

PERFORMANCE EVALUATION OF MIL-STD 188-220 INTEROPERABILITY STANDARD FOR DIGITAL MESSAGE TRANSFER DEVICE SUBSYSTEMS

R. Gregory Burch, Jr. Paul D. Amer
University of Delaware
Newark, DE 19716

Samuel C. Chamberlain
Army Research Laboratory
Aberdeen Proving Ground, MD 21005

Abstract

MIL-STD 188-220 specifies the procedures, protocols, and parameters needed for interoperability of Digital Message Transfer Devices (DMTD's) and applicable control systems. One of 188-220's Link Layer mechanisms is the Net Access Control Algorithm which in turn defines the Net Access Delay (NAD). MIL-STD 188-220 provides several NAD algorithms to handle "the time a station with a message to send shall wait to send a frame after the TP timer has expired." Supported by an ARL contract, the University of Delaware compared the performance of Randomized (R-NAD) and Prioritized (P-NAD) NAD algorithms using the OPNET simulation package. Simulation results comparing R-NAD and P-NAD under varying traffic conditions (offered loads, priority distributions, etc.) are presented. Results show that R-NAD outperforms P-NAD in both access delay and collision statistics. More importantly, this simulation exposes 188-220's ambiguous definition of the P-NAD algorithm, an ambiguity that continues even in the later versions of the standard. This ambiguity can result in a P-NAD implementation that actually encourages packet collisions.

Introduction

MIL-STD 188-220 was designed to allow efficient communication between stations sharing a low-bandwidth broadcast channel. Because of the inherent physical limitations of networks of this type, network access control is crucial to their overall performance. In particular, some control mechanisms are needed to assure a reasonable combination of "fairness" to all transmitting stations without sacrificing network performance. Multiple stations accessing the channel at the same time results in packet collisions, causing losses (and perhaps retransmissions). However, schemes which

avoid collisions by allowing a station to monopolize the channel can drastically affect performance at the other nodes in the network.

The next three sections provide (1) background information on NAD calculation in MIL-STD 188-220, (2) a simplified OPNET model of the 188-220 Link Layer and the assumptions present in that model, and (3) the simulation performance results of the model under various network loading conditions. This work is based on the May 1993 version of the standard.

1 MIL-STD 188-220

The presence of multiple nodes on a single communication medium necessitates a means of controlling which node may access the net at any given point in time. MIL-STD 188-220 provides a distributed Net Access Control mechanism which defines several different schemes to allow stations to transmit. These schemes are intended to minimize the number of instances when more than one station is transmitting simultaneously, i.e., colliding.

Appendix C of the May 1993 version of 188-220 defines three algorithms for net access.¹ All three are based on four primitive functions: net-busy sensing, response hold delay, timeout period, and net access delay (NAD). Net-busy sensing determines the presence of a signal on the channel. Response hold delay defines the time a station delays prior to returning an acknowledgment. Timeout period is the time all stations must wait before they may schedule their NAD. These three functions are identical for all nodes, regardless of which NAD algorithm is in use.

The fourth function is the NAD calculation itself. All three NAD algorithms divide time between station synchronization points into slots. A station wishing to transmit computes (via the NAD algorithm) a particular slot as its NAD slot. When that slot time arrives, if the channel is free, the station transmits; if the channel is occupied, the station waits for the next

Supported, in part, by the Advanced Computational and Information Science Directorate under the auspices of the US Army Research Office Scientific Services Program administered by Batelle (DAAL03-91-C-0034), and by the Army Research Office (DAAL03-92-G-0700).

¹The current version defines five algorithms.

synchronization point and repeats the process. New synchronization points occur when the current transmission ends or when the NAD Period (another Link Layer parameter) expires.

This paper focuses on two specific NAD algorithms. In Random Net Access Delay (R-NAD), a station chooses its slot as follows. At each station, a random number between 0 and $3/4 * NS$ (the number of stations on the network) is generated. That number is the slot number for the current transmission (the first calculable slot is number 0, the second is number 1, etc.). All nodes are treated equally, and all transmissions have the same priority. However, collisions are possible when multiple stations select identical random values.

