

Spectrum Management: The Killer Application of Ad Hoc and Mesh Networking

John A Stine

The MITRE Corporation
McLean, VA 22102-7508
jstine@mitre.org

Abstract— The traditional approach to managing radio frequency (RF) spectrum concludes with a spectrum management administration identifying how a band of spectrum may be used in a region it administers and who may use it. This approach to spectrum management is encumbered with political process that is challenged to keep pace with the rapid advances in technologies that exploit RF spectrum. Recognizing that this process can stifle innovation, the FCC has proposed new spectrum management models and the use of a measure of interference temperature. Unfortunately, they too result in picking winners and losers and can have detrimental effects on legacy users. In this paper we review these concepts and then describe how ad hoc networks can offer new solutions. Three different spectrum management ideas are described. The Synchronous Collision Resolution* (SCR) MAC protocol enables a strict arbitration of spectrum access based on spectrum rights thus enabling a hierarchy of networks in the same spectrum that always guarantees the primary rights holder precedence. Second, it autonomously manages the use of an arbitrary number of channels in the same network all of which support the network achieving a higher capacity. The third and most exciting idea is a new Fast Command and Control model for spectrum management. An underlying ad hoc network built using the Nodes State Routing* (NSR) protocol is used to track and manage the use of spectrum of attached RF emitters. NSR tracks the state of the network by collecting and disseminating the states of the nodes. These states can include relevant information on the spectrum these nodes are using and are observing others use. Thus the network supports tracking and monitoring spectrum use spatially in near real time. Spectrum management utilities built on top of the network could allow users and spectrum managers to rapidly negotiate the use of spectrum for short periods of time in small regions. This finer resolution command and control model supports rapid prototyping of new services while simultaneously giving spectrum managers the ability to identify and turn off emitters that cause harmful interference to higher precedence users. We conclude with proposed standardization and regulatory changes to make this feasible.

Keywords—spectrum management, ad hoc networking, MANET, synchronous collision resolution, SCR, node state routing, NSR, fast command and control model, FCCM

I. INTRODUCTION

RADIO frequency (RF) spectrum is a critical resource for many services that people across the world rely on for their safety, employment, and entertainment. Technological advances are making further uses of RF spectrum possible

increasing demand for it and fueling competition among government, public, and commercial sectors for access. In the interest of all it is important to make the use of RF spectrum efficient. Spectrum efficiency is normally quantified in how well a particular technology uses spectrum, e.g. the rate bits are communicated per band of spectrum. This method of measurement is becoming less relevant as people realize that many users of spectrum receive dedicated bands of spectrum for their use but their applications may only require sporadic access leaving opportunities for others to use it. Creating efficiency in these circumstances requires more sophisticated spectrum management. Multiple spectrum management approaches have been proposed but are not getting traction for various reasons. In this paper we describe these multiple spectrum management mechanisms and propose a new approach to spectrum management which uses an underlying ad hoc network to coordinate spectrum use allowing short term licenses with enforcement. It would be built on top of the Synchronous Collision Resolution (SCR) and Node State Routing (NSR) protocols that are particularly well suited for this task.

We begin this paper with a brief overview of spectrum management and the different proposals for improvement. We follow with a discussion of the radio technologies, software defined radio (SDR), spectrum detection, and cognitive radio (CR), and the networking protocol technologies, SCR and NSR, that we exploit in our proposed spectrum management approaches. Our focus is on the benefits that the network and its protocols bring to the spectrum management task and so we follow with a description of the spectrum management mechanisms that are enabled by SCR and then NSR and then describe a broad spectrum management approach that would integrate these mechanisms into a fast command and control model (FCCM) for spectrum management. We describe the necessary changes in network standardization and spectrum regulation to make FCCM possible.

II. SPECTRUM MANAGEMENT

Spectrum is a renewable resource that is finite in any instant of time but through its different dimensions of use: space, time, frequency and bandwidth, can be distributed to many users simultaneously. The process of distributing spectrum to users is spectrum management. Traditionally this function has been performed globally through international agreements and nationally by government administrations. Bands of spectrum

*Patent pending

are divided into allocations that are designated to support particular services. The allocations are subdivided into allotments that may be used by administrations in specified geographic areas. National administrations may further allot the spectrum into channels, specify the conditions of their use, and assign (a.k.a. license) them to users. Historically, the growth in spectrum requirements was accommodated through technology making the higher frequency bands available for use. Little unassigned spectrum remains and so now spectrum management is the business of reallocating, re-allotting and reassigning spectrum. This places government, public, and commercial interests in tension as each has a perceived need for spectrum access and operational and financial stakes in the decisions that are made. In 2002 the Federal Communications Commission (FCC) established a Spectrum Policy Task Force (SPTF) to provide recommendations on how to evolve spectrum policy into an “integrated, market oriented approach that provides greater regulatory certainty while minimizing regulatory intervention.” [1] In November of that year it produced a report [2] that as its most significant recommendations proposed that the FCC move more spectrum from the command and control management model to the exclusive use and commons models and to employ an interference temperature measure as a new paradigm for interference protection. Below we review these different spectrum management models and the interference temperature concept and identify some of the pros and cons of each. We conclude with a brief description of the intent of our spectrum management approach and why it differs from these.

A. *Command and Control*

The command and control model is the legacy model where an administration licenses spectrum to users under specific conditions. Changing uses of spectrum is a deliberative process that involves study and opportunities for public comment. The major complaints against this approach are that it is very slow to adapt, it is unfriendly to commercial interests, and it results in inefficient use of spectrum. Nevertheless, the command and control model is still necessary to protect public interests that are not market-driven such as public safety, scientific research, and government operations, and to conform to treaty obligations. Even with the use of the other spectrum management models and the interference temperature, the command and control model will remain the overarching spectrum management model, the difference being that parts of the spectrum will have more liberal rules that allow commercial development and changing uses without the administrative proceedings.

B. *Exclusive Use*

The exclusive use model is a licensing approach in which the licensee has exclusive rights to a band of spectrum within a defined geographic region. The licensee has flexibility to implement different technologies and can transfer the use rights. The best example of the exclusive use model in practice is cellular telephony. The licensees develop the technologies, infrastructure and services and transfer spectrum use to subscribers of those services. There are great incentives to promote this model especially for the most desirable spectrum because licensees bid for the spectrum which brings revenue to

governments and creates the incentive that the licensees apply the spectrum for its best valued use. This model favors service providers.

