 – automated radio resource management for automated deployment. Ideally, we would like for channel tuning portion to be as effective and as simple as possible – perhaps as simple as when Abram ensured that his flock would not interfere with Lot's with the promise, "*If you go to the left, I'll go to the right; if you go to the right, I'll go to the left.*"

Leveraging the insights gained over the preceding chapters, this chapter proposes low-complexity algorithms for autonomous dynamic frequency selection (DFS) for interference minimization among secondary users[3]. These algorithm are suitable for implementation in 802.22 and 802.11h networks, require no direct collaboration between devices and are easily implemented. Section 7.1 introduces a general model of DFS adaptation algorithms and presents related work. Section 7.2 formally introduces the algorithms, proves important results related to steady-states, the desirability of those steady-states, convergence, and stability and verifies these results via simulation. Section

[3] Dynamic frequency selection is proposed in 802.11h and 802.22 primarily as a means to minimize interference with primary spectrum users – military radars and television broadcasts, respectively – with minimization of interference to other 802.11h and 802.22 devices a secondary consideration.

7.3 studies the proposed algorithms under various realistic conditions – policy variations, asynchronous timing, local frequency preference, noise, and the impact of differing power levels.

7.1 Background

This section introduces a model of distributed DFS algorithms, presents related work, and briefly reviews the IRN design framework.

7.1.1 Modeling DFS Algorithms

Modifying the notation of Chapter 2 to be DFS specific, we can model a network of cognitive radios (or any goal-driven adaptive radios) by the tuple, $\langle N, F, \{u_i\}, \{d_i\}, T\rangle$ where N represents the set of n cognitive radios, F is the frequency space formed as $F = F_1 \times \cdots \times F_n$ where F_i specifies the frequencies available to cognitive radio $i \in N$, $\{u_i\}$, $u_i : F \to \mathbb{R}$, is the set of goals that inform the cognitive radios' decision processes, $d_i : F \to F_i$, which are implemented at the decision timings contained in T.

For our DFS algorithm we model the goal of the radios as minimizing perceived interference as shown in (7.1)

$$u_i(f) = -I_i(f) = -\sum_{k \in N \setminus i} g_{ki} p_k s(f_i, f_k) \tag{7.1}$$

where σ measures the fractional interference, i.e., $s(f_i, f_k) = \max\{B - |f_i - f_k|, 0\}/B$, f_i is the frequency of cognitive radio i's RTS/CTS signal, p_k is the transmission power of radio k's waveform, and g_{ki} is the link gain from the transmission source of radio k's signal to the point where radio i measures its interference. It is assumed that the network design objective function is to minimize the sum network interference, $\Phi(f)$, as shown in (7.2).

$$\Phi(f) = \sum_{i \in N} \sum_{k \in N \setminus i} g_{ki} p_k s(f_k, f_i) \tag{7.2}$$

7.1.2 Related Work

Many authors have attacked the problem of DFS, or more generally dynamic spectrum access (DSA), by requiring assuming a centralized decision maker. After noting that finding the optimal frequency allocation is a NP-complete problem, [Leung_03] proposes a heuristic centralized algorithm based on a local search algorithm with random restart to search through the possible frequency combinations. As part of a solution to network formation problem [Steenstrup_05] utilizes a central controller to assign frequencies to each link in the network according to the abbreviated algorithm shown in Figure 7.1.

For each $v_i \in V$ do
  $F_i^R \leftarrow F_i - P_i$
  If $F_i^R = \emptyset$, no feasible assignment exists
  $F_i^T \leftarrow F_i^R$
For each $v_i \in V$ do
  For each v_j s.t. $(v_i, v_j) \in E$ do
    $F_i^T \leftarrow F_i^T \cap F_j^R$
For each $V_i \in V$ do
  $B \leftarrow \emptyset$ (indices of boundary nodes)
  For each v_j s.t. $(v_i, v_j) \in E$ do
    If $F_i^T \cap F_j^T = \emptyset$
      If $F_i^T \cap F_j^R \neq \emptyset$ and $F_j^T \cap F_i^R \neq \emptyset$
        $B \leftarrow B \cup \{i\} \cup \{j\}$
      Else, no feasible assignment exists
For each $b \in B$ do
  $S \leftarrow \{v_b\}$ (nodes with common broadcast frequencies)
  $C \leftarrow F_b^T$ (common broadcast frequencies)
  $U \leftarrow \{v_b\}$
  For each $v_i \in U$ do
    For each v_j s.t. $(v_i, v_j) \in E$ do
      If $F_i^T \cap F_j^T \neq \emptyset$
        $S \leftarrow S \cup \{v_j\}$
        $C \leftarrow C \cap F_j^T$
        $U \leftarrow U \cup \{v_j\}$
        If $v_j \in B$
          $B \leftarrow B - \{v_j\}$
    $U \leftarrow U - \{v_i\}$
  $c \leftarrow$ any frequency in C
  For each $v_i \in S$ do
    $c_i \leftarrow c$ (selected broadcast frequency for v_i)

Figure 7.1: Frequency Assignment Algorithm in [Steenstrup_05]

Other authors do not assume a central controller, but instead assume extensive message passing between the devices so each radio can effectively calculate the same solution.
For instance, [Zhao_05] considers a network of orthogonal channels where adaptive secondary users coordinate their adaptations via a common channel. [Etkin_05] considers a system wherein optimal frequency/power allocations are achieved by employing punishment strategies similar to the ones considered in Chapter 4 but applied to DFS

where the optimal strategy is known *a priori.* [Nie_05] considers a DSA scheme wherein radios must share information over a common channel to compute the interference levels each radio would induce to other radios in order to evaluate its goal (U2 in [5]). While this has the virtue of being both an exact potential game and an IRN, it requires significant overhead to distribute the information needed to evaluate the goal and requires that decisions are made sequentially. For DSA systems where spreading codes adapted (viewed in the context of signal space representations, spreading code adaptation algorithms could be directly applied to DFS problems), [Sung_03a] presents an algorithm where each radio's goal incorporates the interference measurements of all other radios in the system. [Xing_06] considers a *Homo Egualis* ("fair man") implementation where each access point chooses frequencies so as to maximize (7.3) where x_j is the usable spectrum for user j and $a_i, b_i \in \mathbb{R}$. Thus each access point attempts to ensure that every access point is receiving approximately the same amount of spectrum.

$$u_i = x_i - \frac{a_i}{n-1}\sum_{x_j > x_i}\left(x_j - x_i\right) - \frac{b_i}{n-1}\sum_{x_j < x_i}\left(x_i - x_j\right) \tag{7.3}$$

[Villegas_05] considers a distributed graph coloring algorithm where edges are formed between interfering access points. Each access node recursively distributes frequency and interference measurements and selects the frequency it believes will result in the least interference.

