A LINK SCHEDULING AND AD HOC NETWORKING APPROACH USING DIRECTIONAL ANTENNAS

by
J. Bibb Cain, Tom Billhartz, and Larry Foore
Harris Corporation
Edwin Althouse and John Schlorff
Naval Research Laboratory

ABSTRACT[1]

There is strong interest within DoD to utilize high-gain, directional antennas at both the transmitting and receiving end of the link in a dynamic, ad hoc network environment. However, the application of directional antennas (e.g., phased-array or sectorized antennas) in a dynamic network of mobile nodes requires coordination of antenna steering at both the receiver and transmitter ends of the link. Our solution is to apply adaptive, link-state routing (be performed by the OLSR ad hoc routing protocol [1]) supported by a distributed, adaptive Time Division Multiple Access (TDMA) scheduler, which determines schedules based on cooperative decisions between each pair of neighbor nodes. The architecture that has been developed contains a high rate mission data channel with an adaptive TDMA link scheduling protocol designed to take advantage of high-gain directional antennas. Time slots on this channel are adaptively scheduled to meet dynamic traffic demand requirements and to avoid interference from adjacent transmitting nodes. In addition, the link scheduling protocol must adapt to changes in node neighborhood topology caused by node mobility and link obstructions. This architecture will be tested and evaluated in two ways: by OPNET simulations and by field demonstrations.

INTRODUCTION

Most current efforts to apply ad hoc networking to tactical military systems have used simple RF physical layer solutions with omni-directional antennas. Accompanying that simplicity are the following problems: 1) low-data rates, 2) inadequate link-budget margins at required stand-off ranges, 3) spatially-uncontrollable RF signature, 4) extreme susceptibility to jamming and denial of service, and 5) channel-sharing inefficiency. Consequently, there is strong interest within DoD to utilize high-gain, directional antennas at both the transmitting and receiving end of the link to successfully mitigate all of the problems cited above. However, the application of directional antennas (e.g., phased-array or sectorized antennas) in a dynamic network of mobile nodes requires coordination of antenna steering at both the receiver and transmitter ends of the link. An algorithm or protocol is needed to control how a number of aircraft and terrestrial nodes with directional antennas communicate in order to support a highly responsive and dynamic network. The Naval Research Laboratory (NRL) and Harris Corp. have embarked upon a multi-year project under ONR sponsorship to develop protocols to support networking with directional antennas with the objective of supporting Network-Centric Operations via Littoral extension of the network from navy ships to the forces ashore.[1] The link scheduler consists of algorithms, protocols, and software that will adaptively schedule connections among pairs of directional antennas in a mesh network in order to provide complete network connectivity, meet aggregate traffic demands, and support highly variable traffic exchanges with neighbor nodes. Shortly after the start of this program, a DARPA program, identified as Future Combat Systems – Communications (FCS-C), was initiated with similar objectives. The major differences in technical approach between the two programs reside in the design of the link-scheduling and routing algorithms. This paper focuses exclusively on the ONR/NRL/Harris algorithms that are intended to support ship-to-forces-ashore networking in the context of Network Centric Operations in the Littoral Battlespace. However, the algorithms are also appropriate to support Army Future Combat Systems scenarios.

[1] This work was supported by the Naval Research Laboratory under Contract No. N00014-96-C-2063.

