Protocol Design Challenges for Multi-hop Dynamic Spectrum Access Networks*

Pradeep Kyasanur
Dept. of Computer Science, and
Coordinated Science Laboratory,
University of Illinois at Urbana-Champaign
Email: kyasanur@uiuc.edu

Nitin H. Vaidya
Dept. of Electrical and Computer Engineering, and
Coordinated Science Laboratory,
University of Illinois at Urbana-Champaign
Email: nhv@uiuc.edu

***Abstract*— Driven by the need to improve network capacity, there is a growing interest for dynamically utilizing spectrum over a wide range of frequency bands. The available spectrum is typically divided into multiple channels. Past work on designing protocols for multi-channel wireless networks has assumed that all channels are homogeneous. However, channels that are located in widely separated frequency bands exhibit considerable heterogeneity in transmission ranges, data rates, etc. In this paper, we identify the impact of channel heterogeneity on network performance, and motivate the need to account for channel heterogeneity while designing higher layer protocols. We present some approaches for managing heterogeneity, and propose hiding most of the channel heterogeneity from higher layers by designing suitable channel abstractions.**

I. Introduction

Recent years have seen significant increase in the use of wireless networks that operate in unlicensed spectrum using standardized wireless technologies (e.g., IEEE 802.11 [1] based networks). A common mode of communication in many wireless networks involves single-hop operation with infrastructure support, though multi-hop operation is being considered in emerging network architectures, such as mesh networks. However, a fundamental impediment to building large multi-hop networks has been the insufficient network capacity, due to the *limited spectrum available for unlicensed use*.

One proposal for alleviating the scarcity of spectrum is by *dynamically utilizing existing licensed spectrum* with spectrum agile "cognitive radios" [2]. Cognitive radios are being envisioned that allow secondary users to co-exist with the primary users of the licensed spectrum. Secondary users may use the spectrum only when such use does not interfere with primary user's communication. This requirement implies that the spectrum that is available for use varies with time (i.e., dynamically changes), and is dependent on the load imposed by the primary users. There are several research initiatives for building cognitive radios that can opportunistically identify free spectrum, and designing medium access control and link layer protocols that utilize the available spectrum (e.g., see [3], [4]). When dynamic spectrum access networks operate in a multi-hop fashion, several other challenges arise at higher layers of the protocol stack that need to be carefully considered before the full performance benefits can be realized. The focus of this paper is to highlight some of the research challenges from the *perspective of higher layer protocol design.*

The available spectrum is typically divided into multiple channels [1]. There are several proposals for using multiple channels in wireless networks (c.f., [5]–[14]). However, existing proposals for using multiple channels typically make several assumptions that may fail to hold in dynamic spectrum access networks. For example, most proposals assume that the set of available channels is static, i.e., the channels available for use is fixed at the time of network initiation. Since cognitive radios may allow the spectrum available to change dynamically, the *set of channels may dynamically change as well*. Furthermore, existing proposals often assume that the available channels are "homogeneous", i.e., different channels have similar range and support similar data rates. The homogeneity assumptions are broken when *different channels may be located on widely separated slices of frequency spectrum* with different bandwidths, and different propagation characteristics. Thus, there is a need to design higher layer protocols that suitably manage the (possibly) "heterogeneous" channels supported by a cognitive radio.

The rest of the paper is organized as follows. In Section II, we identify the impact of dynamic spectrum access on higher layer protocols, and motivate the need for developing simple abstractions for hiding heterogeneity in Section III. In Section IV, we provide an example to highlight the possibility of benefiting from channel heterogeneity, and conclude in Section V.