Other authors have considered single cell adaptations without the need for communication beyond reporting measurements from a common receiver. As discussed in Chapter 6, [Sung_03b], [Hicks_04], and [Ulukus_04] consider spreading code adaptations where each access node is isolated in frequency and spreading codes are chosen so as to minimize the interference of clients/mobiles are – a situation analogous in signal space to DFS applied to the clients in a single isolated cluster.

[Nie_05] also proposes another goal (or utility function) for DSA (U1) that is identical to the goal used in this paper (equation (7.1)). However, because [Nie_5] places no restrictions on the observation mechanism, [Nie_05] is unable to show that system forms an exact potential game which would permit the use of a simple distributed and autonomous algorithm. Instead [Nie_05] employs a no-regret learning algorithm wherein the radios autonomously try every possible frequency and then adapt to frequencies that yield the best weighted cumulative utility and show that the algorithm converges to a mixed-strategy equilibrium – a less than optimal result as mixed strategies in frequency selection imply continuous probabilistic adaptation.

[Luo_04] considers a closely related algorithm applied to a regular 10x10 grid of access points where each radio is guided by (7.4)

$$u_i(f_k) = \frac{M_i}{\sum_{M_j \in S_k} M_j} f\left(\sum_{M_j \in S_k} M_j\right) \tag{7.4}$$

where M_i is the number of users attached to access node *i*, S_k is the set of nodes operating on f_k and *f* evaluates the throughput if for the users in the argument. Each access node then chooses the channel that maximizes its throughput and switches to it with a random probability.

does not have to infer what other radios are experiencing, does not does not assume the existence of a centralized decision maker proposes a low-complexity autonomous distributed DFS algorithm suitable for use in ad-hoc 802.11h networks. After briefly defining the concepts of interference reducing networks and exact potential games and defining the proposed algorithm, this paper shows via analysis and simulation that this algorithm results in a frequency allocation that is a minimizer of sum network interference even when different policies are applied to different channels, asynchronous decision timings are used, access nodes exhibit private frequency preferences, and spectral signals are imperfectly estimated. Additionally, the impact of combining transmit power control (TPC) with our DFS algorithm is explored.

7.1.3 Interference Reducing Networks

Then a cognitive radio network is said to be an *interference reducing network* (IRN) if all adaptations decrease the value of the sum of observed interference levels $\Phi(f) = \sum_{i \in N} I_i(f)$

where $I_i(f)$ is the interference observed by cognitive radio *i* when the frequency vector $f \in F$ is implemented by *N*.

Chapter 6 states that an IRN can be realized in a distributed and autonomous fashion by selfish interference minimizing radios if adaptations are made by only one radio at a time under if the condition of *bilateral symmetric interference* (BSI) holds which happens if $g_{ki} p_k s(f_i, f_k) = g_{ik} p_i s(f_k, f_i) \forall f_k \in F_k, \forall f_i \in F_i$. BSI implies that a network is an IRN for unilateral adaptations because BSI implies that is an *exact potential game*. An exact potential game is a normal form game for which there exists a function, called the *exact potential function*, $V : F \to \mathbb{R}$, such that $u_i(\hat{f}_i, f_{-i}) - u_i(f_i, f_{-i}) = V(\hat{f}_i, f_{-i}) - V(f_i, f_{-i}) \forall i \in N$ where f_{-i} refers to the *n*-1 dimensional vector formed by excluding the contribution of *i* from *f*. By examining this definition, it is apparent that when profitable unilateral adaptations are made in an exact potential game, *V* constitutes a monotonically increasing sequence. Since when BSI holds, $\Phi(f) = -2V(f)$ Chapter 6, a monotonically increasing *V* implies a monotonically decreasing $\Phi(f)$ and an IRN is realized. This monotonicity property can then be used to prove the convergence of all selfish decision rules with unilateral timings Chapter 6.

7.2 An IRN DFS Algorithm

As opposed to algorithms with a centralized decision maker, or a single cell network, or a network that requires significant message passing just for the radios to evaluate their own goal, this chapter presents a DFS algorithm proposed in [Neel_06] which is completely distributed and requires no message passing between clusters (presumably an access point has to signal its users when changing frequencies). Further it does this without requiring complex computations or observations – the radio merely has to measure the received power of the RTS/CTS messages sent by other access nodes in the network and then choose a channel that the radio believes will reduce its interference.

This simple algorithm works because it satisfies the Interference Reducing Network (IRN) framework, and in particular, the Bilateral Symmetric Interference condition. This implies that the network's goals and frequency space form an exact potential game with a potential function which is a scalar multiple of the negation of the network interference function, $\Phi(f)$. The following provides detailed information about the proposed algorithm.

7.2.1 Algorithm Details

The proposed algorithm can be described as follows. Suppose each access node maintains a table with $|F_i|$ entries initialized to zero corresponding to the $|F_i|$ channels available to the network. Whenever access node i detects an RTS/CTS signal from another access node, j, i adds the power received from j to the table entry corresponding to the channel used by j. If j had been previously observed, then its previous received power is subtracted from its previous entry so that an access node only impacts a single table entry.

Now let us make the following assumptions about the network.

(A1) All RTS/CTS messages are transmitted at the same power level – a reasonable assumption as these messages are typically transmitted at maximum power to clear out hidden nodes.

(A2) The access nodes are not mobile so that $g_{ij} = g_{ji}$ or that power measurements are averaged over a long period so that $g_{ij} = g_{ji}$ is approximately true.

(A3) All channels have the same bandwidth.

(A4) A single access node adapts at a time.

(A1)-(A3) imply that $g_{ji} p_j s\left(f_i, f_j\right) = g_{ij} p_i s\left(f_j, f_i\right)$ which means that the bilateral symmetric interference of Chapter 6 holds. By the methods of Chapter 6, if each adaptation reduces the access node's observed interference, then the network is an IRN if one access node adapts at a time which is assured by (A4). While (A4) is a difficult condition to assure, if the frequency space is finite, then it can be shown that asynchronous adaptations converge to the same set of steady-states. As the network is an IRN, it is expected that this distributed algorithm will converge to a low interference steady-state which would be stable if the equilibria are isolated.

