Performance Evaluation of a Modified Aloha Protocol for Tactical Radio Communications

Mike Markowski
US Army Research Laboratory
Aberdeen Proving Ground, MD 21005

April 5, 1996

Abstract

A performance evaluation of a modified Aloha protocol for the Media Access Control (MAC) layer is presented. The modifications are that frames are not constrained to be fixed length, and frames are assigned a waiting time for their first attempt rather than only after collisions. Equations are derived that predict the throughput and average delay for a finite user system with a Poisson arrival process. These expressions can be used in a higher layer protocol so that channel use is maximized while still meeting application layer requirements. Next, an analysis for Logical Link Control or higher protocol layers is presented for the case where several priority queues are used. The MAC layer receives the frames at intervals from these queue which are emptied from highest to lowest priority. It is shown that if either channel utilization of traffic rate is held constant, there is no benefit through the use of priority queues.

1 Introduction

Modems are used in conjunction with the SINCGARS or SINCGARS-SIP VHF radio to transmit digital information. Aloha or variants of it are often used because of simplicity of implementation and robustness in the event of noisy channels. The analysis in this paper is for a variant of Aloha. Analysis of the classic Aloha protocol assumes that all frames are the same length. Furthermore, when a node receives a frame, it immediately transmits it. If there is a collision, the node is assigned a random waiting interval before its next transmission attempt. The variant considered here assumes frames have a maximum length but can be smaller. Also, even initial transmissions are assigned a random waiting time.

For a concrete example, the Harris 3490A Tactical Data Buffers (TDB) is considered and evaluated using the technique presented. It transmits data by using two internal protocols: immediate transmit if the channel has been idle for 10 s or more, or a typical backoff protocol if the channel has been busy within the last 10 s. Though waiting intervals are assigned to frames, the random backoff times picked are chosen in a slightly more complex manner than in pure Aloha. The specific backoff protocol is described in detail in the later sections. The 3490A TDB has an internal buffer of 24 kbytes and provides signaling to indicate when the buffer is empty or not. This specific implementation, however, makes yet another modification to the Aloha protocol: there is no collision detection even though there is carrier sensing.

When several such devices make up a single hop network, throughput and average delay are the performance measures of concern. It is not, however, readily apparent what these values will be for different frame sizes and traffic rates. Because of the limited bandwidth available, e.g., 2133 data bits/s after error encoding an actual 16 kb/s rate, it is critical to maximize channel utilization. Average delay is important only insomuch as the application's delay requirements are met. Ideally, as the application's requirements change, or network properties such as channel utilization change, a protocol could adapt by modifying traffic rate, frame size, or utilization so that the best use is made of the low bandwidth channel.

After presenting an overall analysis of the MAC layer, the next item of interest is performance of higher layer protocols which make use of Aloha-type protocols for the MAC layer. Interest is specifically focused on the use of several priority queues, whose packets are concatenated to form a larger frame, with the frame then passed to the MAC layer for transmission. It is shown that in certain instances, it is not beneficial to use priority queues.

Equations are derived which characterize network performance as various design parameters vary. These equations can be easily incorporated in future protocols which adapt to the changing network.

2 The RF-3490A Protocol

The MAC layer implemented in the Harris RF-3490A TDB is described by Kaste in [Kas90] and summarized here. For lack of a given name, this paper refers to the protocol as the 3490A backoff protocol. An important factor with the 3490A backoff protocol is that it offers carrier sensing but not collision detection. As a result, it is not possible to directly design an inference seeking algorithm, i.e., one which makes use of channel state history. It is possible, however, to make use of acks (acknowledgements) to simulate collision/noncollision channel feedback. This will decrease usable bandwidth however. Like all MAC protocols, reliable message delivery is not guaranteed. If reliable message delivery is required, it must be implemented at a higher level.

The 3490A backoff protocol can be thought of as serving a queue of length one where the single frame can be up to 24 kbytes in length. The 3490A indicates whether or not a frame is currently awaiting service. Once service is complete, the 3490A indicates to the higher layer its readiness to accept another frame. See [Kas90] for the details of the physical interfaces.

An arriving frame is assigned a delay time so that the departure process is made independent of the arrival process. After the assigned delay has expired, the radio channel is sensed. If idle for 10 s or more, transmission commences. If busy, the frame is reassigned a wait time. If after three assignments of a "standard" waiting time and the frame is still not transmitted, it is assigned a "priority" wait time. Priority waits are shorter than standard waits as shown in Table 1.

The assignments are made by generating a random number between 0 and 255 and then performing a table lookup. A range of numbers is mapped to a specific wait time. The random number generator in the 3490A will not repeat a random number until all 256 numbers are used. Table 1 indicates the actual times being assigned and their probabilities. For standard wait times, delays are assigned with a slight bias towards the middle values. It should be noted that while delays are assigned in half-second increments, this is not a globally slotted system, and so node slot times between nodes are not synchronized. Also, if the channel has been idle for 10 s or longer, an arriving frame is sent immediately without being assigned a delay. For a very lightly loaded net, then, the MAC delay reduces to the channel service rate.