0-7803-8140-0/03$17.00 (C) 2003 IEEE

BACKGROUND/PROBLEM DISCUSSION

We focus on using directional antennas at both ends of the high-rate data link (i.e., for transmitting and receiving) because only then do we create conditions that provide satisfactory solutions to all of the five problem areas cited above. Directional gain at both ends of the link provides greater link power-budget for supporting higher data rates and longer stand-off ranges required for aircraft-to-aircraft and aircraft-to-ground links. It strongly reduces RF signature in unnecessary directions and reduces receiver jamming susceptibility (although this is somewhat dependent on antenna side-lobe characteristics). Finally, because of spatial containment of radiated power by the directive gain of the antenna, it is possible for multiple pairs of nodes to occupy the same space-time-frequency domain simultaneously without interference. Our media-access (MAC)-layer protocols capitalize on this property and attempt to maximize the number of simultaneous uses of the RF channel. In principle our protocol can support multiple simultaneous beams and RF channels at each node, thereby creating a richly-connected network with lower latency than could be obtained with a single beam and single RF channel per node. However, our networking protocols will initially be required to support only a single-beam per node and a single-RF-channel per node. Following successful demonstration of this capability, these restrictions will be later removed to allow addition of multiple beams and RF channels when affordable.

The process of coordinating the antenna pointing at both ends of the link requires identity of nearest "reachable" neighbors (neighbor-discovery algorithm) and the corresponding position or steering direction. Information exchange that is required to coordinate pointing schedules is exceedingly difficult and time consuming to achieve using highly directional antennas; consequently, we follow the common procedure to use an omni-directional channel to pass control information required to establish the directional network. Robustness of the omni control channel is achieved by minimizing the quantity of control data and frequency of exchange with concurrent maximal waveform spreading and processing gain. The dynamics of node motion dictate the rate of control information exchange. Ideally, the communication range of the omni control channel should be about equal to that of the directional links. This condition can be arranged by increasing the rate on the directive link until the link power budgets on both links are comparable.

Network dynamics require protocol adaptation at both the link and routing layers. Adaptive routing can find alternate paths to destinations as link topology changes. We have chosen the combination of an adaptive TDMA protocol for link scheduling and a proactive link-state protocol for routing. Adaptability to highly varying traffic demands is achieved in the TDMA link schedule by pooling the slots in the TDMA frame into a small number of semi-permanent slots and a larger group of on-demand slots. Generally, the number of semi-permanent slots will equal the number of neighbor nodes and they will be used to maintain solid links of minimal capacity to those nodes. The on-demand slots will be dynamically reallocated, as required to meet existing traffic demands, by adding to the capacity of links initiated with semi-permanent slots. For example, the pool could be allocated to achieve equal capacity among all links or to achieve dominant capacity on one of the links. The ability to achieve the latter condition improves as the frame size (number of slots per frame) increases at the cost of increased frame-to-frame latency.

TECHNICAL APPROACH

The directional antenna scheduling problem can be represented as a graph coloring problem. A related graph coloring formulation was used in a paper by Ma and Lloyd [2] to represent the multihop packet radio problem where nodes used omni-directional antennas. In that paper the omni antenna constraint leads to a solution whereby the algorithm must color the nodes (i.e., assign time slots to the nodes for transmission since a transmission from a node will be heard by all of its neighbors). The coloring or time slots must be assigned such that no two nodes are assigned the same color if they are distance-2 neighbors.

The difference in problem formulation between the Ma and Lloyd paper and the directional antenna scheduling problem is the following. If we make the assumption that the antennas have either very narrow or zero beamwidth, then a transmission by a node will be received only by the neighbor node to which it is transmitting (this assumption will be relaxed later). This allows a less restrictive set of constraints on the coloring solution. In this case the links will be labeled with colors, which represent a set of Tx and Rx time slots assigned to the link. The new constraint is that *no node may have more than one link labeled with the same color*. This implies that a color assigned by a node to one of its links is constrained by the previous colors assigned by that node to its other links and the colors previously assigned by that neighbor node to each of its links. Thus, nodes may determine schedules through a cooperative decision process with their one-hop neighbors.

The graph coloring problem formulation for phased array networking is illustrated in Figure 1, which shows a network with neighbor connectivity and a graph coloring solution. The constraints listed above are satisfied in this graph in which the links are colored such that no node uses the same color on two of its links (colors are represented by the number labels on the links in the figure).