7.2.2 An 802.11h Application

As Chapter 6 asserts, since the only requirement on the decision process of cognitive radio is that adaptations increase (7.1) in order to decrease (7.2), great variation in the implementation of the decision process is permissible. Here we assume that each access node implements the following protocol:

1) Observe the spectral energy of the RTS/CTS messages of all observable access nodes.
2) At the time of its choosing, choose the channel on which the node observed the least amount of energy.

Consider a network of 802.11h access nodes (and presumably their client devices, but as the client devices are not involved in the decision process, they are irrelevant to the interactive decision problem). Suppose the access nodes are policy constrained to operate in the eleven channels available in the 5.47-5.725 GHz European band (channels 100-140) so that the assumption that "all RTS/CTS are transmitted at the same power level" holds for all channels (in this case, 1 W). Further, let us assume each radio has an equal probability of being the only radio allowed to adapt at each instance. As this is just a direct application of the general RTS/CTS DFS algorithm in Chapter 6 (where σ is now a binary function and discrete channels are used), we expect that the network will automatically sort itself into a low-interference frequency reuse pattern and that each adaptation will reduce the sum of perceived interferences in the network.

These expectations are confirmed in simulation of thirty access nodes randomly distributed over 1 km^2 operating in an environment with a path loss exponent of 3 with random placements and random initial channels and noise floors of -90 dBm. The geographic distribution of devices and their final operating frequencies are shown in Figure 7.2 where a circle notes the position of an access node with its final channel id labeled just below and to the right of the circle. Figure 7.3 depicts the operational channels for each access node (top), perceived interference levels by the access nodes (middle), and the sum of perceived interference levels (bottom) for the simulated network. Note that $\Phi(f)$ (bottom) decreases with each adaptation thereby satisfying the definition of an interference reducing network even though there are instances of interference increasing for individual access nodes (middle). Thus as is the case for all IRNs, self-interested adaptations led to a socially desirable outcome (at least when socially desirable is defined as the sum of observed network interference levels). As this algorithm converges to a minimum of $\Phi(f)$ (though not necessarily the global minimum), the algorithm performs at least as good as the centralized local search algorithm of [Leung_03] if no restarts are employed. So somewhat remarkably, this scalable distributed low complexity algorithm yields results as good as the high complexity centralized algorithm – a rare case of a "free lunch" in an engineering application.

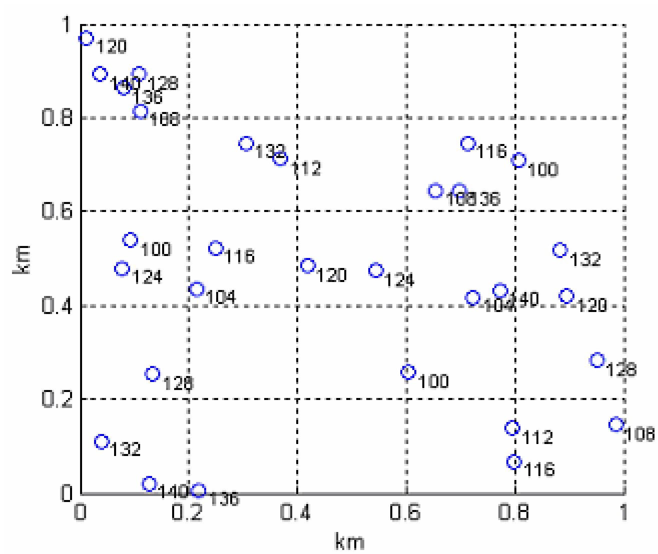

Figure 7.2: Steady-state Channels Selected for a Random Distribution of Access Nodes with Random Initial Channels in the 5.47-5.725 GHz Band.

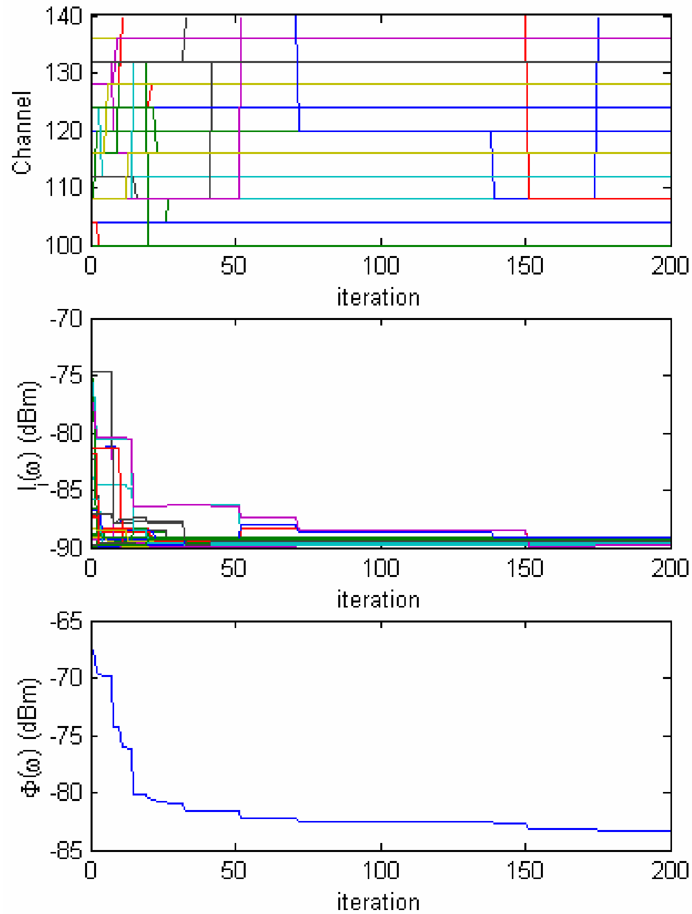

Figure 7.3: Instantaneous Statistics for Network in Figure 7.2.

7.3 Algorithm under Non-Ideal Conditions

In the preceding, we made a number of assumptions to make the network be an ideal IRN. As the following progressively relaxes these assumptions, it is seen that the proposed algorithm retains its desirable properties.

7.3.1 Policy Variations

If we permit the radios to choose permissible channels beyond channels 100-140, the assumption that all RTS-CTS messages are transmitted at the same power level fails as the lower and middle UNII bands (channels 36-64) limit transmission power levels to 200 mW [Etkin_05]. This violates (A1) ($p_k = p_i \ \forall i,k \in N$). However, for non-overlapping signals, $\sigma(f_i,f_k) = \sigma(f_k,f_i) = 0$, so BSI still holds and the network is still an IRN. Repeating the previous simulation and changing only the permissible channels and reflecting the transmission power policy variation we get the instantaneous statistics shown in Figure 7.4 where it is evident that the network continues to be an IRN.