	Time (sec)	*Range*	*Probability*
Priority	1.0	0–63	64/256 = 25%
	1.5	64–127	64/256 = 25%
	2.0	128–191	64/256 = 25%
	2.5	128–191	64/256 = 25%
Standard	3.0	0–24	25/256 = 9.76%
	3.5	25–49	25/256 = 9.76%
	4.0	50–74	25/256 = 9.76%
	4.5	75–103	29/256 = 11.33%
	5.0	104–132	29/256 = 11.33%
	5.5	133–161	29/256 = 11.33%
	6.0	162–190	29/256 = 11.33%
	6.5	191–215	25/256 = 9.76%
	7.0	216–240	25/256 = 9.76%
	7.5	241–255	15/256 = 5.86%

Table 1: Packet Delay Assignments.

3 Tactical Radio Protocol Stacks

The 3490A, as a MAC layer protocol, would presumably be the foundation for a Logical Link Control (LLC) layer. Generally, the MAC layer of any protocol stack provides unreliable frame transmission. An LLC layer immediately above this would convert the unreliable service into reliable links between service access points at each node or nodes. A typical LLC, e.g., the IEEE 802.2, provides three types of services: acknowledged connectionless, connectionless, and connection-oriented. These two layers roughly correspond to the OSI layer 2, the Data Link Layer (see Tanenbaum [Tan89] for a discussion of the Open Systems Interconnection model).

Ideally, an LLC protocol would be used in conjunction with the 3490A. However, this is currently not the case. The protocol stack of interest in this paper, as illustrated in Figure 1, is that of the Distributed Fact Base (DFB). In the DFB, the 3490A protocol would be used to support the combined LLC, network and transport layers of the Fact Exchange Protocol (FEP). Because the 3490A protocol does not offer reliable transmission, it is the FEP which must provide this service if needed by the application. Working under the assumption that the applications will be the set of DFB programs, Figure 1 shows the complete protocol stack required for a full implementation of the DFB. As can be seen, the FEP supports two physical interfaces, 802.3 and RS-232. The diagram shows the FEP disregarding OSI protocol layering recommendations by spanning several layers, because the code was written when there was no IP (Internet Protocol) support for multicast addressing and when no LLC for the Harris TDB existed. Until the needed protocols were developed, the FEP was written to implement necessary, but at that time unavailable, services of various protocol layers.

The 802.3 connection is used to interconnect stationary nodes at the higher 10 Mb/s Ethernet bandwidth while the RS-232 serial interface can be used in a given network topology one of two ways. The less common use is to directly connect two computers running DFB by using an RS-232 cable. This is illustrated in Figure 2.

The 3490A protocol is not required for a DFB implementation since the DFB model is built only on the assumption that some unreliable data link layer exists. However, the Harris 3490A TDB is the most readily available equipment and should be studied to understand what services it can and cannot offer to higher protocol layers. With this in mind, the more

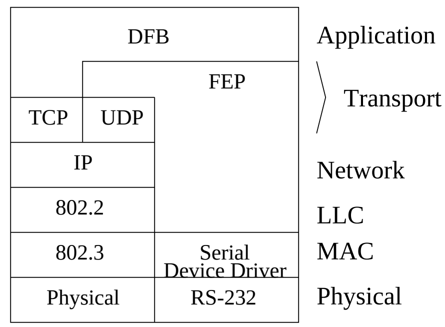

Figure 1: Full DFB protocol stack.

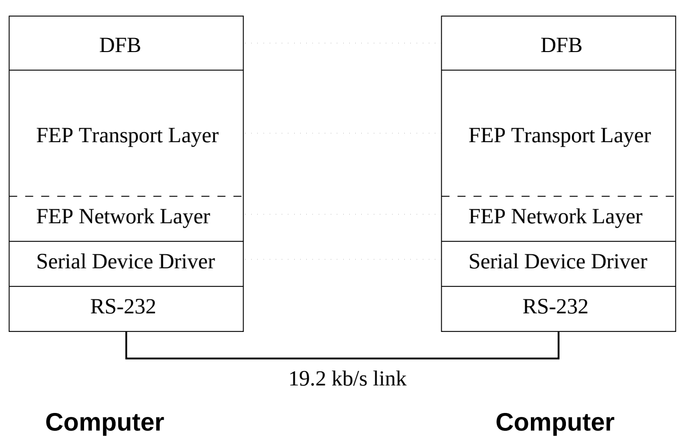

Figure 2: Partial DFB protocol stack with RS-232 interconnection.

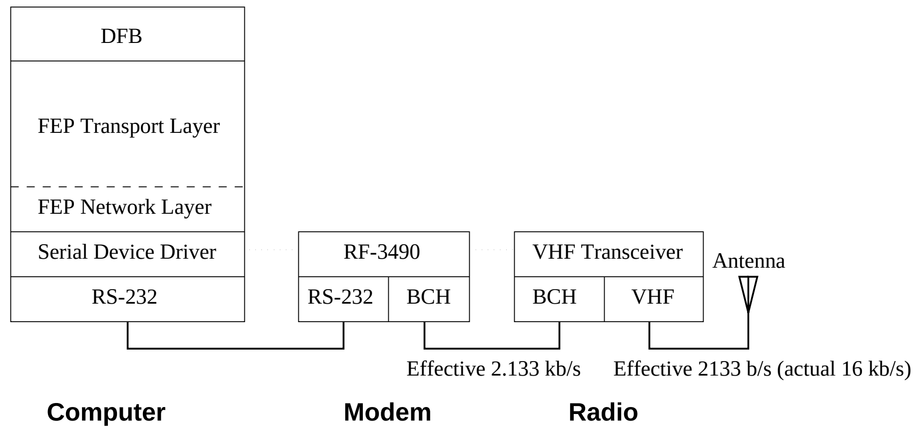

Figure 3: Partial DFB protocol stack with radio interconnection.

