Spatial Backoff Contention Resolution for Wireless Networks

Technical Report

March 2006

Xue Yang
Communications Technology Lab
Corporate Technology Group
Intel Corporation
2111 N.E. 25th Ave
Hillsboro, OR 97124
Email: xue.yang@intel.com

Nitin Vaidya
University of Illinois at Urbana-Champaign
Department of Electrical and Computer Engineering,
and Coordinated Science Laboratory
1308 West Main Street
Urbana, IL 61801
Email: nhv@uiuc.edu

***Abstract*— Traditional medium access control (MAC) protocols utilize temporal mechanisms such as *access probability* or *backoff interval* adaptation for contention resolution. They typically take the set of competing nodes as a given, and address the problem of adapting each node's channel access behavior to the given channel contention level. This is a temporal approach for contention resolution, which aims to separate transmissions from different nodes in time to achieve successful transmissions.**

We explore an alternative approach for wireless networks—named "spatial backoff"—that adapts the "space" occupied by the transmissions. Each transmission in a wireless network competes for a certain space. By adapting the space occupied by transmissions, the set of "locally" competing nodes, and thus, the channel contention level, can be adjusted to reach a suitable level. There are different ways to realize spatial backoff. In this paper, we propose a dynamic spatial backoff algorithm using the joint control of carrier sense threshold and transmission rate. Our results suggest that spatial backoff can lead to a substantial gain in channel utilization.

I. Introduction

Past studies on contention-based medium access control (MAC) protocols have often taken a temporal approach. That is, when two nodes are competing for a common channel, their channel accesses are separated in time to ensure successful transmissions. For example, nodes A and B in Figure 1 compete for the channel access. Since their transmissions interfere with each other, A and B's transmissions are separated in time using temporal contention resolution. Such a temporal approach typically takes the set of competing nodes as a given, and addresses the problem of adapting each node's channel access behavior (e.g., channel access probability, or backoff interval) to the given channel contention level.

In this paper, we explore an alternative contention resolution approach for wireless networks—named "spatial backoff"—that adapts the space occupied by transmissions. Observe that transmissions in wireless networks compete for space as well, since wireless nodes communicate over air and there is significant interference among nodes that are spatially close to each other. We consider each transmission as occupying certain part of the space. Intuitively, node B is said to be within the space occupied by a transmission from node A, if a concurrent transmission from node B will prevent reliable reception of A's transmission. Thus, the space occupied by a transmission depends on the signal level at the intended receiver as well as interference that may be posed by other nodes. It should be noted that our proposed protocol does not rely on knowledge of the space occupied by a transmission.

As shown in Figure 1, for the sake of illustration, we use a shaded circular area to represent the space occupied by a transmission. In practice, however, the occupied space is not necessarily circular.[1] Nodes A and B's transmissions in Figure 1 cannot overlap in time because they are within each other's occupied space. On the other hand, if small enough spaces are occupied by nodes A and B's transmissions, as shown in Figure 2, the two transmissions can be separated in space and they both can proceed successfully at the same time.

We refer to the space occupied by a transmission as the "contending region" around the transmitter. It should be noted that contending region is a property of a transmission, although we may often associate the contending region with the transmitter for the sake of brevity. Figure 3 illustrates the contending region ω around S for a transmission from node S to node D.[2] Transmissions by nodes located within the contending region of S (e.g., $S4$, $S5$, $S6$, $S7$) will cause erroneous reception of the transmissions from S to D. Therefore, node S has to compete for the channel *in time dimension* with the nodes in its contending region ω. We refer to such contention as "local channel contention" of node S.

[1] In general, the interference that may be posed by other nodes depends on the protocol parameters. The interference tolerance level of a transmission depends on the transmission power, rate, and etc.

[2] Again, the area is shown circular only for the sake of illustration.

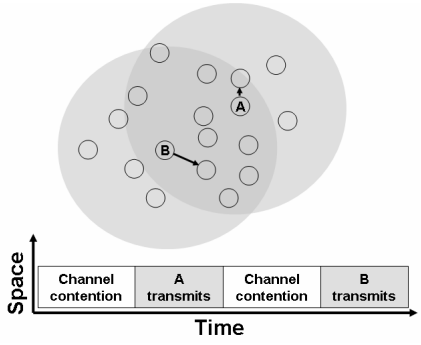

Fig. 1. Temporal contention resolution.

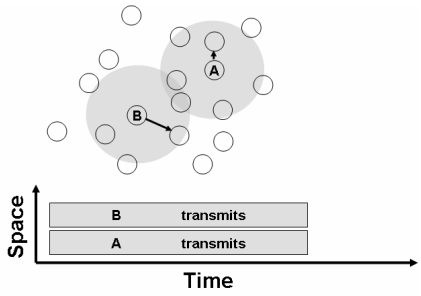

Fig. 2. Spatial contention resolution.

Nodes outside of the contending region of S (e.g., $S1$, $S2$, $S3$) may have their transmissions overlapped in time with node S as long as the SINR requirements of these transmissions can be satisfied. By using spatial backoff, the size of contending region ω can be adapted (as elaborated later). As a result, the local channel contention of each node can be adjusted to a suitable level, such that the local contention can be resolved efficiently using temporal contention resolution mechanisms. At the same time, by possibly allowing more concurrent transmissions, the space can also be utilized more efficiently. Consequently, spatial backoff can help to improve the channel utilization, and thus, the network aggregate throughput.

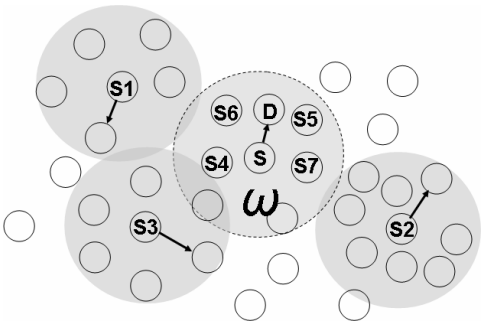

Fig. 3. Contending region.

The rest of this paper is organized as follows. In Section II, we discuss how spatial backoff affects throughput. Proposed dynamic spatial backoff algorithm is discussed in Section III, and its performance is evaluated in Section IV. Section V summarizes the related work. We present conclusions and future directions in Section VI.

II. Spatial Backoff Affects Throughput

A. Adapting CS threshold and transmission rate

The space occupied by a node while competing for channel access depends on many factors such as its transmission power, transmission rate, and also the interference caused by other transmissions. Different approaches can be designed to adjust the contending region ω. We consider a MAC protocol based on Carrier Sense Multiple Access (CSMA) as an example. Carrier sensing refers to listening to the physical medium to detect ongoing transmissions. Only if the received signal strength detected at a node is below a *Carrier Sense (CS) Threshold* CS_{th} may the node access the wireless channel. Given a fixed transmission power used by other nodes, a node will transmit more aggressively using a higher CS threshold, as the example in Figure 4 illustrates. The horizontal axis in Figure 4 represents the distance from node A; the vertical axis represents the signal strength of A's transmissions; and the curve plots the received signal strength versus distance for A's transmissions. When node D uses CS threshold $CS1$, D is required to defer its transmissions whenever A is transmitting, which implies that D has to compete for the channel access *in time* with node A. On the other hand, when a higher CS threshold $CS2$ is used, node D is allowed to transmit to C at the same time when A is transmitting to B.

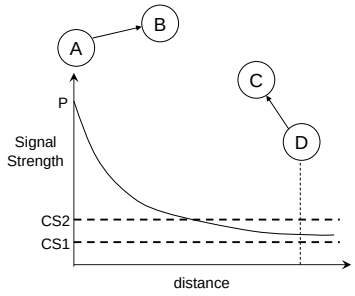

Fig. 4. Larger CS threshold and lower rate lead to smaller contending region.

