53-32
N93-17513
P-13
DS 8533217

A Simple, Effective Media Access Protocol System for Integrated, High Data Rate Networks

by

E.C. Foudriat, K. Maly, C.M. Overstreet, S. Khanna, L. Zhang
Department of Computer Science
Old Dominion University
Norfolk, VA 23529
foudr_e@cs.odu.edu

Submitted to the Journal on Selected Areas in Communications

Abstract

The paper describes the operation and performance of a dual media access protocol for integrated, gigabit networks. Unlike other dual protocols, each protocol supports a different class of traffic. The Carrier Sensed Multiple Access - Ring Network (CSMA/RN) protocol and the Circulating Reservation Packet (CRP) protocol support asynchronous and synchronous traffic, respectively. The two protocols operate with minimal impact upon each other. Performance information presented in the paper demonstrate that the they support a complete range of integrated traffic loads, do not require call setup/termination or a special node for synchronous traffic control, and provide effective pre-use and recovery, The CRP also provides guaranteed access and fairness control for the asynchronous system. The paper demonstrates that the CSMA-CRP system fulfills many of the requirements for gigabit LAN-MAN networks most effectively and simply. To accomplish this, CSMA-CRP features are compared against similar ring and bus systems, such as Cambridge Fast Ring, Metaring, Cyclic Reservation Multiple Access and DQDB.

I. Introduction

The direction for networking has been to increase basic network data rates toward the gigabit range. For gigabit networks, the most of the recent development is concentrated on circuit-based, packet networks like ATM which are applicable generally to WAN and MAN systems [1]. Since circuits and switching add to the complexity of ATM based systems, there is a place, especially in the LAN - MAN area for fixed connectivity networks based upon a ring-and bus architecture. Here, no connection needs to be established and information can flow from any-to-any/many simply by access to the network and insuring uncorrupted data to the receiver(s).

Gigabit ring and bus network media access protocols have been described in recent literature. These include Metaring [2-4]; the Cyclic-Reservation Multiple-Reservation system [5-8], most recently CRMA-II; the Cambridge Ring [9, 10]; DQDB [11]; and the Carrier Sensed Multiple Access - Circulating Reservation Packet system, CSMA/CRP, [12, 13]. The CSMA - CRP protocol is a dual, non-interfering protocol which supports both asynchronous and synchronous traffic. It is the intent of this paper to describe the CSMA-CRP system and its performance and then to compare its features with the other systems above based upon their published results.

There are a number of requirements for high data rate network systems. Since data is arriving at the node at 1000 bits/msec., logic decisions related with network access must be high speed. Thus, a major requirement for high data rate media access is that the protocol be simple, i.e., it leads to simple logic circuitry. If complex data formulation and logic operations are required they should be to a large extend pre-computable, and the event times to which the operations relate be highly predictable. This permits these complex operations to be executed quickly based upon boolean indications of status.

A second major requirement is for these high data rates systems to handle integrated traffic efficiently. In the future, high data rate networks will be the backbone upon which distributed multimedia operations are supported. At times, much of the network load may be voice/video traffic while at others almost all of the traffic may be asynchronous packets [16]. Also, load can vary from full load to no load over short periods of time. Therefore, the protocol must support highly dynamic load and load type ranges without renegotiation of its operational parameters.

A third requirement is for these systems to support maximum capacity since even for simple protocols, equipment will be expensive. For example, dual counter rotat-

[1]Research support has been provided by DARPA, N00174-C-91-0119. NASA. Langley Research Center grant NAG-1-87263 and Virginia Center for Innovative Technology grant 596045.

ing ring configurations can, in some cases, provide as much as 8X the capacity with the same equipment count a dual bus configuration [2]. Also, within the scope of effective equipment utilization, the network structure and operations should be easily reconfigured in case of node or link failure. This feature provides added reliability.

In this paper, we describe a system, CSMA - CRP, with most adequately supports these requirements. The next section describes the system operation. This is followed by development of an analytical model and then a description of system performance based upon both analysis and simulation modelling. In the remaining sections, the operational and performance features of CSMA-CRP are compared to the four other systems listed above based upon descriptions and performance information found in the open literature. To accomplish this task each system's operation is described briefly. This is followed by a comparison of capabilities related to the features discussed above.

II. The Carrier Sensed Multiple Access - Circulating Reservation Packets Media (CSMA-CRP) Access Protocol

The CSMA-CRP system is a dual media access protocol for high data rate networks. Unlike other dual systems which support asynchronous traffic over different load ranges, the CSMA and CRP protocols support asynchronous and synchronous traffic, respectively. The following sections describe each protocol.

A. Asynchronous Protocol System

The CSMA/RN system has been described in reference [12]. Figure 1 shows the basic logic operation which occurs. The physical level media access logic interface can be "almost" all optical since it conforms closely to the logic operations in reference [17]. The logic must make the decision as to whether the incoming data or data queued at the node should be transmitted on the outgoing line. Queued data may be transmitted under the condition that incoming block is idle or that the incoming block is destined for this node, i.e., it uses destination removal (spatial reuse). At low data rates, the idle condition is equivalent to the fact that the total network is free, as for example in the CSMA/CD Ethernet protocol. At high data rates, like 1 Gbps, the access conditions are local to a node, which allows simultaneous access by multiple, geographically separated nodes. Most high data rate protocols [2, 5, 9, 12, 14] including some of those listed above use simultaneous access and destination removal.

Destination removal has demonstrated that it can significantly enhance peak performance for high data rate networks. In a single ring under uniform load conditions, the nominal capacity of a network is twice the basic data rate and for dual counter rotating rings capacity can increase to 8 times. The CSMA-CRP system is designed so that it can use this feature. It will be shown later in the

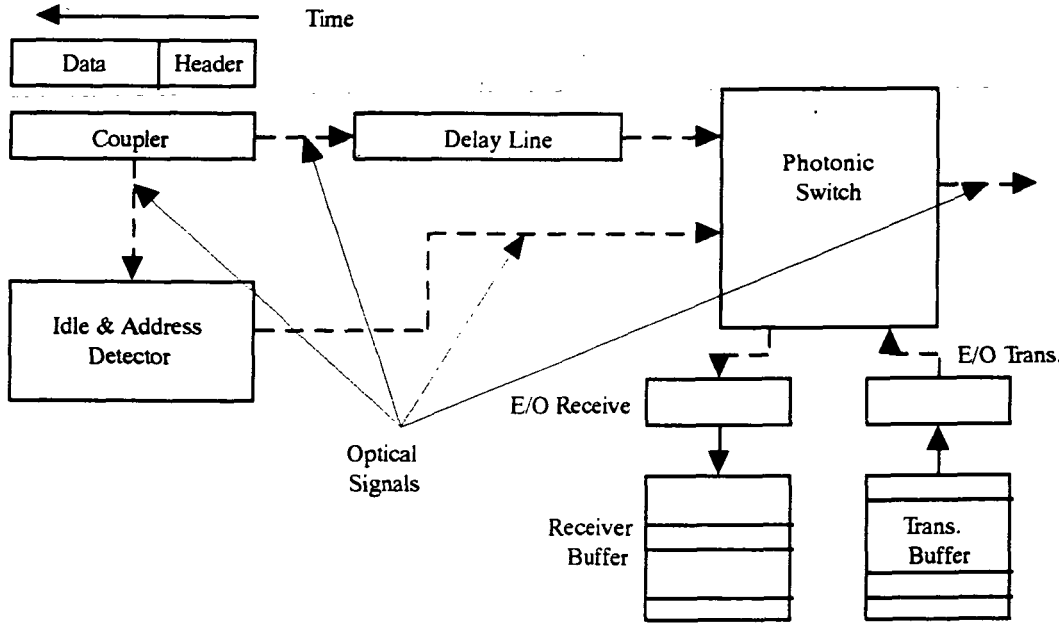

Figure 1 Node Inteface Logic for CSMA-CRP System

paper, that destination removal also influences fairness access performance at high load factors.

