Mesh Networking Protocols to Exploit Physical Layer Capabilities*

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

Xue Yang Nitin H. Vaidya
Dept. of Electrical and Computer Eng., and
Coordinated Science Laboratory,
University of Illinois at Urbana-Champaign
Email: {xueyang,nhv}@uiuc.edu

Abstract— **Wireless mesh networks are becoming increasingly popular as a low-cost alternative to wired networks for providing high-speed last mile connectivity. A key challenge in mesh networks is the need for sufficient network capacity to meet the requirements of applications, especially when network density increases over time, and newer applications require higher throughputs. In this paper, we consider some approaches for improving network capacity by exploiting various physical layer capabilities. Specifically, we consider the use of multiple wireless channels, and improving utilization of any given channel by introducing "spatial backoff", for improving network capacity. Through these example scenarios, we highlight the challenges involved and benefits possible by exploiting physical layer capabilities in mesh networks.**

I. INTRODUCTION

In recent years, mesh networks have been advocated as a cost-effective approach for providing high-speed last mile connectivity. In addition, mesh networks may also spur the growth of new neighborhood-specific applications, such as video sharing among community members, that require high throughput. To build cost-effective mesh networks, it is desirable to operate the network in a *multi-hop* fashion using commodity wireless hardware. However, wireless medium is a shared resource, and the capacity of a multi-hop network quickly degrades as the node density and network diameter increases. As a result, there is a need to develop solutions to enhance the network capacity to meet the needs of mesh network applications.

Existing commodity hardware [1], [2] often allows users to control several physical layer parameters, such as transmission power, data rate, frequency of operation, etc. Such support can be utilized to improve the network capacity. In this paper, we present two examples of improving network capacity by using physical layer support. The first example considers utilizing multiple frequency-separated channels, which are often provisioned for in wireless standards, to increase network capacity. The second example considers an approach to improve the capacity of any given channel by introducing "spatial backoff" mechanisms. In the rest of this section, we motivate the need for new protocols in the above-mentioned two example scenarios for exploiting physical layer capabilities.

The rest of the paper is organized as follows. We motivate the need for new mesh networking protocols in Section II. We present approaches for utilizing multiple channels in Section III. "Spatial backoff" mechanisms are discussed in Section IV, and we conclude in Section V.

II. MOTIVATION

In this section, we motivate the need for mesh networking protocols that exploit physical layer capabilities to utilize multiple channels, and to improve channel utilization using "spatial backoff".

A. Need for protocols to exploit multiple channels

There has been significant recent interest in designing protocols for using multiple channels in wireless networks [3]–[20]. The (unlicensed) spectrum that is available for use in mesh networks is typically divided into multiple channels [21]. Commodity radio interfaces expose the list of available channels to higher layers, while hiding the physical layer details of supporting channels. For example, IEEE 802.11a standard defines 12 channels in the 5 GHz band (in US). Commercially available IEEE 802.11 interfaces [1], [2] allow higher layer protocols to select the channel of operation, and switch channels when necessary. In spite of this support from the physical layer for using multiple channels with a single interface, typical multi-hop mesh networks have used a single channel for the whole network. Two adjacent wireless nodes can communicate with each other only if they have at least one interface on a common channel. This condition is met in typical mesh network configurations by ensuring all nodes use a *single common channel*, thereby allowing any pair of adjacent nodes to communicate with each other.

Reducing hardware costs have made it feasible to equip wireless nodes with multiple radio interfaces. However, it may be difficult to equip nodes with one interface for each channel, especially when the number of available channels is large. Interfaces typically allow channels to be switched when necessary, though the switching of channels may incur a non-negligible delay. In our research, we attempt to fully utilize

---

[*] This work was supported in part by the US Army Research Office under grant W911NF-05-1-0246, and National Science Foundation under grant ANI-0125859.

all the available channels even if the number of interfaces is smaller than the number of channels. For example, when c channels are available, but each node has only $m < c$ interfaces, one possibility is to keep the m interfaces fixed on some m common channels [10], thereby not using the remaining $c - m$ channel. Although this approach simplifies co-ordination among neighboring nodes, it may offer only a m-fold (potentially $m \ll c$) increase in the network capacity. In contrast, we aim to utilize all the available channels and achieve close to a c-fold increase in capacity, by switching the available m interfaces among the c channels. Our past analysis has shown that in theory [22], significant performance gains can be obtained even if only a few interfaces are available. In this paper, we will briefly outline a practical architecture for realizing the theoretically predicted gains using commodity hardware.