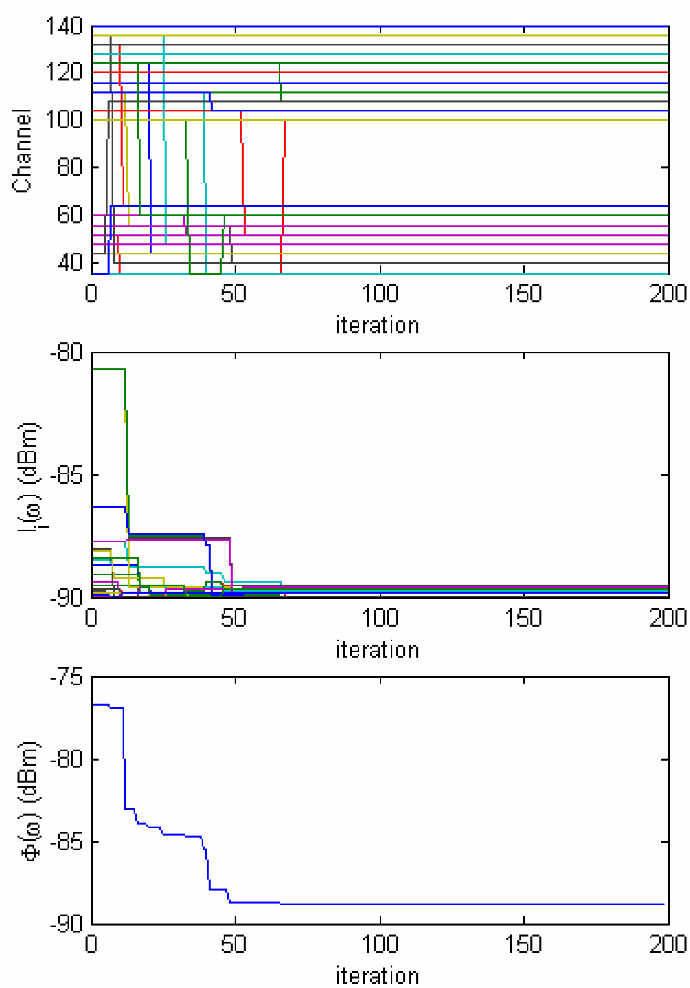

Figure 7.4: Instantaneous Statistics with Policy Variations

7.3.2 Asynchronous Timing

In the preceding, we assumed that one and only one access node adapted at any instance in time (A4). However, because adaptations and observation processes do not occur in infinitesimal periods of time it is likely that multiple access nodes will occasionally adapt simultaneously – a trend that becomes more likely as the number of access nodes in the network increase. So now continue to have the policy variations of the previous section and now assume that (A4) does not holds and instead of assume that each access has an opportunity to adapt at each iteration of the algorithm with non-zero probability.

Following the algorithm considered in this paper and the relaxed timing constraint two radios which are operating in the same channel and in close proximity to each other could simultaneously choose to adapt to another channel where a distant radio is operating. In this case, $\Phi(f)$ would increase even though each radio chose the channel which the radio had measured as having the least interference. Thus with (A4) relaxed, the proposed algorithm cannot be guaranteed to yield the strict monotonicity required by the definition of an IRN.

Yet this network will still converge to a steady-state with that is a minimizer of $\Phi(f)$. This again is a result of $<N,F,\{u_i\}>$ forming an finite exact potential game which implies FIP. As it is an exact potential game, minimizers of $\Phi(f)$ are Nash equilibria and the game has FIP which means that from any starting state, every sequence of self-interested unilateral adaptations must terminate in a minimizer of $\Phi(f)$. Due to these two properties, the network can be modeled as an absorbing Markov chain where minimizers of $\Phi(f)$ are the absorbing states of the chain. By virtue of being a minimizer, there can be no unilateral deviations that reduce interference; thus minimizers are absorbing states. By virtue of the finite improvement path property, there always exists a sequence of adaptations of non-zero probability that terminate in a minimizer as long as the probability of a unilateral deviation is always nonzero. Thus even with (A4) relaxed to asynchronous timings for adaptations, the network will still converge to a minimizer of $\Phi(f)$.

To verify this assertion, we modified the preceding simulation so that at each "iteration" each access node had an opportunity to adapt with probability 0.02. The instantaneous statistics for this simulation are shown in Figure 7.5. While $\Phi(f)$ still trends down, it longer does so monotonically. Nonetheless, because this system forms an absorbing Markov chain, it eventually converges to a frequency vector that is a minimizer of $\Phi(f)$.

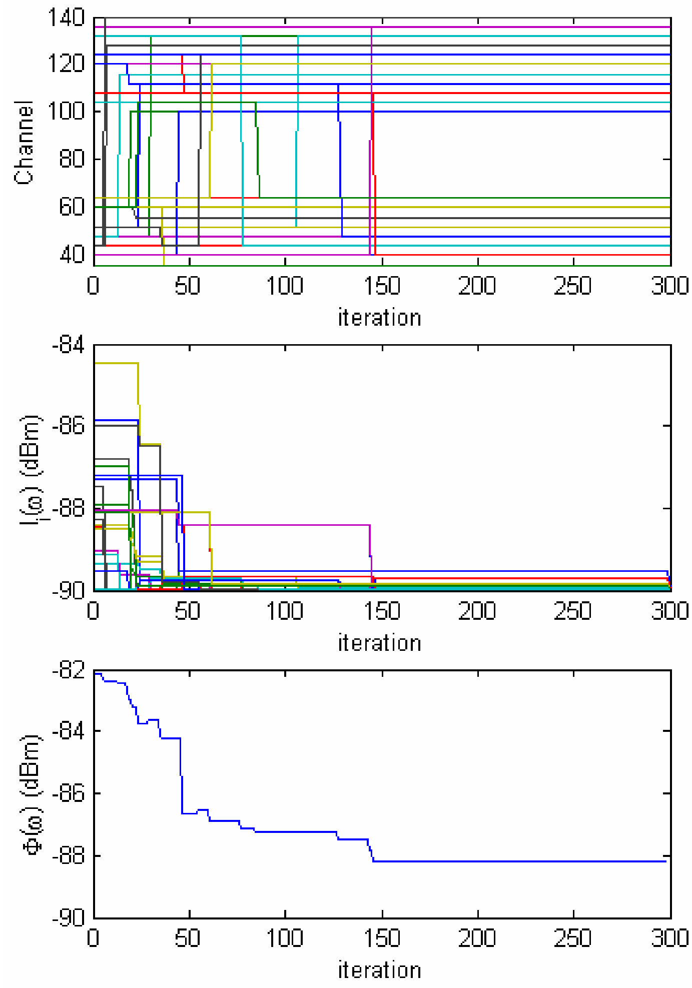

Figure 7.5: Impact of Asynchronous Decision Timings

7.3.3 Private Frequency Preferences

Throughout this discussion we have assumed via (A2) that each access node only intends to minimize the interference it perceives from other adaptive access nodes. However, because of the presence of interferers or because of local channel conditions, different access nodes may also exhibit different preferences for different frequencies. If we denote the frequency preferences of access node i as $S_i(f_i)$, these preferences might be incorporated as shown in (7.5).

