Integrating the Physical and Link Layers in Modeling the Wireless Ad Hoc Networking MAC Protocol, Synchronous Collision Resolution

John A. Stine
The MITRE Corporation
7515 Colshire Drive
McLean, Virginia 22102-7508
Email: jstine@mitre.org

Abstract

Providing a set of protocols that will support the demands of mobile wireless ad hoc networking is fraught with many challenges. Despite over 30 years of research and effort, ad hoc networking still remains a field of research rather than a field of practice. Many of these challenges, however, are created by the use of an inappropriate paradigm to define the network's behavior. To date, the bulk of the research in this area has used a wired paradigm. We propose a new wireless paradigm that we believe is more suitable. The significance of the paradigm is that the wired paradigm tends to promote research in routing and higher layer protocols as a result of over simplifying the behavior of the link and physical layers. The wireless paradigm reveals that most of the networking functionality needs to be implemented at the link and physical layers. In this paper, we contrast the different paradigms and explain how the protocol, Synchronous Collision Resolution (SCR)[1], provides a framework for implementing algorithms that will solve the challenges of ad hoc networking according to the wireless paradigm. To do research in this area and to ultimately answer critics' concerns about the impact of the unreliable physical layer on the performance of SCR, has required us to make a sophisticated model of the physical layer. We describe how we implemented such a model using OPNET's Modeler environment.

I. Introduction

Mobile ad hoc networks are proposed as the preferred wireless networking solution when nodes are mobile and infrastructure is unavailable. The unique characteristics of the wireless channel make them very different than the ubiquitous wired networks. Nevertheless, the wired paradigm has been and remains the paradigm under which most try to understand ad hoc networks and in turn design protocols to operate these networks. In this paper, we describe a new paradigm that better describes the behavior of wireless ad hoc networks. We describe how the medium access control (MAC) protocol, Synchronous Collision Resolution (SCR), provides the framework to operate an ad hoc network according to this paradigm. With this new paradigm and the use of SCR there is greater emphasis on integration of the physical and link layers. Modeling this integration is necessary to validate the suitability of SCR. We describe how we model the link and physical layers using OPNET's Modeler environment.

We start our discussion in Section II by describing both the wired and wireless paradigm and then comparing how they influence the study of ad hoc networks. We explain why the wired paradigm is inadequate for a wireless network. Then in Section III we describe SCR and demonstrate how it provides the framework to understand and to manage a wireless network according to our new paradigm. In Section IV we describe the most common objections to the use of SCR and make the case that accurate modeling of the integration of the physical and link layers is necessary to validate its capabilities. In Section V we describe how we accomplished this integration in a simulation model. Section VI presents some simulation results using our model that demonstrates the utility of have an integrated physical and MAC layer model. Section VI concludes the paper.

II. Comparison of the Wired and Wireless Paradigms

Wired networks are built using the paradigm of routers and links. The critical resource of a router is its buffers and the critical resource of a link is its capacity. The router contains the logic to control its buffers and the use of capacity on any of its outgoing links. The design problem in wired networks is to use as much of each resource as possible without compromising performance. This involves preventing any convergence of flows that may cause buffer overflow or an excessive backlog at a link. When applied to a wireless network, the wired paradigm is slightly modified to acknowledge that links are temporal, they come and go based on the movement of nodes. Thus the design problem is expanded to include more advanced protocols that can respond to a changing topology.

The wired networking paradigm above fails to adequately describe wireless networks because of its oversimplification of the nature of wireless links. Its abstraction of wireless links as having two states, available and not available, and being independent of each other is incorrect. In wireless networks, links are not physical entities. Connections may be created between pairs of nodes but these are states. The state of a connection between two nodes exists through the cooperation of their neighbors. Consider Figure 1. In this example, there are 46 potential directed connections between pairs of the 10 nodes. Say node 4 establishes a connection to node 7. To ensure no interference from adjacent nodes with this connection and vice versa would require none of the other 45 potential directed connections be used. We see that the resource of interest is not a link but a wireless channel in a geographic space. Our wireless paradigm is based on this observation and its ramifications.

Several observations form our wireless paradigm:

1. The critical resource is a wireless channel in a geographical space.
2. No single node controls the resource.

[1] Patent pending

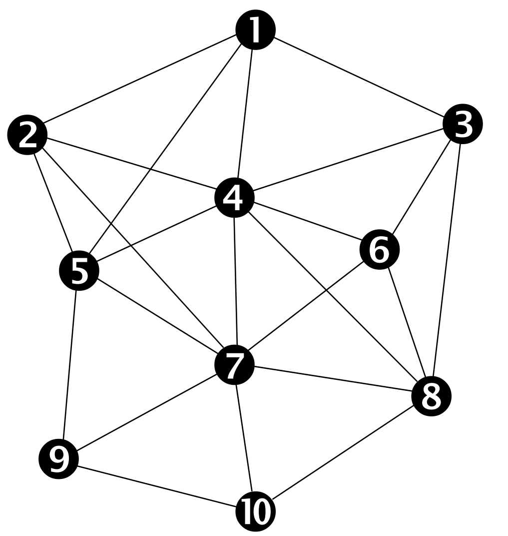

Figure 1. An example of link interdependence in wireless networks: When the directed connection from 4 to 7 is used, no other directed connection may be used.

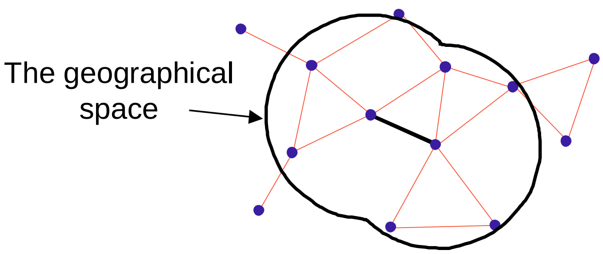

Figure 2. An example of a channel space required for a connection between two nodes. The extent of the channel space determines the membership of the distributed queue, in this case all nodes.