Note that increasing the CS threshold (with a fixed transmission power) allows transmitters to be closer to each other and causes more interference. In the example of Figure 4, if D uses CS threshold $CS2$ instead of $CS1$, the transmission from A to B will encounter a larger interference due to the concurrent transmission from node D. As we know, the quality of a communication link depends on the interference at the receiver caused by other transmissions; the higher the signal-to-interference-and-noise-ratio (SINR), the higher the rate at which packets can be transmitted reliably. To account for the increase of interference when using a larger CS threshold, the transmission rate often needs to be reduced. In essence, as the above discussion suggests, a larger CS threshold and a lower transmission rate lead to a smaller contending region. Thus,

spatial adaptation of contending region can be achieved using the joint control of CS threshold and transmission rate.

For the ease of argument, we assume for now that all nodes in the network use the same CS threshold and transmission rate. However, our protocol allows each node to choose these parameters independently, as elaborated later. For a given transmission from some node S to another node D, for future reference, we define *concurrent* transmissions as other transmissions that can overlap in time with the S-to-D transmission without causing unreliable reception at D. We also define *simultaneous* transmissions as those transmissions from nodes within the contending region of node S, which start shortly before or after the start of S-to-D transmission and cause erroneous reception at D. Notice that carrier sensing cannot help prevent simultaneous transmissions that start within a short time interval, due to the delay required for carrier sensing (which includes the propagation delay). From the perspective of MAC layer, the aggregate throughput for a given network depends on the medium access efficiency in resolving the local channel contention, the number of concurrent transmissions in the network, and the transmission rate between each transmitter/receiver pair. It is straightforward to see that a smaller contending region allows more concurrent transmissions in the network, at the price of lowering transmission rates. A less straightforward, yet important, observation is that a smaller contending region can also help to improve the medium access efficiency in resolving the local channel contention. We elaborate on this observation below.

- **A smaller contending region can reduce the collision probability when resolving the local contention.**
  Consider the transmissions from S to D in Figure 3. There are two primary sources of interference for receiver D. One type of interference comes from other concurrent transmitters outside the contending region ω (e.g., transmitters $S1$, $S2$, and $S3$), which may transmit even when S is transmitting because the detected signal strength from S's transmission is below their CS threshold. The other type of interference comes from what we usually refer to as *collisions*, when the simultaneous transmission attempts from transmitters inside the contending region occur. Such events can happen, for example, when nodes $S4$, $S5$, $S6$, $S7$ start their transmissions close enough to the start time of node S's transmission.
  By reducing the transmission rate when using a smaller contending region, the interference from concurrent transmissions outside the contending region can be taken into account. However, such adaptation cannot address the interference from simultaneous transmissions inside the contending region effectively. This is because simultaneous transmissions may be from a node that is arbitrarily close to the receiver; for instance, node $S5$ is very close to the receiver node D in our example in Figure 3. Therefore, local channel contention inside the contending region has to be resolved in time domain. When using temporal MAC protocols (e.g., IEEE 802.11 DCF) to resolve the local channel contention, the probability of simultaneous transmissions (i.e., collision probability) often increases rapidly with the number of competing nodes, which, in turn, is a non-decreasing function of the contending region. As such, a smaller contending region can result in a smaller number of locally competing nodes and can lower the collision probability. The tradeoff here is that transmission rate decreases due to the increased interference from outside contending region when the contending region is reduced.
- **A smaller contending region can reduce the "rate-independent" MAC overhead when resolving the local contention.**
  We define *rate-independent overhead* of the MAC protocol as the overhead, by which the channel time consumed is independent of the transmission rate used for data packets [1]. For example, the duration of inter-frame spaces (DIFS, SIFS, EIFS, etc.) in IEEE 802.11 DCF is fixed regardless of the transmission rate used for data; hence, they are rate-independent overheads. Let P_l (in bits) be the packet payload size, T (in seconds) be the channel time consumed by the rate-independent overhead associated with each transmission, and R (in bits per second) be the transmission rate. Adopting a simplified model, it can be readily shown that $\frac{TR}{P_l+TR}$ fraction of channel capacity is wasted in the rate-independent overhead. Therefore, the smaller the rate R, the smaller the channel wastage in rate-independent overhead. As we have discussed before, when using a smaller contending region, the transmission rate often needs to be reduced to account for the increased interference from outside the contending region. As a result of the lowered transmission rate, the channel wastage in rate-independent overhead can be reduced. This observation is also made by Yang et al. in [2], although that paper does not present a protocol utilizing the observation.

To summarize, a larger CS threshold and a lower transmission rate lead to a smaller contending region. By using a larger CS threshold to bring concurrent transmitters closer to each other, the MAC efficiency in resolving the local channel contention can be improved, due to the reduced number of locally competing nodes and the reduced rate-independent overhead. At the same time, since a lower transmission rate can tolerate more interference at the receiver given a received signal strength, more concurrent transmissions can proceed reliably. The price paid to gain above benefits is the reduced transmission rate. Such a tradeoff implies that there exists optimal CS threshold and transmission rate which can maximize the aggregate throughput for a given network. Intuitively, a network with a larger transmitter density will prefer a smaller contending region for the transmissions. Here the transmitter density is defined as the number of transmitter/receiver pairs in the area covered by the maximum transmission range. With the increase of transmitter density, the benefit of reduced local channel contention resulting from a smaller contending

region becomes more significant, and more transmitters are available to exploit the improved spatial reuse. Hence larger CS threshold and lower rate are generally preferable for networks with denser traffic patterns.

To provide a better understanding of the impact of CS threshold and transmission rate on the aggregate throughput, we now present some simulation results for random networks, obtained using a modified ns-2 simulator. We modified the interference model in ns-2 version 2.26 such that the interference from all concurrent/simultaneous transmissions are accumulated to properly evaluate SINR at a receiver. The physical layer characteristics follow the specifications of IEEE 802.11a, with transmission rates at 54, 36, 18, and 9 Mbps. We assume that transmissions at a certain rate are successful if the corresponding SINR threshold is met. The SINR threshold used are listed in Table I [3]. The MAC protocol follows the specifications of IEEE 802.11 DCF, but with the contention window size CW fixed at 31 (i.e., exponential backoff is disabled). RTS and CTS transmissions are disabled so that we can focus on the impact of CS threshold. Two-ray ground radio propagation model, which models the large scale path loss, is used in the simulations.

TABLE I

Rates (Mbps)	SINR (dB)	Rates (Mbps)	SINR (dB)
54	24.56	36	18.80
18	10.79	9	7.78

In addition to adequate SINR, the received signal needs to be above a certain *received signal threshold* RX_{th} for reliable reception to occur. The value of RX_{th} thus limits the maximum transmission range. In all our simulations, we use identical RX_{th} and transmit power levels at all nodes such that the receive power level at 35 meter distance is equal to RX_{th} under the two-ray ground radio propagation model. Four randomly generated networks with increasing transmitter density are simulated. More specifically, in a 300 m $\times$ 300 m area, 8, 11, 16 and 40 transmitter/receiver pairs are randomly placed, as shown in Figures 5(a), (b), (c), and (d), respectively. As we are interested in the maximum achievable aggregate throughput, all flows are constantly backlogged. The payload packet size is 512 bytes. All results presented in Section II are averaged over 20 simulation runs, and the 99% confidence interval for the presented results are less than 1% of the mean values.