$$\tilde{u}_i(f) = \sum_{k \in N \setminus i} g_{ki} p_k s(f_i, f_k) - S_i(f_i) \tag{7.5}$$

Note that $S_i(f_i)$ indicates that this component for access node i is only a function of access node i's choice of frequency and makes the most sense express additively as in (7.3) where $S_i(f_i)$ models the influence of static interferers.

Under the assumption that $S_i(f_i)$ models static interferers in the environment (7.2) no longer reflects the sum network interference. Instead sum network interference with frequency preferences is given by (7.6) .

$$\Phi^S(f) = \sum_{i\in N}\left(S_i(f_i) + \sum_{k\in N\setminus i} g_{ki}p_k s(f_k, f_i)\right) \tag{7.6}$$

This inclusion of additional interferers or jammers or local channel conditions may also impact bilateral symmetric interference as the interferers may not be transmitting at the same power level as the cognitive radios or may be operating with differing bandwidths.

Regardless of the loss of bilateral symmetric interference due to variances in the static interferers, $\langle N, \Omega, \{u_i\}\rangle$ remains an exact potential game but with an exact potential function given by (7.7).

$$V^S(w) = -\sum_{i=1}^{n}\left(S_i(f_i) + \sum_{k=i+1}^{n} g_{ki}p_k s(f_k, f_i)\right) \tag{7.7}$$

Note that the differences between (7.6) and (7.7) imply that the network is not strictly an IRN. Consider the scenario where a unilateral adaptation is made from a channel that is originally only occupied by the adapting access node i and a static interferer to a channel that is occupied only by access node k such that (7.8) holds.

$$g_{ki}p_k s(f_i, f_k) < S_i(f_i) < 2g_{ki}p_k s(f_i, f_k) \tag{7.8}$$

This adaptation would increase (7.5) – thereby satisfying the proposed algorithm – but (7.6) would also increase – violating the definition of an IRN. However, the exact potential in (7.7) will always increase, ensuring the algorithm's convergence. And when the only maximizers of (7.7) are those for which $S_i(f_i)=0$ $\forall i\in N$, the algorithm will converge to a minimizer of (7.6) as under this condition $\Phi^S(f) = -2V(f)$. Even though it is trivial to construct two-access node, two channel, single interferer scenario with non-random geographic and channel distributions where (7.8) is satisfied, repeated trials of our randomly placed, random initial channel simulation have not yielded an adaptation that satisfies (7.8), which indicates the condition might be rare in practical settings. For example, modifying the policy variation simulation so it includes five static interferers operate in both channels 132 and 136, but distributed randomly geographically yield the simulation shown in Figure 7.6.

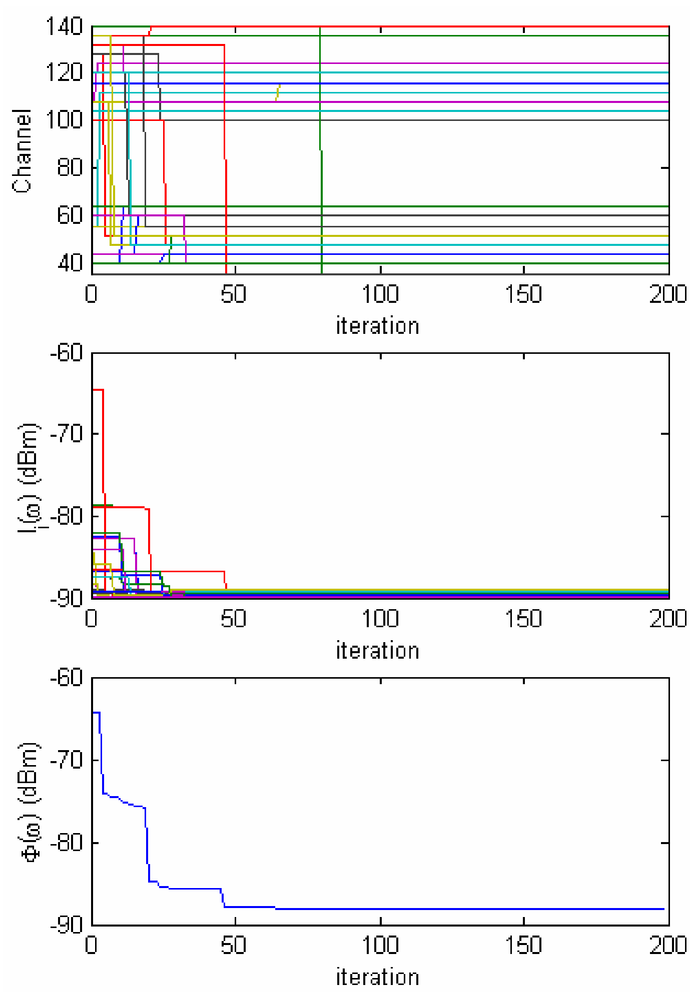

Figure 7.6: Algorithm with Private Frequency Preferences

7.3.4 Effect of Estimations

Throughout the preceding, we have implicitly assumed that the access nodes are perfectly measuring the signal strength of the RTS/CTS signals. However, in a practical setting, measurements of interference levels in differing channels would be corrupted by receiver noise, non RTS/CTS signals, and RTS/CTS signals too weak to decode and recognize as from an access node. Thus measurements of received power will at best be corrupted estimations. In such a scenario, the access nodes' goals would again take the form as shown in (7.5) but with $S_i(f_i)$ a stochastic variable.