II. Protocol design challenges

Cognitive radios are being designed that opportunistically utilize spectrum across a wide range of frequency bands, thereby providing for a large set of frequency-separated channels. The set of available channels in such dynamic spectrum access networks may be "heterogeneous", i.e., different channels may support different transmission ranges, data rates, delay characteristics, etc. Furthermore, the set of available channels itself may change with time, depending on traffic

*This work was supported in part by the US Army Research Office under grant W911NF-05-1-0246.

imposed by the primary user. In the rest of this section, we illustrate the impact of heterogeneous transmission ranges and changing channel set, on higher layer protocols.

A. Heterogeneous transmission ranges

Transmission range of a channel is informally the maximum distance up to which a packet transmitted by a node on that channel may be successfully received. The exact region over which a transmission can be received may have a complex shape (although circular regions are often used as approximations) that depends on channel propagation characteristics, obstructions, etc. Wireless transmissions on different frequency-separated channels can suffer varying amounts of frequency-dependent path loss, multi-path effects and attenuation [15]. Since channels supported by a cognitive radio may be located on widely separated slices of frequency, different channels may experience significantly different propagation characteristics. In addition, FCC regulations may specify different limits on the maximum allowed transmission power for different frequency bands. As a result, *different channels may support different transmission ranges*. Furthermore, on account of variable attenuation and multi-path effects, a channel with longer (average) transmission range may not cover all the area covered by a channel with a shorter transmission range.

Unequal transmission range may affect the performance of many existing higher layer protocols. Therefore, it may be beneficial if the range of different channels is equalized. One way to equalize the transmission range of different channels is to reduce the transmission power on channels with longer range, such that all channels have the same range as the channel with the shortest range. However, this approach may be excessively conservative, and limiting the transmission range of all channels to that of the shortest range channel may break network connectivity. Furthermore, choosing appropriate transmission powers for equalizing the transmission range of different channels may not be feasible when there are time varying differences in propagation characteristics of different channels brought about by fading, multi-path effects, etc.

Using different modulation schemes on different channels may be another approach toward equalizing the range of different channels. Different modulation schemes require different Signal-to-Noise ratios (SNR) for successfully decoding a packet. Therefore, on any given channel, while using a fixed transmission power, the distance over which a packet can be successfully decoded is dependent on the modulation scheme used. However, even if range is equalized using this approach, the transmission range of the whole network will reduce to the range of the shortest range channel (with the channel with the shortest range using the lowest possible modulation rate to maximize its range). Furthermore, using different modulation schemes on different channels results in data rate differences across channels, which will also complicate protocol design.

Existing multi-channel protocols often assume that a node has a *common neighbor set on each channel*, where informally, neighbors of a node X on some channel c are all nodes that X can directly communicate with on channel c. In a homogeneous multi-channel network, where all channels have the similar propagation characteristics, a node X can reach the same set of neighbors on any of its channels. However, in a heterogeneous multi-channel network, a node may be able to communicate with different (potentially overlapping) set of neighbors on different channels, and therefore, *whether a pair of nodes can communicate with each other is dependent on the channel that will be used for the communication*.

Distributed multi-channel protocols often need to exchange control information, such as routing information or channel usage information, among all the neighbors. A node may be able to send data to a neighbor only if certain control information (such as a neighbor discovery packet) has been previously exchanged with that neighbor. Typically, such control information is often sent out as broadcasts, which can then be received by all neighboring nodes.

In networks using cognitive radios, the number of available channels can potentially be large, and a single radio may only be able to operate over one channel at a time. Since each node typically has few radios, a broadcast packet sent by a node is received by its neighbor only if the packet was sent on one of the channels on which the neighbor was listening to. To ensure every neighbor receives a broadcast packet, one possibility is to send a copy of the broadcast packet on every channel. However, sending a packet on every channel may be quite expensive when the number of available channels is larger than the number of available radios [11]. In this scenario, when all channels have the same range, a commonly used optimization is to exchange broadcast packets on one common broadcast channel, and all nodes, by design, are required to always listen to the broadcast channel (any channel can be used as the broadcast channel because all channels have the same range). However, when different channels have different ranges and different (possibly overlapping) neighbor sets, it may be necessary to exchange broadcast packets on *all, or a large set* of channels to ensure every neighbor receives a copy. This can significantly increase the cost of broadcast and has to be carefully accounted for in protocol design.