In addition to using spectrum already provisioned for unlicensed use, there is a growing interest in exploiting additional spectrum by *dynamically utilizing existing licensed spectrum* with spectrum agile "cognitive radios" [23]. Cognitive radios are designed to allow secondary users to co-exist with the primary users of the licensed spectrum, and cognitive radio technology can significantly increase the total spectrum that is available for use. Cognitive radios may be an excellent fit for mesh networks that will be deployed in dense urban areas with possibility of significant contention. However, realizing the performance gains promised by cognitive technology requires several protocol design challenges to be addressed. In this paper, we will also highlight the challenges of using cognitive radio technology, and outline how the capabilities of a cognitive radio may be used to enhance network capacity.

B. Need for protocols to improve channel utilization using "spatial backoff"

Wireless channel is a shared medium, and medium access control (MAC) protocols are used to regulate the channel access among multiple competing stations. Taking the set of competing stations as given, prior research on MAC protocols proposed numerous ways for each station to adjust its channel access behavior (e.g., using temporal backoff), so that transmissions from different stations may be separated in time to achieve successful transmissions. This is a temporal approach to resolve channel contention. Since the given set of competing stations may vary significantly depending on the network load, it remains a major challenge to design MAC protocols that can function efficiently under various network loads.

We propose an alternative approach for wireless networks – named "spatial backoff" – that adapts the "space" occupied by the transmissions. Wireless nodes communicate over the air and there is significant interference among nodes that are spatially close to each other. On the other hand, due to radio signal attenuation, nodes that are sufficiently apart from each other are able to reuse the channel spectrum and transmit at the same time. In other words, for a node S, one can visualize the channel contention by means of contending area ω around S, where nodes located within this area compete for the channel with node S and nodes outside of this area (e.g., S1, S2, S3) may transmit concurrently with node S, as illustrated in Figure 1[1]. By spatially adjusting the contending area ω, the set of competing stations can be controlled. In this paper, we will discuss the benefits of spatial backoff, and suggest different ways to realize it by exploiting physical layer capabilities such as transmission power control, rate control, adjusting carrier sense threshold, etc.

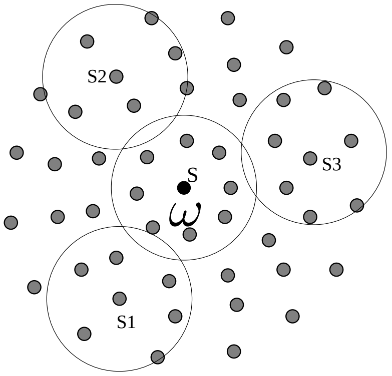

Fig. 1. Contending Area

III. UTILIZING MULTIPLE CHANNELS

In this section, we will first briefly describe an architecture for exploiting multiple wireless channels using currently available commodity hardware. We then discuss challenges in exploiting multiple channels with smarter "cognitive" radios that may be available in the near future.

A. Architecture for utilizing multiple channels

We present an architecture for utilizing multiple channels in multi-hop wireless networks using commodity hardware. Our solution requires at least two wireless interfaces at each node. We exploit the ability of the interfaces to switch the channel of operation under the control of a higher layer protocol. In our solution (which has been discussed in detail in [17], [18]), we develop a link layer protocol to manage the use of multiple interfaces, and a routing protocol that interacts with the link layer protocol to select good routes. Such a separation of functionality is used to simplify protocol design (theoretical results in [22] have shown that even with such a separation, network capacity is not degraded). Interface switching can occur on the time scales of a few packet transmissions; hence it is beneficial to incorporate interface management at the link layer, as part of the kernel. Route selection happens on larger time scales (often hundreds of packet transmissions or more), and it is beneficial to implement it separately, possibly as a user space daemon.