3. There is a distributed buffer for each space formed by the buffers of those nodes that have interest in each space.
4. The geographical spaces overlap.
5. Nodes participate in several distributed buffers.

Figure 2 provides an example scenario that illustrates these observations. The figure illustrates a connection between two nodes (thick line) and the geographic space that is consumed by an exchange between the two. This space is uniquely associated with this pair. The figure also illustrates the potential connectivity between all pairs of nodes (thin lines) in the network. All nodes in this example participate in the distributed queue of the connected nodes due to either being within the channel space or having destinations that are within the channel space. Understanding this scenario is made even more complex when one considers the role of the physical layer. Physical layer characteristics can be used to enable more than this single connection to exist in this example.

Paradigm has had a profound influence on the research and development of ad hoc networks. The wired paradigm, with its discrete view of the nature of links, has bifurcated research. Research and development at the link and physical layers has as its objective the perpetuation of this abstraction of discrete links. Variation of topology and issues of quality of service are then handled by routing and other higher layer protocols that work with this abstraction. The deficiency of this paradigm is that the true nature of the wireless environment disables higher layer protocol mechanisms based on the discrete link abstraction. Higher layer protocols usually assume they can control how traffic is offered to the communications medium. However, the contentious and volatile nature of the wireless medium results in the use of buffers at the link layer. The higher layer protocols can only control how traffic is offered to these buffers. A suggested improvement is to allow cross layer communications but this also comes up short since access to the medium is distributed. To empower higher layer protocols as desired requires MAC and physical layer mechanisms not only to isolate links but also to arbitrate access based on the contents of the distributed buffers. To accomplish this goal requires these mechanisms to be designed using the wireless paradigm. The wireless paradigm increases the emphasis of design at the physical and link layers.

Understanding and developing an intuition of the interactions involved in ad hoc networks using the wireless paradigm is very difficult. The amorphous form and interdependence of connection dependent geographic spaces and their distributed buffers is complex. Fortunately, the structure and mechanisms of the SCR MAC protocol provide an intuitive framework that can be used to both understand and to manage a wireless network according to the wireless paradigm.

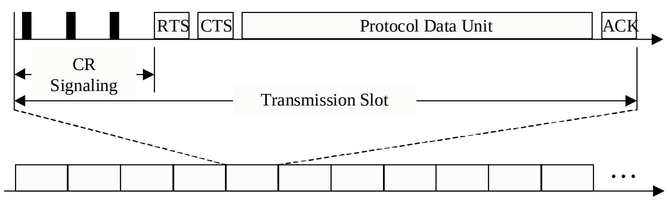

Figure 3. The Synchronous Collision Resolution Protocol

III. Synchronous Collision Resolution (SCR)

Synchronous Collision Resolution is a broad MAC definition and is better viewed as a paradigm for access. SCR has four key characteristics:

1. The communications channel is time slotted.
2. Nodes with packets to send contend in every slot. There are no backoff mechanisms.
3. Signaling is used to arbitrate contentions.
4. Packet transmissions occur simultaneously.

Figure 3 illustrates the general concept. Specific details such as the design of the signaling, whether or not to use the RTS/CTS exchange, or whether to execute the collision resolution signaling (CRS) every slot are choices that are made considering the physical layer and how the network will be used.

SCR's characteristics are what make it suited for the wireless paradigm. The synchronizing of access attempts and the use of an interactive contention arbitration mechanism, CRS, enables SCR to seek out the best collection of nodes to exchange traffic at the beginning of each transmission slot. At the conclusion of the signaling, the set of exchanging nodes is frozen without risk of mid-transmission collisions. The definition of the "best" collection of nodes that CRS arbitrates is dependent on how the signaling mechanism is designed. At minimum, it arrives at a relatively dense set of nodes that can exchange traffic simultaneously. CRS may also be designed to arbitrate access