CSMA/RN can be viewed as a slotted system with a slot length of 1 bit. If a node is transmitting data during a period when the incoming traffic changes from idle to data, the conflict is alleviated by truncating the node's transmission of its data. Under this operation all packets on the network have preemptive rights. This decision logic requires that the delay line for incoming packets be long enough so that the truncation operation can take place, i.e., in the range of 100 bits. The delay line is not operated as an insertion buffer as in some access protocols but only so that the logic operations have sufficient time to complete. Hence, the delay serves only to increase the total ring length. The truncation operation does not require breaking messages into fixed slot lengths as in slotted systems, but at higher loads and for larger messages, fracture does occur.

Figure 2 shows the CSMA/RN packet structure. Information in the header provides for destination removal or source removal for multicast and broadcast messages. Control information is used to signal request for acknowledge, to indicate network control packets and to enable removal control for packets where address is corrupted. Trailer information assists a receiver in reassembly since it provides message initiation and termination information as well as providing for error detection using a check sum. A similar packet structure is used for the CRPs.

B. Synchronous Protocol System

Synchronous access uses the Circulating Reservation Packet (CRP) system. Its operation is described in reference [13]. It uses a special, small packet which circulates continuously around the ring. Synchronous traffic is attached behind this packet. On the cycle prior to a node's need to submit synchronous traffic, the node using the CRP, requests a block of space following the packet be freed upon its return. When the CRP returns, the node frees the packet and places the traffic in the free space behind the packet. Multiple CRP can be used on a net and they can be assigned in a round robin fashion or used in a free access mode. The separation of circulating reservation packets will limit the size of regular data packets as well as the size of synchronous message block a node can send at one time so the number and location of reservation packets on the system should be regulated. Since CRPs circulate without interference or delay, there occurrence at a node is strictly periodic and completely predictable.

As the CRP request circulates around the network, the space behind the CRP is freed as the packets arrive at their designated destination. *Pre-use*[13, i-] is available since the space being freed can be used by a node which has a message for another node between and up to the node that set the request. Recovery of synchronous traffic space occurs automatically since the node controlling the circulating packet does not request space to a greater extent than needed. For example, a telephone call which goes into a "non-talk" mode requires no or only a minimal reserved space during this phase. In addition, there is no need for a node having new synchronous traffic to establish a *call*, since when the node gets its CRP, it can request additional reserved space within the limit of the space available following the CRP.

In addition to handling synchronous traffic, CRPs support guaranteed access. Access is guaranteed one revolution after receipt of a free CRP. To guarantee that nodes have access to free packets requires that a node cannot use a packet on a number of successive revolutions and/or that nodes are assigned CRP use in a round robin fashion. Fairness control operates in a similar manner. Since CRP circulation is strictly periodic and CRP use for synchronous traffic highly regular, the node can easily predict when access will occur. Finally, CRPs can be used for guaranteed time to acknowledge. Since the node sending the acknowledge request can predict the arrival of the packet at the destination node, the time the next CRP arrives at the destination node and the arrival of the returned acknowledge packet, it knows how long it must store the data to replace a lost packet.

The CRP operate without interference or interfering with the CSMA/RN asynchronous access operations. To the CSMA/RN system, it is just another packet since all information on the ring has preemptive rights. When the CRP is set with a request, each node must monitor the block requested. With regular use, the node knows when a CRP is expected to arrive and the node using it. Hence, it can prepare its traffic to take advantage of *pre-use* effectively.

The CRP system causes minimal impact on asynchronous access. It should be noted because of this minimal

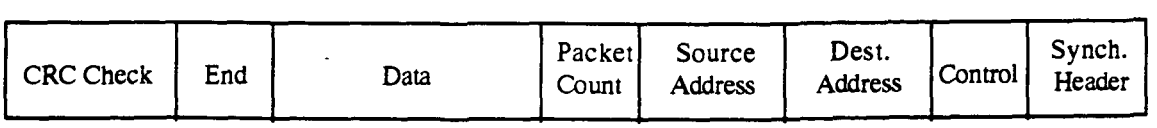

Figure 2 CSMA/RN Packet Structure

impact, the CRP protocol concept is applicable to other networks besides CSMA/RN. These include slotted rings and dual bus configurations if the end stations transfer CRPs from one unidirectional bus to the other.

III. CSMA - CRP Access Performance

A. Analytical Model

An analytical model for describing the CSMA/RN performance was developed in reference [12]. In that paper it was noted that a number of models including those involving contention over the total ring length were evaluated. Since the capability for node access is based upon local conditions, i.e., the condition of the immediate upstream arriving information, and since once placed, the message can not be preempted, a simple priority queuing model suffices.

The model developed has been extended to cover the conditions where load factors are greater than the service rate, i.e., the condition in reference [13, 15] where r > 1. The equation for arrival is:

$$Pr(x = \text{available} = p \quad \begin{matrix} = 1 - lf + lf/(n-1) & lf < 1 \\ = 1/(n-1) & lf > 1 \end{matrix} \quad (1)$$

where lf = load factor;
n = number of nodes; and

where an available packet is describe as one which is either empty or one whose destination is the node under consideration. Note that lf > 1 constitutes a network completely filled with packets.

The remainder of the analysis follows that of reference [12]. Figure 3 compares the calculated and simulator service times for the conditions in equation (1). Note that for lf > 1, the service time, the time it takes a message to be sent once reaches the head of the queue, approaches a constant. This occurs with uniform destination address and destination removal because each node receives 1/(n-1) of the packets on the network. Even though the service time is finite, the wait and response times for this condition are unstable, as is evident based upon the Pollaczek-Kintchine formula [15] and the equation for response time in reference [12].

B. A Question of Fairness

Under the conditions of multiple access and destination removal, if destination addresses are uniformly distributed then node starvation *does not* occur. Further, if each node generates messages in proportion to those it receives then the network is eminently fair, i.e., all nodes will have identical wait time conditions. These conditions exist even if the input message rate exceeds the network capacity. While this can be easily proved, a better demonstration was obtained from the simulator. A file server system condition was modelled. In the worst case, all other nodes (with identical arrival rates) send data to the file server and the file server send to all other nodes uniformly. The load conditions were increased to where the network data rates exceeded capacity. All nodes still had equal wait and service times.