C. *Commons*

The commons model opens bands of spectrum for unlicensed use with etiquettes that allow as much coexistence among different applications and users as feasible. An example of spectrum bands that are managed in this way are the industrial, scientific, and medical (ISM) bands. The 2.4 GHz ISM band has been very successful being used for wireless LAN, personal area networks, microwave ovens, cordless telephones, and other consumer products. This model favors manufacturers of consumer products.

D. *Interference Temperature*

Interference temperature is a measure of the combined interference and noise per unit bandwidth at receivers. The concept of interference temperature was conceived as a way to allow unlicensed use in licensed bands. Unlicensed users would be permitted to use spectrum so long as they did not exceed a particular interference temperature at nearby receivers. Adopting this type of measure is clearly intended to provide opportunities to manufacturers of ultra wide band devices since specifying broad enough bands under any of the spectrum management models would be impractical.

Implementation of the interference temperature concept requires measurement, a means to distribute those measurements, and some means to control the unlicensed users to prevent their violation of the interference temperature thresholds. Part of our contribution is the underlying fabric that enables the distribution of interference temperature measurements and control messages.

Inevitably, primary users have concerns about anyone being able to use their spectrum in this way. Their acceptance will depend on how well technology can be used to prevent harmful interference. Opposition is also motivated by the fact that this scheme creates the conditions for encroachment where secondary users effectively steal away the primary users’ spectrum rights. As an example, consider the bands shared by the primary government user and secondary users with Part 15 devices. In concept, the secondary user must accept interference from and cause no interference to the primary user. However, secondary users are often unaware of their secondary status and perceive the significance of their use to require protection. The combination of public perception and political process effectively steals the primary user’s access rights. This is seen repeatedly with Part 15 consumer products such as baby monitors and garage door openers that are interfered with by primary government users, normally near military bases, where the primary user is expected to avoid exercising their primary rights.

E. *Fast Command and Control*

This review of the management models and interference temperature demonstrates that although the goal is to make spectrum more useful, each technique still picks winners and losers and potentially creates the conditions that legitimate

users will have their spectrum rights violated without ability for recourse. In almost all approaches listed above, from the losers’ perspective, the results are no different than those of the “command and control” model. Decisions are slow and once made are fait accompli. The alternative we provide is a timelier command and control model with greater resolution. In this model, an ad hoc network provides an underlying control mechanism to license spectrum to users for short periods of time in smaller spaces. The network also monitors compliance with the existing policy. The great benefit provided by such a model is that it affords individuals all the way up to large commercial enterprises the opportunity to use spectrum as it is available and as they need it and can exploit it. It enables secondary markets.

III. DYNAMIC SPECTRUM USE TECHNOLOGIES

Here we provide background material on the technologies that can be exploited by our spectrum management approach.

A. Software Defined Radio

The FCC defines a software defined radio (SDR) as “A radio that includes a transmitter in which the operation parameters of frequency range, modulation type or maximum output power (either radiated or conducted), or the circumstances under which the transmitter operates in accordance with Commission rules, can be altered by making a change in software without making any changes to hardware components that affect the radio frequency emissions.” [11] The significance of such a radio is that its presence in spectrum can be changed through software modifications. It can be made to operate wherever spectrum is available. It enables the improvement of waveforms and spectrum use without the requirement for users to buy new devices. However, from the regulatory perspective, this is problematic if the radio can be made to operate where it does not have license to do so. The traditional way of regulating spectrum use has been to license equipment since the spectrum of operation was dictated at the time of manufacture. The SDR changes this paradigm and requires a protocol to control software changes. The contribution of this paper is to propose a network enabled method for a spectrum manager to authorize a change and to verify an appropriate use of spectrum.

B. Spectrum Detection

A benefit of an SDR is that it does an analog to digital conversion of the RF signal. This conversion allows the use of the Fast Fourier Transform (FFT) to identify spectral content. Over time this measurement can indicate the pattern of spectrum use. The sharing of this information among a set of physically distributed radios can be used to understand the spatial use of spectrum. The contributions of the protocols in this paper are a mechanism to distribute this information and an approach to incorporate it in the decisions of the network to assign spectrum for use.

C. Cognitive Radio

A cognitive radio (CR) is a radio that is able to modify its behavior based on external factors. For example, a source CR

may change the characteristics of its transmissions based on its observation of spectrum use. A typical instantiation of a CR would combine an SDR and its ability to detect spectrum use with some sort of logic to choose a modulation scheme that operates in the spectrum that is perceived idle and a method to assist receiver adaptation. The challenge in implementation and in regulation is that such a radio has limited ability to ascertain the ramifications of its choice. At the very least, a radio needs to coordinate spectrum use with the distant end but using spectrum use observations at two ends still does not guarantee that there are no receivers in the region that will be affected. A contribution of using cognitive radios in tandem with a network is that the network supports a spatial assessment of spectrum use and thus a better view to assess the ramification of using a band of spectrum.

IV. NETWORKING PROTOCOLS

The spectrum management methods that we propose are a direct result of the capabilities of the SCR and NSR protocols and are not be possible with any protocols that we are aware. The critical feature that makes these protocols uniquely qualified to support spectrum management is that they have been designed to manage spectrum in space rather than to manage link capacity. Several papers have been written on these protocols. [5], [6] Here we provide an overview of the features relevant to the spectrum management problem.

A. Synchronous Collision Resolution (SCR)

Synchronous Collision Resolution is a framework for MAC protocol design that has four key characteristics:

1. The wireless channel is slotted.
2. All nodes with packets to transmit attempt to gain access to every transmission slot.
3. Contending nodes use signaling to arbitrate their access.
4. All packet transmissions that occur during a transmission slot are sent simultaneously.

Fig. 1 illustrates the basic implementation of SCR. The transmission slot consists of three activities, collision resolution signaling (CRS) to select a subset of all possible contending nodes, a request-to-send (RTS) – clear-to-send (CTS) handshake used to verify capture and to assist physical layer adaptation, and finally the data transmission and acknowledgement (ACK).