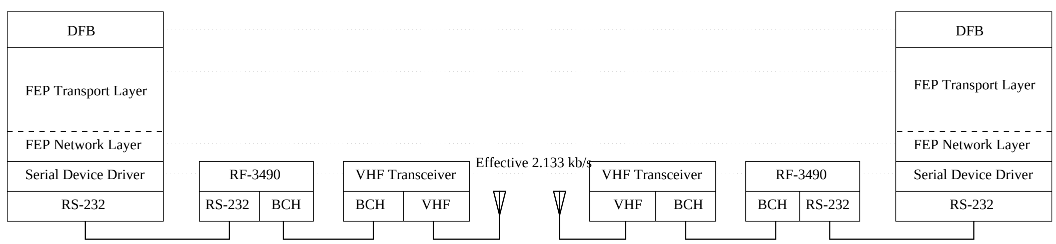

Figure 4: Peer-to-peer DFB Communication via Radio Link.

typical DFB scenario is connection to a Harris RF-3490A TDB which, in turn, is connected to a SINCGARS VHF transceiver as is shown in Figure 3. In this arrangement, peer-to-peer communication between various layers of the DFB can be envisioned as is shown in Figure 4. However, for the sake of simulation, all components need not be considered. For example, the workings of the radio itself, other than delays imposed, etc., are not important to the model. Similarly, the existence of the Harris RF-3490A TDB as a separate and distinct device can be ignored. The resulting simplification of the protocol stack is shown in Figure 5 and the corresponding peer-to-peer communication in Figure 6. As a final illustration of the protocol stacks, Figure 7 shows the stacks as simulated and analyzed in this paper. The topmost layer, entitled "Abridged FEP" in the diagram, will be described in more detail in subsequent sections. Briefly, however, it is little more than a set of priority queues. When the 3490A protocol indicates that it can accept a new frame, packets are concatenated from the priority queues to generate a frame which, in turn, is passed to the 3490A.

4 Analysis

Typical performance measures of random access protocols are throughput and delay. Determining these values is often made easier by making assumptions that frames are fixed size, there are an infinite number of users, and so on. Because we are trying not only to char-

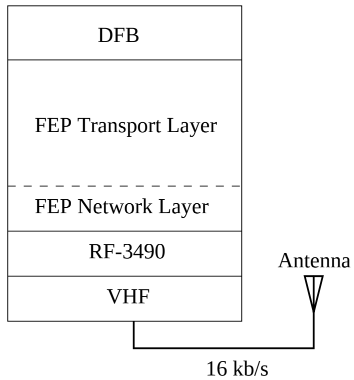

Figure 5: Simplified DFB stack.

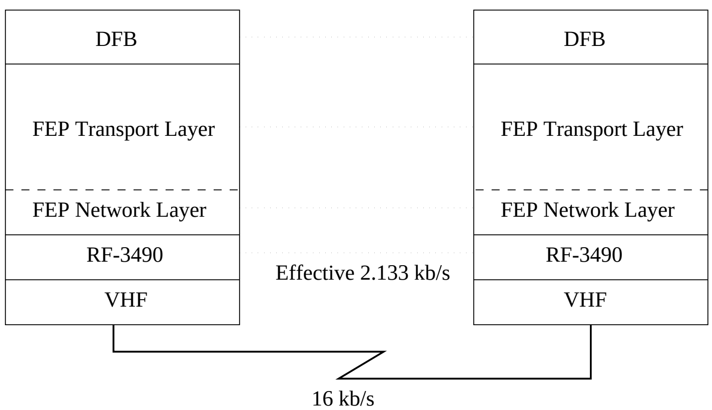

Figure 6: Peer-to-peer DFB communication with simplified stacks.

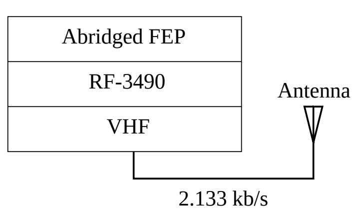

Figure 7: Stack used in simulations.

acterize an existing and therefore finite system, as well as optimize average frame lengths, we cannot make the usual assumptions.[1] However, as will be discussed, a system designed for certain types of environments can often make other assumptions not available in general cases.

4.1 MAC Layer Performance

In a finite user population, total system utilization is greater than that seen by an individual node, because a given node does not see its own channel use. In the following calculations, ρ refers to system-wide utilization while ρ_i refers to the system utilization as observed by node i. For a system made up of similar nodes, it is easy to see that

$$\rho_i = \frac{n-1}{n}\rho$$

and that as $n \to \infty$, $\rho_i = \rho$ as expected.