In the simulations of this section, all transmitters use the same static CS threshold and transmission rate. We refer to this as "static scheme". Various combinations of CS threshold and transmission rate are evaluated. The static scheme helps determine the optimal performance when using a static and identical combination of rate and CS threshold at every node. The aggregate throughput over all flows is presented in Figure 6. In each plot of Figure 6, four different curves correspond to four different transmission rates used; the horizontal axis represents $\beta = \frac{CS_{th}}{RX_{th}}$ in dB (i.e., $10 \log \beta$) and is proportional to

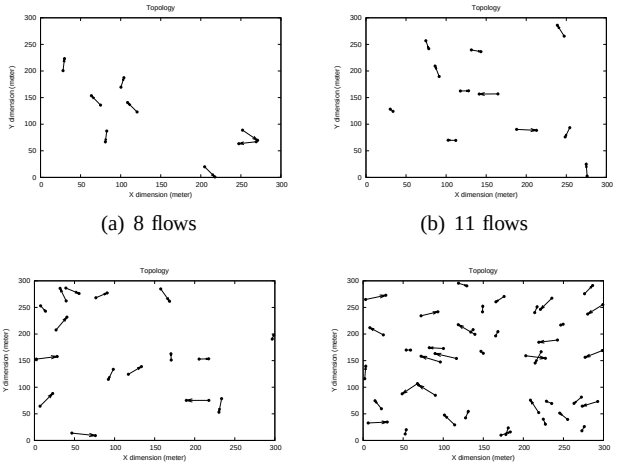

(a) 8 flows (b) 11 flows

(c) 16 flows (d) 40 flows

Fig. 5. Random topologies with increasing transmitter density.

the value of CS threshold (CS_{th}); the vertical axis represents the aggregate throughput of all the flows in the unit of Mbps. We refer to $\beta = \frac{CS_{th}}{RX_{th}}$ as the "normalized" CS threshold.

As we can see in Figure 6(a), when there are only eight flows, it is best to transmit at the highest available rate of 54 Mbps and use $\beta = -22$ dB. However, when the number of flows increases to 11 in Figure 6(b), the maximum aggregate throughput is obtained when the rate is set to 36 Mbps and $\beta = -14$ dB. Observe that in Figure 6(b), the peak throughput obtained using rate 54 Mbps is 28% lower than the maximum throughput. Thus, the rate that was optimal in the previous case is significantly sub-optimal for 11 flows. Further increasing the number of flows to 16, as shown in Figure 6(c), the maximum throughput is now achieved when the transmission rate is set to 18 Mbps and $\beta = -10$ dB; the peak throughput obtained when using both rate 36 Mbps and rate 54 Mbps are now much worse than the maximum throughput. When increasing the number of flows to 40 in Figure 6(d), the maximum throughput point remains at rate 18 Mbps and $\beta = -10$ dB; this is because the SINR thresholds of rates 18 Mbps and 9 Mbps are relatively close (see table I). In the topologies we have simulated, the slightly improved spatial reuse when using 9 Mbps is not sufficient to compensate the halved transmission rate (from 18 Mbps to 9 Mbps).

The optimal CS threshold and transmission rate depend on the transmitter density in the network, and more generally, on the traffic patterns in the network. It is important to use the appropriate values for both CS threshold and transmission rate, otherwise the aggregate throughput may suffer a significant loss. Additionally, note that in above simulations, we have forced *all* transmitters to use the same CS threshold and transmission rate. In general, different source nodes may view the network conditions differently, depending on their neighborhood. A good choice of CS threshold for one source node may not be good for others. Ideally, we would like each node to make its own decision on the values of CS

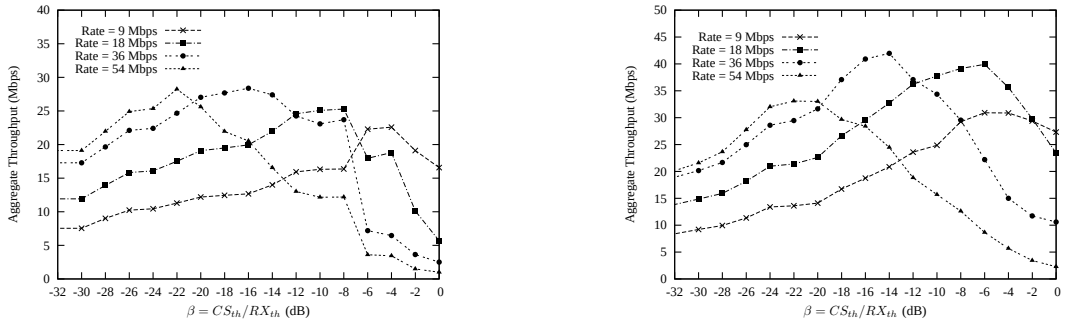

(a) Throughput with 8 flows

(b) Throughput with 11 flows

(c) Throughput with 16 flows

(d) Throughput with 40 flows

Fig. 6. Aggregate throughput in random topologies (payload size: 512 bytes).

threshold and transmission rate to be used. Allowing each node to independently choose its parameters also alleviates the overhead occurred in global coordination. Because of the above reasons, distributed algorithms that allow each node to dynamically search for the appropriate CS threshold and transmission rate are desired. One such algorithm, which we name as *dynamic spatial backoff* algorithm, is presented in Section III.

B. Adapting transmission power and rate

The joint adaptation of CS threshold and transmission rate can help control the size of the contending region, and thus, realize spatial backoff as discussed above. There exist other ways as well. For example, it is possible to use the joint adaptation of transmission power and rate, as the example in Figure 7 shows. Assume node A is transmitting to B using power $P2$. With a fixed CS threshold CS, node D has to defer its transmissions whenever A is transmitting, since the signal strength from A's transmissions at D is higher than CS. As a result, nodes A and D have to compete for the channel access in time. On the other hand, if A uses a lower transmission power $P1$, node D can transmit to C at the same time when A is transmitting, although they may both have to transmit at lower rates. Therefore, a lower transmission power and a lower transmission rate lead to a smaller contending region; and spatial backoff can also be realized using the joint adaptation of transmission power and rate. In this paper, we focus on investigating spatial backoff algorithms that control the CS threshold and transmission rate, assuming that transmission power is fixed. Note that reducing transmission power at MAC layer may affect the connectivity, and the resulting interaction between MAC and routing layers adds complexity to the spatial backoff problem. We delay the joint control of transmission power, rate and CS threshold to our future work.

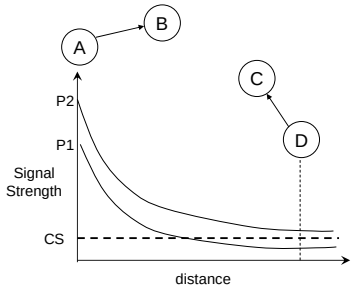

Fig. 7. Lower power and lower rate lead to smaller contending region.

III. Dynamic Spatial Backoff Algorithm

The goal of the spatial backoff algorithm is to allow each node to search the two-dimensional space defined by the *CS threshold* and the *transmission rate* to determine the suitable values of these parameters. Note that when multiple flows originate from a single source node, this node may use a different combination of rate and CS threshold for each of its receivers. In this paper, we often associate these parameters with the source node for simplifying the description. However, it should be noted that these parameters are per flow basis. Ideally, we

would like to find the optimal operating point for each node such that the network aggregate throughput can be maximized. Note that the optimal operating point for different nodes can be different, depending on their neighborhood. The difficulty in the optimization problem is that aggregate throughput is a global metric. Aggregate throughput not only depends on the local contention resolution efficiency experienced by each individual node, but also depends on the transmission rates used by various transmitter/receiver pairs and the total number of concurrent transmissions in the network. In other words, for a transmitter to reach the optimal operating point that maximizes the aggregate throughput, it has to gather information from other transmitters in the network. Due to the dynamic nature of wireless networks and the substantial cost associated with obtaining global information, we are more interested in devising a distributed algorithm that enables each node to make decisions based on its local information. On the other hand, by using only local information, it is not always possible to find the optimal operating point, as we demonstrate using the example below.

Consider the first scenario in Figure 8(a). There are two flows, one from node 1 to 1R and the other from node 2 to 2R. The topology is symmetric for nodes 1 and 2. Assume that there are two "non-trivially different" CS threshold values, cs_1 and cs_2, that nodes 1 and 2 can use. By "non-trivially different," we mean that the interference will change whenever CS threshold changes. Without loss of generality, assume $cs_1 < cs_2$. By using cs_1, suppose that node 1 will defer when node 2 is transmitting, and vice-versa. As a result, nodes 1 and 2 can transmit alternately at the same rate, say, R_1, if they both use cs_1. On the other hand, when using cs_2, suppose that node 1 will transmit even though node 2 is transmitting. As the network is symmetric, nodes 1 and 2 can both transmit concurrently at some rate R_2 if they both use cs_2. Assuming $R_2 < \frac{R_1}{2}$, in this scenario, the optimal CS threshold for node 1 is cs_1 in order to maximize the aggregate throughput.