1) Collision Resolution Signaling (CRS)

The goal of CRS is to select a subset of contenders from among all contending nodes in the network so that the nodes in the subset are physically separated from each other by at least

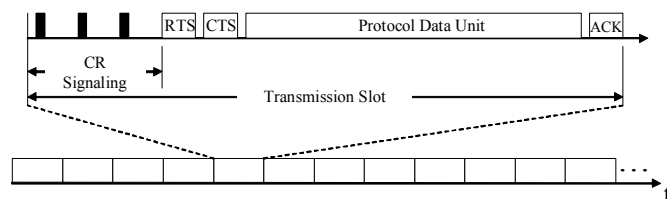


Fig. 1. Basic implementation of the Synchronous Collision Resolution MAC protocol

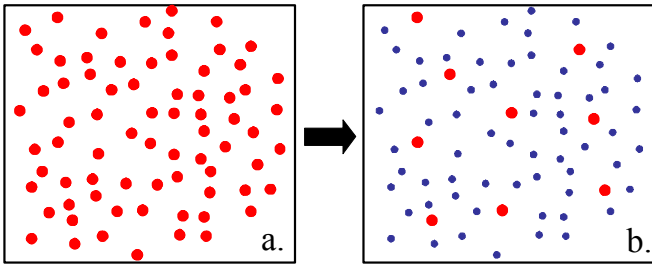


Fig. 2. The effects of signaling. All nodes are contenders in panel a and then signaling resolves a subset of these contenders in panel b, where all the surviving contenders are separated from each other by at least the range of their signals. Large nodes are contenders.

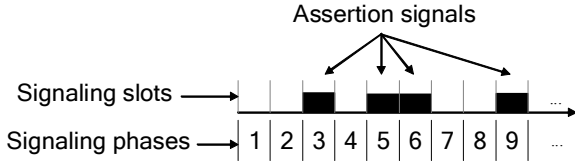


Fig. 3. Collision Resolution Signaling using single slot phases the range of their radios. Fig. 2 illustrates the starting and ending condition of this process.

CRS consists of a series of signaling slots organized into groups of slots called phases in which contending nodes may send very short signals¹. The simplest and generally most effective at arbitrating contention is illustrated in Fig. 3, and consists of one signaling slot per phase. In this design, a probability is assigned to each signaling slot and a contending node will signal in that slot with that probability. The rules of signaling in this design are as follows.

1. At the beginning of each signaling phase a contending node determines if it will signal. It will signal with the probability assigned to the slot of that phase.
2. A contender survives a phase by signaling in a slot or by not signaling and not hearing another contender's signal. A contender that does not signal and hears another contender's signal loses the contention and defers from contending any further in that transmission slot.
3. Nodes that survive all phases win the contention.

Signaling performance is a function of design and can be made better than 99% effective at arbitrating contentions locally and separating surviving contenders by at least the range of their signals. Details about the design of signaling to cause physical separation of contenders can be found in [5].

The separation caused by the basic CRS does not prevent collisions. This is intentional so the protocol can benefit from using physical layer techniques (e.g. channelization [6], [7] and smart antennas [8]) to improve capacity. In some cases; however, contenders can still block each other from gaining access. This is detected by observing repeated successful contentions but then failed handshakes. Signaling can increase separation and resolve blocking through the use of echoing.

¹ The size of the signaling slots and the duration of the signals are selected to prevent ambiguity as to when signals are sent that may result from propagation delays or potential inaccuracies in synchronization.

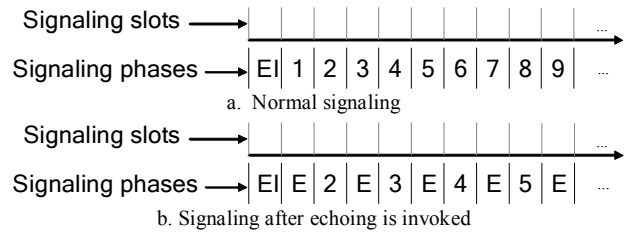


Fig. 4. A signaling design to selectively use echoing: In most contentions, nodes use the signaling design shown in a. If the source detects a blocking condition, knows the source to be an exposed node, or wants to broadcast a packet, it may invoke echoing. If a node signals in the echo invoke (EI) slot then that node and all of its neighbors use the echoing design of b.

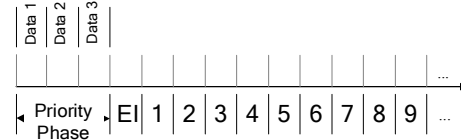


Fig. 5. Modified CRS for providing priority access. This design provides four levels of priority, three levels, one associated with each slot in the priority phase and one level associated with not signaling in the priority phase.

Echo signaling phases consist of two slots. Non-contenders that hear a contender's signal in the first slot echo that signal in the second slot thus extending the effect of a contender's signal two hops. Signaling can be designed to conditionally use echoing. Fig. 4 illustrates a 9 "single slot" phase design that can be dynamically converted to a 4 phase echoing design. If a contender detects the condition that a possible block is occurring it invokes echoing by signaling in the EI slot. The signaling design in Fig. 4b is the design used by all nodes that hear the EI signal.

2) Prioritized Access

Priority access is easily added to the CRS mechanism. In Fig. 5, we add a multi-slot priority signaling phase to the front end of the CRS. In multi-slot phases the node that signals first wins the phase. Here, each slot in the phase is mapped to a different priority with highest priority first. Contenders use the slot that corresponds to the priority of the packet they are contending to send. If a node has a higher priority packet than its neighbors, it will signal first causing those neighbors to defer from contending. The remainder of CRS resolves the contention amongst nodes using the same priority.

3) Channelization

Channelization is used to increase capacity; however, in ad hoc networks, broadcasting on a common channel is necessary for the discovery of network topology. Any channelization scheme must enable nodes to broadcast on a common channel but then to pull peer-to-peer transmissions to separate, preferably orthogonal, channels. This is challenging since destinations normally do not know on which channel to listen.