This creates an interesting question with regard to high data rate network design using multiple access and destination removal. Node starvation can occur only when two conditions exist simultaneously. They are: (1) the network is heavily loaded, i.e., greater then 75% capacity and (2) one or more nodes are sent a significantly reduced message count. Since the simultaneous occurrence of this

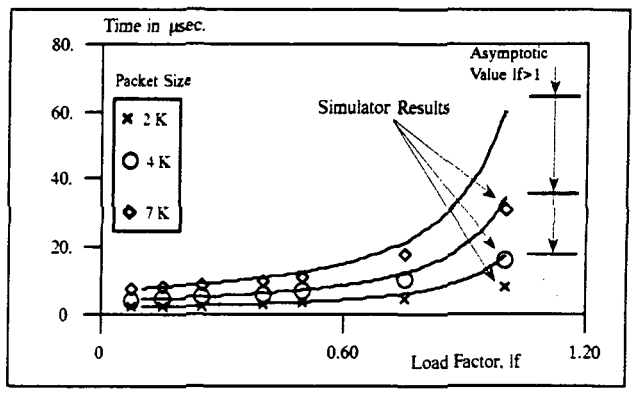

Figure 3 Service TimeComparison Between Analytical and Simulation Models
Conditions - Data Rate = 1 Gbps
Node Count = 10
Ring Length = 10 km

condition is probably rare, to what extent should high data rate networks go in order to provide access fairness? If, for example,the fairness system reduces capacity by 25%, it may be doing more harm than good. If the fairness system is complex, is it worth the cost since it is rarely needed?

C. CSMA/RN Performance

The simulation determined CSMA/RN performance under a range of system parameters [12]. Three conditions of major interest are message length, ring length and node count. Runs were made to examine the effects of these parameters. Figures 4a - 4d present data for a 1 Gbps, 10 km, 10 node ring for message lengths ranging from 2K to

20 Kbits. Note that load fraction is shown as a function of basic network bit rate and that because of destination removal the capacity of the network is actually 2 Gbps. We see from Figures 4 that average performance characteristics for CSMA/RN are not detrimentally altered by message length. Both wait and service time are similar to that predicted by the analysis. Mean response time is greater for the larger packets, primarily because service time is greater. Finally average message fracture ratio does not change significantly as packet length increases, indicating that message fracturing does not materially increase, at least, when all messages are of the same size and nodes are uniformly distributed around the ring.

A similar set of curves is plotted in Figure 5 to show the effect of ring lengths from 2 km to 1000 km. Only response time is shown because the other performance measures are only minimally affected. It shows significant dependence upon ring length, mainly due to the travel time necessary from source to destination. For longer length rings, travel time dominates, so service and wait time become less significant. However, the latter two are the only controllable factors in the response performance. The CSMA/RN protocol provides excellent operations over a range of LAN and MAN conditions. Additional information provided in reference [12] demonstrates that the protocol is applicable to WANs as well.

Figures 6a - 6b show the simulation results when node count is varied from 10 to 200 nodes for a 50 km ring; node spacing range from 0.25 km to 5 km. Service time and packet fracture ratio are shown because at the higher node count they are the factors most affected. At the larger node count and high load factor, both service time and message fracture show a definite increase. Under these conditions, the CSMA/RN protocol would have it worst operational problems as the packets on the ring would have the greatest tendency to fracture. However, performance is completely acceptable for loads up to 150%.

Reference [12] also included performance information for data rates from 100 Mbps to 1 Gbps. At the lowest speed, short length networks, for example, 2K, tend to suffer because multiple simultaneous access no longer exists. At data rates above 1 Gbps CSMA/RN performance only gets better. Also, it shows that CSMA/RN can handle overloads without significant capacity decrease, and recovery, after load is reduced, depends upon queue buildup during overload.

CSMA/RN provides a number of excellent performance features for supporting asynchronous network traffic including:

1. immediate access for an idle network since there is no waiting for a token or slot;
2. no automatic message breakup due to slotting;

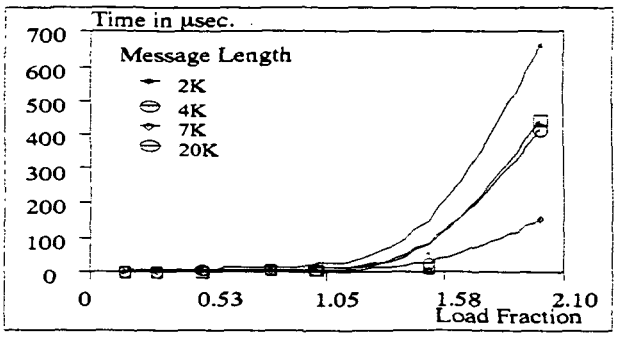

a) Wait Time

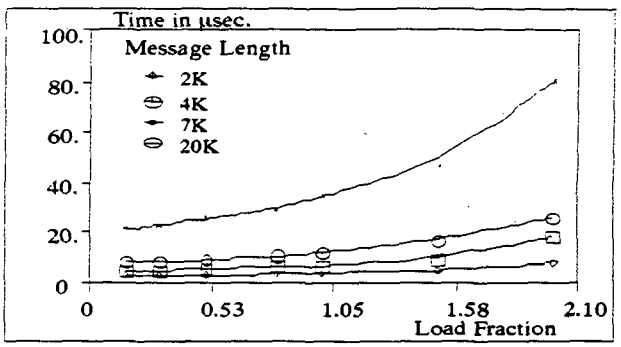

b) Service Time

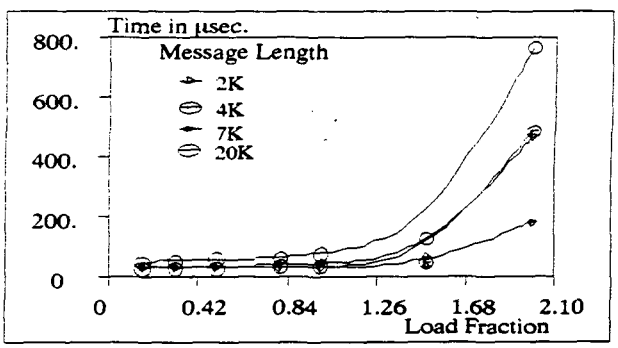

c) Response Time

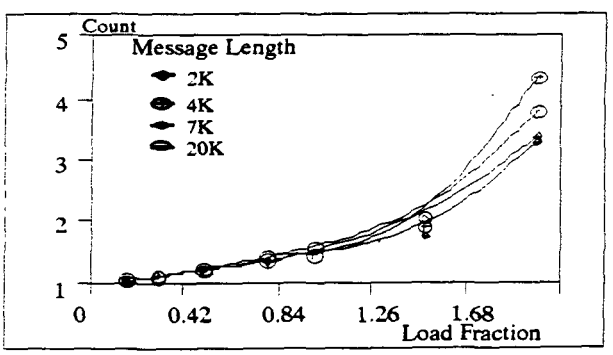

d) Message Fracture

Figure 4. CSMA/RN Performance for Various Message Length
Conditions 10 km, 10 nodes
1 Gbps