---

[1] The area is shown circular only for the sake of illustration. In general, whether a station may be allowed to transmit or not depends on the protocol design. By "transmit currently", we mean the transmissions from more than one stations overlap in time

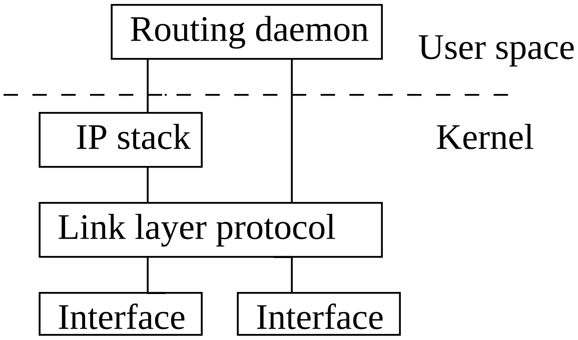

Fig. 2. Proposed multi-channel architecture

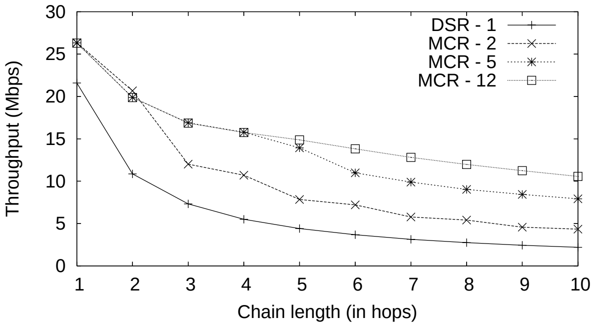

Fig. 3. Performance of single FTP flow: Even two interfaces are sufficient to utilize 12 channels.

Figure 2 outlines the proposed architecture. A key benefit of this approach is that even existing routing protocols can be used without any modifications, since the link layer protocol completely hides the complexity of managing multiple channels and interfaces from the higher layers. We have already implemented most components of this architecture in a Linux-based testbed as a user space daemon complemented with a kernel module, and our ongoing implementation experience has indicated that the architecture is indeed practical to implement.

Simulation experiments to evaluate the efficacy of the architecture have been encouraging. We will present a sample result to illustrate the benefits obtained by using multiple channels and multiple interfaces. In Figure 3, we evaluate the performance of the proposed architecture in simple chain topologies. The length of a chain is varied from 1 to 10 hops. A FTP flow is setup from the first node to the last node of the chain. Figure 3 compares the flow throughput with DSR protocol operating in a 1 channel network (labeled “DSR - 1”), and the flow throughput with the proposed architecture when using our link and multi-channel routing protocols (MCR) with varying number of channels (labeled “MCR - x” where x is the number of channels available). The experiment assumes that our architecture uses two interfaces. As we can see from the figure, the FTP throughput in single channel networks rapidly degrades when the number of hops along a chain increases (this behavior is well-known).

When multiple channels and multiple interfaces are used in the proposed architecture, the link layer protocol assigns the fixed channel of successive nodes along the chain to different channels. Also, when an intermediate node is receiving data using one interface, it can simultaneously forward data to the next node using the second interface. Consequently, MCR offers higher throughput by *using different channels on successive hops*, and by *using the two interfaces to receive and send data in parallel*.

The key observation from Figure 3 is that *multiple channels can significantly improve throughput of a flow* in multi-hop scenarios. Furthermore, even with only a few interfaces (2 in this example), having large number of channels (up to 12 channels in this example) is beneficial. Hence, by exploiting the physical layer capability of channel switching that is exposed by an interface, a large number of channels can be utilized with few interfaces.