Now we consider the second scenario in Figure 8(b), which only differs from Figure 8(a) in node 2R's location. In this scenario, node 2R is moved away from the interfering node 1. As a result, node 2 can transmit to node 2R at a higher rate, say R_3 ($R_3 > R_2$), when nodes 1 and 2 transmit concurrently. At the same time, the transmission rate of node 1 remains at R_2 when nodes 1 and 2 transmit concurrently. Node 2R can be placed close enough to node 2 such that the condition $R_3 + R_2 > R_1$ is satisfied, where R_1 is the aggregate throughput when nodes 1 and 2 transmit alternately using cs_1. As such, in Figure 8(b), the optimal CS threshold for node 1 to maximize the aggregate throughput is cs_2.

The two scenarios in Figures 8(a) and 8(b) cannot be differentiated from the local viewpoint of node 1. In both cases, node 1 gets to transmit at rate R_1 when using cs_1, and gets to transmit at rate R_2 when using cs_2. Limited by the local information, it is not possible for node 1 to find the optimal CS threshold for both scenarios.

Using the above example, we argue that, limited by the local information only, it is not always possible for an individual source node to find the optimal operating point to maximize the aggregate throughput. As such, the goal of our work is to design a simple mechanism that can be easily incorporated into existing MAC protocols to take advantage of spatial backoff and to improve aggregate throughput. We do not make claims regarding the optimality of the proposed dynamic spatial backoff algorithm in this paper. As the above example shows, no local optimal algorithm based only on local information exists. However, as our results in the next section suggest, our protocol is able to achieve good performance in typical network topologies.

A. Protocol Description

Our dynamic spatial backoff algorithm allows each source node to search for an operating point with appropriate values of CS threshold and transmission rate for itself, in order to improve the throughput. Given the possible range of CS threshold[3] and the multiple levels of transmission rate supported by a wireless transceiver, a naive exhaustive search in the two dimensional space can lead to poor performance. This is because of the potentially large number of operating points, a significant fraction of them not being desired ones. Therefore, the design of our dynamic spatial backoff algorithm includes reducing the search space, and then incorporating suitable search rules.

1) Reducing the search space: Suppose that each node may use one of available K rates. We represent them using an array $Rate[]$, where $Rate[j] > Rate[i]$ if $j > i$ ($i, j \in [1, ..., K]$). In order to reduce the cost of searching over the two-dimensional space of rate and CS threshold values, for each available transmission rate, we identify the smallest CS threshold that may be used in conjunction with that rate. In particular, let $CS[i]$ be the smallest CS threshold that may be used in conjunction with $Rate[i]$. How should we determine the suitable value for $CS[i]$? While many different approaches may be devised for this, we use the approach described below.

When some node S transmits to another node D using a *fixed* transmission rate, two reasons (other than the collisions due to simultaneous transmissions from nodes within the contending region) can cause an erroneous reception at node D:

(i) The first reason is that CS threshold used by node S may be too large. The interference at D is proportional to the detected signal strength at node S (this is because the signal detected at node S can also propagate to node D but possibly with a different channel gain). Thus, increasing the CS threshold used by node S will allow node S to start transmitting even if node D is experiencing a higher level of interference when the transmission from S begins, and vice-versa. Thus, transmission failure occurring with larger CS threshold used by node S may imply that S has over-estimated the interference tolerance level of node D.

(ii) The second reason for error may be due to the large interference from other transmitters, which begin transmitting

[3]The minimum CS threshold is constrained by the radio sensitivity.

(a) The optimal CS threshold for node 1 is cs_1

(b) The optimal CS threshold for node 1 is cs_2

Fig. 8. The optimal CS threshold for node 1 is different for two different network settings.

after node S has started its transmission. Carrier sensing at node S cannot help to detect this type of interference, since carrier sensing is performed before a node begins its transmission.

We consider an interference limited environment. As defined previously, RX_{th} is the lower bound on the received power level for the receiver to be able to decode the signal. Thus, if $SINR[i]$ is the SINR threshold for rate $Rate[i]$, interference less than or equal to $\frac{RX_{th}}{SINR[i]}$ shall not affect the correct reception at rate $Rate[i]$, if the received power level is more than RX_{th}. Secondly, observe that when node S uses a CS threshold, say, P_{CS}, it will transmit when the detected signal strength is not higher than P_{CS}. Although the interference at receiver D will not be identical to that at S, we use P_{CS} as an approximate estimate on the interference posed to node D when the transmission from S starts. We define $CS[i]$ as follows.

$$CS[i] = \frac{RX_{th}}{SINR[i]} \tag{1}$$

Above discussions suggest that, when using a CS threshold less than or equal to $CS[i]$ to transmit at rate $Rate[i]$, it is not likely that node S will over-estimate the interference tolerance level at node D. In other words, if a transmission at rate $Rate[i]$ using a CS threshold larger than $CS[i]$ fails, then the cause may be that the used CS threshold is too large. On the other hand, if a transmission at rate $Rate[i]$ using $CS[i]$ fails, there is not much benefit in reducing the CS threshold further. Since for a higher rate, the SINR threshold is higher, it follows that $CS[j] > CS[i]$ if $j < i$ ($i, j \in [1, ..., K]$). As shown in Figure 10, our dynamic spatial backoff algorithm searches the subspace above or on the diagonal line for the best combination of $Rate[]$ and $CS[]$ for each node (we will elaborate on the search process). That is, our algorithm limits each node to use one of the available K rates, and for each such rate $Rate[i]$, the choice of CS threshold is limited to $CS[j],\ j \leq i$. It should be noted that the search rules described below is independent of the manner in which the values of $CS[i]$ above are determined.

2) Search rules: With K transmission rates available, we define "rate level" as a number from 1 to K such that rate level 1 is the lowest transmission rate $Rate[1]$, and rate level K is the highest rate $Rate[K]$. In addition to rate array $Rate[]$ and CS threshold array $CS[]$, each node also maintains an index variable R_{index} and an index array $CS_{index}[]$. R_{index} represents the array index of transmission rate (i.e., transmission rate equals to $Rate[R_{index}]$), where $R_{index} \in [1, ..., K]$. Given the R_{index}, there is an associated CS threshold array index $CS_{index}[R_{index}]$ that is dynamically adjusted. At any given time, a node will transmit at rate $Rate[R_{index}]$ using the associated CS threshold $CS[CS_{index}[R_{index}]]$. The data structure is illustrated in Figure 9.

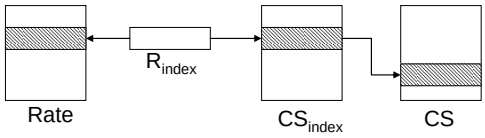

Fig. 9. Data structure of dynamic spatial backoff.

Initially, at each node, $R_{index} = 1$ and $CS_{index}[i] = i$ ($1 \leq i \leq K$). Thus, each node starts from using the transmission rate $Rate[1]$ and the CS threshold $CS[CS_{index}[1]] = CS[1]$. A node will take different actions, depending on whether its past transmissions have been *successes* or *failures*. We delay our discussions on the conditions, under which a node considers its transmissions as *successes* or *failures*, to the latter part of this section. The actions of each node on *successes* or *failures* are governed by the 4 rules below. For each rule, we also provide our motivation behind the rule. Note that, in the following text, "diagonal line" refers to the diagonal in Figure 10.

- **Rule 1: When transmissions are successful**, the node will increase its transmission rate by one level unless the highest transmission rate is already in use, i.e., $R_{index} = min(R_{index} + 1, K)$. If this results in a rate increase, the CS threshold associated with the old rate will be associated with the new rate, i.e., using the new R_{index} value, we perform $CS_{index}[R_{index}] := CS_{index}[R_{index} - 1]$.
  Motivation for Rule 1: Successful transmissions indicate that interference is small enough that the SINR threshold of current rate is met. Following two actions may possibly be taken by the node, each of which attempts to increase throughput by exploiting any remaining interference margin at the receiver. One possibility is to transmit at a higher rate, and the other is to transmit more aggressively by increasing the CS threshold. However, if we allow a node with successful transmissions to increase its CS threshold, the successful node will become more and more aggressive and consume more and more channel resource, potentially starving other nodes. In view of this, in our dynamic spatial backoff algorithm, a node increases its transmission rate when transmissions have been successful, as elaborated earlier in this paragraph.