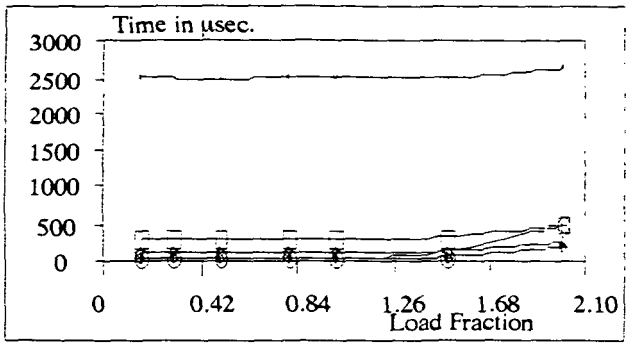

Figure 5 Response Time For Ring Lengths of 2, 10, 50, 200 & 1000 km

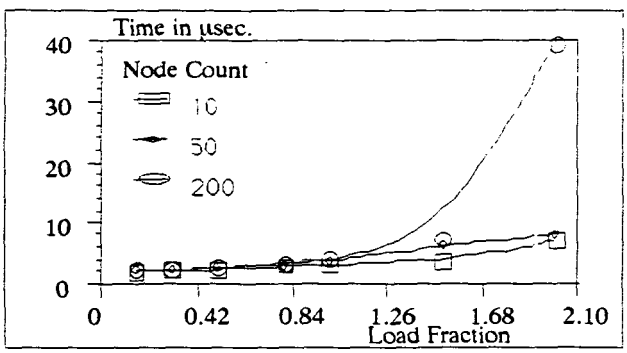

a) Service Time

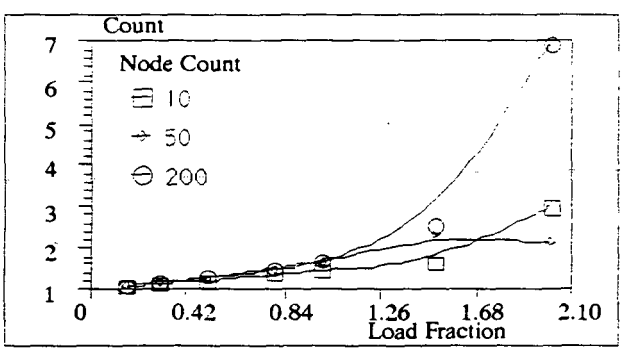

b) Fracture Ratio

Figure 6 Network Performance for Node Counts
Conditions - Data Rate = 1 Gbps
Ring Length = 50 km
Packet Size = 2 Kbits

3. no wasted capacity at high loads due to collisions which abort transmissions;
4. applicable to a wide range of network conditions including:

lengths	10 - 5000 km;
nodes	10 - 1000;
data rates	>100 Mbps; and

5. simple node operations since nodes react to local traffic conditions only.

D. CSMA/RN Performance with CRP Supporting Integrated Traffic

1. CRP Implementation

CRP performance in its own right, is fully deterministic, since CRP travel is completely predictable. Here, synchronous traffic is served strictly periodically and both the CRP and its attached traffic are not affected in any way by the CSMA/RN operations. Any study of the combined protocols need only determine the effects of the CRP system and its synchronous traffic on CSMA/RN operations.

In the simulator, the CRP system and synchronous traffic conditions were established at initialization time. CRPs (one or more depending upon run conditions) were assigned in a simple round robin fashion. They were spaced at equal distances around the network and nodes were assigned a single CRP. Assignments were interlaced, e.g., for 3 CRPs, CRP #1 was assigned nodes 1, 4, •••; CRP #2 nodes 2, 5, •••, etc. No complex assignments were used to mitigate particular load assignments or alleviate starvation and no sharing of CRPs between adjacent nodes was implemented. When the CRP arrives at the node for which it is scheduled, synchronous traffic is checked, and the total space needed is calculated and loaded in to the CRP data structure. When the CRP arrives back at the starting node, the synchronous traffic is placed using a last-out, first-in structure, i.e., messages traveling farthest are placed first in the free space. This format is used because as messages are removed, the free space available may be contiguous and better suited for CSMA/RN operations.

Interference effects were studied for fixed synchronous traffic for each run. Except for fairness runs, source and destination for both synchronous and asynchronous traffic was assigned uniformly. Asynchronous traffic arrival was Poisson. For synchronous traffic, two message types, standard TV and digital telephone, were used; the former requiring a bandwidth per call of 10 Mbps and the latter 64 Kbps. TV blocks were sent at a minimum interval of 1 msec. (>10,000 bits per access); telephone at a 10 msec. (>64 bits per access). All bits accumulated since the last CRP arrival were sent. Synchronous traffic mix could be varied between all TV or all telephone, but for most runs TV and telephone traffic load was 90 % and 10%, respectively.

After fixing synchronous traffic load, asynchronous traffic was varied over the remaining range. Run combinations used fixed synchronous traffic load while asynchronous traffic varied over the remaining range. For example, if synchronous traffic was 25%, i.e., 250 Mbps, then

Table I. Total Load Bit Rate

in BPS x 10 -6

Async. Load Percent	Synchronous Load Percent				
	0%	10%	25%	50%	75%
0%	0	100	250	500	750
50%	500	550	625	750	875
100%	1000	1000	1000	1000	1000
150%	1500	1450	1375	1250	1125

asynchronous traffic loads could range from 0 to 1500 Mbps for a network with 1 Gbps data rate. Under these conditions and asynchronous load of 750 Mbps in conjunction with the 250 Mbps synchronous traffic is considered a 100% load. Table I shows bit loads and splits under various load conditions.

2. CSMA Performance Under CRP Operational Conditions

Table II illustrates the CRP access time capabilities for some typical LAN - MAN network configurations. By selecting the number of CRP based upon the ring length and the number of nodes, a maximum access time of 1 msec is feasible in cases where the nominal ring travel time is somewhat less than 1 msec. For WAN length rings, guaranteed access after arrival of the CRP below 1 msec. can not be obtained. However, with multiple CRPs, their arrival can be made to occurs at less than 1 msec. intervals.

Table II. Interarrival Time for Multiple Use CRPs

Time in sec.

CRP Count	Node/Lng(km)				
	4/10	10/10	20/50	40/100	50/200
1	2.50E-04	5.50E-04	5.25E-03	2.05E-02	5.10E-02
2	1.50E-04	3.00E-04	2.75E-03	1.05E-02	2.60E-02
5	N/A	1.50E-04	1.25E-03	4.50E-03	1.10E-02
10	N/A	5.00E-05	7.50E-04	2.50E-03	6.00E-03
20	N/A	N/A	2.50E-04	1.50E-03	3.50E-03
50	N/A	N/A	N/A	N/A	1.00E-03

Here, nodes would have to use substitution or reserve extra space apriori if critical message time constraints exist.

Figure 7a - 7d show the wait time, service time, response time and fracture ratio for a LAN configuration of 10 km - 10 nodes using 2 CRPs. It can be concluded that CRPs, by themselves, have very little affect on CSMA/RN operations with the exception that the fracture ratio at low asynchronous loads is raised from 1.1 to 1.4. Here, the curves to compare are those with No CRPs and 2 CRPs with 0% synchronous traffic conditions.

Examination shows that asynchronous and synchronous traffic are very similar in their affect on CSMA/RN operations. As synchronous load increases from 0 to 75%, both wait and service time for low asynchronous traffic loads are increased, from 0 to 6 msec. and from 7 to 16.5 msec., respectively. This increase is nearly identical to that which occurs when asynchronous traffic load increases from 0% to 75% with no synchronous load.

Reference [13] discusses the affect of CRP count. Larger CRP counts actually improve network operations especially at high traffic loads. The reasons are that the network load is more evenly distributed since the CRPs tend to be in different stages at any particular time and that no extremely large reserved block of space is being requested which can shut off many nodes simultaneously for a long period of time. Here with minimal CRP count, a node receives the CRP at longer intervals, as shown in Table II, so it may request an extremely large block of space, effectively shutting down other nodes for a long period as the space is first cleared, then filled and finally cleared during CRP use cycle.