B. Exploiting capabilities of cognitive radio interfaces

Earlier in this section, we presented an architecture for utilizing multiple channels using commodity radios. The protocols used in the architecture assumed that all channels have similar range and support similar data rates (i.e., “homogeneous” channels). Such an assumption is reasonable when all the channels are part of a common frequency band (e.g., the 12 channels in 5 GHz band in IEEE 802.11a). However, with improving radio technology, in the near future, it may be possible to use cognitive radios [23]–[25] that opportunistically utilize spectrum across a wide range of frequency bands.

Channels that are present in widely separated frequency bands may be “heterogeneous”, i.e., different channels may support different transmission ranges, data rates, delay characteristics, etc. In the rest of this section, we describe the impact of heterogeneous channels on higher layer protocols. We will also outline some approaches for addressing these challenges.

1) Heterogeneity in transmission range: 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 from a node can be received may have a complex shape 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 [26]. The frequency-dependent variations is larger with channels that are operating farther apart in frequency.

Figure 4 shows the unlicensed spectrum bands that are used by popular wireless technologies, such as IEEE 802.11. In addition to the already available unlicensed spectrum, FCC is considering allowing additional spectrum below 3 GHz (for example in the 700 MHz band) for use by nodes equipped with cognitive radios. Therefore, the channels supported by a cognitive radio may be located on widely separated slices of

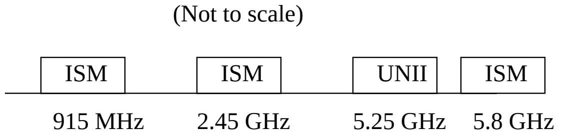

Fig. 4. Chart of commonly used unlicensed spectrum. Spectrum in the 2.4 GHz band is being used by IEEE 802.11b/g devices, and spectrum in the 5 GHz band is being used by IEEE 802.11a devices.

the frequency spectrum, with different propagation characteristics. Furthermore, 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, multi-path effects, and interference from other devices, a channel with longer transmission range may not cover all the area covered by a channel with a shorter transmission range.

As we discuss later in this section, 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 (as we will discuss in Section III-B.2). Therefore, equalizing the range of all channels to allow the use of existing higher layer protocols may often involve several tradeoffs such as loss in connectivity, which have to be taken into consideration.

We next discuss the impact range heterogeneity has on 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 [17]. 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

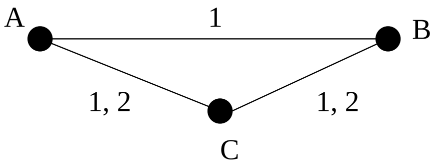

Fig. 5. Impact of range heterogeneity: The route chosen between A and B may depend on the channel used to exchange routing information.

of this approach is the possibility of not using certain communication links which could otherwise have improved network performance.

We illustrate the impact of range heterogeneity with a simple example. Figure 5 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 the route discovery (i.e., exchanging control information through broadcasts), 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 used 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. Thus, new trade-offs arise with range heterogeneity between reducing protocol overheads and maximizing performance.

2) Heterogeneity in channel performance: The rate at which data can be exchanged between a pair of nodes on a given channel may depend on the modulation scheme used (as well as the redundancy in the error correcting code, etc.). On a given channel, one approach for supporting multiple data rates is by appropriately changing the modulation scheme based on the channel quality between the sender and the receiver. Different pairs of nodes may use different data rates on the same channel depending on their observed channel quality. This feature is commonly called “auto-rate” selection, and is supported in commodity IEEE 802.11 hardware. Such data rate heterogeneity on any given channel has been accounted for, during route selection, in some existing multi-channel protocols [10], [18]. However, in networks with cognitive radios, the supported data rates *across channels* may also significantly vary, in addition to different data rates being supported on a *given channel*. The variation in data rates across channels may in part be due to the variations in the bandwidth of channels. For example, less spectrum is available in lower frequency bands, and therefore, channels in the lower frequency bands may have smaller bandwidth. As a result, the range of data rates supported by different channels may vary and has to be accounted for by higher layer protocols.