Because a frame is assigned either a standard delay or a priority delay, both must be considered when calculating the expected MAC delay. Looking at the probable delay assignments in Table 1, it is easy to determine the expected (average) standard delay, $d_s = 5.1612$ s, and priority delay, $d_p = 1.75$ s. In conjunction with these values, the expected delay also depends on how often carrier sensing finds the channel busy. On average, the probability of finding the channel busy is the channel utilization seen by the sensing node, or ρ_i. The total delay is determined by an infinite sum representing the probability of success after one carrier sense plus that of success after two sensings, and so on. Therefore, the expected MAC delay, D_i, for node i can be expressed as

$$\begin{aligned} E\{D_i\} &= P(1 \text{ sensing}) + P(2 \text{ sensings}) + P(3 \text{ sensings}) + \cdots \\ &= (1-\rho_i)d_s + \rho_i(1-\rho_i)2d_s + \rho_i^2(1-\rho_i)3d_s + \rho_i^3(1-\rho_i)(3d_s+d_p) \\ &\quad +\rho_i^4(1-\rho_i)(3d_s+2d_p) + \cdots + \rho_i^n(1-\rho_i)\big(3d_s+(n-2)d_p\big) + \cdots \\ &= (1-\rho_i)d_s\sum_{k=0}^{2}(k+1)\rho_i^k + \rho_i^3(1-\rho_i)\Big[3d_s\sum_{k=0}^{\infty}\rho_i^k + d_p\sum_{k=0}^{\infty}(k+1)\rho_i^k\Big]. \end{aligned}$$

Note that after three "busy" sensings, all subsequent delays are the shorter priority delays. Reducing the previous expression above to a closed form yields an expected delay function

$$D_{\rho_i} = E\{D_i\} = d_s(1+\rho_i+\rho_i^2) + (2d_s+d_p)\rho_i^3 + d_p\frac{\rho_i^4}{1-\rho_i}. \tag{1}$$

The ρ_i subscript on D indicates that this is the expected MAC delay for a frame at node i when node i perceives a channel utilization of $\rho_i \le \rho$.

Plotting Equation 1 in Figure 8 shows that net delays increase dramatically as channel utilization approaches 0.9 offering where average delay is about 35 s. Looking at the range of values only up to a utilization of 0.9 shows that for a more reasonable net utilization of 0.5, less than 20 s MAC delay can be expected. It is also clear that even if the utilization drops close to zero, a 5 s wait will be encountered.

Generally, the application will determine what average net access delay, D_{MAC}, can be tolerated. Note, however, that D_{MAC} is dependent on channel utilization ρ which is ultimately dependent on the aggregate traffic rate, average frame length and channel service rate. The aggregate traffic rate is further dependent on the number of nodes in the net. Figure 9 shows only channel utilization; the actual breakdown of utilization into the other variables is still necessary for a network design.

[1] While Aloha performs optimally using fixed size frames, it is often more important to transmit a small frame quickly than it is to wait for a frame to reach full size.

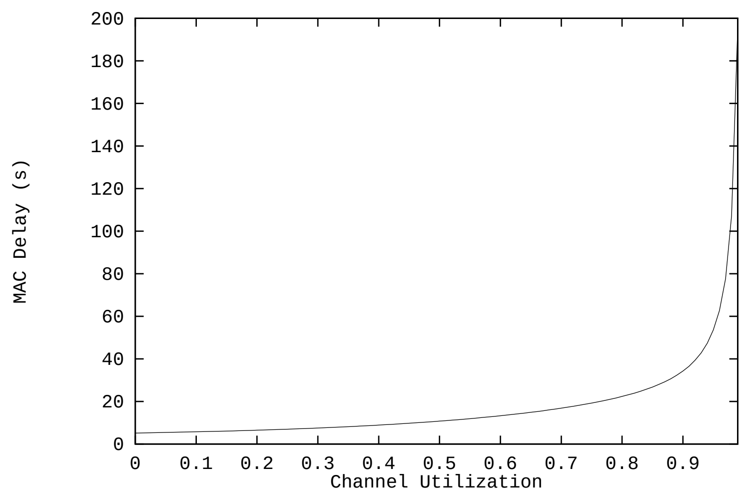

Figure 8: MAC delay vs. net utilization

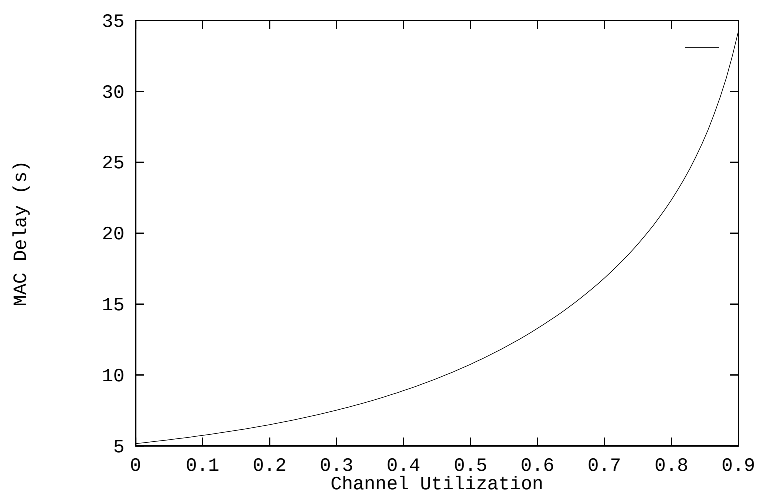

Figure 9: Net utilization ≤ 0.9.