Figures 8 shows the adaptability of CRPs to MAN network conditions for a ring of 200 km and 40 nodes. Note that for low loads, minimal response time. At higher synchronous and asynchronous loads on the MAN, travel time still dominates, while in a LAN, the response time, as illustrated in Figures 7, is more dependent on protocol access changes.

It can be concluded that CRPs, do indeed, operate with minimal impact on the underlying asynchronous CSMA/RN access protocol and that the combined protocols provide an extremely effective system for supporting integrated traffic over a wide range of network conditions. The results illustrate that CRPs are effective over a wide range of network conditions and configurations.

3. Pre-Use Effectiveness

Pre-use is the capability for a node to use requested reserved space when it has traffic destined for nodes between itself and the reserving node. Pre-use was used for the runs presented in the previous figures. To illustrate

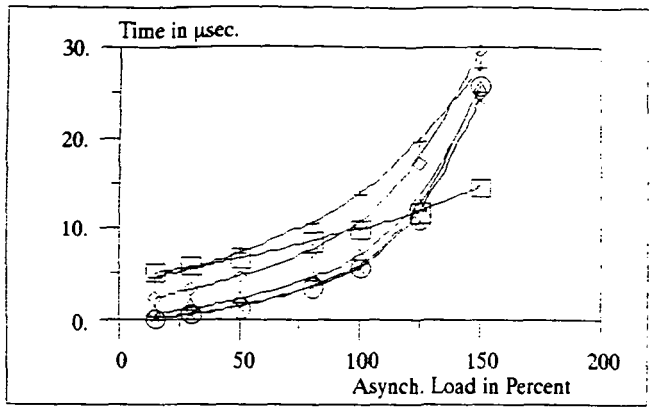

a) Wait Time

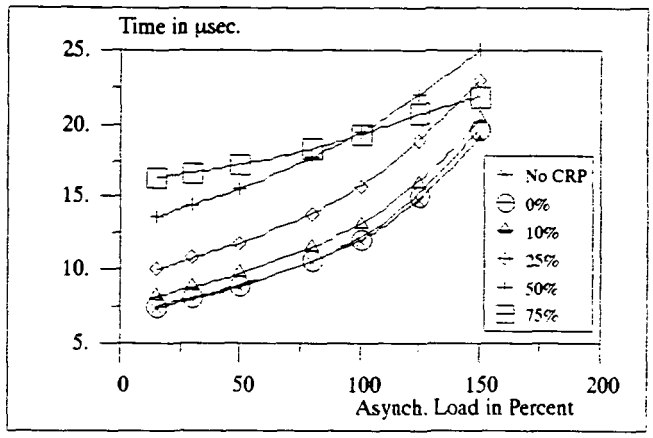

b) Service Time

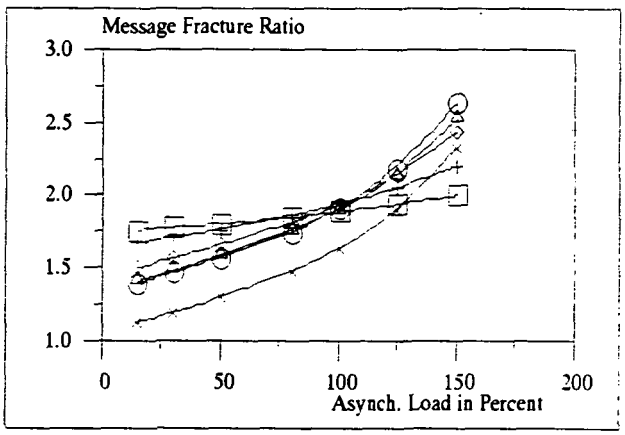

c) Message Fracture

Figure 7. Asynch. Message Performance for Dual Protocol System

Conditions: Ring Length - 10 Km; Node Count - 10; CRP Count - 2; Data Rate - 1 Gbps; Message Lngth - 7 K bits.

Note: Synchronous Load Key in Figure 7b).

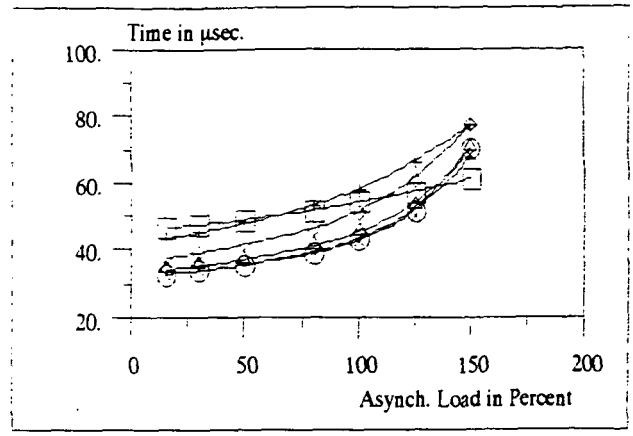

d) Response Time

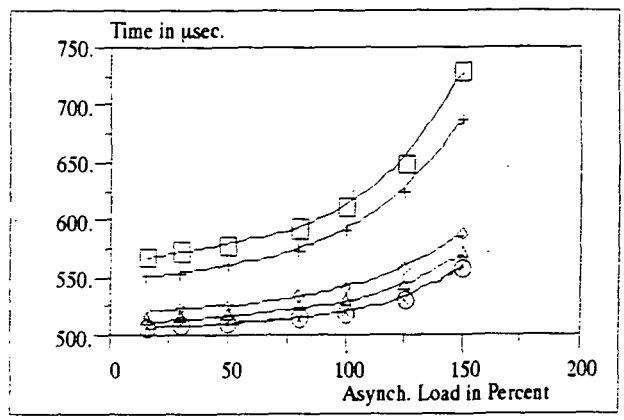

Figure 8. Dual Protocol Response Time for MAN Conditions

Conditions: Ring Length - 200 Km; Node Count - 40; CRP Count - 8; Data Rate - 1 Gbps.

Note: Synchronous Load Key in Figure 7b)

its effectiveness, runs were made for the LAN with a 50% synchronous load condition both with and without pre-use. The comparison, Figure 9, shows that pre-use significantly improves asynchronous access for operations especially under high load conditions. For larger synchronous loads conditions, pre-use is even more effective.

4. Fairness Using CRPs

As noted previously, CRPs can be used to support fairness and guaranteed access as well as synchronous traffic. Since unfairness situations in multiple access, destination removal, ring networks are undoubtedly complex, our approach in this paper is restricted to demonstrating the CRP's effectiveness, not developing a comprehensive fairness system. With this objective, the simulator

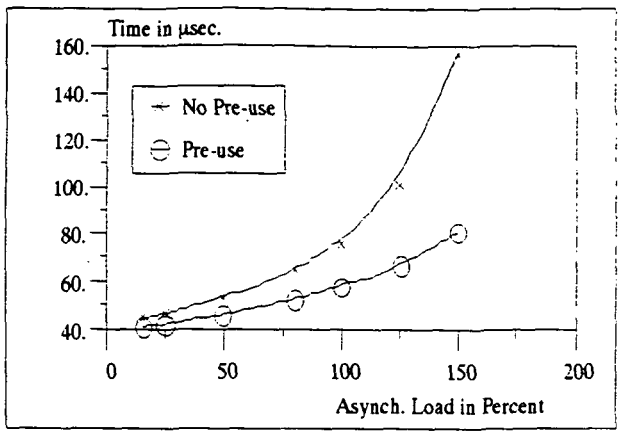

Figure 9. Wait Time Showing Pre-Use Effectiveness
Conditions: Ring Length - 50 Km; Node Count - 20; Data Rate - 1 Gbps; Synch Load - 50%; CRP Count - 5.