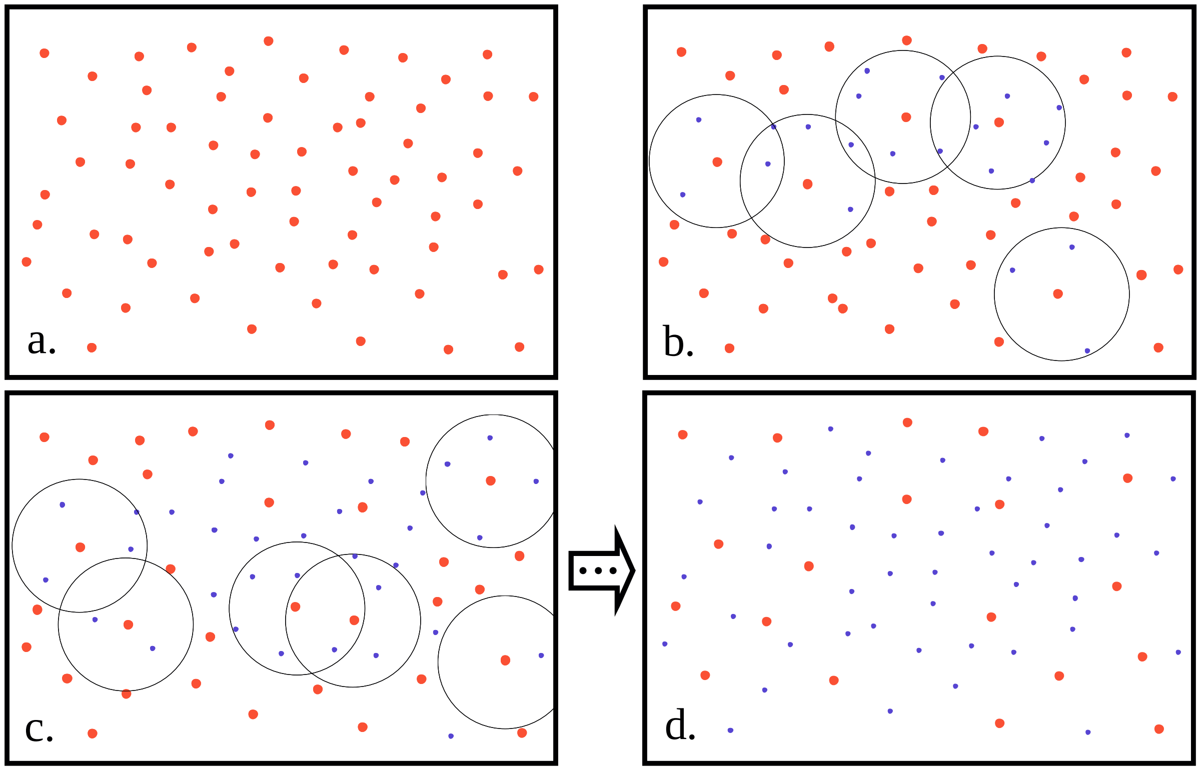

Figure 4. An example of the effect of Collision Resolution Signaling: All nodes start off as contenders (large dots) in Panel a. Then, through a series of signals, two sets of which are illustrated in Panels b and c, a final subset of contenders is selected in Panel d. The large circles represent the range of the signals. The small dots represent nodes that have deferred from contending.

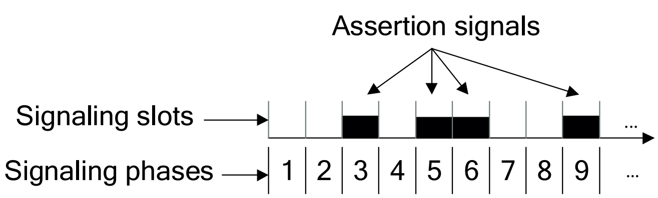

Figure 5. Collision Resolution Signaling using single slot phases

giving preference to nodes with the highest priority packets in their queues or to coordinate physical layer characteristics. In this paper, we only discuss the basic arbitration mechanism.

Collision resolution signaling consists of series of short signaling slots organized into groups of slots called phases in which contending nodes may send very short signals. There are numerous ways to design signaling. The simplest and generally most effective at arbitrating contention is illustrated in Figure 5, which illustrates a series of slots where each slot constitutes a phase. These signaling slots should not be confused with the longer transmission slots of Figure 3. Rather, they occur within a transmission slot during a short period at the very beginning. In Figure 3, we labeled this short period "CR Signaling." In this design using single slot phases, a probability is assigned to each slot. This value is the probability that a contending node will signal in that slot. The rules of signaling in this design are as follows.

1. At the beginning of each signaling phase a contending node determines if it will signal. In single slot phases a contending node will signal with the probability assigned to the slot of that phase.

2. A contender survives a phase by signaling in the slot or by not signaling and not hearing another contender's signal. A contender that does not signal and hears another contender's signal loses the contention and defers from contending any further in that transmission slot.

3. Nodes that survive all phases win the contention.

Figure 4 illustrates the process. In Panel 4a, we illustrate a scenario where all nodes in the network start as contenders, and then, through the series of signals, two sets of which are illustrated in Panels 4b and 4c, reduce these contenders to the final subset of contenders illustrated in Panel 4d. The large dots are the nodes that view themselves as contenders, the small dots are nodes that view themselves as having lost the contention, and the large circles represent the range of the signals. The desired outcome of CRS is to arrive at a subset of contenders that are separated from each other by at least the range of their signals.

We discuss the design and analysis of CRS in [1], and present the more interesting results in Figures 6 and 7. This type of contention has three measures of effectiveness: the probability that CRS will isolate a single survivor amongst nodes in range of each other, the average separation distance of a surviving contender and its nearest surviving neighbor, and the average density of surviving contenders. Figure 6 illustrates the ability of CRS to arbitrate contention. These results were derived

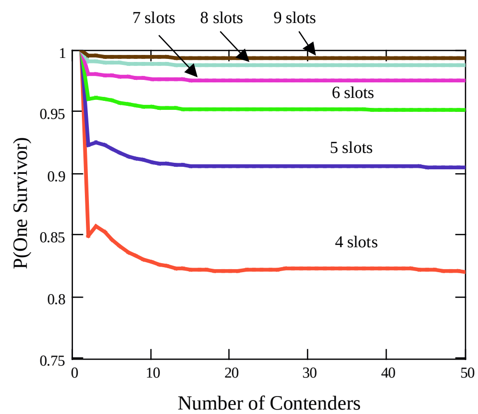

Figure 6. Contention resolution performance of different signaling designs

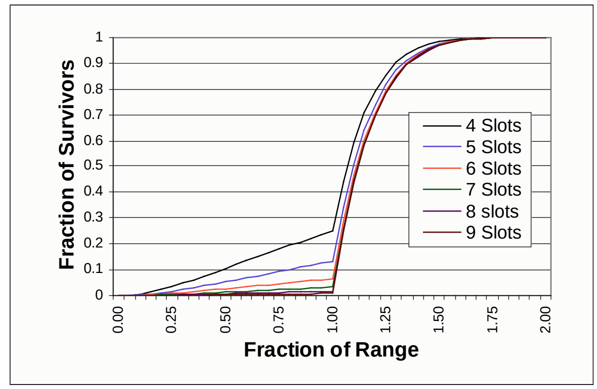

Figure 7. Cumulative distribution of range from a survivor to its nearest surviving neighbor