As shown in the preceding section, a goal of the form of (7.5) implies that while $<N,F,\{u_i\}>$ is still an exact potential game, the network will not necessarily remain an IRN for all possible realizations. Further, for channels with very low interference levels, $S_i(f_i)$ may be a dominant term and its natural time variation may spawn unnecessary adaptations.

For example consider a modification of the preceding simulation where the -90 dBm noise floor is implemented as a Gaussian stochastic variable. The results of this simulation are shown in Figure 7.7. While the algorithm still yields an almost 15 dB reduction in interference levels from the initial random distribution, $\Phi(f)$ is no longer monotonic, overall performance is decreased and significant bandwidth would be wasted signaling all of these adaptations. However, by modifying the algorithm so the access nodes only adapt if the improvement in performance is predicted to be more than a small threshold (-85 dBm), the system behaves as shown in Figure 7.8 – generally like a convergent IRN, but with the caveat that there exists the small probability that either an adaptation may increase sum interference. Based on the discussion of Chapter 5, we recognize this as an ε-better response decision.

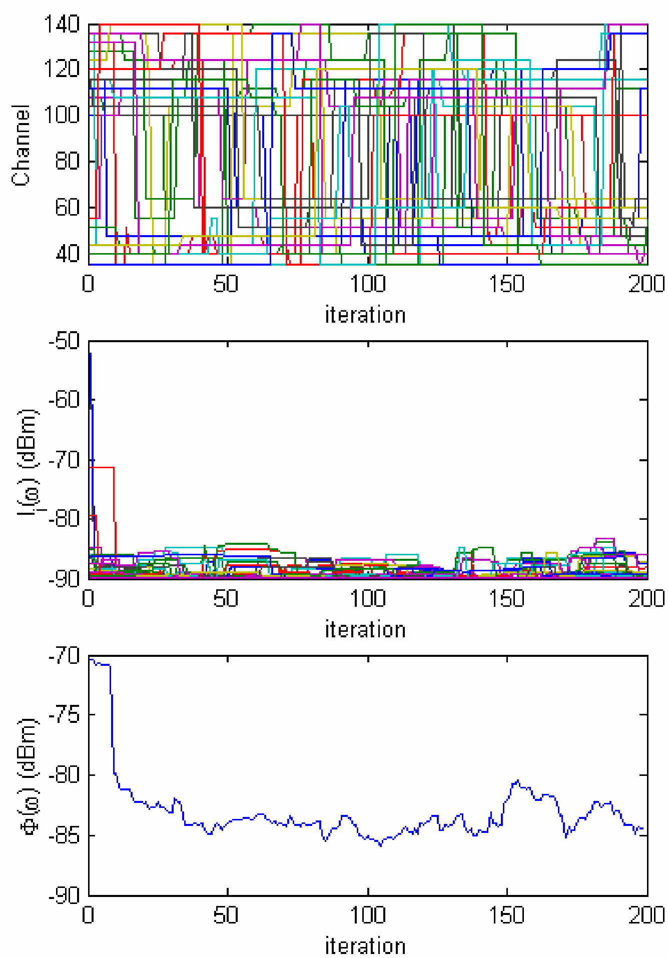

Figure 7.7: Algorithm with Stochastic Estimations.

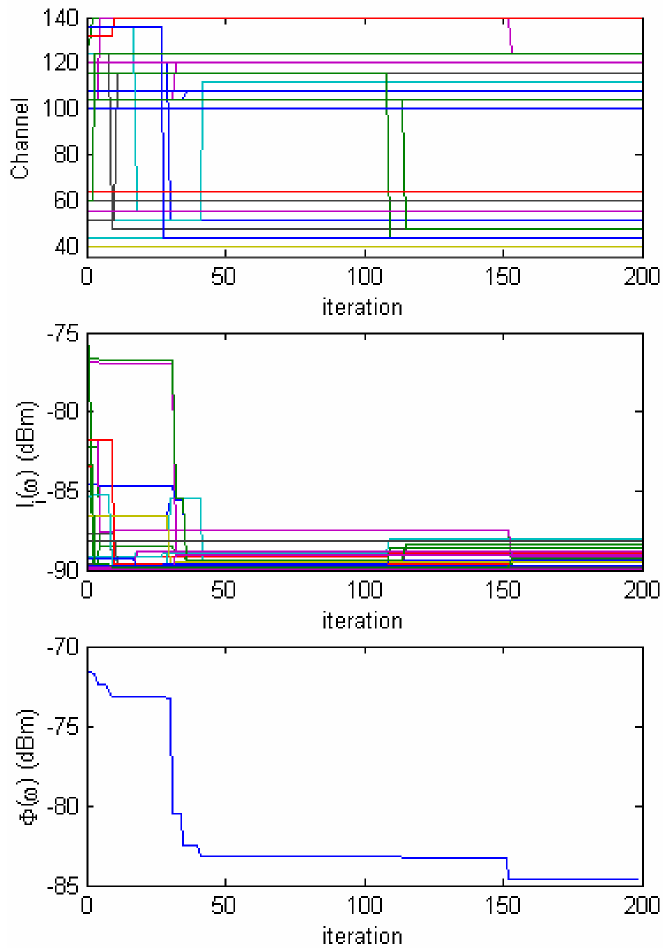

Figure 7.8: Algorithm with Stochastic Estimations and a small adaptation threshold (-85 dBm)

While this modified decision rule is stable, it is somewhat at odds with how we have been treating the stability of decision rules. Specifically, we have been primarily operating under the assumption that each decision rule can be characterized as $d_i : A \to A_i$ and characterizing stability as for every ε>0 there exists some δ>0 such that if the system is perturbed off the steady state, a^* by a distance no greater than $\|a^* - a\| < d$ then the system will remain no further away from a^* than $\|a^* - a\| < e$ absence further perturbances. Because of the finiteness of the action space and the lack of isolated steady-states, this formulation is impossible.

Yet the system is stable as shown in the simulation. Recall that in the model in Chapter 2 we said that expressing the decision rule as a function of the action space was really only a useful analytic conceit for decision rules that are a function of the outcome space and that the radios were really observing and reacting to the outcome space. So rather than $d_i : A \rightarrow A_i$ the radios are actually implementing $d_i : O \rightarrow A_i$. When we introduce noise to the observations of the radio, this has the effect of perturbing the observed outcomes in *O*. So we are actually referring to the stability of $d_i : O \rightarrow A_i$. With this in mind, it is relatively easy to show that the threshold causes d_i to be stable. Because of the threshold in the decision rule and assuming the system is at an NE, it now takes a perturbation in the outcome space at least as great as the threshold to induce an adaptation. Or more formally, for any arbitrarily small $\varepsilon>0$ and assuming that the system is operating at an NE, for all $\tau>\delta>0$ the system remains within an ε of the steady-state, specifically, the system remains on the original steady-state. Note that the thresholded decision rule induces many more equilibria in the system (specifically a number of ε-NE) so different steady-states may have smaller permissible values for δ.