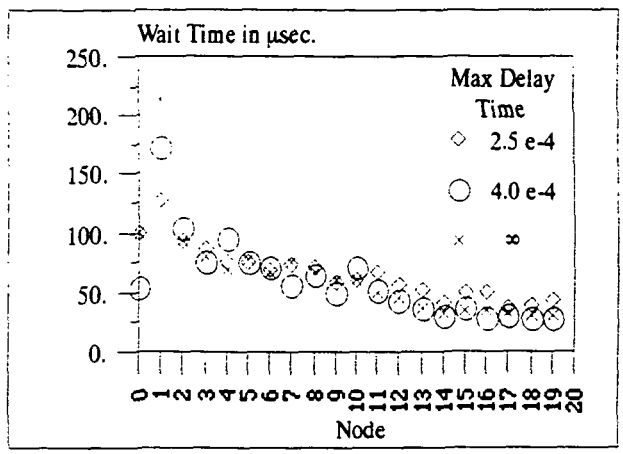

Figure 10. Mean Wait Time Reduction Using CRPs for Fairness Control
Conditions: Ring Length - 50 Km; Node Count - 20; Data Rate - 1 Gbps; CRP Count - 4 (fairness) - 4 (Synch Load); Synch. Load - 10%; Asynch. Load - 175%.

message generation system was modified to incorporate both non-uniform message arrival load and to effect node starvation through non-uniform destination selection. The first condition represents network loads typical of file server systems. As noted previously, no unfairness was observed. The second can create unfairness conditions because the model assumes uniform message arrival addresses but destination addressing reduces the reception of messages at designated nodes. For example, a starving node would receive only x% of the messages which it would normally receive under uniform loading. Messages not received by starving nodes were distributed to non-starving nodes uniformly. Starving nodes could be selected at any location, in any combination and at any percentage of starvation. However, for the runs used to show starvation and the CRP's effectiveness for fairness control, two adjacent nodes were selected to starve. Runs were made under varying load conditions and it was found that starvation was most severe for asynchronous loads above 150% when synchronous traffic load was small.

A simple fairness control scheme was used. Fairness CRPs where added to the synchronous traffic CRPs. They were used only to provide fairness access and were not used for synchronous traffic. Alternatively, synchronous traffic CRPs could not be used to alleviate starvation although in a realistic situation this is easily implemented. If a message waits for greater than Y time, the node uses the next free fairness CRP to request space to service this message. Since it takes the CRP one rotation before it can be used, there is the possibility that the delayed message will be serviced by the time the CRP returns with free space. When the CRP returned, the present message at the head of the queue is tested to see if its wait time interval is greater than Y, regardless of whether this is the message in the queue at the time that the CRP request was made. If so, this message is sent in the CRP space available.

Figure 10 shows node fairness control results for nodes 0 and 1 starving at 25% of their expected traffic. Node starvation affects not only the targeted nodes (in this case nodes 0 and 1), but also nodes downstream. Hence, the fairness control CRPs may be used by nodes other than those actually starving. Figure 10 shows that the fairness system using a delay time of Y = 250 μsec. significantly reduces the mean wait time for node 1 which is starving badly while the mean wait times for the non-starving nodes is not altered significantly. Runs at other network conditions and loads showed similar results using fairness CRPs.

IV. Comparison of CSMA-CRP with Similar Media Access Protocols

In this section we will briefly describe alternative media access protocols designed to operate in a manner similar to the CSMA-CRP system. After this description, we will compare the features of each.

A. Cambridge Fast Ring (CFR)

As originally conceived, the Cambridge Fast Ring (CFR) [9, 10] is a slotted ring with source removal. Access logic for asynchronous traffic uses a free/busy bit but since source removal is required logic decisions must be made on address also. Others have redesigned the ring to use destination removal [10]. Performance of asynchronous

slotted rings is well known [10].

Synchronous access is by one of two mechanisms - slot reservation or a cycle time scheme [10]. Slots in the former system are either normal or channel; channel slots are reused by the source but normal slots must be passed on. For the reservation system, nodes must capture sufficient channel slots for their synchronous traffic. In some configurations, slot rotation may not fit efficiently with synchronous traffic rates, so slots may not be efficiently used. Pre-use is possible but it is impossible to recover unused slots for messages in "silent" phase.

Cycle time, similar to an implicit token rotation time, establishes a time which is sufficient for nodes to send all their synchronous traffic plus some asynchronous traffic. Nodes cannot send newly arrived asynchronous traffic until a new round has started. Using cycle time leads to some difficult choices with regard to setting parameters. Poorly set cycle time parameters can lead to wasted capacity and changes in load may require renegotiation of parameters [10].

B. Metaring

The Metaring is being developed by IBM [2 - 4]. The ring is described is a dual, counter rotating, slotted ring with insertion buffers and destination removal. Nodes use the insertion buffers to avoid conflict between incoming and data being transmitted by the node.

Asynchronous traffic is controlled by a fairness algorithm using a SATisfied control packet. The SAT packet circulates in the opposite direction. Nodes between SAT arrivals can send up to a fixed amount of traffic. If a node, when it gets a SAT packet is not satisfied, it holds the packet until it is. SAT packets, which use uncoded data words, have preemptive priority so that a SAT packet must be serviced even within the middle of data packet operations. Results [2, 4] show that the SAT system can reduce peak capacity by about 12.5% but that significant fairness can be achieved for nodes which are starving. However, as noted previously true starvation can be a rare occurrence. While multiple SATs are mentioned it is not clear that multiple SATs can be used effectively in longer distance rings, without an additional mechanism for regulating SAT spacing.

Synchronous traffic access is supported by an ASYNC-EN packet with rotates in the opposite direction to the SAT packet. The ASYNC-EN system is equivalent to the timed token function in FDDI. The reservation protocol accompanying the system requires that a node send a request to all intervening nodes between source and destination indicating how many synchronous traffic data units it needs. If all nodes have less than a maximum synch traffic then the each intervening node adds that traffic to its reserve. Like the cycle time system above, synchronous traffic is transmitted before asynchronous traffic after receiving the ASYNC-EN packet. Normally, this packet circulate around the ring in at a fixed rate. Holding the ASYNC-EN packet, which is done by a node when its synchronous traffic queue builds up, signals nodes on the ring to stop sending their asynchronous traffic. No information is given showing either synchronous or integrated traffic performance in the open literature. To what extent delaying the ASYNC-EN packet throttles asynchronous packets is not clear but it would appear holding the token will reduce asynchronous access for at least that ring cycle. It is not clear that either SAT nor SYNC-EN operations which support Metaring are amenable to larger ring distance.

C. Cyclic Reservation Multiple Access-II

CRMA-II is also an IBM development [5 - 8]. It is a reservation based protocol where access to slots is either restricted by a scheduler or free gratis if the slot is not reserved and is free. The concept is applicable to either a ring or bus architecture but we will only consider its operation in a ring environment because, in a bus configuration, it is similar to DQDB.