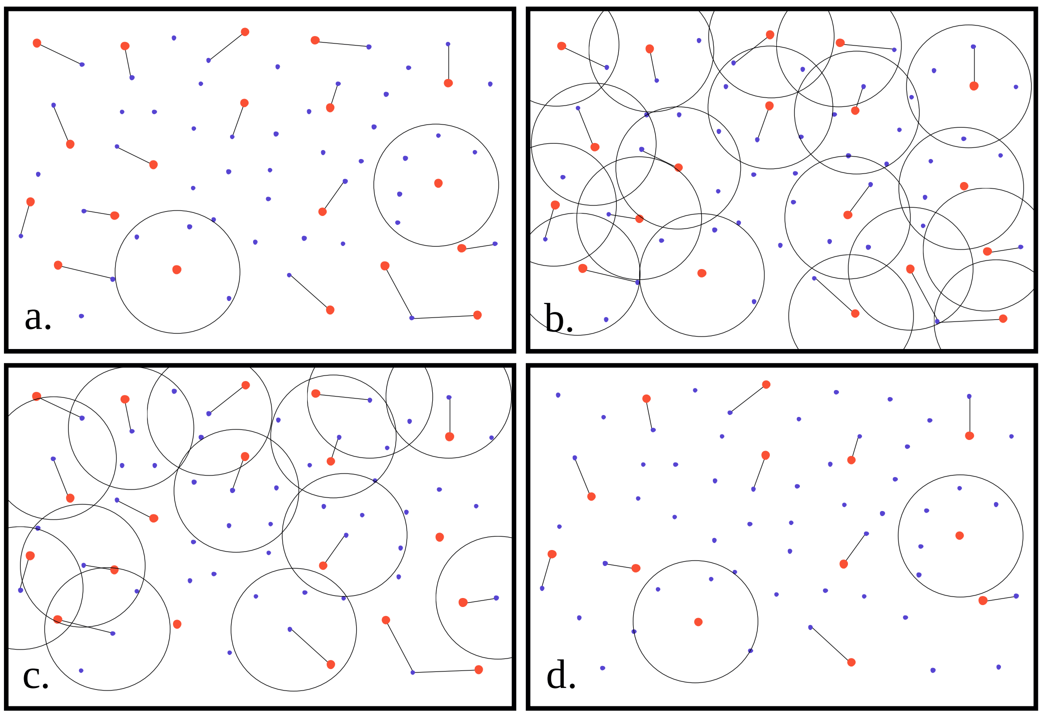

Figure 8. Example of the RTS-CTS handshake finalizing the set of nodes to exchange packets: Panel a illustrates the set of contenders that survived signaling and their intended destinations (circles indicate intended broadcasts). Panel b illustrates the contenders' simultaneous RTS transmissions. Panel c illustrates the destinations' simultaneous response of CTS packets. Panel d illustrates the winners of the contention.

analytically. The horizontal axis is the number of contenders and the vertical axis is the probability that there will be just one survivor at the conclusion of the signaling. The performance improves with the number of slots and with as few as 9 slots CRS is better than 99% effective at isolating a single contender with as many as 450 nodes contending. There is one especially nice feature. Collision resolution performance can be designed to be consistent across a large range of contenders making capacity a function of geographic space rather than a function of the density of contenders. This is the desired result under the wireless paradigm. Figure 7 further illustrates the suitability of CRS. These are the cumulative distributions of separation distances of survivors and their nearest surviving neighbors using the signaling designs whose contention resolution performance is illustrated in Figure 6. The probability that a survivor's nearest surviving neighbor will be within the range of the radio matches that predicted by Figure 6. Most nearest surviving neighbors lie at a distance from 1 to 1.5 times the range of the radio. These results came from geometric simulations executed in MathCAD where nodes were arranged on a toroidally wrapped surface according to a Poisson point process. The results were consistent for all node densities tested, degrees[2] of 5, 10, 15, and 25.

At the conclusion of signaling, survivors are separated but this is not necessarily true for their destinations where we are most concerned about interference. The purpose of the RTS/CTS handshake that follows CRS is to verify that sources and destinations can close a connection and to provide a feedback mechanism to adapt the transmission characteristics for a more

[2] Node degree is the average number of neighbors within range of each node.

reliable and energy efficient PDU exchange. Figure 8 illustrates the process. In Panel 8a, the large nodes are the signaling survivors. We have drawn lines from signaling survivors to their intended destinations. Circles are drawn around nodes that are broadcasting a packet. Panel 8b reveals those nodes that transmit RTS packets. The large circles are the ranges of their RTS transmissions. If a destination receives a RTS packet, it responds with a CTS packet, see Panel 8c. These CTS packets are also sent simultaneously. Recipients of the RTS packets for broadcasts do not respond, since the source would not be able to distinguish CTS packets from multiple destinations. In the end, all broadcasting nodes and those nodes that have received a CTS from their destination transmit PDUs, see Panel 8d. The RTS and CTS packet transmissions create the worst-case mutual interference. Since contenders may defer if they do not receive a CTS, mutual interference conditions will only improve for the subsequent PDU and ACK packets.

As has been explained, SCR exploits the properties of the physical layer to arbitrate contention and to achieve high capacity. Specific features that are very important are the attenuation of signals that result from propagation and terrain effects and the capture of signals in the presence of interference.

IV. Common Objections

The most commonly voiced objections to SCR are concerning its reliance on synchronization and the quantity of overhead there appears to be with signaling. The significance of synchronization and overhead are both dependent on the capabilities of the physical layer.

The role of synchronization in SCR is to prevent ambiguity as to when signals are sent during signaling. In this role, SCR relies on synchronization only to the extent that it affects the protocol's efficiency. Figure 9 compares the factors that determine the size of a backoff slot in the IEEE 802.11 MAC protocol and those that affect the sizing of a signaling slot in SCR. The only difference is additional time to accommodate the level of synchronization, here measured as the maximum difference in synchronization that may exist between neighboring nodes. With good synchronization, the overhead of signaling is comparable and most likely better than that of the 802.11 MAC carrier sensing and collision avoidance mechanism. Table 1 lists some values of interest for these factors. The significance of this information is that other physical factors are likely to have a more significant role in determining the quantity of overhead than synchronization.