7.3.5 TPC and DFS

[Etkin_05] states that TPC is intended to support variations in policy and adaptations based on "a range of information including path loss and link margin." As we showed in the Policy Variations section, as long as it is applied consistently across a channel policy variations do not impact the IRN features of the algorithm. However, if the RTS/CTS power levels are set at varying levels by the differing access nodes operating in the same channel, then it is likely that (A1) will be violated in situations where $s(f_i, f_k) \neq 0$ which means the BSI condition will not be satisfied. For instance, consider a modification of the original policy variation simulation where the transmit power that each access node applies to its RTS/CTS signals has been scaled by a factor randomly drawn from a clipped Gaussian distribution (clipped so as to rule out negative power levels) whose results are shown in Figure 7.9. Note that $\Phi(f)$ does not decrease monotonically in this simulation, though it does trend downwards fairly consistently and converges for all simulations to date. When TPC is applied to the RTS/CTS messages, it is observed that the system still converges to an interference minimizer. Currently, we do not have a firm analytic explanation for this phenomenon, though it is known that for relatively small variations in transmit power levels, $\Phi(f)$ will be an ordinal potential function for $<N, F, \{u_i\}>$ so for many realizations of TPC applied to RTS/CTS signals the network will still behave as an IRN. However, without a firm analytical basis for stating why desirable behavior results we are unable to rule out unforeseen pitfalls from the interactions. So we recommend that application of the proposed algorithm be limited to scenarios where TPC is applied only to the DATA and ACK messages. While this assumption would still enable improved battery life and would be consistent with the RTS/CTS messages original intent for clearing out hidden nodes, it would limit the gains seen from frequency reuse. However, all localized TPC schemes face a functionality tradeoff of clearing out hidden nodes versus maximizing frequency reuse. By reducing transmit power on the RTS/CTS message, a higher cluster density can be achieved, but this comes at a cost of increasing the probability that a hidden node will miss the RTS/CTS signal and subsequently interfere with the data transfer, particularly where TPC is guided by local decisions instead of policy.

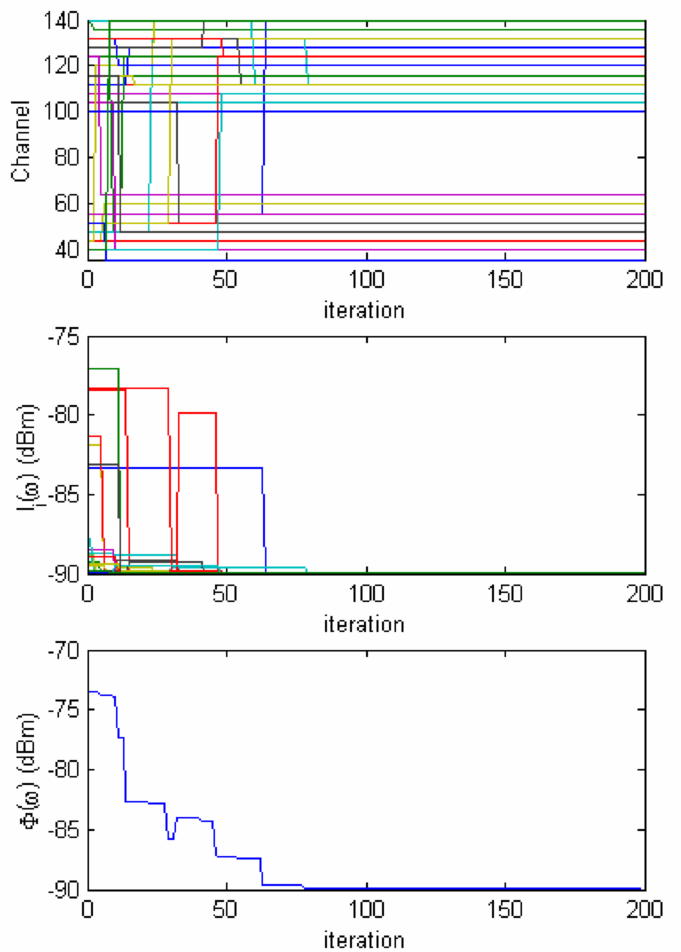

Figure 7.9: Algorithm with TPC Applied to RTS/CTS.

7.4 Summary and Conclusions

This chapter presented a novel algorithm for performing DFS which does not require the use of a centralized controller, specialized network topologies, or even any message passing between nodes but still achieves performance as good as could be expected from a centralized heuristic algorithm, e.g., [Leung_03]. The remainder of this section summarizes the results of the DFS algorithm, describes how game theory aided the design, and describes further extensions that can be made to this algorithm.

7.4.1 Algorithm Summary

By leveraging the framework of interference reducing networks, this chapter proposed a low complexity autonomous distributed ad-hoc DFS algorithm whose adaptations converge to a minimizer of the sum of observed interference levels by minimizing their own perceived interference measured from the RTS/CTS signals of other access nodes. We showed that this non-cooperative non-collaborative algorithm is robust to policy variations, timing variations, the presence of interferers, and noisy estimations of signal strengths when a simple adaptation threshold is applied to the algorithm. Though empirically convergent, when TPC is applied to the RTS/CTS signals, the algorithm fails to satisfy the IRN framework. However, the assumption of TPC applied to RTS/CTS signals may not be realistic as it necessarily increases susceptibility to hidden nodes. While all simulations implemented a best-response dynamic, any self-interested decision rule – including an ontological reasoning engine – will converge by virtue of being an exact potential game and an IRN.

7.4.2 How Game Theory Aided the Design

This chapter demonstrated what can be gained by leveraging the techniques of the previous chapters. We knew that if a cognitive radio network could be designed as a potential game, then unilateral deviations would converge and the maximizers of the system's potential function would be steady-states. The one hole in applying potential games is that convergence is not guaranteed to be to a desirable steady-state. As noted in Chapter 5, this problem can be solved by designing networks where the potential function is the design objective function. Chapter 6 gave the framework for doing this wherein the concept of bilateral symmetric interference guaranteed that self-interested interference minimizing adaptations yield an interference reducing network. The condition of bilateral symmetric interference was an application of bilateral symmetric interaction exact potential game to cognitive radio networks. Asynchronous convergence was then assured by the FIP convergence analysis of Chapter 4 which was in turn aided by the Markov models of Chapter 3. We also knew that the network would be stable – a concept introduced in Chapter 3 – because of the consideration of the stability of ε-better responses in finite potential games.