SCR uses receiver directed channelization. This means, in addition to a shared broadcast channel, all nodes will have a separate channel that they will use to receive peer-to-peer packets. Nodes broadcasting packets use the broadcast channel and nodes sending peer-to-peer packets use their destination's receive channel. We enable destinations to determine the

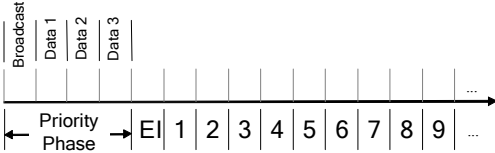


Fig. 6. Modified CRS for providing channelization. When the broadcast priority is used destinations listen to the broadcast channel; otherwise, they listen to their own peer-to-peer channel.

channel to listen to through the addition of a broadcast signaling slot to the priority phase as illustrated in Fig. 6. This slot is used by nodes wanting to broadcast a packet. Not only does it provide a higher priority to broadcast packets over other best effort packets it also serves to indicate to which channel a destination should listen. All nodes will know which priority was used to gain access at the conclusion of the CRS. Nodes that do not survive CRS listen to the broadcast channel if they hear the broadcast priority used, otherwise they listen to their own peer-to-peer channel. Support for the selection and dissemination of receiver channels is provided by the Node State Routing mechanism.

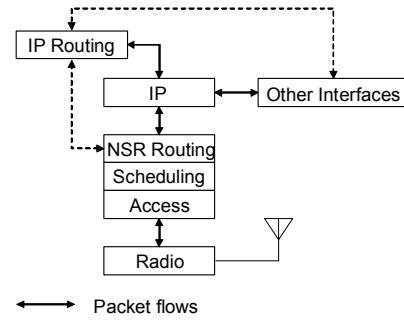
B. Node State Routing

Node State Routing (NSR) is an alternative to the standard link driven approaches to routing. The distinction is that rather than discovering and explicitly disseminating connectivity in terms of links, node states are disseminated and connectivity is inferred from their pairwise use. Articulating network state information in node states allows NSR to support other functions such as quality of service [6], multicasting [9], and as we describe here spectrum management. NSR is implemented beneath IP and is very much a part of the link layer. It is intended for a homogeneous wireless network. Fig. 7 illustrates that additional routing functionality above IP is needed for heterogeneous networks. IP routing exchanges information with NSR routing and does not offer load to the wireless network.

1) Overview

In lieu of links, there are two different routing constructs used in NSR, a node and a wormhole. The node construct is modeled as a point in space and is assumed to have connectivity with other nodes through the use of wireless connections. In many cases nodes may be connected using a dedicated link such as a cable. To use these links within the node state routing protocol we define a second routing construct called a wormhole. We define our wormhole construct as a directed path between two points in the network. The basic algorithm used to select which routing constructs to use in a route considers the cost of sending a packet to a construct and the cost of using the construct. These costs are derived from the states of the nodes and the wormholes.

NSR requires two capabilities: location awareness and the ability to measure signal strength. With this information, each node creates a pathloss map. Location and the pathloss maps of all nodes and wormhole endpoints provide sufficient information to determine connectivity between the constructs and then the overall topology.



←→ Packet flows
 - - - - - Cross layer communication

Fig. 7. NSR's multilayer routing functionality

Table 1 Proposed node states that are useful for spectrum management

STATE	DESCRIPTION
Address	MAC address of the node or the wormhole. In the case of the wormhole, the address is associated with the node at the front end.
1-meter Path loss	Pathloss of the first meter of propagation used with the log distance path loss model.
Propagation map	Propagation conditions can vary based on the location of nodes and the direction of propagation. To accommodate this concern we propose nodes measure and estimate a path loss exponent for the path loss model. We require each node that broadcasts a packet to announce the power level it is using. We assume that each destination node that hears a broadcast can determine the power level of the received signal and can then estimate a path loss exponent using the attenuation of the signal and the separation distance from the source. When propagation characteristics vary to different destinations, these states can be broken up into different sectors that account for these differences.
Channel	The channel the node uses to receive a peer-to-peer packet. This state complements the channelization capability of SCR.
IP Addresses	IP addresses that are used by the node. It includes multicast addresses.
Voice nets	Voice net IDs that the node subscribes to.
Configuration	A description of the quantity and types of radio interfaces at the node.
Frequency use	A listing of the current channels used by ganged radios
Direction	Current direction of movement of the node. Used to predict future topology
Location	The location defines where the node or where the wormhole's endpoints physically exist in the network. Node state routing requires location awareness.
Spectrum use detection map	A data structure articulating the measured level of energy in different bands of spectrum
Time Stamp	This is the time that the reported state was measured. We assume time is absolute and synchronized throughout the network.
Velocity	Current velocity as measured by the node. Used to predict future topology.

NSR consists of three processes: propagation map discovery, node state dissemination, and a route calculation. On a periodic basis, each node in the network transmits node state update packets. These transmissions are used to discover propagation conditions and to disseminate the node states. Either on a periodic basis or as required, nodes use these states to determine topology.

2) Node States

The node states used in NSR may describe any type of state information for a node. As a minimum, it provides the node's location, the propagation conditions about the node, and a mapping between IP and MAC addresses. Table I lists some possible states required to implement basic routing and then additional states that are used in our story on how NSR supports spectrum management. Other possible states not listed here can support energy conservation [10], quality of service [6], and multicasting [9]. Details on propagation maps can be found in [6].

3) Topology Determination

Given a set of node states, each node determines topology in three steps. First, connectivity between constructs is inferred using their propagation maps and locations. Second, for all inferred links a metric is assigned. These metrics are formed from the node states and include the cost of transmitting the packet and using the destination construct. Finally, Dijkstra's algorithm is used with the weighted set of inferred links to find the shortest paths to all destinations. The power of this approach is that a whole assortment of filters and weighting techniques can be used to affect the routing tables that are calculated without having to change the state dissemination mechanism. In our case this mechanism disseminates the information necessary for spectrum management.

4) Node State Dissemination

Nodes distribute the node states using a diffusion mechanism. On a periodic basis a node will broadcast a node state packet (NSP) which will include its own state and other states in its list restricted in number by the maximum packet size. The states that are included in these updates are selected by two criteria, a threshold that indicates whether an update is needed and a prioritization criterion to enable selection amongst several states that meet the update threshold. In the diffusion process, the update threshold depends on the distance between the node that owns the state and the node doing the rebroadcast.