Prioritized NAD (P-NAD) avoids collisions (supposedly) and allows both station and message priorities to improve performance for important DTMD's and transmissions. In P-NAD, the slot number calculation is deterministic. Each station is given a unique Subscriber Precedence number, perhaps allocated arbitrarily, perhaps allocated according to the importance or military rank of the station's user. Packets are divided into three priorities: urgent (high), priority (medium), and routine (low). The unique [sic] slot number calculation is as follows:

- Parameter MP is calculated based on the message priority. For urgent messages, $MP = 0$. For priority messages, $MP = (NS + 1)$. For routine messages, $MP = 2 * (NS + 1)$.
- For initial transmissions: $Slot = Subscriber\ Precedence + MP$;
- For subsequent transmissions: $Slot = NS + MP$;

Unfortunately, a definition of "initial" and "subsequent" is not included in the May 1993 version on which this simulation is based.² This ambiguity can result in implementations that cause collisions instead of eliminating them. Further discussion of the "initial vs. subsequent" factor in choosing the P-NAD slot is included in Section 3.1.

2 OPNET Model

A model network was created using OPNET, a simulation and analysis tool. The model focuses solely on the performance effects of the NAD algorithm being used; thus the model implements a simplified 188-220 Link Layer. The simplifications and assumptions are

as follows:

- Only 188-220 "Type 1" link layer (connectionless) operation is modeled. Packets are transmitted once and are not acknowledged (and hence not retransmitted if lost via collision).
- The transmission medium is considered perfect (no bit errors introduced), so losses occur only as the result of collision.
- The propagation delay between stations is assumed to be zero. For the sake of transmission "correctness," this assumption simply means that carrier sensing is considered perfectly accurate at all times for all stations.
- The network contains 16 stations, only 4 of which actively submit traffic.
- Arbitrary values for the Keytime Delay and CRYPTO Preamble Transmission Time (two Link Layer parameters) were chosen. These two parameters are the only factors which control the NAD slot size. The chosen values resulted in a slot size of 1.0 second.

These assumptions allow an inspection of phenomena which result from the NAD algorithm in use, and not from other aspects of the protocol. The only assumption which may bias any conclusions about the NAD algorithms themselves is the assumption that the propagation delay is zero (i.e., that net-busy detection could not yield the wrong answer). In the case of non-zero propagation delay, stations could either transmit when they should not (resulting in a higher collision rate), or vice versa (resulting in longer net access delays). For a study of the actual values of collision rates and delays under particular NAD schemes, this error could be problematic. However, for a comparative study of the NAD algorithms, the error is not a concern. It is also noted that the expected small network bandwidth and relatively large packets result in packet transmission times of 8-28 seconds, thus making microsecond propagation times insignificant.

The OPNET model is repeated for the following combinations of independent variables and their values:

- NAD Scheme: R-NAD, P-NAD
- Packet Priority (only in P-NAD): 1/3 urgent, 1/3 priority, 1/3 routine
- Total Offered Load - total amount of traffic all nodes try to send over the network. This value is a fraction normalized to the network's maximum usable bandwidth: .1, .2, .3, .4, .5, .7, .85, 1.0, 1.5, 2.0

²An attempt to clarify these terms appears in the Feb. 1995 version. However, that clarification, while an improvement, remains ambiguous.

- Offered Load Distribution - the distribution of offered loads among the four nodes:
 balanced: 1/4, 1/4, 1/4, 1/4
 unbalanced: 1/2, 1/6, 1/6, 1/6
- Packet Size - in a given simulation trial, all packets by all nodes are assumed to be of one size: large, medium, small. Relative to the bandwidth, the three packets required roughly 28, 18, and 8 seconds, respectively to transmit. (Remember, this is a LOW bandwidth network.) We do not consider the effects of different packet size distributions within the same trial. Given a packet size and desired offered load, the simulation model computes the packet interarrival rate at each of the 4 nodes taking into account a balanced vs. unbalanced loading.

For each simulation trial, we measure network access delay and collision statistics. Net access delay begins when a packet to be transmitted reaches the head of the queue at the sending node and ends when the first bit of that packet is transmitted. This delay does not include the delay a packet incurs while waiting for previous packets at the same node to be transmitted. Thus even for offered loads greater than 1.0, although the sending queues will necessarily grow without bound, the net access delay need not necessarily also grow without bound. This definition for net access delay is appropriate because this is the delay that is most affected by the NAD scheme in use (and not by other protocol transmission features). Also this delay is independent of whether or not a transmission results in a collision. Retransmissions, if any, must be handled at a higher layer. Obviously, with this definition, a small average access delay is only good when simultaneously there are few collisions.

3 Performance Results

The primary results of the OPNET study are summarized in Figures 1 through 8.