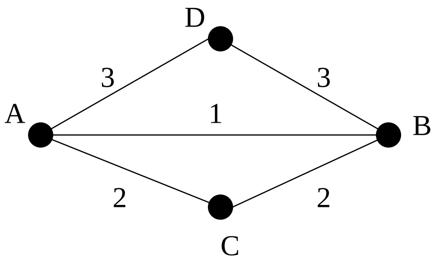

Fig. 6. Impact of variable data rate and variable access time: The choice of route to use between A and B depends on the channel data rates and access latencies.

Figure 6 provides an example to illustrate the impact of variable data rate on performance. Suppose node A is setting up a route to node B, with three possible routes A-B, A-C-B and A-D-B, with the routes using channels 1, 2 and 3 respectively. Now, if all three channels support the same data rate, then route A-B is preferable as it uses smaller number of hops. However, when different channels support different data rates, route A-C-B or A-D-B may be preferable over route A-B if channels 2 and 3 support significantly higher data rate than channel 1. Now, suppose channels 2 and 3 support the same data rate. Even then, whether route A-C-B or A-D-B is better may depend on whether channel 2 or channel 3 has lower channel access time, respectively (as performance of certain protocols such as TCP depends on end-to-end latency). Thus, in general, higher layer protocols need to be aware of the properties of the heterogeneous channels.

To summarize, cognitive radios may enable access to larger amount of spectrum, but at the same time require more complex protocols to utilize the spectrum. Therefore, there is a need for new mesh networking protocols that are aware of the physical layer capabilities offered by the radio, and carefully utilize the available capabilities to maximize performance.

IV. SPATIAL BACKOFF

In this section, we consider improving the utilization of any given channel by introducing spatial backoff. The “space” occupied while contending for channel access depends in part on the transmission power used, data rate of transmission, carrier sense threshold of nodes, etc. Different choices can be made to adjust the contending area. To do so, we often need to explore the interactions between MAC and physical layers.

In particular, let us consider MAC protocols based on Carrier Sense Multiple Access (CSMA). Carrier sense refers to listening to the physical medium to detect any ongoing transmissions. Only if the radio signal strength detected at a station is below a *Carrier Sense Threshold* CS_{th}, may the attempt of the station to access the channel proceed. Given a fixed transmission power used by other stations, a node will transmit more aggressively using higher carrier sense threshold values. For example, in Figure 7, station A is transmitting to

B. The curve represents the signal strength versus distance for A's transmission. When station D uses carrier sense threshold CS1, D has to compete the channel access with station A. Whenever A is transmitting, D is required to defer its transmissions. On the other hand, when carrier sense threshold CS2 is used, D is allowed to transmit concurrently when A is transmitting. Therefore, a higher carrier sense threshold will lead to a smaller contending area. Similarly, given a carrier sense threshold used by other stations, a lower transmission power will lead to a smaller contending area.

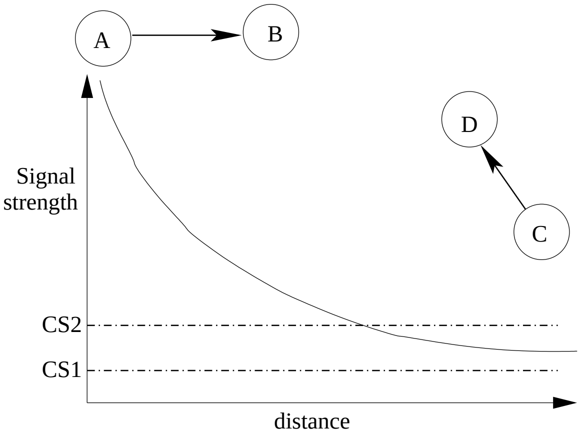

Fig. 7. Larger carrier sense threshold leads to smaller contending area

Notice that, when adjusting the contending area, the interference present in the network varies. For example, increasing the carrier-sense threshold (with a fixed transmission power) allows transmitters to be near each other and causes more interference. Typically, the quality of a communication link depends on the interference at the receiver caused by other concurrent transmissions; the higher the Signal-to-Interference-and-Noise-Ratio (SINR), the higher the rate that packets can be transmitted reliably. To account for the change of interference, the transmission rate often needs to be adjusted along with the contending area. A smaller contending area often reduces the channel contention at the cost of poorer link quality.