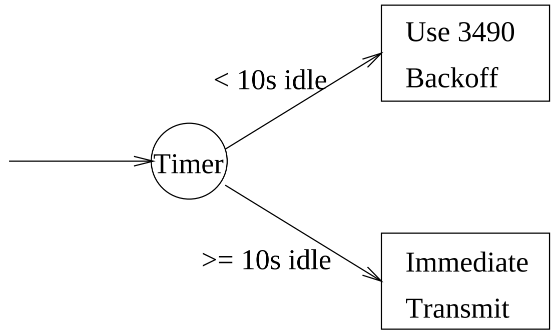

Figure 10: The two protocols in the TDB 3490A.

The calculations of Equation 1 do not take into account that in the 3490A protocol, a channel idle for 10 s or more allows a frame to be immediately transmitted. This tends to decrease average MAC delay and is illustrated by the block diagram of Figure 10.

For a Poisson arrival process, the probability of experiencing a 10 s or longer idle period is just $e^{-10\lambda}$. As an illustration, this is 0.05 for $\lambda = 0.3$ frames/s. To take protocol switching into account, the total expected delay is now:

$$D_{\text{MAC}} = D_{\rho_i}(1 - e^{-10\lambda}). \tag{2}$$

This is important at low traffic rates. For instance, suppose the application requires an average MAC delay of 10 s, Figure 11 shows that until a traffic rate of 0.6 frames/s is reached, the expected MAC delay is lower than 10 s.

Once the expected delay D_{MAC} is determined by the application, then the bound on the outgoing data rate at a node is the maximum size of a frame divided by the sum of the MAC delay and the channel service rate. If the maximum transmittable unit (MTU) is m bits and the average frame length is $L = m/2$ bits, then the *frame* service rate by the channel is $\mu = C/L$. And the system traffic rate in frames per second is bounded by

$$\lambda\left(\frac{1}{\mu} + D_{\text{MAC}}\right) < 1. \tag{3}$$

Equation 3 must be met to ensure system stability. Because of the standard and priority delays built into D_{MAC}, the simpler $\rho < 1$ does not guarantee stability. When the system is comprised of n similar nodes, then we can say

$$n\lambda_i\left(\frac{L}{C} + D_{\text{MAC}}\right) < 1 \tag{4}$$

for each node i.

When any two of the three unknowns are given, Equation 4 can be rewritten three ways to solve for the remaining unknown:

$$\lambda_i < \frac{C}{n(L + CD_{\text{MAC}})}, \tag{5}$$

$$n < \frac{C}{\lambda_i(L + CD_{\text{MAC}})}, \quad \text{or} \tag{6}$$

$$L < C\left(\frac{1}{n\lambda_i} - D_{\text{MAC}}\right). \tag{7}$$

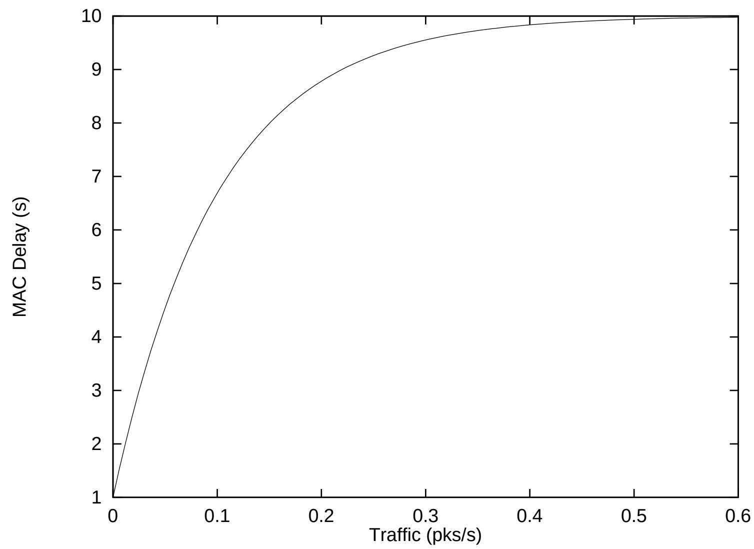

Figure 11: Expected MAC Delay for $D_{\rho_i} = 10$ s vs. traffic rate.

When either choosing or calculating λ_i, it is important that Equation 2 be used to account for reduced MAC delay due to long idle periods.

Typically, the number of nodes in a network is given. Assuming that each node would like to transmit as large a frame as possible as often as possible, it is necessary to find the maximum value for $L\lambda$, or finding the maximum z, where

$$z = \frac{\lambda LC}{(1 - n\lambda D)(C + LD)}. \tag{8}$$

Optimization of z requires nonlinear programming techniques and is beyond the scope of this paper. If, however, the application's requirements determine two of the unknowns, it is simple to find the third. As an example, for given λ and n with application demands of $\rho = 0.3$, the resulting ideal average frame lengths are graphed in Figure 12. (An MTU would be twice that length.) Lengths beyond 24 kbytes are ignored since that is the capacity of the RF-3490A [Cor87]. These results represent a system at full load (worst case) and under the previously mentioned node number and traffic rate constraints. These values are generated for the case when 1/7.5 of the 3490A MTU are data bits [Kas90] which corresponds to a data rate of 2133 bits/s. The plot is intended to show trends rather than to allow study of particular cases. These should be generated individually and in more detail. For instance, a net with five nodes is considered in Figure 13. Suppose, for example, that the application requires that one message be transmitted every 25 s, then the ideally sized MTU is approximately 9 kbits, or 1.125 kbytes. This is intuitively easy to verify: it takes 4.2 s at 2.133 kb/s to transmit 9 kbits. Five nodes using 4.2 s to transmit takes 21 s. This is nearly the 25 s message period with headroom for the backoff algorithm. Also notice that the graphed data in Figure 13 exhibit "bumps" at 0.02 and 0.04 frames/s/node. These are regions where, for the given MTU size, the 3490A switches between the backoff and immediate transmit protocols. For an infinite user population analysis, this would not happen. It does here because the analysis is of a finite system, i.e., n is finite and integer, which must meet the constraints imposed by Equation 7. As a result, there are regions where to maintain stability the optimal MTU size is not the ideal size.