Scaling is forced using a minimum interval between NSP updates, i.e., a node may send one NSP per interval. However, NSP updates are accelerated when routing failures are observed. Loops do not occur in link state routing protocols if all nodes use the same states. In NSR, nodes may have different node state information and loops may occur. The observation of a loop triggers accelerated updates. The goal of these updates is to synchronize the node state tables of all the nodes in the loop so it can be broken. After identifying a looping condition, a node in the loop broadcasts a relevant subset of its node state table that covers the region of interest, recalculates its routing tables and then forwards the packet that was looping. This process is repeated so long as the packet remains in the loop. Ultimately, all nodes in the loop will have a common picture of the network and the packet will progress.

V. SPECTRUM MANAGEMENT MECHANISMS

SCR and NSR provide a rich set of features that can be exploited in many ways to manage spectrum. In this section we describe the specific mechanisms that can arbitrate the use of spectrum.

A. Access Mechanisms

1) Primary and Secondary Access Arbitration

The prioritized access described above can be exploited to arbitrate the primary and secondary use of both the channel and the network. Fig. 8 illustrates the signaling design and describes the process. A separate spectrum management phase is prepended to CRS. This echo phase design is used by primary users to assert their rights over secondary users.

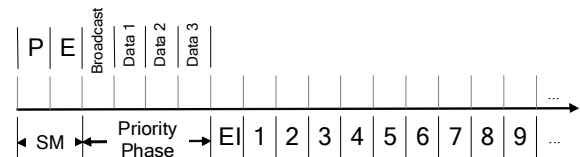


Fig. 8. Modified CRS for primary secondary access arbitration. The SM phase is used by primary users to assert their right. By the protocol, primary users signal in the P slot of the SM phase and all neighboring nodes, primary and secondary, echo that signal in the second slot of the phase. For the remainder of CRS, primary users only contend with other primary users.

Echoing insures these rights extend 2-hops from the primary contender. The SM phase can be designed to support more than two levels of SM access priority by adding more P-E slots.

The ad hoc network that uses this primary secondary access approach may be designed with various rules about how primary and secondary users cooperate with each other. On one extreme, the primary and secondary users may have completely isolated networks and on the other extreme primary and secondary users fully cooperate to form a single network where access rights transfer with packets. The packets originating from primary users would always have precedence over packets of secondary users.

This mechanism supports several spectrum sharing scenarios.

a) Isolated Networks

A network built to support municipal services such as police or emergency dispatch may allow this same spectrum to be used on a secondary basis when there is no demand for the primary use.

b) Secondary Market

The primary user may sell secondary access rights to his spectrum and may even give the user access to his infrastructure through the cooperative networking approach. In this way primary users can get secondary users to support the development and maintenance costs of the infrastructure without sacrificing his access rights.

c) Broadband Development

Primary use is sold in an exclusive use model for the purposes of providing fee based broadband access to a community. As it may not be financially viable for a provider to build infrastructure and support access in some regions the same equipment can be used by local communities and neighborhoods to build their own networks and wireless broadband access. If and when the service provider decides to develop infrastructure and provide services in the region, users can continue using the network on a secondary basis or pay for the primary use and its associated services.

2) Channelization

We described earlier how the signaling mechanism supports knowing which channel to listen to in the ad hoc network. This feature makes it possible to add capacity to a network through the addition of more channels. The channels become part of

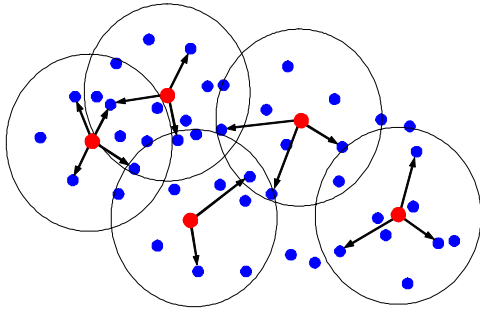


Fig. 9. An example of multi-destination communications. The large circles show the range to which sources can send packets and the arrows show the direction of the downlink. The combined use of SCR, channelization, CDMA and smart antennas can enable this capability.

the pool that nodes can choose from for receiver directed communications. The larger the pool the less likely that neighbors will use the same channel and the less likely that neighboring contenders will interfere with each other at destinations. A large set of channels combined with the construction of nodes that can talk on multiple channels simultaneously can enable one to many communications as illustrated in Fig. 9.

B. Routing Mechanisms

1) Channel Selection and Distribution

Peer-to-peer channel selection is distributed. The nodes initialize the process by randomly selecting a channel from the pool and then advertising their selection in their node state. If there is a conflict with a node's own selection and that of any of its two-hop neighbors, it chooses a new channel. It chooses an unused channel if there is one or, if not, it randomly selects a channel from the least used channels in the pool. It broadcasts its channel selection before using it. We limit the rate at which random changes can be made, e.g. one change every 5 seconds. Due to the physical separation result of the contention there are rarely more than three contenders in range of any destination, so despite the reuse of channels, collisions on the same channel are rare.

2) Interference Temperature Collection and Distribution

Assuming radios have the capability to detect spectrum use and there is a data structure to articulate these observations, NSR can disseminate these observations as an additional state. As a node state they are combined with other state information, specifically location and propagation observation that as a combination provide greater context. When further combined with the observation of the plurality of nodes in the network they can provide a spatial map of spectrum use. These observations can be used to support identifying opportunities where in space spectrum is available or as a monitoring tool to identify where inappropriate uses of spectrum are occurring.

3) Ganged Radio Channel Control

As illustrated in Fig 7, the ad hoc network can be connected to other interfaces. These interfaces may be to other radios. The ganging of radios at nodes in this way would be a node state. The ad hoc network can then serve as an underlying control network for the distribution of the channel, waveform,

and transmission parameters that these radios' use. Through the use of the NSR node state dissemination mechanism a controller (i.e. person or automated process) at a controlling node can monitor the assignment and use of spectrum across the region the network stretches, direct who may use spectrum and under what conditions, and identify when transmitters might be contributing too much interference to higher priority services of neighboring users.