- **Rule 2: When transmissions fail and the operating point is above the diagonal line**, the CS threshold associated with current transmission rate will be decreased by updating the CS_{index}. In particular, we perform $CS_{index}[R_{index}] := min(CS_{index}[R_{index}] + 1, R_{index})$. Note that, for the operating point to be on or above the diagonal in Figure 10, we must have that $CS_{index}[R_{index}] \leq R_{index}$.
  Motivation for Rule 2: Recall that $CS[j] > CS[i]$ if $j < i$ $(i, j \in [1, ..., K])$. As we discussed in Section III-A.1, when transmitting at rate $Rate[R_{index}]$ using CS threshold larger than $CS[R_{index}]$, the cause of transmission failures may be that the transmitter has chosen CS threshold too large, over-estimating the interference tolerance level of the receiver. By reducing the CS threshold, the node will transmit more conservatively, which may help it improve its transmission success probability at the current rate.
- **Rule 3: When transmissions fail and the operating point is on the diagonal line**, the node will decrease its transmission rate by one level unless the lowest rate is already in use, i.e., $R_{index} := max(R_{index} - 1,\ 1)$. At the same time, the CS threshold associated with the reduced rate will be applied (CS threshold will increase); that is, the CS threshold will be $CS[CS_{index}[R_{index}]]$ using the new value of R_{index}.
  Motivation for Rule 3: Suppose that a node's transmission fail at rate $Rate[R_{index}]$ and CS threshold $CS[R_{index}]$, which is one of the diagonal points. In this case, the CS threshold used is already the lowest threshold deemed reasonable for the chosen rate. As the discussion in Section III-A.1 suggests, in this case, the cause of transmission failures is likely to be that the other nodes have been transmitting more aggressively, and excessive interference occurs due to transmissions that begin after the node has started its own transmission. Therefore, continuing to decrease its CS threshold will not help the node to improve its transmission success probability. On the other hand, by reducing the transmission rate and using the larger CS threshold associated with the reduced rate, the node can improve its success probability and gain more chances to access the channel. Recall that a larger CS threshold and a lower rate lead to a smaller contending region, which can result in the improved local contention resolution efficiency and better spatial reuse.
- **Rule 4: Avoid starvation when probing small CS threshold.** Notice that, following the above rules, a node has to suffer transmission failures before its CS threshold can be increased. It could happen that, when using a small CS threshold, a node becomes so conservative in transmitting that it no longer has sufficient chances to access the channel. The transmission successes or failures of a node cannot be observed if the node does not transmit at all; as a result, the unsuitable small CS threshold can be retained and the node can lose its chances to access the channel. To improve on such an undesirable situation, in our proposed algorithm, each backlogged node keeps track of whether or not it has made any transmission attempt during a certain time period $T_{timeout}$. If no such transmission attempts are made, and $R_{index} > 1$, then the transmission rate of the node will be reduced by one level (i.e., $R_{index} = R_{index} - 1$), and CS threshold associated with the reduced rate, namely, $CS[CS_{index}[R_{index}]]$, will be used.

Figure 10 illustrates how the proposed dynamic algorithm probes the two-dimensional space. We elaborate on the search process using an example below. In the example, when we mention that a node operates at the point (x,y), we mean that the node uses CS threshold x and transmits at rate y. Assume that transmissions of a node have been successful starting from the point ($CS[1]$, $Rate[1]$). Following *rule 1* mentioned above, the node increases its rate and moves to the point ($CS[1]$, $Rate[2]$). Assume that the transmissions continue to be successful and the node keeps on increasing its transmission rate until it reaches the point ($CS[1]$, $Rate[K-1]$). The node then encounters transmission failures at ($CS[1]$, $Rate[K-1]$) because $CS[1]$ is too large for rate $Rate[K-1]$. Following *rule 2*, the node decreases its CS threshold and moves to the point ($CS[2]$, $Rate[K-1]$). As the node transmits more conservatively now, let us pretend that its transmissions become successful again. The node again increases its transmission rate per *rule 1* and moves to ($CS[2]$, $Rate[K]$). At this point, the node starts to suffer transmission failures, which causes it to move to ($CS[3]$, $Rate[K]$) according to *rule 2*. Suppose that the node continues to fail and eventually it reaches the point ($CS[K]$, $Rate[K]$). Transmission failures at ($CS[K]$, $Rate[K]$) cause the node to reduces its rate to $Rate[K-1]$, following *rule 3* mentioned above. While reducing the rate to $Rate[K-1]$, since $CS[2]$ is associated with $Rate[K-1]$ from prior search process, the node returns back to the point ($CS[2]$, $Rate[K-1]$) directly.

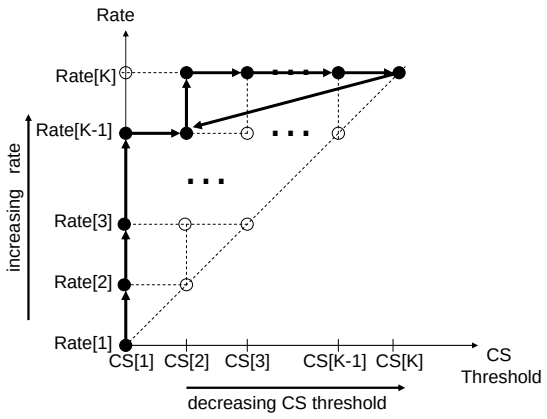

Fig. 10. Search space of the dynamic spatial backoff algorithm.

In our dynamic spatial backoff algorithm, nodes adjust their CS threshold and transmission rate based on whether their past transmissions have been successes or failures. In general, the conditions under which a node considers its transmissions

as successes or failures can be defined in many ways. The condition currently used in our algorithm is as follows. We have two additional arrays $S[i]$ and $F[i]$ ($i \in [1, ..., K]$); and their elements are initialized to $S_{initial}$ and $F_{initial}$, respectively. A node will consider its transmissions at rate R_{index} as successful if it has had $S[R_{index}]$ consecutive successful transmissions. On the other hand, a node will consider its transmissions fail if it suffered $F[R_{index}]$ consecutive transmission failures.[4] To avoid the performance loss due to frequent unsuccessful probing, the values of $S[]$ and $F[]$ can be adjusted dynamically. Intuitively, if a node can transmit at the rate level R_{index} successfully with high probability, but it fails often at the rate level $R_{index}+1$, we would like the node to probe the rate level $R_{index}+1$ less frequently. Therefore, if the total number of successful transmissions for a source node at rate level $R_{index}+1$ is less than a certain threshold S_{th}, $S[R_{index}]$ will be increased by 1 when the node returns from the rate level $R_{index}+1$ to R_{index}; otherwise, $S[R_{index}]$ will be reset to the default initial value $S_{initial}$. From another perspective, if a node can transmit successfully at the rate level R_{index} with high probability, we would like the node to fall back to the lower transmission rates less frequently. Hence, if the total number of successful transmissions at rate level R_{index} is more than a threshold F_{th}, $F[R_{index}]$ will be increased by 1 every time when the node falls back to a lower transmission rate; otherwise, $F[R_{index}]$ will be reset to the default initial value $F_{initial}$.