Figures 1 and 2 show the delays incurred under varying offered loads for balanced and unbalanced offered loads, respectively. In each figure, six lines are plotted. These represent the six combinations of NAD scheme and packet size. As expected, all delays increase with increasing offered loads. Of importance are the following observations:

- The delays for R-NAD are consistently smaller than those of P-NAD for similar packet size independent of total offered load, and independent

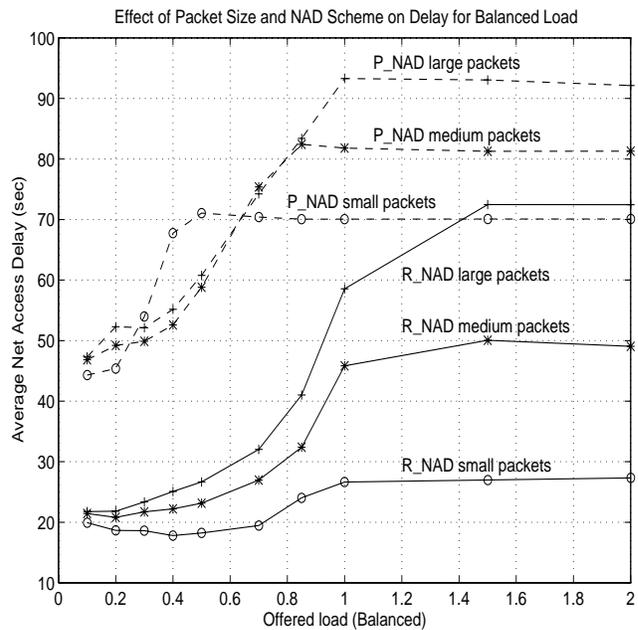


Figure 1:

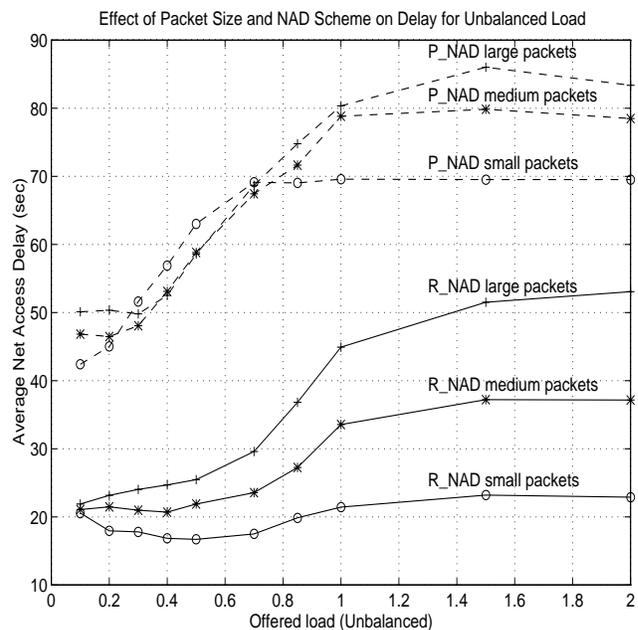


Figure 2:

of whether or not the traffic is balanced. It is good to see that neither NAD scheme is negatively affected by an unbalanced load since this is what is likely to occur in a battlefield situation. One may conclude that the added delays incurred by P-NAD over R-NAD (roughly two to three times the delay on average) represent the implicit cost of providing different priority classes. It is good to see that delays for R-NAD over reasonable operating offered loads (.1-.8) remain

stable.

- In general, a fixed load consisting of larger packets results in greater delays than the same load with a larger number of smaller packets. This may be expected inasmuch as packets which do not succeed in accessing the channel in their first attempt must wait longer for a synchronization point if the packet size is larger. This is simply due to the fact that the next synchronization point will come when the channel becomes free (i.e., one packet transmission time from the access attempt), and the channel will obviously become free later for larger packets.

Similarly, Figures 3 and 4 show the collision statistics under various offered loads, again for balanced and unbalanced traffic, respectively.