VI. THE FAST COMMAND AND CONTROL MODEL

The Fast Command and Control model (FCCM) envisions a near realtime control of spectrum across some set of spectrum bands. It uses a network connected between a controlling entity and users as the means to monitor and assign spectrum use. Users would use ganged radios where one radio would be an ad hoc networking radio as described above with the ability to determine its own location and possibly the ability to sense spectrum use. The additional ganged radios would be controllable over the network. Either the radio would have the ability to adapt its transmission characteristics to those dictated by the controlling node or in cases where the transmission parameters are fixed would have its operation slaved to the permission of the controlling node. The concept in the second case is for the radio to notify the controlling node of its spectrum consuming parameters and then for the controlling node to either grant or not grant it permission to transmit. The controlling node would need to continuously grant permission for this type of ganged radio to continue using the spectrum. The monitoring features and the collection of observations of all nodes provide the data that enables the controlling node to monitor spectrum use and to prevent harmful interference at networked nodes with critical services. Through connectivity to the internet, users in remote areas could obtain access to spectrum by creating the network connections to a controller assigned to the region they operate. In the following we describe several implementations that demonstrate the possible evolution of this concept.

A. Military Voice Nets

Traditionally, voice radio nets have been created in military applications that match the organization. For example a platoon leader and his squads share a common net, a company commander and his platoon leaders share another, and then a battalion commander and his company commanders share yet another. This series of voice nets matches the hierarchical structure of command and control. Channels are allocated to these nets so that they can coexist spatially. Transitioning to a data network that connects all personnel in the organization in a single network does not preclude the need for these type of voice networks. Although it is conceivable to implement multicast within a common network to create this service it is neither practical nor efficient. The value of the ad hoc network is that it provides a ubiquitous way for all users to communicate to each other but it is not the most efficient way for all subsets of those nodes to communicate. In a group of nodes that are generally in close proximity where most traffic is meant for the entire group, sharing a broadcast network like these voice nets remains very efficient and a much better paradigm to deliver the desired performance. What makes this

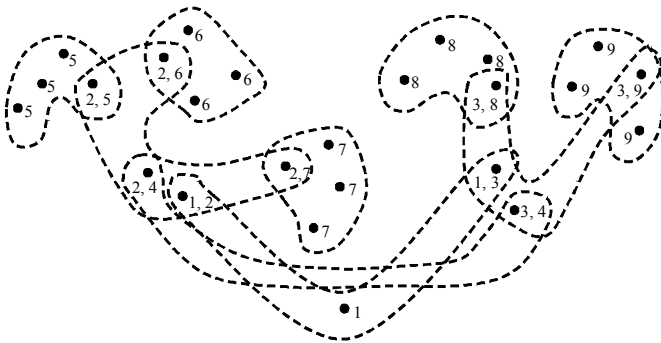


Fig. 10. Example scenario of military voice networks.

approach inefficient is that a priori assignment of the channels requires a large set of channels since the mobility of these nets in military scenarios can cause a large number of voice nets to come in range of each other at different times. The FCCM can alleviate the need for such a large set since it could manage the assignment of channels dynamically based on the actual proximity of the nets. Fig. 10 illustrates the idea. It illustrates a notional layout of a formation. Each circle represents a member of the organization. All have a radio on the common ad hoc networking channel. The numbers adjacent to these circles are the voice nets (i.e. platoon net, company net, battalion net) that each member subscribes. If a member has one number it has two radios, one is a member of the common ad hoc network and the second is available for the voice net. If a member has two numbers it is a member of two voice nets and has three radios total. The dashed lines circumscribe nodes that belong to the same voice nets. The number of the voice net is a logical association; it does not map directly to a radio channel. Rather, it only indicates that all nodes that subscribe to the same net should be on the same channel. The command and control network assigns the channel and perhaps even the waveform and transmit power. Through this approach, voice nets that are separated from each other, for example 5 and 9, could be assigned the same channel to use. This assignment process could be fully automated using a pool of channels some of which may even have restrictions on their use to sub-regions of the maneuver space. Additionally, assignment can take into consideration the unique operational requirements of some voice nets. For example lower frequencies with a suitable modulation may be assigned for use in an urban environment to overcome the harsh fading conditions.

B. Network Engineering

The voice nets above and other specialized connections can be combined to create a more robust network. Most all access mechanisms (we do not know of any where this is not the case) assume the radios have similar ranges. Using different transmit powers to achieve variable range in the same network befuddles the carrier and virtual sensing mechanisms that are the foundation of these protocols. The selection of a common range tries to balance the needs for connectivity and capacity. Inevitably, mobile environments result in a heterogeneous distribution of nodes and load. Overlaying networks (e.g. the voice nets in the previous example) and specialized long range links (i.e. wormholes) on top of the ad hoc network can mitigate these occurrences. The control mechanism that is used for spectrum management can be combined with a network

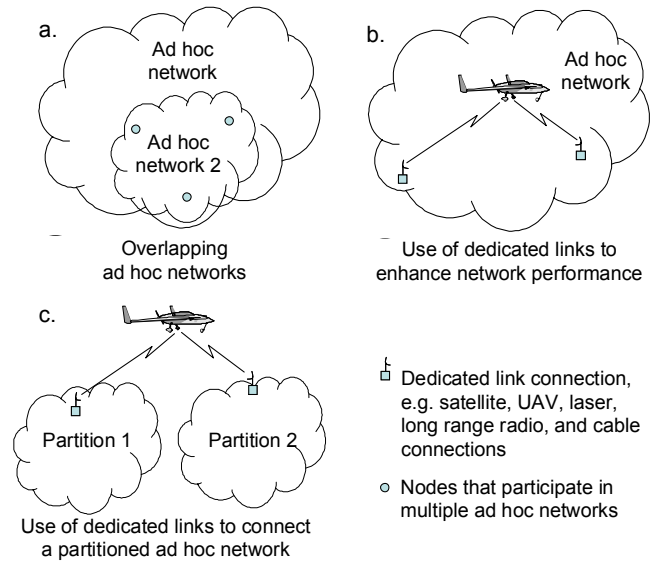


Fig.11. Example network engineering scenarios.

engineering capability where the pool of available spectrum and the radio assets at each node become part of a network engineering design space that can be used to maintain network performance through maneuver that would otherwise create partitions and congested network bottlenecks. The networking engineering function would use a set of command and control messages that assign ganged radios to channels.