Intuitively, using the proposed algorithm, a source node is likely to oscillate around a point where, given the interference present in the network, it can transmit with a high success probability using the highest possible transmission rate and correspondingly the largest suitable CS threshold. We use the example in Figure 11 to illustrate this. Assume that there are 4 available transmission rates (i.e., $K = 4$) and we consider two nodes, node 1 and node 2, in a certain network. Suppose the point ($CS[2]$, $Rate[3]$), which we called point A, is the best operating point for both nodes 1 and 2, and they have stayed there for a relatively long duration. Then, node 1 begins to probe point B after $S[3]$ consecutive successful transmissions. Since point B has higher rate and higher SINR threshold, node 1 may suffer transmission failures, and thus, move to point C (trajectory 1 of Figure 11(a)). When using $CS[3]$, node 1 transmits more conservatively, which in turn, can reduce the interference to node 2's transmissions. As a result, node 2 may also move to point B after it has had $S[3]$ consecutive successful transmissions (trajectory 1 of Figure 11(b)). If node 1 continues to suffer transmission failures at rate $Rate[4]$, or it is starved as a result of using a smaller CS threshold than node 2, node 1 will end up with going back to point A (trajectory 2 of Figure 11(a)). Once node 1 goes back to point A and uses $CS[2]$, node 2 will encounter larger interference. As we previously assume point A is the best operating point for both nodes 1 and 2, node 2 will likely suffer transmission failures at rate $Rate[4]$, and will thus move back to point A, following the trajectory 2 of Figure 11(b). As the above example illustrates, a node may not always stabilize at one particular operating point, but we anticipate that nodes will tend to oscillate close to their appropriate operating points.

[4]Alternatively, we can also define the success or failure conditions in terms of the percentage of successful transmissions over a certain time period.

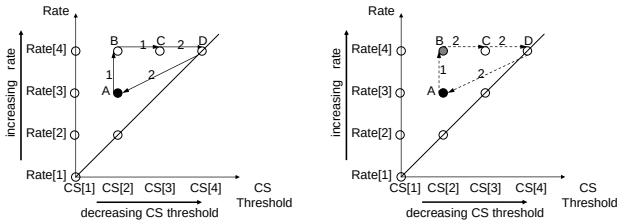

(a) Trajectory of node 1 (b) Trajectory of node 2

Fig. 11. Interactions between two source nodes

IV. Performance evaluation

We simulated our dynamic spatial backoff algorithm using a modified ns-2 simulator. Simulation settings are similar to static simulations presented in Section II. As a reminder, the physical layer follows the specifications of IEEE 802.11a and four transmission rates (i.e., 54, 36, 18, 9 Mbps) are used. To distinguish the effects of spatial backoff from temporal backoff, exponential backoff of IEEE 802.11 DCF is disabled and a constant contention window size (i.e., 31) is applied. All flows are constantly backlogged. The rate array $Rate[]$, and the CS threshold array $CS[]$ derived from equation 1, are listed in Table II. More specifically, given a rate $Rate[i]$, we have $10 \log \frac{CS[i]}{RX_{th}} = -10 \log SINR[i]$, where $SINR[i]$ values are based on Table I. Other parameters used by our dynamic spatial backoff algorithm are as follows: $S_{initial} = 10$, $S_{th} = 20$, $F_{initial} = 3$, and $F_{th} = 100$, $T_{timeout} = 0.1$ second.

TABLE II

$Index$	$Rate[Index]$	$SINR[Index]$	$\frac{CS[Index]}{RX_{th}}$
1	9 Mbps	7.78 dB	−7.78 dB
2	18 Mbps	10.79 dB	−10.79 dB
3	36 Mbps	18.80 dB	−18.80 dB
4	54 Mbps	24.56 dB	−24.56 dB

To show how our spatial backoff algorithm performs differently from the conventional rate control algorithm, we also implemented the Automatic Rate Fallback (ARF) algorithm [4] with a slight modification. The ARF algorithm is implemented as follows. When transmitting at rate level K, a source node will increase its rate upon $S[K]$ consecutive successful transmissions, and will decrease its rate upon $F[K]$ consecutive transmission failures. We allow the values of $S[K]$ and $F[K]$ to be adjusted dynamically, in order to avoid the performance loss due to frequent unsuccessful probing. The $S[]$ and $F[]$ update algorithm we used for ARF is the same as that for the dynamic spatial backoff algorithm, as described in Section III-A. Simply put, if rate level K turns out to be the

most suitable rate level for a node, then $S[K]$ and $F[K]$ of this node will be increased over time. As a result, this node will probe other rate levels less and less frequently.

The four plots in Figure 12 correspond to the four random topologies in Figure 5, respectively. In each plot, the horizontal axis represents the normalized CS threshold $\beta = \frac{CS_{th}}{RX_{th}}$ in dB; the vertical axis represents the aggregate throughput in the unit of Mbps. Results of our dynamic spatial backoff, ARF rate control, and static simulations from Section II are presented. Recall that in static simulations, all source nodes use the same transmission rate and CS threshold. The throughput of ARF is measured when all source nodes use the same CS threshold, but vary their rates independently; various CS threshold values are evaluated for ARF.

We observe the results for the topology with 8 flows in Figure 12(a). Each static simulation curve represents the aggregate throughput when using a specified transmission rate and various CS threshold values. Also shown are the curves for the ARF scheme and our dynamic spatial backoff. The curves for dynamic spatial backoff are flat, since the scheme does not depend on the value of $\beta = CS_{th}/RX_{th}$ on the horizontal axis of the graphs. Note that in each plot in Figure 12, the 99% confidence interval of the dynamic spatial backoff algorithm is shown by the two dashed lines that are above and below the average. As we can see, the performance of ARF approximately follows the envelop of four static simulation curves. In other words, if we know what is the suitable CS threshold for a given network and we use this CS threshold, ARF can perform well by controlling the rate alone. On the other hand, if an inappropriate CS threshold is used, the performance of ARF can be poor. In general, no single CS threshold performs close to optimal for all networks; hence ARF cannot perform well in all networks. Using our dynamic spatial backoff algorithm, on the other hand, each source node dynamically searches for the appropriate combination of transmission rate and CS threshold, which helps to achieve good aggregate throughput for any given network. As shown in Figure 12(a), the maximum aggregate throughput from the static simulations is 28.4 Mbps, while the average aggregate throughput using our dynamic spatial backoff is 28.9 Mbps.

The results for the topology with 11 flows are shown in Figure 12(b). In this topology, the static maximum aggregate throughput is 41.9 Mbps, while the average aggregate throughput of our dynamic spatial backoff is 36.2 Mbps. Although the throughput of dynamic spatial backoff algorithm is 13% lower than the static optimal point, dynamic spatial backoff achieves this performance without a priori knowing the optimal CS threshold and transmission rate. For the topology with 16 flows, as shown in Figure 12(c), the static maximum aggregate throughput is 44.2 Mbps; and the average aggregate throughput of dynamic spatial backoff is 40.6 Mbps. For the topology with 40 flows, the static maximum aggregate throughput is 54.5 Mbps, while the average aggregate throughput of dynamic spatial backoff is 54.8 Mbps, as shown in Figure 12(d). More topologies have been simulated, and the results are not presented here for conciseness. Our results indicate that the proposed dynamic spatial backoff algorithm is able to achieve performance close to the static optimal point. Moreover, unlike the static scheme, the proposed dynamic spatial backoff algorithm achieves this performance without having to a priori determine the optimal CS threshold and transmission rate.

We plot the traces of CS threshold and transmission rate in Figure 13, for one source node in the topology with 40 flows. The horizontal axis represents the simulation time (in seconds); the vertical axis represents the normalized CS threshold β, and transmission rate, in Figures 13(a) and 13(b), respectively. Each dot in the plots indicates a transmission attempt (the transmission may or may not be successful). As we can see, normalized CS threshold of this node oscillates between -7.78 dB and -10.79 dB; transmission rate of this node stays at 18 Mbps most of the time. Notice from Figure 12(d) that, in static simulations, both points (-7.78 dB, 18 Mbps) and (-10.79 dB, 18 Mbps) have throughput close to the static optimal. Using our dynamic spatial backoff, most nodes use -7.78 dB or -10.79 dB as their CS threshold. Some nodes are able to transmit at rate 36 Mbps successfully, while others use 18 Mbps or 9 Mbps, depending on the channel condition around each individual node. Our dynamic spatial backoff algorithm allows each node to choose the appropriate combination of rate and CS threshold based on its own channel condition, while static simulations enforce the same rate and CS threshold for all nodes. This explains why our dynamic spatial backoff algorithm achieves slightly higher throughput than the static optimal point in Figures 12(a) and 12(d).