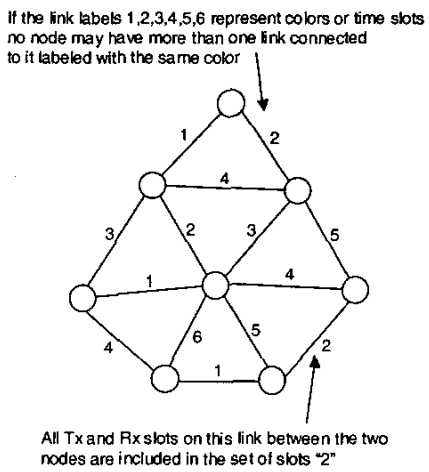

Figure 1. Example network for directional antenna graph coloring analog.

In this example, the graph could be labeled with the minimum possible number of colors (6) because the network was static and a centralized algorithm was used for labeling. This allows a TDMA epoch length of 6 time slots to be used. In solving the real problem, we would like a distributed algorithm solution that can rapidly adapt time slot assignments to changing network topologies with a minimum number of nodes involved in time slot assignment. The approach developed will require only 2 nodes to cooperate in the assignment of time slots. It will also require that the epoch be larger than the minimum possible epoch in order to allow this simple distributed solution. Assume that we allow no more than N directional neighbors at each node. Then any pair of nodes with $(N-1)$ time slots currently assigned neighbor nodes can find a time slot to assign for a new directional link between them as long as the epoch length satisfies

$$2 \cdot N - 1 \leq Epoch_length .$$

This can be done because at most only $(2N-2)$ different time slots were already assigned by the 2 nodes prior to this. If N is assumed to be the maximum number of neighbor nodes that any node might have, then an epoch length satisfying this relation will allow assignment of time slots based on the cooperation of just the two nodes. This is an ideal bound based on very narrow antenna beams not constrained by interference. Several practical issues force us to make the actual epoch length much larger. These include:

- Use of broader beam directional antennas produce interference constraints on allowable scheduling of time slots.
- Time-varying unbalanced traffic loads require the assignment of additional time slots on some links to meet variations in traffic demand.
- The ability to flexibly reallocate time slots is required due to the changing environment caused by node mobility.

The directional link will be operated with TDMA access with dynamic time slot assignment to network node pairs for communication between them as illustrated in Figure 2. The directional link epoch framing will be slaved to a specified Time-of-Day (TOD). The epoch and time slot length are configurable according to application. Nominal numbers for the epoch and time slot time allocations are time slots of 4 ms with 25 time slots per epoch for an epoch length of 100 ms. A node pair is considered to have a directional link between them if they have at least one directional time slot assigned.

Figure 2 illustrates a representative time slot assignment for a directional channel with an epoch of n time slots. In the figure we show use in a half-duplex mode between nodes A and B. The first minislot is used for transmissions from A to B, and the link is reversed in the second minislot. Each minislot also contains a guard time specified to allow for propagation delay uncertainty and other known delays.

The separate omni control channel provides a mechanism for neighbor nodes to discover each other and to exchange network control information without knowledge of the location of the neighbor node and without having to coordinate the information exchange with that neighbor node. Each node is allocated a transmit time slot in the omni epoch. During all other omni time slots, the node is listening for omni transmissions from its neighbor nodes.

There are two types of information that must be exchanged over the omni control channel:

- Node/link status information in Link_HELLO messages is used to inform neighbor nodes of information needed to understand how to best allocate time slots in the directional channel. This list of information includes node ID, location, list of neighbor nodes and the directional time slots used with those neighbors.

- Directional channel allocation control messages

Any neighbor node with which a bi-directional omni link is established is a candidate for connection of a directional link via scheduling a directional time slot.