One possible solution for reducing the cost of broadcasts is to carefully identify a small subset of channels which cover the neighbors of a node on all channels, and use this subset of channels for exchanging broadcast packets. Implementing this solution will require the development of new techniques for efficiently identifying the set of neighbors of a node on any given channel, under the constraint that there are few radios and a large number of channels. Another possible solution is to carefully restrict the set of nodes higher layer protocols communicate with to neighbors on a specific channel, and always send broadcasts only on that channel. This will ensure that nodes with which data communication takes place are only those nodes that can receive the broadcast (control) information. The key drawback of this approach is the possibility of not using certain communication links which could otherwise have improved network performance.

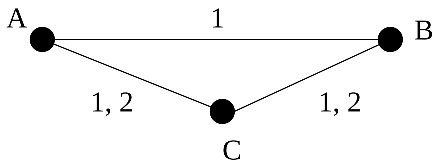

Fig. 1. Impact of range heterogeneity

We illustrate the impact of range heterogeneity with a simple example. Figure 1 considers a network with three nodes A, B, and C. Suppose two channels are available, with channel 1 having a longer range than channel 2. Also, assume that communication on link A-B is possible only on channel 1, while communication on links A-C and B-C is possible over both channel 1 and channel 2. Suppose node A is discovering a route to B. If channel 2 is used for route discovery (i.e., exchanging control information), then the direct route between A and B on channel 1 will not be discovered. On the other hand, if channel 1 is used for route discovery, but channel 2 is preferred for data communication, then data communication is only possible through route A-C-B, but route discovery may select route A-B. Therefore, when different channels have different ranges, restricting control operations (e.g., route discovery) on a specific channel may be sub-optimal. However, exchanging control information on all channels may be quite expensive, especially when the total number of channels is large.

B. Dynamic changes to available channel set

The set of available channels in dynamic spectrum access networks may vary with time, based on the spectrum usage of primary users. It is possible to hide the changing channel set from higher layer protocols by using suitable MAC and link layer protocols. For example, [4] proposes a new architecture that uses a link layer protocol to discover neighbors over a common control channel. Data communication with a neighbor is over an appropriate channel, and the channel used for data communication may vary with time. Therefore, the channel to use for communicating with a neighbor is hidden from the higher layer protocols.

While a link layer approach is quite attractive by hiding the complexity of managing a changing channel set, it may introduce unwanted side-effects that affect performance. Changes to the set of available channels can change the set of neighbors of a node. However, if the channel set information is managed at the link layer, higher layer protocols may not be aware of any change in the reachability of a neighbor. For example, considering Figure 1, suppose that A is communicating directly with B over channel 1. Now, suppose that channel 1 is no longer available (say, the primary user is now using that spectrum), and instead a new channel 3 is available (say, some other primary user has now stopped transmitting). If the MAC protocol hides this change to the available channel set from a higher layer routing protocol, then data traffic may still be routed along A-B. However, if channel 3 has a shorter range than channel 1, then B is no longer directly accessible from A, leading to route breakage. On the other hand, suppose that channel 3 has a longer range than channel 1, but supports a lower data rate. Then, higher throughput may be possible if data is sent from A to B over route A-C-B (if channel 2 supports a significantly higher data rate than channel 3), instead of using route A-B. Therefore, for achieving good performance, there may be a need to expose changes to the available channel set to higher layer protocols.

III. Designing Channel Abstractions

In the previous section, we have argued that wireless networks with heterogeneous channels require the development of new higher layer protocols. However, new protocols may need to obtain new types of information about the channels from lower layers. It is an important research challenge to identify what information has to be exported to higher layers, and developing clean interfaces for exporting the information.