Because of game theory and the results established in the preceding chapters, we knew that a low complexity, scalable DFS algorithm, convergent, stable, and desirable network would result if we could get the network to satisfy the bilateral symmetric interference condition. And because of game theory and the earlier results we knew that this would occur without having to resort to a centralized controller, without message passing between the radios so they could all independently find the same solution, and without resorting to specialized network topologies.

With all this in hand, the only insight needed to design this network was identifying a situation where symmetric link gains and equal transmit powers could reasonably be assumed to be present. And this is satisfied simply by having each access point tabulate the channel and received power of the RTS/CTS messages of other access points which it can observe.

Beyond the single best response decision rule presented in the original paper and covered in this chapter, the FIP results of Chapter 4 and Chapter 5 also assure us that the combination of decision rules and timings whose entries contain a 'Y' in Table 7.1 also converge. Thus with the correct observations, goals, and action space in place many different scenarios are known to converge.

Table 7.1: Other Conditions Guaranteed to Converge to a Low Interference State

	Timings			
Decision Rules	Round-Robin	Random	Synchronous	Asynchronous
Best Response	Y	Y	N	Y
Exhaustive Better Response	Y	Y	N	Y
Random Better Response[(a)]	Y	Y	Y	Y
Random Better Response[(b)]	Y	Y	N	Y
ε-Better Response[(c)]	Y	Y	N	Y
Intelligently Random Better Response	Y	Y	N	Y

(a) Proposed random better response (b) Random better response of [Friedman_01] (c) Convergence to an ε-NE

7.4.3 Further Extensions

This algorithm need not be specifically limited to finite channel sets (though the table entry routine would require modification) nor to nonoverlapping channels as relaxing these assumptions still preserves the bilateral symmetry assumption that $s(f_i, f_k) = s(f_k, f_i) = \max\{B - |f_i - f_k|, 0\}/B$. For instance, the same algorithm could be readily applied to a 2.4 GHz 802.11b network which has 11 channels where at most 3 channels (1, 6, and 11) can be made to not overlap. However, because the bilateral symmetry assumption still holds, the observation of access nodes' RTS/CTS signals will still satisfy bilateral symmetric interference.

Obviously, this algorithm can be extended to the other decision rules and timings listed in Table 7.1. Because the random better response decision rules converge, such an algorithm could be easily inserted into genetic algorithm based cognitive radios and used as part of a rapidly deployed network of radios under the assumption that symmetric link gains and constant observed power levels still hold. Further, because all exhaustive better response algorithms converge, it seems reasonable that if the key observation (RTS/CTS signals of access nodes) is used to drive the decision process of ontologically defined cognitive radios, an interference reducing network will still emerge.

So in the end, we can replace the need for a wireless PhD student on every street with the simple requirements that 1) each access point notes the received power and channels of the RTS/CTS messages from other access points that it can observe and 2) each access point acts to reduce its own perceived interference.

7.5 References

[Etkin_05] R. Etkin, A. Parekh, D. Tse, "Spectrum Sharing for Unlicensed Bands," *DySPAN2005*, pp. 251-258, Nov. 2005.

[Hicks_04] J. Hicks, A. MacKenzie, J. Neel, J. Reed, "A Game Theory Perspective on Interference Avoidance," *IEEE GlobeCom*, vol.1, pp. 257-261, Dec. 2004.

[Leung_03] K. Leung, B. Kim, "Frequency Assignment for 802.11n Wireless Networks," VTC 2003, vol. 3, pp. 1422-1426, Oct. 6-9 2003.

[Liu_05] X. Liu, W. Wang, "On the characteristics of spectrum-agile communication networks," *DySPAN 2005*, pp. 214-223, Nov. 8-11 2005.

[Luo_04] H. Luo, N. Shankaranarayanan, "A distributed dynamic channel allocation technique for throughput improvement in a dense WLAN environment," *ICASSP '04,* vol. 5, pp. 345-348, May 17-21, 2004.

[Neel_06] J. Neel, J. Reed, "Performance of Distributed Dynamic Frequency Selection Schemes for Interference Reducing Networks," Accepted to *Milcom 2006*. Washington D.C. Oct 23-25, 2006.

[Nie_05] N. Nie, C. Comaniciu, "Adaptive channel allocation spectrum etiquette for cognitive radio networks," *DySPAN2005*, Nov. 2005 pp. 269-278.

[Sung_03a] C. Sung, K. Leung , "On the stability of distributed sequence adaptation for cellular asynchronous DS-CDMA systems," *IEEE Transactions on Information Theory*, vol. 49, no. 7, July 2003, pp. 1828-1831.

[Sung_03b] C. Sung, K.W. Shum and K. Leung, "Multi-objective power control and signature sequence adaptation for synchronous CDMA systems - a game-theoretic viewpoint", *Proceedings of the IEEE International Symposium on Information Theory*, July 2003, pp. 335-335.

[Steenstrup_05] M. Steenstrup, **"**Opportunistic use of radio-frequency spectrum: a network perspective," *DySPAN2005*, Nov. 2005 pp. 638-641.

[Ulukus_01] S. Ulukus and R.D. Yates, "Iterative construction of optimum signature sequence sets in synchronous CDMA systems," *IEEE Transactions on Information Theory*, vol. 47, no. 5, July 2001, pp. 1989-1998.

[Villegas_05] E.Garcia Villegas, R. Vidal Ferr0, J. Paradells Aspas. "Implementation of a Distributed Dynamic Channel Assignment Mechanism for IEEE 802.11 Networks," *16th International Symposium on Personal, Indoor and Mobile Radio Communications*, pp. 1458-1462, 2005.

[Xing_06] Y. Xing, R. Chandramouli, S. Mangold, S. Shankar, "Dynamic spectrum access in open spectrum wireless networks," IEEE Journal on Selected Areas in Communications, vol. 24, No. 3, pp. 626-637, March 2006.

[Zhao_05] J. Zhao, H. Zheng, G. Yang, "Distributed Coordination in Dynamic Spectrum Allocation Networks," *DySPAN 2005*, pp. 269-278, Nov. 2005.