Fig. 11 illustrates several examples of the idea. In panel a an overlay ad hoc network is created with a few nodes that are in both networks. Nodes in the underlying ad hoc network can route packets through the overlay network which we assume in this example to have longer range radios. Short range destinations may be reached through the underlying network alone and long range communications could be passed to the overlay network since it can transport the traffic across the combined network in fewer hops. In panels b and c, we see examples of long range links. The link in panel b may have been set up to bypass a region of congestion and in panel c to bridge a partition. We note that the selection of which radios serve in the overlay network or as the end points of the long range links is all done from the network engineers control station. Network engineering responsibilities may be distributed throughout the network.

C. Independent Spectrum Use

In the independent spectrum use case, owners of ganged radios attempt to negotiate use of spectrum for their ganged radios for a purpose of their choosing. The request would provide details on the capability that is desired and if available the spectrum manager can allocate the spectrum and specify the conditions for its use. It is envisioned that these grants could be for short periods of time to support a specific operation. In the management process, users identify radios to support the service, they render a request to the spectrum manager, and then the terms of spectrum use are negotiated. The specific transmission parameters for the ganged radios are communicated from the spectrum manager to the radios through the network thus enabling the radios to be used. In cases where the ganged radio's transmit parameters are fixed,

the user renders a request for their use, the spectrum manager decides if the use is feasible, and then, if feasible, the spectrum manager grants permission.

D. Future Work

Enabling the FCCM requires algorithms and protocols. Algorithms are required to spatially track spectrum use and to determine where spectrum is available and protocols are necessary for the processes of requesting and granting spectrum use. Additionally, mechanisms are necessary within ganged radios to insure they are controllable by the remote spectrum manager. Our future work will start by trying to create a military voice nets capability and then expand the capabilities as proposed.

VII. STANDARDS AND REGULATORY REQUIREMENTS

Implementing the concepts proposed in this paper requires additions to the standardization efforts of ad hoc networking. Following from these changes would be the development of procedures used by administrations to manage spectrum through the internet and to license network controlled devices.

A. Standards

In complex systems, there is the need to balance innovation with standards. Innovation is required where performance is the issue and standards are required where integration is the issue. The genius of the internet was the choice to standardize the Internet Protocol (IP) allowing for innovation both above and below in the protocol stack with IP being the point of integration. The primary decision that IP makes is to which interface to send a packet and the next hop address to use. The ramification of having a point of integration is that it causes a fairly restrictive view of what exists on the other side. The view from above IP is that the network consists of *links* and routers which map to a connected graph of edges and nodes. Below IP, protocols oblige this view, even with shared media, and innovation focuses on providing higher capacity links. Routing protocols logically fall above IP where they collect information to ascertain which next hop addressee is the best router to forward a packet to its final destination. Guided by the objective to be IP-compliant, the major standardization effort to support ad hoc networking, the Internet Engineering Task Force (IETF) Mobile Ad Hoc Networking (MANET) Working Group (WG) chose to embrace this traditional view and to focus its design efforts on routing protocols that are placed above IP. [12] The ramification of this choice is that the protocols that are managing wireless ad hoc networking view the task as tracking the connectivity of the network through the abstraction of a set of ephemeral links. This abstraction occludes many of the interesting possibilities that the wireless media offers to the networking problem and does not capture the spatiotemporal context that is necessary for spectrum management. There are ample examples of cross-layer design where higher layers attempt to control physical layer properties, (e.g. transmit power, data rate, transmit channel, and antenna pointing) to create performance that is not possible otherwise. Kawadia and Kumar make the case that relying on such cross layer communications can result in “spaghetti design, which can stifle further innovation and be difficult to

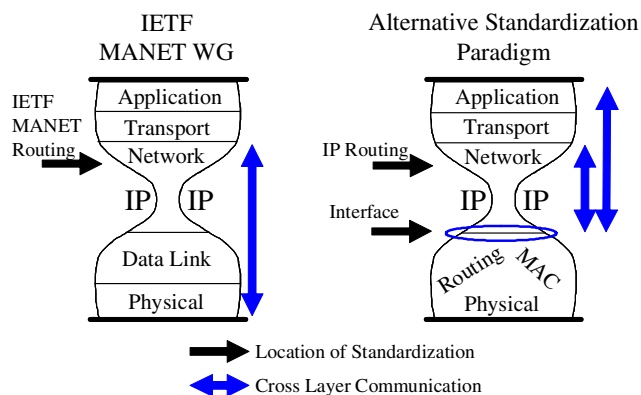


Fig. 12. Comparison of standardization focus.

upkeep.” [13] The major problem is that any cross-layer optimization across IP violates the Internet architecture and the idea of a center point of integration. Spectrum management built upon the standard MANET routing paradigm would be one more example of a cross layer protocol design.

The current trend toward cross-layer design is an indication of the misplacement of function in the protocol stack. Above IP routing is about managing the use of links but links are not the resource in wireless ad hoc networks. The resource is the spatiotemporal use of RF spectrum. Although IP integration will ultimately require abstraction to a link paradigm, we contend that we can prevent cross IP design and fully open-up the wireless design space for innovation by bifurcating the routing function. Routing functionality would be placed on either side of IP with the intent that the below IP portion creates the ad hoc network and that the above IP portion learns the wireless topology through queries to the below IP portion. This approach fully opens the wireless network to innovation where physical layer control can be integrated into the routing logic, it preserve’s IP’s role as the point of integration, and it eliminates the need for the above IP portion to implement its own topology discovery mechanisms. Standardization would focus on the needs of the above IP routing protocol, the interface to the below IP portion, and the set of messages to communicate through IP. Fig. 12 illustrates the differences in approach. In the IETF MANET WG approach the routing function is above the IP waste and efforts to manage the physical layer reach through IP. The methods of communication and the parameters that are controlled are unique to each type of device. In the alternative, the IETF would standardize the communications that pass through IP, the interface to the wireless devices, and the physical layer optimization would remain below the interface. It is envisioned that both the above IP routing protocol and applications would use the messages to learn the wireless device’s vision of the network and its use of spectrum.