One of the features of wireless systems is that it is ideally suited to achieve synchronization. The delay accrued by a signal while propagating from a transmitter to a receiver is predictable based on the separation of the devices. With short range radios, a simple beacon can be used to achieve synchronization of a few microseconds. [2] More advanced systems combine synchronization into position location protocols. The most notable of these systems are GPS and the Enhanced Position Location Reporting System (EPLRS) used by synchronization of a few microseconds. [2] More advanced systems combine synchronization into position location protocols. The most notable of these systems are GPS and the synchronization of a few microseconds. [2] More advanced systems combine

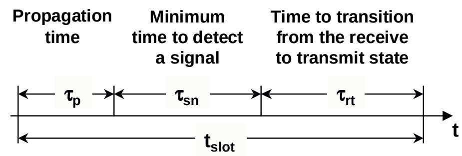

a. Factors that affect the sizing of a backoff slot in IEEE 802.11

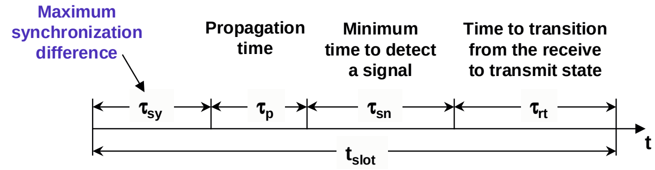

b. Factors that affect the sizing of a signaling slot in SCR

Figure 9. Comparison of factors that affect the sizing of backoff slots of the IEEE 802.11 MAC and signaling slots of SCR

Table 1. Protocol sizing factors

Parameter	802.11 FH	802.11b	802.11a	JTRS WNW*
τ_p	1µs	1µs	1µs	34µs
τ_{SN}**	27 µs	14 µs	<6 µs	?
τ_{rt}	20 µs	5 µs	<2 µs	?

* The Joint Tactical Radio System (JTRS) Wideband Networking Waveform (WNW) is an active program to build an ad hoc network for the U.S. military. Its specification requires that the land radio have a range of 10 kilometers and that it be able to achieve its own synchronization. Otherwise other timing capabilities of the radio are not specified.
** The sensing time has combined the times to make a clear channel assessment and to process this information by the MAC. [3,4,5]

synchronization into position location protocols.The most notable of these systems are GPS and the Enhanced Position Location Reporting System (EPLRS) used by the U.S. military. These systems achieve synchronization of less than 1 µs. Protocols that resolve location and synchronization require destinations to understand when transmissions are sent. Clearly, this is a feature of SCR. There is no ambiguity as to when a node intended to send a signal as there would be in a MAC protocol that uses random backoff to avoid collisions. In military systems, synchronization is necessary regardless of the MAC protocol used since time synchronization is a necessary component for transmission security. Integral synchronization is a specification for military ad hoc networks. We hope to build a position location/ synchronization protocol that works on top of the basic SCR access mechanism.

V. Physical Layer Modeling

We have demonstrated that a significant number of the features necessary to make an ad hoc network work well must be accomplished between the MAC and physical layers of these networks. In Section II, we demonstrated that the paradigm that best defines the nature of ad hoc networks is driven by the factors at the physical layer that affect whether transmissions can be received or detected by another transmitter. In Section III, we demonstrated how the SCR MAC attempts to exploit the physical layer to arbitrate contention and to maximize the utilization of the wireless channel. Then, in Section IV, we demonstrated that the performance of a MAC protocol is directly attributable to the capabilities of physical layer and that the integral use of the MAC and physical layer can provide services such as position location and synchronization. There are several additional issues that also require an accurate portrayal of the physical layer such as modeling whether radios are ready to receive signals and the energy consumption of nodes. The integration of the MAC and physical layers is essential in the modeling of ad hoc networks.

We chose the OPNET Modeler environment to develop models to do our research and development work in this area. Figure 10 illustrates our basic wireless node model before adding higher layer protocols. The feature that we want to highlight is the inclusion of a radio process between the MAC process and the transceiver and receiver modules that are provided by OPNET. OPNET's transceiver and receiver modules provide a means to model most waveform and propagation characteristics. Specifically, it supports the modeling of the waveform (i.e. frequency, bandwidth, modulation, error correction, data rate, and processing gain), propagation effects (i.e. pathloss and propagation delay), and some radio and antenna effects (i.e. transmission power, antenna gain, signal capture and transceiver/receiver isolation at a node). They do not model the states of the radio, the transitions between them, nor the energy consumption in those states. They also require pre-selection of the transceiver parameters (i.e. frequency, bandwidth, transmission power, modulation, data rate, and processing gain) or the use of an outside process to change them. The purpose of the radio process is to model the states of the radio, the energy consumption in those states, and the potential capability of a radio to dynamically change its transceiver parameters.

Figure 11 illustrates our radio process. We have used six states to model the radio: four states associated with transmission and reception, a state associated with transition, and a state associated with inactivation for the purposes of energy conservation. The activity of a radio and its energy consumption is determined by its state.

There are two different states for receiving and two different states for transmitting. This model supports the use of fundamentally different transmissions for signaling and for packet exchanges. In signal reception the goal is for the receiver to simply detect a signal. In packet reception the goal is for the receiver to capture the signal and to recover the bits of the message. Signaling uses the characteristics of where a receiver is looking for a signal together with the energy level of an arriving signal to determine detection. A signal should be designed such that it can be detected very quickly even in the presence of multiple arriving signals. The determining factor of whether a receiver detects a signal is that it is looking for the signal in the right place. Packet reception is much more complex and very dependent on whether a receiver can capture the signal in the presence of interference. So in the OPNET transceiver pipeline stages where capture, interference, and error correction are assessed, signals are treated differently from packet transmissions. The pipeline stages understand what type of transmission should be sought and whether the radio is capable of correctly receiving a transmission by knowing the state of the radio.