Cognitive radios can be used to opportunistically utilize available spectrum, and may offer a wide range of features, ranging from dynamic channel selection to controlling transmit power. Exposing all the details of channels supported by the cognitive radio to higher layer protocols will significantly increase the complexity of higher layer protocols. Therefore, it is necessary to develop simple abstractions of cognitive radios that can represent all the supported features. Such a standard representation is especially useful in developing generic protocols that operate over the cognitive radios.

Many higher layer protocols often only require support for identifying the set of neighbors on each channel, and the "cost" of reaching a neighbor. In general, we will use the term "configuration" to describe one setup of a channel. For example, one configuration may use channel 1 and transmission power of 10 mW. Another configuration may use channel 2 and transmission power of 1 mW. A cognitive radio may support many such "configurations", wherein each configuration is one valid combination of available parameters, such as channel of operation, power to use on the channel, data rate supported by the channel, etc. The notion of configuration offers a generic way of representing a wide and varying set of features that may be supported by a cognitive radio. Higher layer protocols may need information about the set of available configurations, set of neighbors reachable through a configuration, and the cost of using a configuration. The interface between the cognitive radio and higher layers should be designed to allow the higher layers to query for configuration information, as well as set the configuration to be used for transmitting a packet.

Since a configuration is composed of multiple features, it may be difficult to quantitatively compare two different configurations. For example, it may be possible for some node A to communicate with a node B using channel 1 and 10 mW power at 1 Mbps rate, and using channel 2 and 100 mW power at 2 Mbps rate. It is not immediately clear which configuration is better. If the desired goal is to minimize energy consumption, the first alternative may be preferable. On the other hand, if the goal is to minimize transmission time, the second alternative may be preferable. Often, the appropriate choice is application-specific. However, it is still important to provide a small set of metrics to quantify the cost of a configuration. For example, one possibility is to measure the cost of a configuration as some weighted combination of data rate, transmission delay, energy consumed, etc. Higher layer protocols may be allowed to control the choice of weights based on the importance of different parameters.

In this section, we have presented some directions for abstracting channel information from higher layers. It is an avenue for future work to develop mechanisms for quantifying the cost of different configurations, and developing abstractions for cross-layer exchange of configuration information.

IV. Exploiting heterogeneity

Heterogeneity can be managed by designing mechanisms that hide most of the heterogeneity from higher layer protocols (as discussed in Section II). In addition, channel heterogeneity can also be exploited in innovative ways to enhance network performance. In this section, we provide one example to highlight the possibility of benefiting from channel heterogeneity.

When different channels exhibit different characteristics, they can be used to support different operations. For example, channels situated in low frequency bands (where less bandwidth is available) may support low data rates, but have a longer transmission range. On the other hand, channels situated in high frequency bands may support high data rates, but have a shorter transmission range. Now, consider a mobile ad hoc network that can use such heterogeneous channels. In such networks, routes break when nodes forming the end-points of a link move away from each other. Route failure frequently requires a new route to be discovered, and while a new route is being discovered no data communication is possible. Higher layer protocols, such as TCP, are adversely affected by such interruptions to data communication. For example, when a route fails, TCP may timeout and reduce its congestion window, and even after a new route is discovered, there is a delay before TCP starts sending data at a high rate.

To overcome poor TCP performance, the network can maintain two routes for any pair of communicating nodes [16]. One route uses shorter-range higher data rate channels, and is normally used for data communication. The second route, called “backup route”, uses longer-range lower data rate channels, and is used to provide connectivity when the higher data rate route fails. When a higher data rate route breaks, possibly on account of node mobility, it is likely that the lower data rate route is still available, as it uses longer range links (between two nodes, longer range link may still be present when the shorter range link has failed). Until a new higher data rate route is discovered, the lower rate “backup route” can be used for data communication. Since data can still be sent at a lower rate over the backup route, TCP does not timeout, enabling faster TCP recovery when a higher data rate route becomes available. In this fashion, channel heterogeneity can be exploited using specialized protocols to improve TCP performance.