If pursued, this standardization approach can engender further innovation opportunities including the vision of spectrum management proposed in this paper. Fig. 13 illustrates the idea that there could be numerous types of devices connected to an IP network. Interface 1 is a device, wireless or wireline, that presents the standard wireline view, interface 2 is an ad hoc networking device, and interface 3 is a physical device controlled through the internet but not a networking path. The proposed standardization effort would

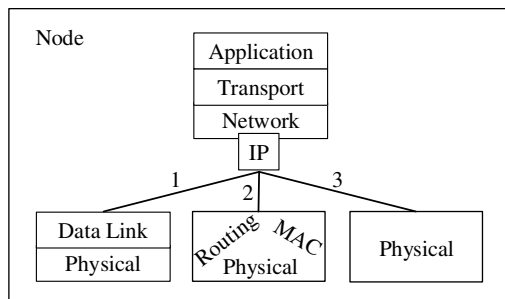


Fig. 13. Below IP devices.

create the communications that allows any protocol or application above IP to discover the physical devices that are connected to IP. Part of this standardization effort would be to establish how devices and ad hoc networks articulate their use of spectrum and the messages that allow applications to direct devices in which spectrum to use. A potential use case could find a non-communication spectrum using device, e.g. a wall penetrating radar, connecting to a network node as a type 3 device. The node would be connected to the internet either through a type 1 or 2 device. The spectrum management utility employing the FCCM would receive requests and direct when the device could be used.

In summary, the goal of any standardization effort that uses the internet should be to retain IP as the point of integration. Unfortunately, efforts to solve the MANET routing problem in protocols above IP stifle innovation by limiting the opportunity to exploit physical layer properties. We have proposed an alternative standardization approach for ad hoc networking that not only enables this control of physical layer properties in ad hoc networks but also enables the direct connection of non-networking RF devices to the network that could also be controlled by the FCCM application.

B. Regulation

In the past, compliance to regulatory use of spectrum was achieved by licensing devices, but the RF properties of these devices could not be changed. Our implementation proposed in Fig. 13 has intentionally provided an approach that allows spectrum management administrations to still license devices. A new requirement however is that licensing would ascertain whether the devices are compliant with the FCCM approach and that there are sufficient controls to prevent rogue use of spectrum. Regulation would focus on the details of FCCM, i.e. how it arbitrates spectrum use, and the licensing of FCCM compliant devices.

VIII. CONCLUSION

In this paper we reviewed spectrum management models and some new technologies that give promise of better use of spectrum. Consistent with the same goals, we have proposed several very dynamic approaches to spectrum management that can be enabled within an ad hoc network. We described how two or more ad hoc networks can coexist in the same spectrum and space where access will always be granted based on the hierarchy of spectrum access rights of those networks. We describe how these coexisting networks can cooperate to

enhance the performance of all the networks. We describe how multiple channels can be exploited in an ad hoc network to increase its capacity. Finally, we propose a new spectrum management model that implements a command and control approach that is made possible by the ad hoc network. We emphasize that the ability to enable this breadth of spectrum management concepts using an ad hoc network does not come from just building an ad hoc network but from using the SCR and NSR protocols in the ad hoc network. SCR provides the mechanisms that enable the arbitration of access based on spectrum rights and solves the hard problem of exploiting multiple channels in an ad hoc network, i.e. enabling destination nodes to know which channel to listen to. NSR provides the network state dissemination mechanism that can allow a spectrum manager to track the use of spectrum across the area the network covers. We described current MANET standardization efforts and identified that they would not support the FCCM we envision and so propose an alternative focus to IP standardization effort. Further work is required to develop the algorithms, protocols, and policies that the spectrum manager would use for this type of dynamic spectrum management. Although much work is required to achieve this vision, with these approaches, the opportunities to innovate are unbounded and the utility of spectrum would increase dramatically.

REFERENCES

- [1] <http://www.fcc.gov/sptf>
- [2] FCC, "Spectrum Policy Task Force Report," ET Docket No. 02-135, November 2002.
- [3] FCC 05-57, "Facilitating opportunities for flexible, efficient, and reliable spectrum use employing cognitive radio technologies," ET Docket 03-108, March 2005.
- [4] FCC 03-289, "Establishment of an interference temperature metric to quantify and manage interference and to expand available unlicensed operation in certain fixed, mobile, and satellite frequency bands," ET Docket No. 03-237, November 2003.
- [5] J. A. Stine, G. de Veciana, K. Grace, and R. Durst, "Orchestrating spatial reuse in wireless ad hoc networks using Synchronous Collision Resolution," *J. of Interconnection Networks*, Vol. 3 No. 3 & 4, September and December 2002, pp. 167 – 195.
- [6] J. A. Stine and G. de Veciana, "A paradigm for quality of service in wireless ad hoc networks using synchronous signaling and node states," *IEEE J. Selected Areas of Communications*, Sep. 2004, pp. 1301-1321.
- [7] J. A. Stine, "Exploiting processing gain in wireless ad hoc networks using synchronous collision resolution medium access control schemes," *Proc. IEEE WCNC*, 2005.
- [8] J. A. Stine, "Exploiting smart antennas in wireless ad hoc and mesh networks," Submitted to *IEEE Wireless Communications Magazine* for consideration for the Feb 2006 issue on "Wireless Mesh Networking: Theories, Protocols, and Systems."
- [9] J. A. Stine, "Node state multicasting in wireless ad hoc networks," *Proc. IEEE MILCOM* 2005.
- [10] J. A. Stine and G. de Veciana, "A comprehensive energy conservation solution for mobile ad hoc networks," *IEEE Int. Communication Conf.*, 2002, pp. 3341 - 3345.
- [11] Title 47, U.S. Code of Federal Regulations.
- [12] S. Corson and J. Macker, IETF RFC 2501, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations," 1999.
- [13] V. Kawadia and P.R. Kumar, "A cautionary perspective on cross-layer design," *IEEE Wireless Communications*, Vol. 12, No. 1, Feb 2005, pp. 3-11