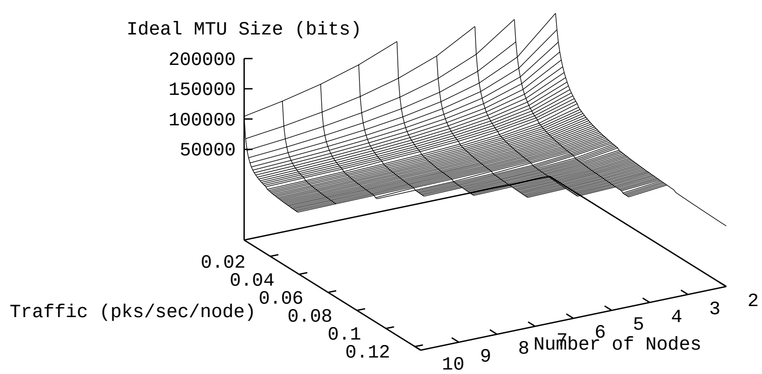

Figure 12: Max MTU size determined by n, λ, and $\rho = 0.3$.

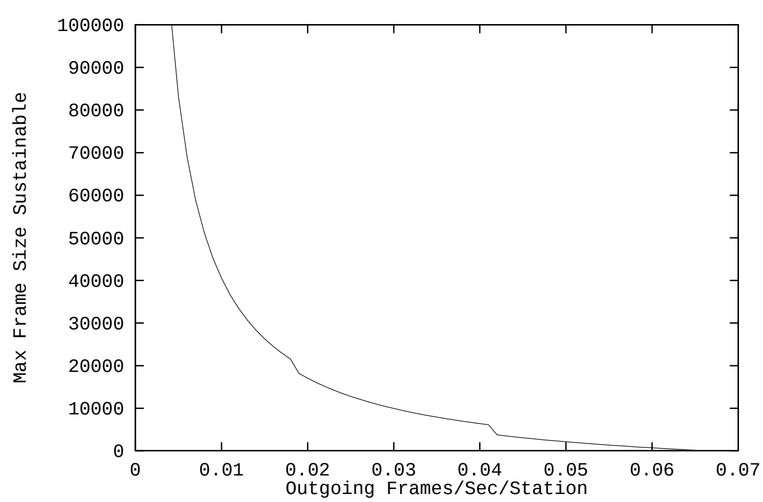

Figure 13: Max MTU sizes for five nodes.

As would be expected, at a given traffic rate, the ideal frame length increases as the number of nodes decrease. Similarly, traffic intensity must remain in a relatively limited range to meet a given utilization value ρ. When utilization moves beyond a maintainable range, values for MTU length, number of nodes, and/or traffic rate must be recalculated. If they are not, the system will operate in a suboptimal state, possibly resulting in an unstable system.

4.2 Performance Above the MAC Layer

The previous section developed an analysis of the RF-3490A independent of the system in which it is used. This section focuses on the performance of the 3490A when used to support the DFB system [Cha90]. While the DFB is a complex software package, its interface to the 3490A is not. Messages between nodes of the DFB communicate via the FEP [Kas90]. The FEP has four priority queues that hold packets smaller than the 3490A's MTU (frame) size. Once per second, the priority queues are concatenated a packet at a time, from highest to lowest priority, until the frame is at most as big as the MTU or until the queues are empty. The frame is then passed to the 3490A, if the 3490A is not busy. The important properties of the FEP are packet size and arrival rate distributions.

It should also be noted that the analysis in this report ignores the possibility of collisions resulting from carrier sensing delay. If β is the amount of time for a signal to propagate and be detected, then any carrier sense performed within β time units after a transmission incorrectly senses an idle net. Its transmission then collides with the one in progress. When the arrival process is Poisson of intensity λ, the probability of a collision is $1 - e^{-\beta\lambda}$. If, for example, traffic rate is 0.4 frames/s, and there is SINCGARS-SIP carrier sense length of approximately $\beta = 100$ ms [SL93], then the probability of collision is 0.039. However, when collisions occur, this affects only protocols at higher layers. Because the 3490A does not take collisions into account in its operation, the occurrence of collisions has no affect on protocol operation. For higher layer characterization of a network using 3490As, though, collisions will certainly have to be considered.