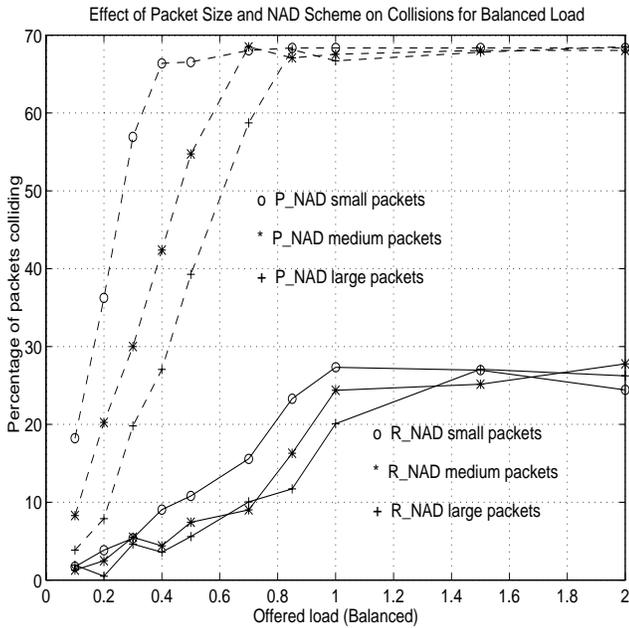


Figure 3:

Once again, we see the expected result that collision probabilities increase with increasing offered loads. These figures lead to the following commentary:

- For small load values, a fixed load consisting of larger packets results in fewer collisions than the same load with a larger number of smaller packets. However, as the load values increase, packet size no longer affects the percentage of collisions. These results are certainly reasonable. At small loads, small packet size (meaning more packets) certainly increases the chances of collision since it increases the chances of two stations even *having* packets to send at the same time. However, at larger loads when stations are more likely to

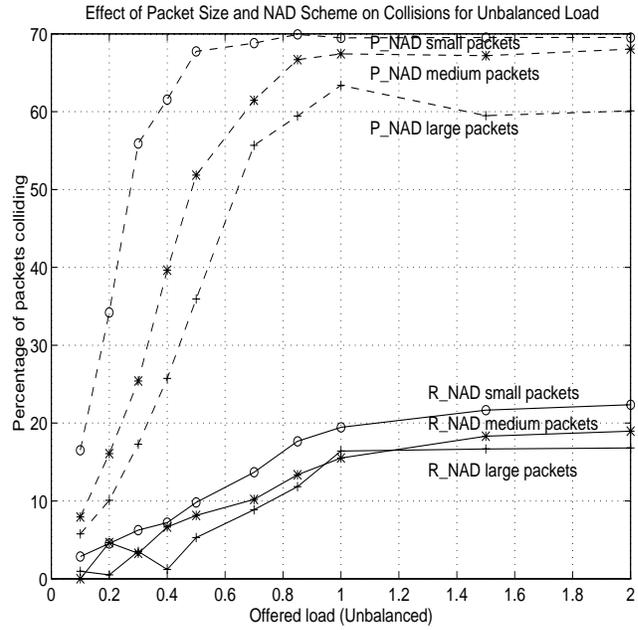


Figure 4:

already have packets to send at the same time (by virtue of the load itself being high), a larger number of packets simply means a larger number queued at the sending node (and hence should not affect collision rates).

- Balanced loads generally lead to more collisions than unbalanced loads (as expected, since balanced loads make it more likely that stations even *have* something to send at the same time).
- R-NAD grossly outperforms P-NAD in collision percentage for all simulation parameters. In fact, the percentage of collisions for P-NAD is roughly 3 times that for R-NAD. High collision rates for P-NAD were *not expected*. A discussion follows.

Figures 5 and 6 further investigate the performance of P-NAD. These figures have four lines plotted, where three of the lines represents the three packet priorities, and the fourth is an aggregate average over all packets.

Figure 5 shows that the P-NAD delay characteristics are exactly what is desired. Urgent packets have the lowest delays, which are significantly better than delays for priority and, in turn, delays for routine. Similar results were observed for different packet sizes and for balanced loads and are not included here.

Figure 6, however, is most startling. This figure shows P-NAD collision percentages at various loads (with small packets) for each of the three priority classes. The first notable result from this graph is that the absolute values of collision rates are significantly

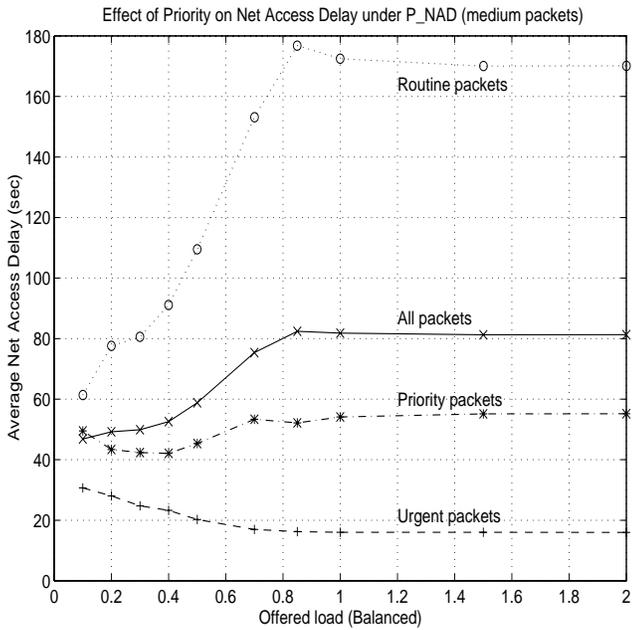


Figure 5:

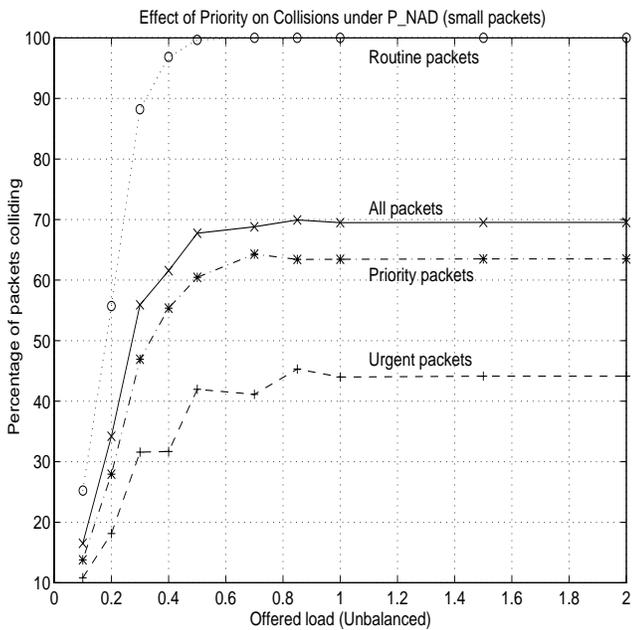


Figure 6:

high. Even though the collision rates themselves depend critically on parameter values, with reasonable choices for these parameters, an aggregate collision rate of 70% is simply unthinkable. Furthermore, at high loads, routine packets are seen to collide 100% of the time! The P-NAD algorithm is supposed to *avoid* collisions.

3.1 P-NAD Specification Ambiguity

According to the May 1993 version of 188-220, for “initial” transmissions, the subscriber precedence is included in the calculation of the P-NAD slot time. However, for “subsequent” transmissions, the subscriber precedence is omitted from the slot calculation. This omission causes the aforementioned collision phenomenon. Without the subscriber precedence, there are only three possible transmission slots: numbers NS , $2 * NS + 1$, and $3 * NS + 2$ for priorities urgent, priority, and routine, respectively. Hence at high loads (i.e., when every station has something to send), routine packets can only be transmitted when all transmitting stations have routine packets, and hence must collide. This analysis is confirmed in the study’s performance results.

Based on these results, the authors began investigating a better definition of the P-NAD algorithm. Then late in the study, the Feb. 1995 version of 188-220 was published. A paragraph is added to clarify the parameter IS , the “initial/subsequent” factor. This paragraph states “Only one station on the net uses the subsequent factor. That is the station that transmitted last on the net.” We now believe the intended meaning of “initial” and “subsequent” is not as we originally interpreted them in our simulation study. Originally, we assumed that a station had one “initial” transmission and that all others were “subsequent.” As our results indicate, this interpretation (which in our opinion is the most reasonable interpretation of English) produces a P-NAD scheme that is highly *inefficient*.

It now seems that “subsequent” is meant to refer only to the *single* station that most recently transmitted. By implication/definition, all others are sending their “initial” transmission!!! This (terrible) interpretation is likely to significantly improve the collision results of P-NAD and is something the authors plan to investigate in the future. Perhaps the most significant result of this OPNET study is reinforcing the fact that inherently ambiguous languages such as English are inappropriate for the specification of complex communication protocols. Different readers will make different interpretations with even the best of writers. Overcoming ambiguity in specification is further discussed in a separate MILCOM ‘95 paper.

3.2 Effect of Unbalanced Loads

Figures 7 and 8 show the delays and collision rates, respectively, incurred by nodes with unbalanced offered

loads. Specifically, these figures show the delay performance for medium packet size. Assuming that further clarification of the P-NAD algorithm takes place, then the graphed P-NAD values will change; however, R-NAD values will not.