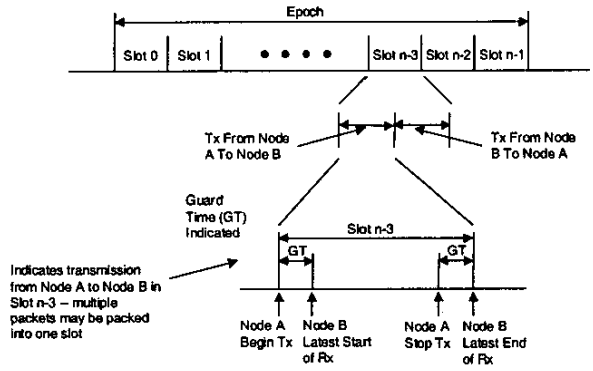

Figure 2. Time slot assignments for the TDMA scheduling of the directional links.

The top-level Concept of Operations (CONOPS) is the following. The omni overhead channel enables neighbor discovery and entry of a new node into the network. A new node entering the network hears the Link_HELLO messages from nearby nodes. It then determines a set of neighbor(s) with which it wishes to establish directional link(s). A series of time slot allocation control messages are exchanged over the omni control channel between this node and a new potential neighbor node to allocate the first directional time slot for the directional link between them. The location information provided in the Link_HELLO messages is necessary for the directional antenna pointing and for avoiding interference from other nodes that may be using the same time slot. The initial time slot assigned to a directional link is referred to as a Semi-Permanent (SP) time slot. This time slot assignment is maintained for as long as possible until it can no longer be used reliably. In addition to the single SP time slot assigned to each link, one or more Demand-Assigned (DA) time slots may be assigned to any link based on the expected traffic demand. DA time slots may also be released if sufficiently underutilized. Sets of message types have been defined for supporting the protocol for both SP and DA time slot allocation.

Some additional observations can be made about this process:

- Selection of the SP time slot allocation between a pair of neighbor nodes requires only coordination between the pair of neighbor nodes. The node requesting the link will send to the neighbor the prioritized list of its time slot choices. The neighbor can then select from this list the time slot it prefers using a specific ranking algorithm and return a reply with this selection. This provides a straightforward, fully distributed algorithm for scheduling the semi-permanent time slots.
- DA time slots will be assigned to best accommodate the traffic load. A node needs only to coordinate the allocation of a DA time slot for a directional link to a neighbor with that neighbor. It will send a request to the neighbor for the time slot assignment and receive either a grant of the assignment or a denial of the request.
- A node requesting a DA time slot allocation from a neighbor will do so based upon a perceived need for additional capacity on that link. This may be prompted by a high link utilization (queue buildup) based on short and long term measurements. The request will contain the number of slots requested and a metric, which indicates the priority to be attached to the request and a prioritized list of its time slot choices. The neighbor selects the time slot(s) it prefers from this list using a specific ranking algorithm and returns a reply.
- A node may have multiple concurrent requests for SP and DA capacity assignments outstanding at any time.

For both SP and DA allocation a set of messages are defined. A REQ message enables a node to request a time slot assignment from a neighbor node. This message includes a list of acceptable candidate time slots, and the neighbor must either accept one of the time slots or reject the request. The acceptance or rejection is communicated via the REPLY message. A CONFIRM message is then returned by the original node indicating that it received the REPLY message. A DELETE_TS message is used to tell a neighbor node that the time slot will no longer be used.

If a link is lost, then the MAC queues for that link are flushed of all packet data for that neighbor node. Additional messages will be needed to add certain features. These have not been completely defined at this point but will include the messages required for rescheduling time slots to avoid and mitigate interference detected or predicted on directional time slots.

The link protocol is agnostic about which routing protocol should be used. In this work we are using the Optimized Link State Routing (OLSR) routing protocol [1]. It can be used with almost no change by making the link layer appear to be a broadcast channel to the network layer. Thus, OLSR will send out HELLO messages for neighbor discovery. These messages will

be sent by the link layer to all directional neighbor nodes. A potential neighbor node will not receive such a message until the link layer has competed neighbor discovery and assigned an SP time slot for the directional link. Only when this is done can OLSR discover that it has a new neighbor for the purposes of routing traffic. However, for all neighbors that are connected on the directional links, OLSR can perform full ad hoc routing and can respond quickly when topology changes occur.