The goal at this layer is to derive expected delays for packets in each priority queue when the system is working at full load. As such, the results from Section 4.1 can be used to calculate the average service rate that the priority queues can expect. With a Poisson arrival process, the first and second moments of the service rates are

$$\overline{X} = E\{X\} = D_{\text{MAC}} + \frac{1}{\mu} = \text{Average service time}$$
$$\overline{X^2} = E\{X^2\} = (D_{\text{MAC}} + \frac{1}{\mu})^2 = \text{Second moment of service time.}$$

Bertsekas and Gallager [BG92] show that for an M/G/1 system the *Pollaczek-Khinchin (P-K) formula* holds:

$$W = \frac{\lambda\overline{X^2}}{2(1-\rho)},$$

where W is the expected customer waiting time. Therefore, the total waiting time of a customer, queue waiting time plus service time, is

$$T = \overline{X} + \frac{\lambda\overline{X^2}}{2(1-\rho)}.$$

Note, however, that we are using the determinstic service of Equation 2, so that the M/G/1 system reduces to an M/D/1 system. As a result, the second moment of the service rate is simply the square of the first moment as shown above. From Bertsekas and Gallager [BG92], the P-K formula is shown to be the basis for priority queuing results in M/G/1 systems. If there are k priority queues, each with their own traffic and service rates, the waiting times in queue k and including service time can be shown to be

$$\begin{aligned} W_k &= \frac{\sum_{i=1}^{n} \lambda_i \overline{X^2}}{2(1-\rho_1-\cdots-\rho_{k-1})(1-\rho_1-\cdots-\rho_k)} \\ T_k &= \frac{1}{\mu_k} + W_k. \end{aligned}$$

Because performance under maximum load is of interest in this analysis, $E\{X\} = 1/\mu + D(\rho)$ is a constant since $1/\mu$ is the radio service rate in MTU/s, and $D(\rho)$ is constant for a given utilization ρ, which is part of the application layer system design specifications. Similarly, $E\{X^2\}$ is a constant. For the sake of illustration, if we additionally assume, probably unrealistically, that each of the four priority queues has the same arrival rate, then an individual queue's arrival rate is $\lambda_k = \lambda_{\max}/4$ where $\lambda_{\max}$ is determined from the analysis in Section 4.1.

Therefore, the priority queue waiting time equations reduce to:

$$W_1 = \frac{2\lambda_k \overline{X^2}}{1-\rho_k}, \tag{9}$$

$$W_2 = \frac{2\lambda_k \overline{X^2}}{(1-\rho_k)(1-2\rho_k)}, \tag{10}$$

$$W_3 = \frac{2\lambda_k \overline{X^2}}{(1-2\rho_k)(1-3\rho_k)}, \tag{11}$$

$$W_4 = \frac{2\lambda_k \overline{X^2}}{(1-3\rho_k)(1-4\rho_k)}. \tag{12}$$

When frame sizes and service rates are fixed, one can expect that as traffic rates increase, queue delay times will become more widely separated with the lowest priority queue experiencing the longest delays. However, by using the analytic results derived to this point, it is possible for the network to adapt to changing situations. As a result, as traffic rates increase, average delay decreases but only because frame size also decreases. There will still be more delay to move a given amount of data on a busier channel. To give fair access to each user, the protocol utilizing these equations would shrink MTU size. There will be some minimum MTU size, below which no data and only packet header information can be passed. This will be the lower bound on performance and the upper bound on sustainable traffic rate. When the application's requirements are used as constants, Equations 5, 6, and 7 can be used to rederive system parameters. For instance, when a node leaves the local net or goes silent, the traffic rate is decreased. This decrease can be used to recalculate a larger MTU which still allows the application dependent maximum delay to be met. Similarly, if packets are noticed to constantly be of smaller length than the current MTU size, the average size can be used as the new MTU length, allowing values for traffic rate and number of nodes that can be supported to be updated.

To illustrate with examples, suppose there are five nodes in a single network. Figure 14 graphs the traffic rate vs. average delay when the utilization is held constant. That is to say, regardless of traffic rate, the MTU size will be adjusted so that the utilization remains constant. In the figure, utilization $\rho = 0.4$ and channel bit rate is 2133 bits/s. Also note that while delays for all four priority queues are plotted, the adaptation performs well enough that the difference in delays is insignificant.

In the next graph in Figure 15, the traffic rate vs. average delay is plotted when the MTU size is held constant. Here, the MTU size is always 1 kbyte with a utilization of

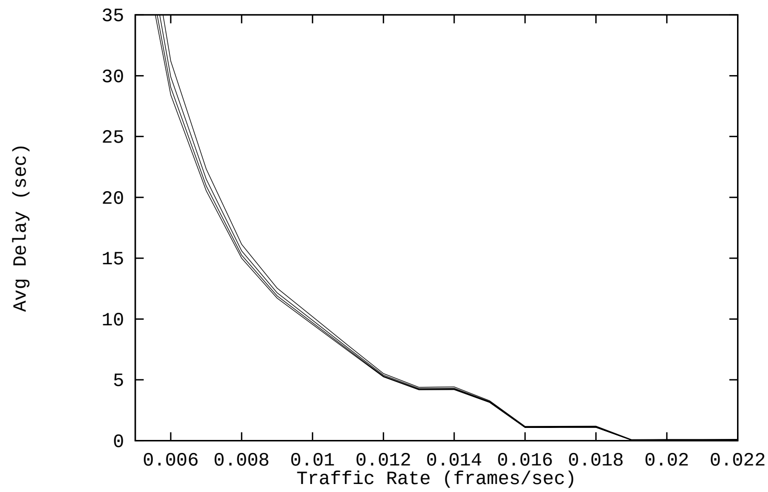

Figure 14: Traffic rate vs. average delay, $\rho = 0.4, C = 2133$ b/s.