Table III lists the throughput of each individual flow from the random topology with 40 flows, for both our dynamic spatial backoff algorithm and the optimal point of static simulations. As we can see, dynamic spatial backoff does not starve any individual flow while improving the aggregate throughput. However, in some cases, we do observe that some flow, like flow 14 in Table III, has relatively fewer chances than others to access the channel successfully. To explain the reason behind this, we plot the traces of CS threshold and transmission rate used by the source node of flow 14 in Figure 14. We observe that this node has been trying very hard to access the channel by using the largest CS threshold and by transmitting at the lowest rate. This node cannot get a chance to access the channel when a smaller CS threshold is used; and it cannot transmit successfully at a higher rate because of the large interference from neighboring nodes. Figure 14 provides an evidence, showing that our dynamic spatial backoff algorithm has helped a node to make the best efforts it can locally do to access the channel and transmit successfully. The reason that flow 14 does not get successful transmissions very often is due to the asymmetric nature of the topology. We use the simple asymmetric topology in Figure 8(b) to explain this behavior. There, since transmissions from node 1 cause little interference at node 2R, node 2 can always transmit to node 2R successfully at high rate using large CS threshold. However, aggressive transmissions from node 2 cause significant interference to the flow from node 1 to node 1R. There is nothing node 1 can locally do to

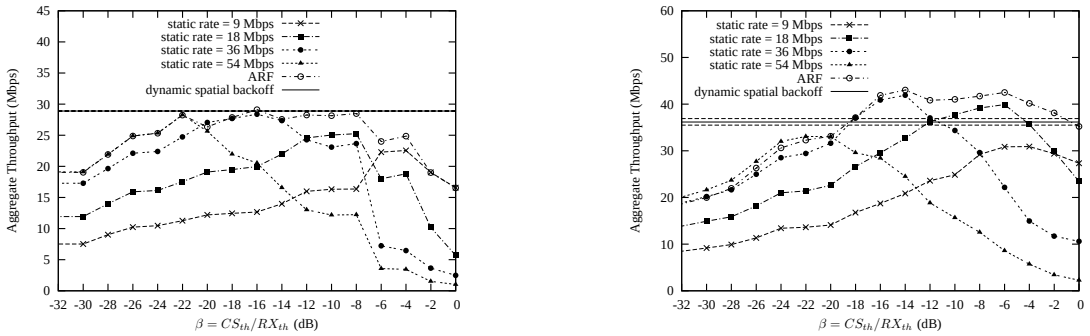

(a) Throughput with 8 flows

(b) Throughput with 11 flows

(c) Throughput with 16 flows

(d) Throughput with 40 flows

Fig. 12. Aggregate throughput of random topologies (payload size: 512 bytes).

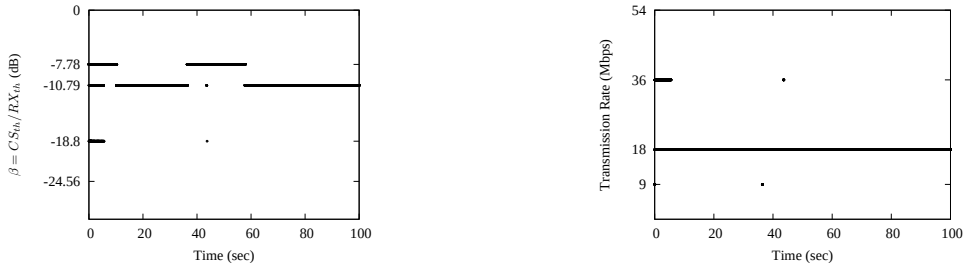

(a) Trace of CS threshold

(b) Trace of transmission rate

Fig. 13. Traces of one typical node in the random topology with 40 flows.

help itself, except by transmitting at a low rate and using a large CS threshold. Flow 14 in Figure 14 encounters a similar situation. The above problem caused by asymmetric topologies can be alleviated by introducing a certain level of coordination among interfering nodes. However, the appropriate level of coordination needs to be carefully investigated to minimize the throughput degradation, which is part of our future work.

V. Related Work

Studies on medium access control to address the channel contention have been conducted extensively in the time domain for the past decades. Temporal contention resolution typically takes the set of competing nodes as a given, and addresses the issue on how to separate transmissions from competing nodes in time to achieve successful transmissions. Numerous methods have been proposed in the past to achieve the temporal separation of transmissions while reducing the overhead introducing by medium access control. Examples of such proposals include [5]–[10]. In this paper, we propose spatial backoff as an alternative approach for wireless networks to address the channel contention. As a result of using spatial backoff, transmissions from competing nodes can also be separated in space to achieve successful transmissions.

Physical carrier sensing provides an effective way to control the interference and the amount of spatial reuse in the network. Guo et al. [11] have noticed the impact of CS threshold on the aggregate throughput. Assuming that the transmission rate is fixed for a given network, Zhu et al. [12], [13] proposed an algorithm that dynamically adjusts the CS threshold to improve spatial reuse and aggregate throughput. Given SINR threshold of the fixed transmission rate, the algorithm in [12], [13]

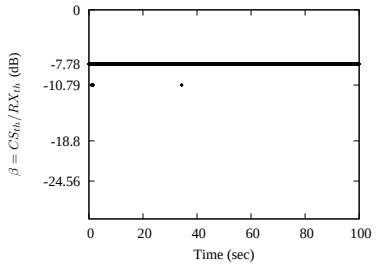

(a) Trace of CS threshold

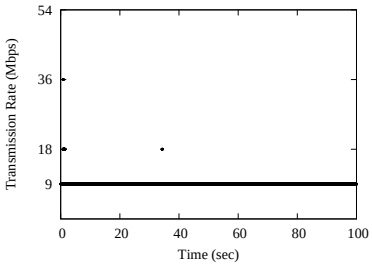

(b) Trace of transmission rate

Fig. 14. Traces of flow 14 in the random topology with 40 flows.

TABLE III
THROUGHPUT OF EACH INDIVIDUAL FLOW IN THE RANDOM TOPOLOGY WITH 40 FLOWS (IN KBPS).

Dynamic spatial backoff				*Static optimal point*			
flow	thr	flow	thr	flow	thr	flow	thr
0	402.5	20	642.4	0	474.6	20	535.3
1	1804.9	21	823.4	1	1957.8	21	626.5
2	1477.1	22	2004.1	2	935.2	22	2096.1
3	1004.5	23	925.2	3	956.7	23	973.0
4	909.7	24	1282.1	4	781.8	24	1302.7
5	2002.1	25	580.8	5	1557.3	25	354.9
6	931.1	26	1490.3	6	949.6	26	1441.7
7	1686.3	27	454.7	7	1426.6	27	1164.7
8	390.5	28	2178.9	8	646.0	28	1971.6
9	1143.3	29	3197.5	9	1353.7	29	2620.7
10	802.3	30	1111.2	10	1232.7	30	607.9
11	1003.0	31	805.5	11	2453.0	31	655.1
12	504.4	32	1709.5	12	715.0	32	1414.2
13	2848.7	33	142.1	13	1618.7	33	329.0
14	74.3	34	2042.9	14	504.7	34	2752.9
15	1338.7	35	729.7	15	1935.4	35	670.1
16	4110.5	36	561.1	16	3762.3	36	743.9
17	4984.3	37	999.6	17	3864.8	37	1144.4
18	716.9	38	1050.8	18	472.7	38	994.7
19	2183.5	39	1737.6	19	1952.2	39	2550.4
Sum	54788.0			Sum	54500.6		

essentially searches the largest CS threshold that can satisfy the SINR threshold for all transmitter/receiver pairs. However, as we discussed in Section II, by reducing the transmission rate and using a even larger CS threshold, we can further improve the spatial reuse and the local channel contention resolution efficiency. The algorithm in [12], [13] is unable to exploit such benefits by adjusting CS threshold alone. Other works that aim to improve the spatial reuse by adjusting CS threshold include [14] by Vasan et al. and [15] by Nadeem et al. The common limitation of the above works is that a fixed transmission rate is assumed. Yang et al. [2] have analytically shown, for uniformly distributed dense networks that, to maximize the aggregate throughput, the optimal CS threshold increases and the optimal transmission rate decreases with the increase of transmitter density. However, [2] does not provide a protocol that utilizes the above observation.