Nodes may access gratis slots. These may occur either by the scheduler issuing free slots between reservation cycles or by destination removal. If a node has packets backed up when a reservation request arrives, it request slots, noting how many gratis slots it has had since the last reservation cycle and how many it needs. It gets a count of slots reserved by the scheduler on arrival of the confirmation packet and uses the subsequent reserved slots as they arrive. If it has been able to use gratis slots in the interim, it marks the reserved slots as free.

One node acts as scheduler. It issues a reservation packet which travels the network. Nodes, as noted above, place their request immediately behind the reservation packet. This requires a limited form of insertion buffer at each node. Upon return of the reservation packet, the scheduler calculates the number of reserved packets it will supply each node and sends this information to each node in a confirmation packet. It then marks sufficient slots to handle all confirmed requests. Note that a reservation access takes one round trip cycle plus the scheduler calculation time plus the slot times of all prior node reservations. Determination of slot status is by bits but since destination removal is used logic based upon address recognition is required.

No asynchronous message performance data is available for CRMA-II or for previous versions of CRMA. As noted previously, equally distributed send and receive load

conditions with destination removal is inherently fair. It is questionable whether the reservation scheme maintains this uniformity since now a large number of reserved slots exist for nodes immediately down stream of the scheduler and a large number of gratis slot immediately up stream. Thus, the cyclic reservation scheme could actually reduce network capacity significantly while enforcing its fairness scheme at high loads. Even Metaring's SAT system, which does not change the symmetrical operational characteristics of the ring had a tendency to reduce capacity although to a limited degree.

No discussion of synchronous traffic access is provided although reference [5] implies that buffer insertion bypass mechanisms are available to support isochronous traffic. With the schedular it would appear that a framing technique similar to FDDI-II [18] and DQDB [14] solutions would be applicable. Here, call setup is an additional task for the scheduler and each node must to keep track of slots reserved for its synchronous use. The framing operation requires frames at fixed periods so that scheduler reserve and confirm operations would have to be integrated it the framing cycle.

D. DQDB

DQDB is being promoted by AT&T [11, 14, 19, 20, 21]. In a unidirectional bus configuration each node must provide two complete sets of hardware and two scheduler nodes are required at each end. With free slot bit access only, hardware is simpler than if destination removal is used but the advantages of additional capacity are not available. In addition, bus configurations are less amenable to multicast and broadcast since now two messages must be sent. Bus configurations also require two scheduler nodes at the bus ends which can be replaced only by immediately adjacent nodes if failure should occur.

Considerable information on DQDB asynchronous performance is available [11, 14, 19-21]. Unfairness is the major issue but to a large extent this has been solved [21, 22]. However, fairness protocol additions tend to increase access time to some extent.

Synchronous traffic is supported in DQDB [14] by a framing system as noted above in the CRMA-II discussion. Call control by the each scheduler and knowledge of nodes reserved for synchronous traffic must be kept by each node. In a WAN version of DQDB [14], both destination removal and pre-use were employed but this adds to the complexity of each access port. Destination removal requires address recognition and for preuse, each node must keep track of the source for each reserved slot. In addition, using framing for synchronous traffic access makes recovery and support for dynamic traffic variation difficult.

E. Comparison of the Media Access Protocols

A comparison of system features is provided in Table 3. The categories include equipment complexity and flexibility, operational and performance for under both asynchronous and synchronous traffic loads, reliability considerations and the expandability to MAN and WAN conditions.

The five systems discussed above are compared. Actually, two Cambridge Fast Ring systems, one based upon slot access and one based upon destination removal, are included.

It can be seen from Table 3 that the CFR systems are the simplest, but they do not support the fairness and guaranteed access requirements for high data rate networks. In addition, the support for synchronous traffic is not effective. Metaring provides excellent asynchronous performance. However, it requires twice as much equipment at each node; synchronous traffic operations are complex with call setup and termination at each node and integrated traffic performance is questionable; and it is not certain whether Metaring can be extended to longer networks. Also, it can not be reconfigured in case of node or link failure. The Cyclic Reservation Multiple Access-II system is most complex. The addition of the scheduler causes this complexity and since the system is no longer symmetrical, it is not capable of supporting increased capacity by using dual counter rotating rings like CFR, Metaring and CSMA-CRP. Also, it is questionable where the reservation capability provided will pay off in fairness performance at high loads. Integrated traffic performance is unknown and since dual ring operation is not supported, reconfiguration is not available. DQDB, like Metaring requires twice as much equipment, but unlike Metaring is not able to use this equipment effectively to increase capacity. Complexity must be added to DQDB to enable fairness through additional reservation control, since, as a bus based system, it does not contain the inherent fairness available in ring based systems. Further complexity is required to support synchronous traffic. Finally, DQDB can not use the additional equipment to enhance reliability.

The CSMA-CRP supports all of the requirements effectively. It can be operated as either a single or dual ring with 2X and 8X capacity, respectively. Complexity is added because of message truncation instead of slots. In most cases packet fracture is not large, and resequencing help is provided in the packet header. However, it is the CRP portion of the dual protocol which really enhances the system. It not only provides an extremely effective mechanism for synchronous traffic with properties such as strictly periodic access with both pre-use and recovery, but also supports fairness for the asynchronous system. It has

Table 3. Operational and Performance Comparison Between High Data Rate Access Protocols.

Rating
+ - good
++ - better
+++ - superior

Network	Equipment Configuration & Flexibility.	Asynchronous Traffic Access Complexity.	Synchronous Traffic Access Complexity	Async. Performance & Fairness	Synch. Performance	Reliability (use secor ring to reconfigure)	Expandable to MAN-WAN
Cambridge Fast Ring (simple)	Connectivity - 1 T/R per node Capacity - 1X data rate. Dual rings - 2 T/R per node. Capacity - 2X data rate	Slot bit access. Message partitioned into slots. On ring preemption. No fairness control. +	Slot Reservation System. +	Access Time - -> 1/2 slot time. Guaranteed Access - none. Fairness - none. Single message multicast. +	Slot Reservation System. Strictly periodic access. No pre-use or recovery. Start - Capture slots. Reduced capacity unknown. Term. - none. ++	Reconfiguration - yes. No special node. ++	Yes with no changes. ++
Cambridge Fast Ring (Destination Removal)	Connectivity - 1 T/R per node. Capacity - 2X data rate. Dual rings - 2 T/R per node. Capacity - 8X data rate +++	Spatial reuse complexity. Message partitioned into slots. On ring preemption. No fairness control. +	Cycle Time. Timers. Renegotiation for parameter setting. +	Access Time - -> 1/2 slot time. Guaranteed Access - none. Fairness - none Single message multicast. +	Cycle Time. Aperiodic access Capacity reduction due to unused slots at end of cycle. Renegotiation time - unknown. ++	Reconfiguration - yes. No special node. ++	Yes with no changes. ++
Metaring	Connectivity - 2 T/R per node. Capacity - ~8X data rate. ++	Spatial reuse complexity. Message partitioned into slots. Register insertion system. Preemption for control pkts. Timer & counter controls. +	Call setup & term. at all nodes. Renegotiation for ASYNC- EN parameter setting. Timers. Preempt operation for ASYNC-EN. +	Access Time - -> 1/2 slot time. Guaranteed Access - Unknown. Fairness - Excellent, ~ 10% capacity reduction. Single message multicast. ++	Aperiodic due to ASYNC-EN rotation time. Start - call set up req. time Term - require time. Capacity reduction - unknown. +	Reconfiguration - no. No special node. ++	Unknown - Fairness and sync. control may not be applicable.
CRMA-II	Connectivity - 1 T/R per node. Capacity - ~1X data rate.	Spatial reuse complexity. Message partitioned into slots. Register insertion system. Reserve/confirm pkt controls. Scheduler complexity. Handling reserve/confirm system with calculations.	Unknown. Same as DQDB?	Access Time - -> 1/2 slot time. Guaranteed Access - Cycle time+sched. calc.+ reserv. slot count. Fairness -Unknown.	Unknown. Same as DQDB?	Reconfiguration - no. Special node.	Undocumented in ring configuration.
CSMA-CRP	Connectivity - 1 T/R per node. Capacity - 2X data rate. Dual ring - 2 T/R per node. Capacity - 8X data rate. +++	Spatial reuse complexity. Message truncation circuitry. On ring preemption. ++	CRP control and access. Control based upon CRP length request. +++	Access Time - -> bit rate. Guaranteed Access - Next CRP + cycle time. Fairness - Excellent, 1% capacity reduction. Single message multicast. +++	Strictly periodic. No interference other than reduced capacity to async. pkts. Pre-use & recovery available. +++	Reconfiguration - yes. No special node. ++	Yes with no changes. ++
DQDB	Connectivity - 2 T/R per node. Capacity - 2X data rate.	Slot bit access. Message partitioned into slots. On bus preemption. Two special end nodes. Counters for reserve and access control. Additional fairness control. +	Call setup & term. at end nodes. Reserved slot access at each node.	Access Time - -> 1/2 slot time. Guaranteed Access - Unknown. Fairness - Excellent, capacity reduction depends upon method used. Two message multicast. ++	Strictly periodic. Start - call set up req. time Term - require time. Capacity reduction - unknown. Pre-use Recovery - none. +	Reconfiguration - no. Two special nodes must be at ends.	Yes with no changes. ++