One possible benefit of spatial backoff is to improve the network aggregate throughput. From MAC protocol point of view, given a network, the aggregate throughput depends on the MAC efficiency in resolving the "local" channel contention, the number of concurrent transmissions in the network, and the transmission rate between each transmitter/receiver pair. Our past study [27] shows that, when transmitter density increases, a smaller contending area is preferred to bring concurrent transmitters closer to each other. By doing this, the MAC efficiency in resolving the "local" channel contention can be improved due to the reduced number of competing stations. At the same time, spatial reuse is improved since more concurrent transmissions are allowed to proceed. Consequently, the aggregate throughput can be higher even though the transmitters may have to transmit at lower rate because of larger interference. Such a benefit of spatial backoff cannot be achieved by existing rate control protocols [28]–[30] because changing transmission rate alone will not improve the spatial reuse. There is some prior work on carrier sense threshold control [31], [32] to maximize the spatial reuse given a predefined transmission rate. However, such schemes are unable to explore possibly more spatial reuse when lower transmission rates are used.

We have investigated different approaches to implement spatial backoff by controlling carrier sense threshold, transmission rate or transmission power based on our prior work [27], [33]. Below, we introduce one spatial backoff algorithm, which controls the carrier sense threshold and transmission rate, assuming that the transmission power is fixed.

The goal of the spatial backoff algorithm is to find a good combination of carrier sense threshold and transmission rate so that the network aggregate throughput may approach the maximum point. Additionally, it is desirable to have a distributed algorithm so that each source station may make decisions based on its local information. To this end, we developed a model to quantify the performance at each individual station. Specifically, let $rate_i$ be the transmission rate and cs_i be the carrier sense threshold used by station i. Let p_{suc_i} denote the percentage of transmitted packets being successful for a certain measuring period, given the chosen cs_i and $rate_i$. We define an utility measure as follows:

$$U_i = rate_i * cs_i^{\frac{2}{\theta}} * p_{suc_i}, \tag{1}$$

where θ is the path loss coefficient and $cs_i^{\frac{2}{\theta}}$ is used to quantify the number of concurrent transmissions that can be possibly allowed in the network, assuming all stations use the same carrier sense threshold cs_i. The utility function defined in Equation 1 has the following desirable properties:

- By introducing p_{suc} into the utility function, we take into account the impact of MAC efficiency on aggregate throughput. If the carrier sense threshold cs_i is chosen to be too small, the contending area around station i will be too large. With inappropriately large number of competing stations, the packet success probability p_{suc} is likely to be low, which results in bad utility measure. In other words, the defined utility function helps to maintain a suitable size for the contending area.
- Given a fixed $rate_i$ and p_{suc_i}, the utility (U_i) is a monotonically increasing function of cs_i, which helps to encourage more spatial reuse. However, if cs_i is inappropriately large, the $SINR$ required by the transmission rate may not be satisfied due to large interference. As a result, p_{suc_i} is likely to be very low and the utility measure will be bad.
- Given a fixed cs_i and p_{suc_i}, the utility (U_i) is a monotonically increasing function of $rate_i$, which encourages the use of the highest transmission rate that can be possibly supported based on the $SINR$ at the receiver. However, if station i chooses an inappropriately high transmission rate, the required $SINR$ can no longer be satisfied. As

a result, the transmission is likely to fail and p_{suc_i} will be very low, leading to bad utility measure.

In essence, the utility function defined in Equation 1 measures the channel utilization per unit area. We can argue that, in dense networks, the carrier sense threshold that maximizes the above utility function approaches the point that maximizes the aggregate throughput.