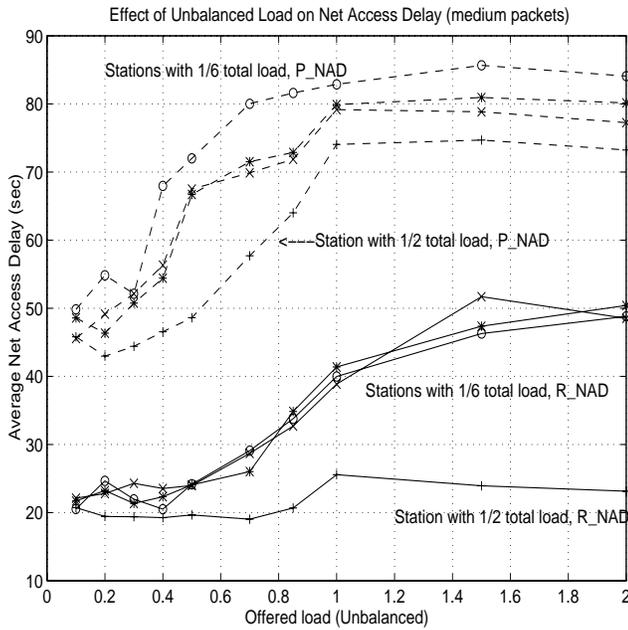


Figure 7:

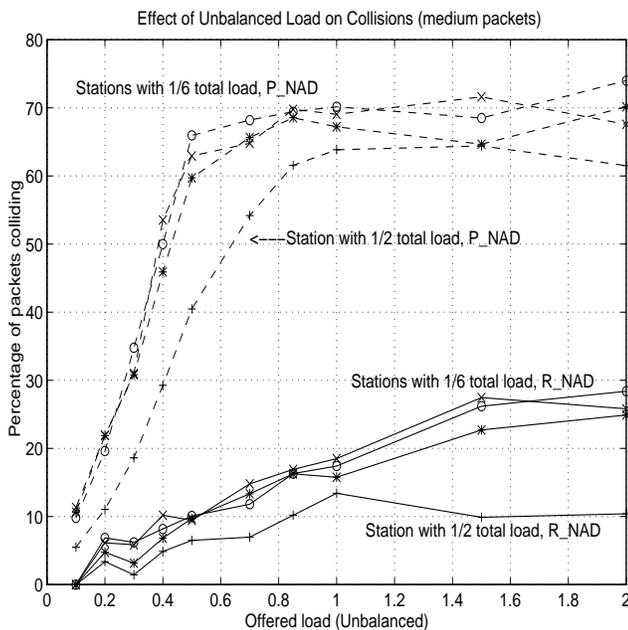


Figure 8:

First, nodes with equal offered loads received roughly the same delay and percentage of collisions. That is, for the three nodes each offering 1/6 of the total load (and independent of packet size, NAD scheme,

and load), each node waited roughly (within 5%) the same amount of time to access the channel and experienced roughly the same collision rate. In addition, for both R-NAD and P-NAD schemes, the delays for the “heavier” node offering 1/2 of the total load were consistently less than the delays of the other three equally balanced nodes. Similarly, the “heavier” node consistently experienced fewer collisions than the other three nodes. These results are reasonable, as the one heavier node faces less competition for the channel from all of the other nodes (1/2 total offered load) than does each of the lighter nodes face from all of the other nodes (5/6 total offered load). Comparable results were obtained for small and large packet size distributions; their figures are not included.

4 Conclusions

In this paper we have

- discussed MIL-STD 188-220 Net Access Algorithm essentials, and reviewed Random Net Access (R-NAD) and Prioritized Net Access (P-NAD) algorithms in some detail,
- presented a simplified OPNET model of the 188-220 Link Layer and explained the assumptions in that model, and
- presented results of simulation of both R-NAD and P-NAD algorithms under a variety of network conditions.

Results show that R-NAD greatly outperforms P-NAD in both net access delay and collision rates, for virtually all other parameters (load balancing, packet size, etc.). Users must decide if the increased delays and collision rates are worth the added functionality of priority classes.

Results also show that interpreting the 188-220 P-NAD algorithm is difficult, particularly in the May 1993 version. The Feb. 1995 version attempts to clarify things, but in the authors’ opinion, the standard remains flawed in its use of the English words “initial” and “subsequent.” This ambiguity can result in implementations that actually encourage collisions among packets, particularly among low-priority packets, and particularly at high loads. The current P-NAD implementation also causes apparently unneeded excessive access delays. We recommend modification to the P-NAD algorithm description to clarify its intended meaning (and hopefully improve its performance).