There also exist some rate control algorithms [4], [16], [17], which aim to adapt the transmission rate based on channel conditions. The major difference between rate control algorithms and our spatial backoff algorithm is as follows. Given the CS threshold used by each node, suppose there are a total of m transmission attempts that overlap in time. Assume that, using a high transmission rate, only m_1 ($m_1 < m$) transmissions can succeed. Also suppose that, using a low transmission rate, all m transmissions can succeed. Moreover, due to the increased interference tolerance level at the low rate, potentially m_2 ($m_2 > m$) transmissions can succeed. Rate control algorithms can only increase the number of successful transmissions from m_1 to m, while spatial backoff algorithm can increase the number from m_1 to m_2 via the joint control of rate and CS threshold.

As discussed in Section II-B, we can also realize the spatial backoff by controlling transmission power and transmission rate. Prior work has proposed power control protocols to improve spatial reuse, by introducing new transmissions without interrupting existing transmissions [18], [19]. Fuemmeler et al. [20] have explored the joint control of transmission power and CS threshold to reduce collisions. Existing work on topology control has addressed the issues on finding the appropriate transmission power each node should use; the objective is to maintain network connectivity while reducing energy consumption and improving network capacity [21]–[26]. However, the problem of realizing spatial backoff using the joint control of power and rate remains open, and will be explored in our future work.

VI. CONCLUSIONS AND FUTURE WORK

The study of medium access control has mostly followed the temporal dimension approach in the past. In this paper, we propose spatial backoff as an alternative approach to address medium access control along the spatial dimension. We propose a dynamic spatial backoff algorithm that jointly controls CS threshold and transmission rate. The results we present in this paper demonstrate that the proposed spatial backoff algorithm can adapt to a give network and achieve a high level of performance. Some topics for further research are as follows: (a) Potentially, all three parameters—transmission power, CS threshold, and transmission rate—can be jointly controlled to adjust the space occupied by transmissions. Work is needed to determine appropriate ways to realize such a joint control. (b) In this paper, we showed the benefits of

spatial backoff without using temporal contention resolution adaptation. Future work is needed to identify appropriate strategies to integrate the temporal and the spatial approaches. (c) The impact of traffic and channel variations over time needs to be further investigated. We have obtained some results on the impact of channel variations, and they suggest that the spatial backoff approach can work well in this case as well.

ACKNOWLEDGMENTS

This work was performed when Xue Yang was with University of Illinois at Urbana-Champaign. This research was supported in part by US Army Research Office grant W911NF-05-1-0246, NSF grant ANI-0125859, and Vodafone Graduate Fellowship.

REFERENCES

[1] T. D. Todd and J. W. Mark, "Capacity allocation in multiple access networks," *IEEE Trans. on Comm.*, pp. 1224–1226, 1985.

[2] X. Yang and N. Vaidya, "On physical carrier sensing in wireless ad hoc networks," in *IEEE INFOCOM*, March 2005.

[3] J. Yee and H. Pezeshki-Esfahani, "Understanding wireless LAN performance trade-offs," *Comm. Systems Design*, Nov. 2002, http://i.cmpnet.com/commsdesign/csd/2002/nov02/feat3-nov02.pdf.

[4] A. Kamerman and L. Monteban, "WaveLAN-II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal*, vol. 2, no. 3, pp. 118–133, Summer 1997.

[5] J. I. Capetanakis, "Tree algorithms for packet broadcast channels," *IEEE Trans. on Info. Theory*, , pp. 505–515, Sept. 1979.

[6] B. Hajek and T. V. Loon, "Decentralized dynamic control of a multiaccess broadcast channel," *IEEE Transactions on Automatic Control*, vol. AC-27, no. 3, pp. 559–569, June 1982.

[7] R. L. Rivest, "Network control by Bayessian broadcast," MIT, Laboratory for Computer Science, Cambridge MA, Tech. Rep. Report MIT/LCS/TM-285, 1985.

[8] F. Cali, M. Conti, and E. Gregori, "IEEE 802.11 protocol: Design and performance evaluation of an adaptive backoff mechanism," *IEEE JSAC*, pp. 1774–1786, Sept. 2000.

[9] H. Kim and J. C. Hou, "Improving protocol capacity with model-based frame scheduling in IEEE 802.11-operated WLANs," in *ACM MobiCom*, Sept. 2003, pp. 190–204.

[10] M. C. Yuang, B. C. Lo, and J.-Y. Chen, "Hexanary-feedback contention access with PDF-based multiuser estimation for wireless access networks," *IEEE Transactions on Wireless Communications*, vol. 3, no. 1, pp. 278–289, Jan. 2004.

[11] X. Guo, S. Roy, and W. S. Conner, "Spatial reuse in wireless ad-hoc networks," in *Proceedings of IEEE 58th Vehicular Technology Conference (VTC 2003-fall)*, vol. 3, Oct. 2003, pp. 1437–1442.

[12] J. Zhu, X. Guo, L. L. Yang, W. S. Conner, S. Roy, and M. 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.

[13] J. Zhu, B. Metzler, X. Guo, and Y. Liu, "Adaptive csma for scalable network capacity in high-density wlan: a hardwire prototyping approach," in *IEEE INFOCOM*, April 2006.

[14] A. Vasan, R. Ramjee, and T. Woo, "ECHOS—enhanced capacity 802.11 hotspots," in *IEEE INFOCOM*, March 2005.

[15] T. Nadeem, L. Ji, A. Agrawala, and J. Agre, "Location enhancement to IEEE 802.11 DCF," in *IEEE INFOCOM*, March 2005.

[16] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *ACM MobiCom*, July 2001.

[17] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *ACM MobiCom* Sept. 2002.

[18] J. P. Monks, V. Bharghavan, and W. mei W. Hwu, "A power controlled multiple access protocol for wireless packet networks," in *IEEE INFOCOM*, April 2001.

[19] A. Muqattash and M. Krunz, "Power controlled dual channel (PCDC) medium access protocol for wireless ad hoc networks," in *Proceedings of IEEE INFOCOM*, April 2003.

[20] J. Fuemmeler, N. Vaidya, and V. V. Veeravalli, "Selecting transmit powers and carrier sense thresholds for CSMA protocols," Oct. 2004, http://www.crhc.uiuc.edu/wireless/papers/PowerCarrierSense04.pdf.

[21] V. Rodoplu and T. H. Meng, "Minimum energy mobile wireless networks," *IEEE JSAC*, pp. 1333–1344, Aug. 1999.

[22] R. Ramanathan and R. Rosales-Hain, "Topology control of multihop wireless networks using transmit power adjustment," in *IEEE INFOCOM*, March 2000.

[23] L. Li, J. Y. Halpern, P. Bahl, Y.-M. Wang, and R. Wattenhofer, "Analysis of a cone-based distributed topology control algorithm for wireless multi-hop networks," in *ACM Symp. on Principles of Distr. Comp.*, Aug. 2001.

[24] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, "Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of the COMPOW protocol," in *Proceedings of European Wireless 2002, Next Generation Wireless Networks: Technologies, Protocols, Services and Applications*, Feb. 2002, pp. 156–162.

[25] N. Li, J. C. Hou, and L. Sha, "Design and analysis of an MST-based topology control algorithm," in *IEEE INFOCOM*, April 2003.

[26] W.-Z. Song, Y. Wang, X.-Y. Li, and O. Frieder, "Localized algorithms for energy efficient topology in wireless ad hoc networks," in *ACM MobiHoc*, May 2004.