One of the benefits of this TDMA approach to link scheduling is that it may provide better QoS performance than a contention-based MAC. In addition, we are providing priority queues at the link layer for the link to each neighbor node so that high priority traffic is not delayed if lower priority queued traffic is waiting for a transmission opportunity.

A key issue that is addressed with the link scheduling approach is interference avoidance and mitigation. Each node understands the geometry of its interference environment through information received from neighbor nodes over the omni control channel. For each neighbor node this includes the time slots and the directions during which the neighbor node transmits. This allows nodes to schedule time slots in a manner that avoids interference from neighbor node transmissions during initial scheduling. It also provides a mechanism for warning of impending new interference in a time slot due to node mobility. The link quality for each directional time slot is also monitored so that interference that is not avoided may be detected and remedied through rescheduling that time slot. The degree to which interference will impact performance and the amount of channel reuse that can be obtained will depend on the directional antenna design/capabilities. Very narrow beamwidths with a large amount of attenuation in the sidelobes eliminate much of the interference problem and can make a large degree of time slot reuse possible. In addition, these antenna characteristics make the interference avoidance and mitigation processing simpler. For cost reasons we will be field testing our protocol with relatively broad beam antennas (45 degree mainlobe beamwidth) which puts much more of a burden on the interference avoidance and mitigation approach to insure good performance.

PLANNED EXPERIMENTS AND TESTING

To support the development of the link scheduling and ad hoc networking protocols for networks with directional antennas, Harris has developed a detailed OPNET simulation model of these protocols and an environment in which to test them. This model is currently working and supports the full TCP/IP stack, the OLSR routing protocol, the distributed adaptive TDMA link scheduling protocol to support directional antennas, the RF transceiver pipeline, and an 8-sector directional antenna model. OPNET simulations are being performed to determine the degree of channel reuse achieved, the response to mobility induced interference and topology dynamics, and the response to traffic dynamics.

Figure 3 shows the result of a simulation to determine the time slot reuse factor for one scenario. In this test low load CBR traffic streams are sent between neighbor nodes, and another CBR stream is sent from the bottom node to the top node. Initially, the simulation starts up and each node obtains an SP time slot to each of its neighbor nodes within the 16 time slot epoch. The traffic generators are turned on at various times over the first 15 seconds of the simulation, and new DA time slots are assigned to nodes as the demand for them occurs. The time slot reuse factor is defined as the number of time slots assigned (both SP and DA time slots) to all nodes per epoch divided by the number of time slots in the epoch. The reuse of 1.6 to 1.7 was achieved. This is an indication of the degree to which the directional antenna combined with the link scheduling protocol can provide greater network throughput on a channel. This is a relatively simple antenna with a 45 degree beamwidth so interference mitigation is necessary in order to assign time slots in an intelligent manner to achieve reuse without scheduling conflicts.

Field tests will be conducted in late 2003 with 15 nodes to obtain quantitative measures of network performance under conditions of extreme mobility and terrain-induced link occlusions from foliage and buildings. The intent will be to create conditions that result in frequent changes in link and routing state and to measure the ability of the system to adapt. The tests will be conducted with reproducible data flows and quantitative post-experiment analysis using NRL's Nettion (http://nettion.pf.itd.nrl.navy.mil) suite of test tools. These tools will allow scripting of desired data flows and collection of performance statistics. In addition, we developed an ancillary post-experiment analysis tool called DAZLE that combines the collected performance statistics and produces extensive performance reports. These reports include send rate, average receive rate, percent successful packet delivery, and latency histogram for each flow. Mobility will be achieved by hosting each mobile node on an automobile, which will follow a prescribed test course defined by a set of visually marked waypoints and a predetermined advancement schedule. To insure reproducible

experiments, VHF voice radios will be used to coordinate node advancement to the next waypoint.