minimal impact on the asynchronous traffic access. In doing so, it still provides enhanced traffic capacity with addition of a second ring and support added reliability through reconfiguration. Further, it is shown to be effective for LAN through WAN configurations.

V. Conclusions

It is shown that ring networks with destination removal require less equipment per node and can have significant capacity advantages and fairness advantages over bus architectures with identical basic data rates. However, to maintain these advantages for rings, it may be necessary that "operational symmetry" exist, i.e., there is not a special node that does scheduling or slot marking in such a manner that all nodes do not experience identical accessibility in a statistical sense.

Three of the ring networks described in this paper provide "symmetry" - CFR, Metaring and CSMA-CRP. Of these CFR is simplest but it does not provide good integrated traffic access and does not provide any mechanism for fairness or guaranteed access. Both Metaring and CSMA-CRP provide fairness. However, CSMA-CRP configuration and operation is simpler than Metaring and its performance capability for integrated traffic over a wide range of network parameters and conditions is demonstrated. Of the systems compared, CSMA-CRP supports integrated high data rate network operations in the most simple and effective manner.

VI. References

1. Minzer, S.E.: "Broadband ISDN and Asynchronous Transfer Mode (ATM)," *IEEE Commun. Magazine*; Sept 1989; pp. 16-25.
2. Cidon, Y.; Ofek, Y.: "Metaring - A Full-Duplex Ring with Fairness and Spatial Reuse," IEEE INFOCOM'90; San Francisco, CA.; June 1990; pp. 969-981.
3. Simha, R.; Ofek, Y.: "A Starvation-free Access Protocol for a Full-duplex Buffer Insertion Ring Local Area Network," IEEE Paper CH2799-5/90/0000/0531. May 1990.
4. Chen, J.; Ahmadi, H.; Ofek, Y.: "Performance Study of the MetaRing with Gb/s Links," Proc. of 16th LCN; Minn. MN.; Oct. 1991; pp. 137-147.
5. Van As, H.R.; Lemppenau, W.W.; Pitro; Z., Zurfluh, E.A.: "CRMA-II: A Gbit/s MAC Protocol for Ring and Bus Networks with Immediate Access Capability," EFOC/LAN '91; London, UK; June 1991.
6. Muller, H.R.; et. al.: "DQMA and CRMA: New Access Schemes for Gbit/s LANs and MANs," Proc of IEEE INFOCOM'90; San Francisco, CA.; June 1990; pp. 185-190.
7. Nassehi, M.M.: "CRMA: An Access Scheme for High-Speed LANs and MANs," Supercomm/ICC'90; Atlanta, GA; April 1990.
8. Heinzmann, P. Muller, H. Wilson, K.: "Configuration Control for Bus Networks," Proc. of 15th LCN; Minn. MN.; Oct. 1990.
9. Greaves, D.J.; Lioupis, D.; Hopper, A.: "The Cambridge Backbone Ring," Proc. of INFOCOM '90; pp 8-14.
10. Zafirovic-Vukotic, M; Niemegeers, I.G.; Valk, D.S.: "Performance Analysis of Slotted Ring Protocols in HSLAN's," Jour. on Selected Areas in Communications; Vol 6; No 6; July 1988; pp 1011-1023.
11. Newman, R.M.; Budrikis, Z.L.; Hullett; J.L.: "The QPSX Man," IEEE Communications Magazine; Vol 26, No 4; April 1988; pp 20-28.
12. Foudriat, E.C.; Maly, K.; Overstreet, C.M.; Khanna, S.; Paterra, F. : "A Carrier Sensed Multiple Access Protocol for High Data Rate Ring Networks," Comp. Comm. Review; Vol. 21; No. 2; April 1991; pp. 59-70.
13. Foudriat, E.C.; Maly, K. Overstreet, C.M.; Khanna; S.; Zhang, L.; Sun, W.: "Combining Two Media Access Protocols to Support Integrated Traffic on High Data Rate Networks," Proc. of 16th LCN; Minn. MN.; Oct. 1991.
14. Edmond, W.; Seo, K.; Leib, M.; Topolcic, C.: "The DARPA Wideband Network Dual Bus Protocol," Proc. of SIGCOMM'90; pp. 79-89.
15. Jaiswal, N.K.: Priority Queues; Academic Press; NY; 1968.
16. Little, T.D.C.; Ghafoor, A.: "Network Considerations for Distributed Multimedia Object Composition and Communications," IEEE Network Magazine; Nov. 1990; pp. 32-49.
17. Haas, Z.: "Optical Distribution Channel: An "Almost-all" Optical LAN based on the Field-coding Technique," Proc. of 16th LCN; Minn. MN.; Oct. 1991, pp. 128-136.
18. Boston, T.: "FDDI-II: An Overview of Call Setup for Isochronous Traffic," British Telecom - TA1131, Aug. 24, 1987.
19. Potter, P.G.; Zuckerman, M.: "Cyclic Request Control for Provision of Guaranteed Bandwidth within the DQDB Framework," Proc. ISS'90, Stockholm, 1990.
20. IEEE 802.6: "Distributed Queue, Dual Bus (DQDB) Subnetwork of a Metropolitan Area Network(MAN)," Draft 15, 1990.
21. Hahne, E.L.; Maxemchuk, N.F.: "Improving the Fairness of DQDB Networks," Proc. of INFOCOM'90, San Francisco, CA, 1990.
22. Maly, K.; Zhange, L.; Game, D.: "Fairness Problems at the Media Access Level for High-Speed Networks," ODU CS Tech. Rpt. TR #90-15; March 1990.