Movement between states requires traversing the transition state. The transition state models the time it takes to move between states and the time to change the parameters of the transceiver

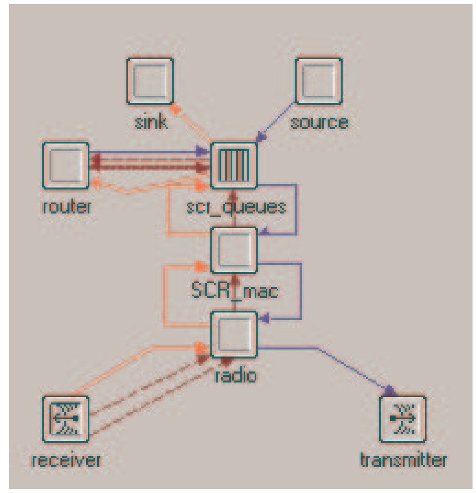

Figure 10. The SCR wireless node model

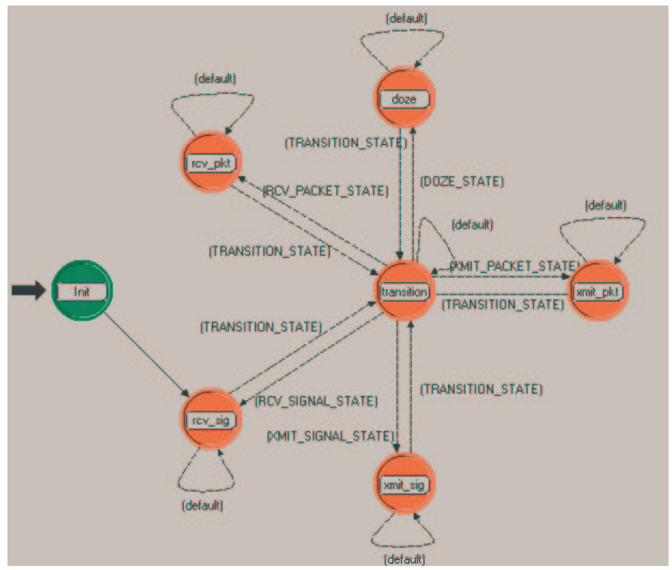

Figure 11. State diagram of the radio process

and possibly of an antenna. We assume that during this state a radio cannot receive nor transmit.

The doze state models a period of inactivation where nodes can conserve energy. Nodes cannot receive nor transmit while in this state. We have modeled the transition and doze states to accommodate multiple different low energy states that differ in the time to transition and the rate of energy consumption while in these two states. For example, the level of energy conserved in the doze state may be correlated with the time to transition into that state where the lower the energy consumption the longer it takes to transition into and out of the state. A protocol may choose a particular low energy state based on how long it intends to doze and whether it is practical to enter the state considering the time to transition and the energy consumption while transitioning. For a discussion of the use of different dozing states with SCR see [6].

Our model of a radio process generates many questions that need to be answered by radio designers. It emphasizes the role that the physical layer plays in the performance of an ad hoc network. Since the design and performance of the SCR MAC is very dependent on these capabilities we use a radio capabilities header file in which to define the critical characteristics of the radio. It includes multiple performance parameters: the time it takes radios to transition between states; the rates of energy consumption; the time it takes a radio to sense a channel before assessing the presence of a signal; and some important characteristics for the pipeline stages such as the threshold levels of received power for signal and packet detection. Portions of this information is used by all processes of our basic wireless node model so that the MAC can be tuned to the radio's capabilities just as the IEEE 802.11 MAC is tuned to each of its different physical layers.

In operation, our SCR MAC protocol acquires information on the reception of signals concerning the strength of a received signal, the number of bit errors in a packet, the time of arrival of a transmission and potentially the direction of arrival of a transmission. This information is passed from the receiver module through the radio to the MAC. It is then used for connection adaptation and network state awareness. Network state awareness information is used to enhance routing, to integrate the use of directional antennas with the SCR MAC, and as described earlier for the integration of a position location and synchronization algorithm with the SCR MAC.

VI. Example Uses of the Model

We provide two examples that demonstrate the exploitation of the radio model to answer questions not typically considered when analyzing the performance of ad hoc networking protocols.

The RTS-CTS mechanism can be avoided by changes in the signaling approach.[3] One of the advantages of the RTS-CTS handshake, however, is that it supports the conservation of energy. Nodes not transmitting or receiving packets can enter a low energy state during PDU exchanges. Without the handshake, nodes cannot be certain that they are not destinations until after receiving a PDU at which time it is too late to enter a low energy state. With the handshake they can enter a low energy state after the RTS. We designed an experiment to study this advantage. We executed simulations with 25 nodes, all within range of each other. In this topology, the RTS-CTS handshake has no function other than being a feedback mechanism. We compared two different signaling designs differing only in whether they used the RTS-CTS handshake. Both used the 9 slot signaling design whose performance is illustrated in Figure 6. The packet overhead was similar to that for the IEEE 802.11 MAC: 192 bits for physical layer overhead, 320 bits for the PDU header, and then 160, 112, and 112 bits for the RTS, CTS, and ACK packets respectively. Using a 506 byte payload size, the design using the RTS-CTS handshake had a payload to transmission slot ratio of 0.66 while the design without the handshake had a ratio of 0.745. Using a 1 Mbps channel, this corresponds to 163 transmission slots per second for the design using the RTS-CTS handshake and 184 transmission slots per second for the design without the handshake. We assumed energy consumption to be 2.75 watts when transmitting, 1 watt when receiving, and 0.1 watt when dozing. We assume the time to transition to the dozing state to be 10 μsec and the time to transition out of the doze state to be 1 msec. We modeled the arrival of packets as a Poisson process at each node. Each packet could be sent in a transmission slot. Packet destinations were selected randomly.