0.47. To maximize net utilization while not surpassing $\rho = 0.47$, the traffic rate λ can vary so that the utilization remains constant. This effects average delay in each queue. In this case, there is a significant difference in the delays experienced packets in each of the priority queues with the highest priority queue seeing the lowest delay.

Finally, if the traffic rate is held constant at $\lambda = 0.03$ frames/s, Figure 16 shows that, like in Figure 14, there is no appreciable difference between delays experienced by the different priority queues. The figure graphs utilization vs. average delay.

5 Conclusions and Future Work

Results from the analysis presented in Section 4 show that simple equations are able to describe properties like throughput and delay of Aloha-type protocols as exemplified the Harris 3490A TDB. While the straightforward use of the equations allows for prediction of system throughput and delay, the more interesting use would be in a new adaptive protocol. By incorporating these equations into the protocol, as net conditions such as available bandwidth, number of nodes, acceptable delay, and so on, change, the protocol could calculate new ideal values, thus quickly and easily adapting to changing net conditions.

Before that can happen, however, higher layer model is needed which accounts for collisions and their affect on traffic rates. Other problems should be considered as well; topologically related problems such as the hidden node problem, for instance. It would also be interesting to analyze a model for average case performance rather than worst case, heavily loaded instances.

The results presented in this paper, however, offer new insights towards the design of an efficient system using a modified Aloha protocol for the MAC layer.

References

[BG92] Dimitri Bertsekas and Robert Gallager. *Data Networks.* Prentice Hall, second

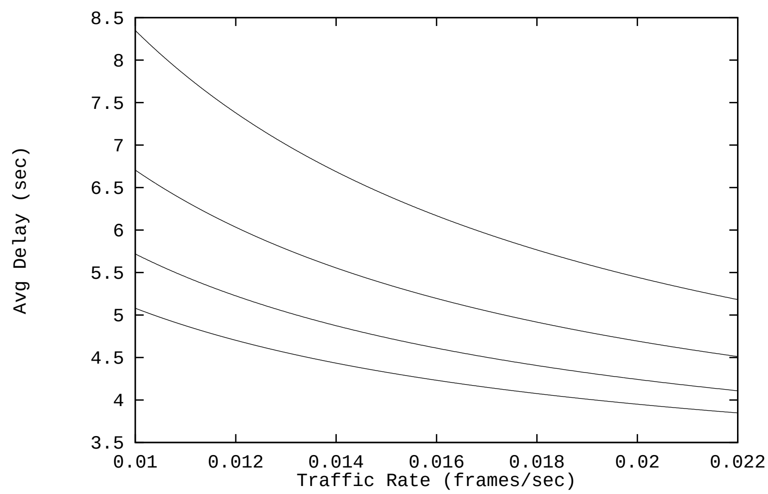

Figure 15: Traffic rate vs. average delay, $\rho = 0.47, MTU = 8192$ bits, $C = 2133$ b/s.

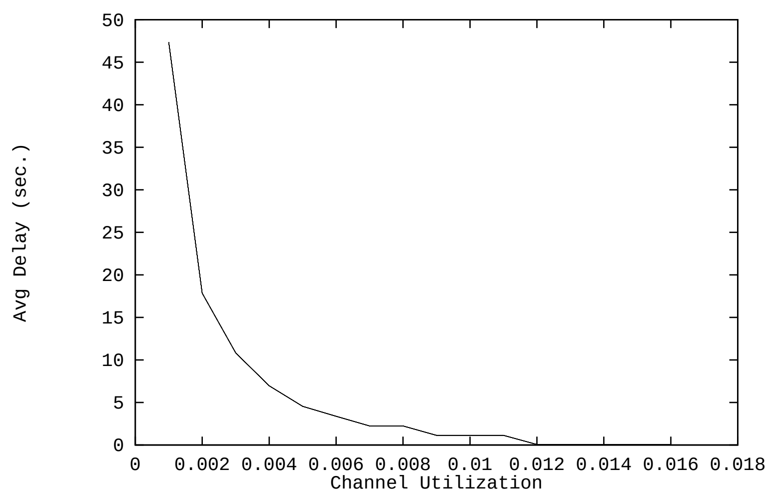

Figure 16: Utilization vs. average delay, $\lambda = 0.03, C = 2133$ b/s.

edition, 1992.

[Cha90] Samuel C. Chamberlain. The information distribution system: IDS–an overview. Technical Report BRL-TR-3114, US Army Laboratory Command, July 1990.

[Cor87] Harris Corporation. *RF-3490 Digital Data Buffer/Processor Instruction Manual.* Harris Corporation, Rochester, NY, September 1987. Publication Number: 10012-0123.

[Kas90] Virginia A. T. Kaste. The information distribution system: The fact exchange protocol, a tactical communications protocol. Technical Report BRL-MR-3856, US Army Laboratory Command, 1990.

[SL93] John Siliato and Jay Latham. Combat net radio (CNR) protocols: A means for battlefield digitization. Technical Report *[Number not given]*, *[Institution not given]*, November 1993.

[Tan89] Andrew S. Tanenbaum. *Computer Networks.* Prentice Hall, second edition, 1989.