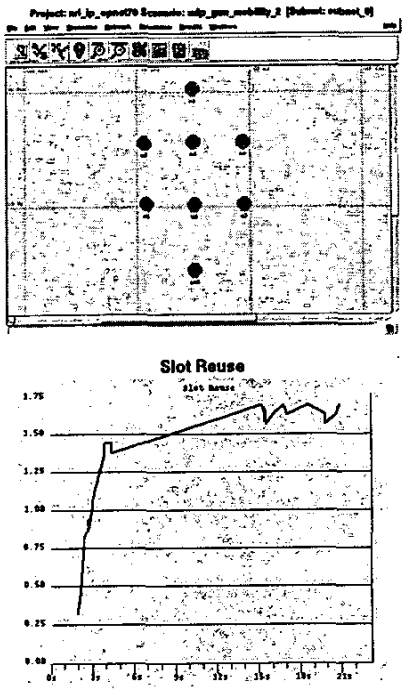

Figure 3. OPNET simulation demonstrating time slot reuse factor at low loads.

A block diagram of the field demonstration system implementation is illustrated in Figure 4. This implementation will require a PC running a real-time Linux operating system in order to support certain link layer functions that might otherwise be supported in hardware in a final configuration of a deployed system. Because of cost and availability, the demo system uses 802.11b radios operating in the 2.4 GHz ISM band. The 802.11b radio MAC firmware (Prism 2.5) is modified to allow it to be used as a burst TDMA modem with all contention-based channel access mechanisms turned off. For implementation convenience and to eliminate the cosite interference problem, a single radio is used for both the omni-directional control channel and for the high rate mission data transmitted over the directional antenna. The TDMA epochs for the two channels are overlayed so that only one channel is active in any time slot. The omni transmissions are deterministic within the time slot schedule thereby providing a logical omni channel. Timing for the epochs is generated from the 1 PPS pulse from a GPS device at each node. An 11 Mb/s data rate will be used on both channels for the field test and attenuators will be used if necessary to balance the omni and directional links and obtain the same effective range for both links. The Power Amplifier (PA) is 1 W. Selection of the omni or directional antenna (and the sector of the antenna) is via RF Switches (SW) controlled from the laptop according to the TDMA schedule established by the distributed TDMA link scheduling protocol

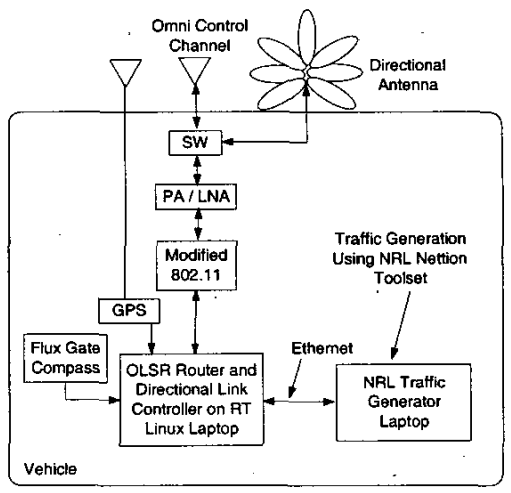

Figure 4. Node block diagram for Harris/NRL testing of networking with directional antennas.

SUMMARY

This paper has described a new distributed, adaptive Time Division Multiple Access (TDMA) link scheduling protocol, which determines schedules based on cooperative decisions between each pair of neighbor nodes. This protocol was developed to take advantage of high-gain directional antennas to allow greater channel reuse and higher network throughputs. The field tests will use relatively broad antenna beamwidths (45 degrees) which is a more challenging environment for the link protocol. Much narrower directional antenna beamwidths and even nulling of interferers can produce much better performance and potentially much higher network throughput.

REFERENCES

[1] T. Clausen et al, "Optimized Link State Routing Protocol," IETF MANET Working Group Internet Draft, draft-ietf-manet-olsr-06.txt, September 1, 2001.

[2] X. Ma and E.L. Lloyd "Adaptive Algorithms for Broadcast Scheduling in Multihop Packet Radio Networks".