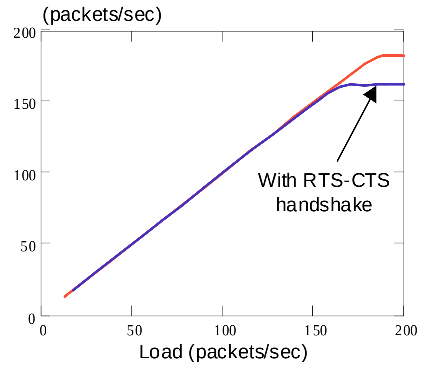

Figure 12. Throughput

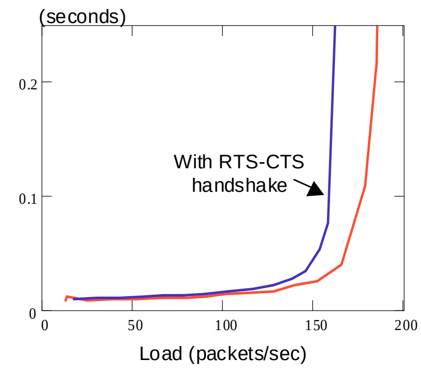

Figure 13. Packet delay

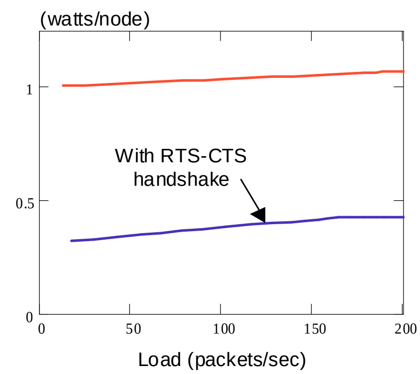

Figure 14. Average energy consumption rate

Figures 12 through 14 illustrate the results of our experiment and the tradeoffs. Figure 12 is the throughput. The capacity matches the load until about 98% of the capacity is reached. There is no congestion collapse. The difference in capacity is as predicted. Figure 13 illustrates the corresponding delay. The average delay is the same up until the load reaches about 120 packets/second. Figure 14 illustrates the average energy consumption per node. We see that in all cases the average energy consumption of nodes when the RTS-CTS handshake is used is less than half that of when it is not used. The analysis follows readily. Unless load is expected to be high, the signaling with the RTS-CTS mechanism

[3] The Synchronized Unscheduled Multiple Access (SUMA) [7] signaling design uses echoing to achieve two hop separation of contenders thus eliminating the need for an RTS-CTS handshake to verify signal capture.

provides equivalent throughput performance while allowing batteries to last more than twice as long on average.

The radio model allows this experiment to be repeated easily with different physical layer parameters that affect the data transmission rate, the energy consumption rates, and the transition times.

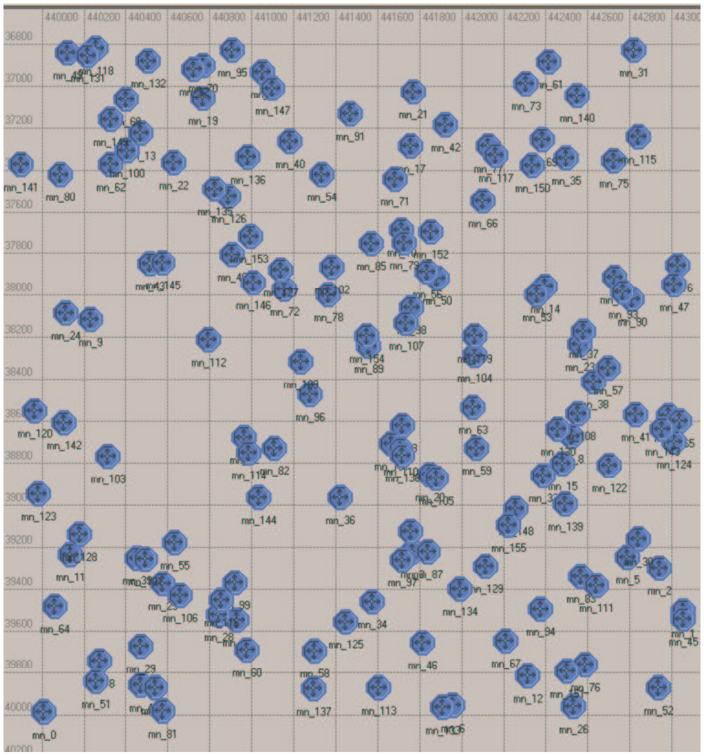

Figure 15. The ad hoc networking scenario

Our second example explores the role of processing gain in a multihop ad hoc network. For our analysis, we created a square toroidally wrapped simulation area, 7 transmission ranges on a side.[4] We define transmission range as the threshold distance to where a received signal has a 10 dB signal to noise ratio. In our simulation, all transmitters used the same transmit power and we modeled pathloss using the two-ray propagation model. We used a constant noise level to account for thermal and ambient sources. Since we were studying physical and link layer issues, we implemented a “perfect” routing protocol and kept all nodes stationary to avoid confounding the results with routing protocol effects. The router was omniscient of all pathloss and considered a connection possible if a received signal could obtain the 10 dB SNR threshold. We randomly placed 156 nodes on the simulation area to achieve an average degree of 10. Figure 15 illustrates our scenario.

We conducted two experiments where one used no processing gain and the second used 10 dB of processing gain.[5] We modeled traffic arrivals as Poisson and randomly selected a