Based on the utility measure, we designed a protocol [34] for each individual station to search for the appropriate carrier sense threshold and transmission rate operating points. In Figure 8, we present the ns-2 simulation results for our spatial backoff algorithm in a circular topology, in which 32 transmitters (always back-logged) are evenly distributed along a circle with a radius of 350 meters. The receiver corresponding to a transmitter is located on the line from the transmitter to the center of the circle, and is 35 meters away from the transmitter. In our simulations, the physical layer follows the specifications of IEEE 802.11a, and MAC layer follows the specifications of IEEE 802.11 DCF, except that a fixed contention window size is used.

In Figure 8, horizontal axis represents the ratio between carrier sense threshold and receiving signal threshold (in dB), vertical axis represents the aggregate throughput. We first obtain the aggregate throughput for different combinations of transmission rate and carrier sense threshold. As we can see, in this example, the maximum aggregate throughput is achieved when all stations choose transmission rate as 18 Mbps and normalized carrier sense threshold as -6 dB. Our spatial backoff algorithm indeed finds the optimal point and approaches the maximum aggregate throughput, as the arrow in the figure points out. We have also evaluated above spatial backoff algorithm in random topologies, and the results are very encouraging.

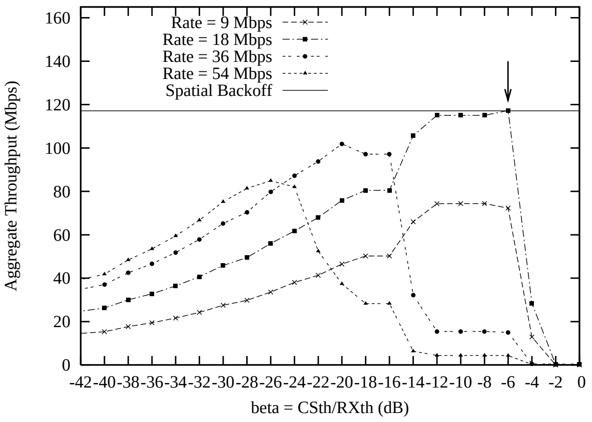

Fig. 8. Aggregate throughput for a circular topology with 32 transmitters

In general, spatial backoff approaches can be combined with temporal contention resolution approaches. For example, spatial backoff can be used to control the number of competing stations to be within a certain range in which the temporal approach is able to resolve the channel contention most efficiently. Essentially, by combining the temporal and spatial approaches, we are expanding our decision space when determining the strategies for resolving the channel contention: strategies decide when is the right time to access the channel, what carrier sense threshold and transmission power to use, and which rate to transmit at. By combining the spatial and temporal contention resolution appropriately, we expect further performance gain.

V. CONCLUSION

Although existing commodity radios offer several features like transmission power control, rate control, channel selection, etc., very few real world protocols actually utilize these features. Furthermore, newer cognitive radio technologies may offer greater flexibility in accessing channel resources. In this paper, we have highlighted the benefits of using physical layer capabilities offered by radio interfaces for enhancing performance. We have focused on two specific approaches to improving performance – using multiple channels to increase total available bandwidth, and improving the utilization of any given channel by applying spatial backoff mechanisms. It is part of our ongoing work to incorporate some of these solutions in a prototype mesh network testbed.

REFERENCES

[1] “Atheros Inc,” http://www.atheros.com.

[2] “Maxim 2.4 GHz 802.11b Zero-IF Transceivers,” http://pdfserv.maxim-ic.com/en/ds/MAX2820-MAX2821.pdf.

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

[4] A. Nasipuri and S.R. Das, “Multichannel CSMA with Signal Power-based Channel Selection for Multihop Wireless Networks,” in *VTC*, Sept 2000.

[5] S.-L. Wu, C.-Y. Lin, Y.-C. Tseng, and J.-P. Sheu, “A New Multi-Channel MAC Protocol with On-Demand Channel Assignment for Multi-Hop Mobile Ad Hoc Networks,” in *International Symposium on Parallel Architectures, Algorithms and Networks (ISPAN)*, 2000.