In general, although channel heterogeneity introduces significant protocol design challenges, it may also offer many opportunities to develop innovative solutions that enhance network performance.

V. Conclusion

In this paper, we have highlighted some of the higher layer protocol design challenges with dynamic spectrum access networks. The opportunistic utilization of available spectrum leads to various forms of channel heterogeneity, and managing this heterogeneity is a key research challenge. We described how channel heterogeneity can be addressed by a combination of solutions that may be distributed between lower and higher layers of the protocol stack. We motivated the need for developing simple channel abstractions to allow higher layers to manage channel heterogeneity. We also illustrated the possibility of using channel heterogeneity in innovative ways to improve network performance.

References

[1] *IEEE Standard for Wireless LAN-Medium Access Control and Physical Layer Specification, P802.11*, 1999.

[2] J. Mitola III, *Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio*, Ph.D. thesis, KTH Royal Institute of Technology, 2000.

[3] C. J. Rieser, T. W. Rondeau, C. W. Bostian, and T. M. Gallagher, “Cognitive Radio Testbed: Further Details and Testing of a Distributed Genetic Algorithm Based Cognitive Engine for Programmable Radios,” in *Milcom*, 2004.

[4] R. W. Brodersen, A. Wolisz, D. Cabric, S. Mishra, and D. Willkomm, “CORVUS: A cognitive radio approach for usage of virtual unlicensed spectrum,” White paper at http://bwrc.eecs.berkeley.edu/Research/MCMA/.

[5] A. Nasipuri, J. Zhuang, and S.R. Das, “A Multichannel CSMA MAC Protocol for Multihop Wireless Networks,” in *WCNC*, 1999.

[6] R. Draves, J. Padhye, and B. Zill, “Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks,” in *ACM Mobicom*, 2004.

[7] J. So and N. H. Vaidya, “A Routing Protocol for Utilizing Multiple Channels in Multi-Hop Wireless Networks with a Single Transceiver,” Tech. Rep., University of Illinois at Urbana-Champaign, October 2004.

[8] J. So and N. H. Vaidya, “Multi-channel MAC for Ad Hoc Networks: Handling Multi-Channel Hidden Terminals using a Single Transceiver,” in *Mobihoc*, 2004.

[9] P. Bahl, R. Chandra, and J. Dunagan, “SSCH: Slotted Seeded Channel Hopping for Capacity Improvement in IEEE 802.11 Ad-Hoc Wireless Networks,” in *ACM Mobicom*, 2004.

[10] A. Raniwala and T. Chiueh, “Architecture and Algorithms for an IEEE 802.11-Based Multi-Channel Wireless Mesh Network,” in *Infocom*, 2005.

[11] P. Kyasanur and N. H. Vaidya, “Routing and Interface Assignment in Multi-Channel Multi-Interface Wireless Networks,” in *WCNC*, 2005.

[12] P. Kyasanur and N. H. Vaidya, “Routing and Link-layer Protocols for Multi-Channel Multi-Interface Ad hoc Wireless Networks,” Tech. Rep., University of Illinois at Urbana-Champaign, May 2005.

[13] P. Kyasanur, J. Padhye, and P. Bahl, “A Study of an 802.11-like Control Channel-Based MAC,” in *IEEE International Conference on Broadband Networks (Broadnets)*, 2005.

[14] P. Kyasanur and N. H. Vaidya, “Capacity of Multi-Channel Wireless Networks: Impact of Number of Channels and Interfaces,” in *ACM Mobicom*, 2005.

[15] Theodore Rappaport, *Wireless Communications Principles and Practice*, Prentice Hall, 2002.

[16] N. H. Vaidya, “Exploiting Diversity in Mesh Networks,” Presentation online at: http://research.microsoft.com/meshsummit/techprogram.aspx.