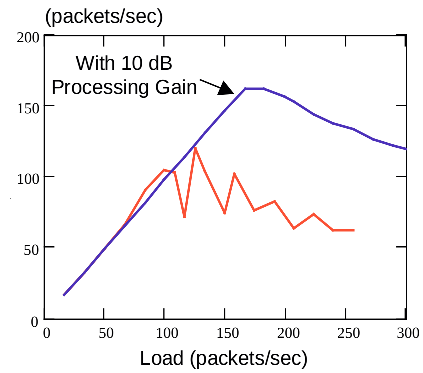

Figure 16. End-to-end throughput

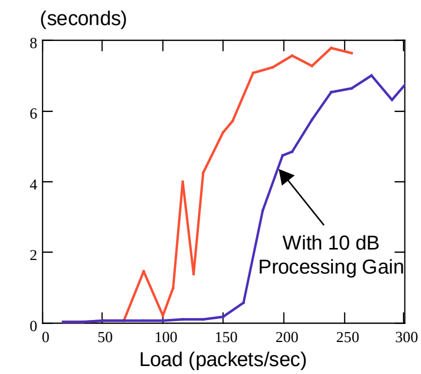

Figure 17. End-to-end delay

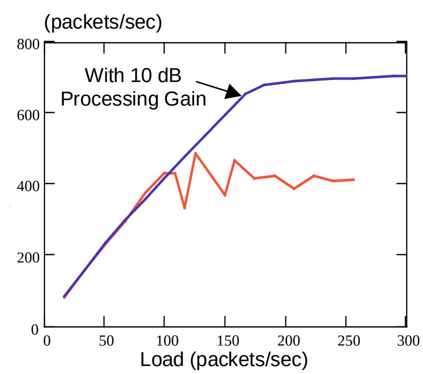

Figure 18. MAC protocol throughput

destination for each. We used an independent random number generator for this traffic so that arrivals were the same for both scenarios when common arrival rates were used. We assumed unlimited buffer space but that packets had an 8 second lifetime and were dropped when it expired. Packets were queued in ascending expiration time order. Figures 16 through 18 illustrate the results. These figures illustrate the end-to-end delivery rate of packets, the end-to-end delivery delay, and the 1 hop exchange rate, all as a function of the packet arrival rate to the network. The results demonstrate that this small amount of processing gain nearly doubles the capacity of the network. Two additional interesting observations come from these results. First, we see the significance of tuning the routing metric. Ad hoc networks are interference, not noise, limited. Routing protocols that seek minimum hop routes, as we are using, will favor longer hops where interference can be a greater problem. The jitter in the results for the network not using processing gain can be attributed to the scenario effect of getting stuck on a bad route where interference continuously blocks the exchange of a

[4] The purpose of toroidally wrapping a simulation area is to remove edge effects. On a toroidally wrapped surface, transmissions can reach across borders and be received on the opposite side of the surface. Nodes close to the border can exchange packets across to the opposite side and nodes near corners can exchange packets across to the opposite corner.

[5] The 10 dB processing gain approximates that of the IEEE 802.11 1 Mbps DSSS physical layer.

particular packet. The second interesting observation is that congestion collapse occurs in the end-to-end throughput but does not occur at the MAC level. Thus, we can attribute the collapse to the queue pruning policy rather than to the access mechanism as would be the case for Aloha or CSMA type protocols. Congestion collapse occurs in networks when capacity is used to move packets through a network but they expire before final delivery. A more aggressive queue pruning policy that drops packets that are unlikely to be delivered prior to timing-out could improve average performance. We used our models to test this hypothesis. We used the following pruning criteria,

$$t > t_x - l\left(h-1\right)^n t_{ts} ,$$

where t_x is the expiration time, l is a linear effect, h is the number of hops to the destination, n is an exponential effect, and t_{ts} is the duration of a transmission slot. We executed two sets of experiments, one in which we used just a linear pruning effect, $n = 1$ and $l = 10$, and a second using an exponential pruning effect, $n = 2$ and $l = 10$. Figure 19 compares the results of using these pruning techniques in the 10 dB processing gain scenario. Indeed, aggressive pruning mitigates congestion collapse and improves the throughput of a congested network.

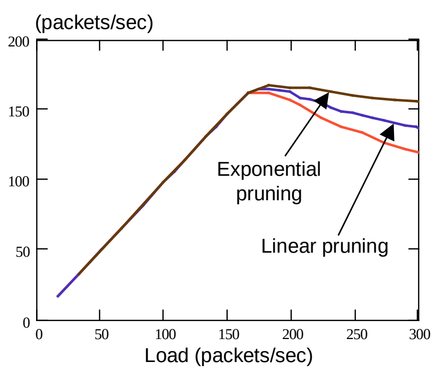

Figure 19. End-to-end throughput

As demonstrated, there are interesting interactions between the link and physical layers of ad hoc networks that can dramatically affect the performance of those networks. We intend to use our models to further investigate other types of these interactions.

VII. Conclusion

In this paper we have made the case that wireless ad hoc networking requires a different paradigm than that used for wired networks and proposed an appropriate alternative. This new wireless paradigm emphasizes the role of the physical layer in determining the performance of an ad hoc network. We demonstrated that the access protocol, SCR, is exceptionally suited to manage a wireless network under this paradigm. However, we also demonstrated that the suitability of this protocol is dependent on the capabilities of the radio itself. For these reasons, we develop a specialized radio process in the OPNET Modeler and enhanced the OPNET radio pipeline stages to capture all the critical performance capabilities of radios. It is our intent to use this model to quickly integrate SCR with different physical layers. We intend to exploit the features of this model to do additional research in the development of connection adaptation, antenna control, routing, energy conservation, and position location/ synchronization algorithms.

VIII. References

[1] J. Stine, G. de Veciana, K. Grace, and R. Durst, “Orchestrating Spatial Reuse in Wireless Ad Hoc Networks using Synchronous Collision Resolution,” *Journal of Interconnection Networks*, Vol. 3 No. 3 & 4, September and December 2002, pp. 167 – 195.

[2] J. Elson, L. Girod, and D. Estrin, “Fine-Grained Network Time Synchronization using Reference Broadcasts,” http://lecs.cs.ucla.edu/~estrin/papers/wsn-ts.pdf, August 2002.

[3] ANSI/IEEE Std 802.11, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Press, 1999.

[4] ANSI/IEEE Std 802.11-a, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-speed Physical Layer in the 5 GHz Band, IEEE Press, 1999.

[5] ANSI/IEEE Std 802.11-a, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-speed Physical Layer in the 2.4 GHz Band, IEEE Press, 1999.

[6] J. Stine and G. de Veciana, “A Comprehensive Energy Conservation Solution for Mobile Ad Hoc Networks,” *IEEE International Conference on Communications*, 2002.

[7] K. Grace, “SUMA – The Synchronized Unscheduled Multiple Access Protocol for Mobile Ad Hoc Networks,” *Proceedings Eleventh International Conference on Computer Communications and Networks, IC3N ‘2002*, Oct. 2002.