[6] N. Jain, S. Das, and A. Nasipuri, “A Multichannel CSMA MAC Protocol with Receiver-Based Channel Selection for Multihop Wireless Networks,” in *IC3N*, October 2001.

[7] W.-C. Hung, K.L.Eddie Law, and A. Leon-Garcia, “A Dynamic Multi-Channel MAC for Ad Hoc LAN,” in *21st Biennial Symposium on Communications*, Kingston, Canada, June 2002, pp. 31–35.

[8] Y. Li, H. Wu, D. Perkins, N.-F. Tzeng, and M. Bayoumi, “MAC-SCC:Medium Access Control with a Separate Control Channel for Multihop Wireless Networks,” in *IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW’03)*, 2003.

[9] J. Li, Z. J. Haas, M. Sheng, and Yanhui Chen, “Performance Evaluation of Modified IEEE 802.11 MAC for Multi-Channel Multi-Hop Ad Hoc Network,” in *IEEE International Conference on Advanced Information Networking and Applications (AINA)*, 2003.

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

[11] 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.

[12] 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.

[13] A. Raniwala, K. Gopalan, and T. Chiueh, “Centralized Channel Assignment and Routing Algorithms for Multi-Channel Wireless Mesh Networks,” *Mobile Computing and Communications Review*, vol. 8, no. 2, pp. 50–65, April 2004.

[14] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, “A Multi-Radio Unification Protocol for IEEE 802.11 Wireless Networks,” in *IEEE International Conference on Broadband Networks (Broadnets)*, 2004.

[15] 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.

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

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

[18] 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.

[19] M. X. Gong and S. F. Midkiff, “Distributed Channel Assignment Protocols: A Cross-Layer Approach,” in *WCNC*, 2005.

[20] 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.

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

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

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

[24] 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*, Monterey, CA, 2004.

[25] Robert W. Brodersen, Adam Wolisz, Danijela Cabric, Shridhar Mishra, and Daniel Willkomm, “CORVUS: A cognitive radio approach for usage of virtual unlicensed spectrum,” White paper, Available for download from http://bwrc.eecs.berkeley.edu/Research/MCMA/.

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

[27] Xue Yang and Nitin H. Vaidya, “On physical carrier sensing in wireless ad hoc networks,” in *IEEE Infocom*, 2005.

[28] A. Kamerman and L. Monteban, “WaveLAN-II: A high-performance wireless LAN for the unlicensed band,” *Bell Labs Technical Journal*, pp. 118–133, Summer 1997.

[29] Gavin Holland, Nitin Vaidya, and Paramvir Bahl, “A Rate-Adaptive MAC Protocol for Multi-hop Wireless Networks,” in *ACM International Conference on Mobile Computing and Networking (MobiCom)*, July. 2001.

[30] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, “Opportunistic Media Access for Multirate Ad Hoc Networks,” in *ACM International Conference on Mobile Computing and Networking (MobiCom)*, Sep. 2002.

[31] Jing Zhu, Xingang Guo, L. Lily Yang, W. Steven Conner, Sumit Roy, and Mousumi M. Hazra, “Adapting physical carrier sensing to maximize spatial reuse in 802.11 mesh networks,” *Wireless Communications and Mobile Computing*, vol. 4, no. 8, pp. 933–946, Dec. 2004.

[32] Arunchandar Vasan, Ramachandran Ramjee, and Thomas Woo, “ECHOS - Enhanced Capacity 802.11 Hotspots,” in *IEEE Infocom*, 2005.

[33] Jason Fuemmeler, Nitin Vaidya, and Venugopal V. Veeravalli, “Selecting Transmit Powers and Carrier Sense Thresholds for CSMA Protocols,” Tech. Rep., University of Illinois at Urbana-Champaign, October 2004.

[34] Xue Yang, *Efficient Packet Scheduling in Wireless Multi-hop Networks*, Ph.D. thesis, University of Illinois at Urbana-Champaign, 2005, Under preparation.