



Calhoun: The NPS Institutional Archive
DSpace Repository

NPS Scholarship

Conferences

2005

The Impact of Contention Resolution verses a priori Channel Allocation on Latency in a Delay Constrained Network

Gibson, John; Coelho, Jose; Diaz-Gonzalez, Leopoldo; Xie, Geoffrey G.

<https://hdl.handle.net/10945/34778>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

The Impact of Contention Resolution verses *a priori* Channel Allocation on Latency in a Delay Constrained Network

John Gibson, Geoffrey G. Xie, José Coelho, Leopoldo Diaz-Gonzalez
Computer Science Department
Naval Postgraduate School

ABSTRACT

The predominant mechanism used to control access to the underwater acoustic channel is a contention-based collision-avoidance scheme to reserve the shared media, on-demand, before sending data to avoid retransmission costs incurred with collisions. This paper reports analysis of such a mechanism for a simple linear, or backbone, type topology as compared to the use of an *a priori* allocation of capacity to each host, thereby eliminating the need for such access coordination. The former incurs RTS/CTS propagation overhead while the latter incurs a transmission penalty due to the reduced transmission rate of each allocated channel. Therefore, the selection of an appropriate access mechanism, given consideration of the message size and the propagation distance, may have a significant impact on the overall latency of a message. In particular, given the topology modeled for the study, this paper suggests that relatively small messages or long propagation distances benefit from an *a priori* channel allocation, while large messages or short distances are more effectively served by an access coordination method, such as RTS-CTS exchanges. Further, as the number of frames in a message or the number of hops over which the message is relayed increases the relative latency performance of the *a priori* allocation improves over that of the coordinated access. The results establish *a priori* channel allocation as an appealing media access scheme in delay constrained networks and should provide an impetus for further research so as to maximize the performance potential of such schemes.

I. INTRODUCTION

One of the fundamental characteristics of shared-medium systems, such as contention-bus or wireless networks, is the manner in which the systems manage contention. Contention for the media may result in collisions, which waste the available network capacity. Methods to resolve the contention may be reactive, requiring resolution as collisions occur, or proactive, where the occurrence of collisions is either minimized or eliminated. These measures are embodied in the media access control mechanism employed by the particular network. Various measures are available for controlling access, including: completely uncoordinated access, a form of the Aloha protocol developed by Abramson at the University of Hawaii for packet radio networks; collision avoidance schemes, adapted from aerial wireless networks; token passing [Kurose]; and contention-free *a priori* channelization, that also enable full duplex communications. Uncoordinated access schemes simply respond to collisions, typically by exponential back-off and retransmission of the frame, while the coordinated access schemes, whether distributed, such as the various carrier sense multiple access methods, or centralized, such as polling or token passing, seek to avoid collisions and in so doing mitigate waste of network capacity.

Each of these schemes has been implemented in wired or aerial wireless networks. In both of these cases, per-hop propagation delay is small enough that it is often considered negligible. For

example, typical Ethernet hops have a one-way propagation delay, over 200 meters, of less than 5 microseconds, allowing 4 microseconds for the delay through a hub or switch. For the longest proposed 10 Gigabit Ethernet hop, a 40 kilometer fiber optic link [Stallings], a round trip propagation delay of 400 microseconds, is more than 3000 times smaller than the one-way propagation delay over a 1 kilometer underwater acoustic link. However, a class of networks, referred to as Delay Tolerant Networks (DTNs), exhibits extreme end-to-end delays either due to intermittent periods of link discontinuities or extreme propagation delays, due to low propagation rates or very long propagation distances [Fall].

This paper addresses the impact of media access control for a subclass of DTNs that are constrained by extreme one-hop propagation delays, rather than those that experience planned or unplanned link drop-outs. Of that subclass of delay constrained networks, we are most concerned with networks established using underwater acoustic links, which exhibit a propagation rate of approximately 1500 meters per second and have single-hop distances of up to 2 kilometers. In some cases, the single-hop distances reach up to 4 kilometers or more [Hartfield]. The limited capacity and extreme propagation delays inherent in underwater acoustic networks (UANs) make it imperative that the implementation of such networks consider the impact of the media access methodology on network throughput and traffic latency. An ill-considered or uninformed decision may result in unnecessarily increased traffic latency.

While all coordinated access schemes seek to avoid collisions, a subset of such schemes, commonly referred to as Collision Avoidance (CA), is derived from the Carrier Sense Multiple Access (CSMA) methodology. They use a form of channel reservation to minimize the contention interval. By limiting the contention interval to the time required to exchange short reservation coordination messages, these schemes resolve contention by eliminating the opportunity for data frames to collide. As only reservation requests may collide, collision resolution is limited to using a random back-off delay to limit the likelihood of sequential reservation request collisions. As compared to the uncontrolled access, where a collision may require a larger data frame be retransmitted, the collision avoidance reservation scheme seeks to reduce total retransmission time. Such CA schemes, operating in a half-duplex mode of communications and providing Stop-and-Wait flow control with Automatic Repeat reQuest (ARQ) error recovery, are the dominant method of controlling access in UANs [Freitag].

One of the key issues regarding the use of CSMA schemes with respect to wireless networks is the possibility of network hosts, or nodes, not being within reception range of all of the other hosts. This leads to the hidden node problem for which the channel reservation scheme was designed to address. The wireless contention-based local area network standard, IEEE 802.11, as an option, allows for the use of Request-to-Send (RTS) and Clear-to-Send (CTS) exchanges to reserve media access for sufficiently long frames. The introduction of such a collision avoidance scheme to resolve contention is in recognition of the increased likelihood of collisions given large frame size and the attendant cost of retransmitting such frames. By reserving the channel prior to data transmission, by way of a RTS-CTS exchange, the likelihood of large frame collisions can be eliminated. Since the RTS and CTS messages are small compared to the data frames for which they are coordinating access, and since the propagation delays in aerial wireless communications are on the order of a few tens of microseconds, their incorporation in an aerial network has minimal impact on data frame latency. However, if the assumption of negligible

propagation delays in not valid, as in the case of acoustic communications where the propagation delay associated with a single RTS-CTS exchange may be on the order of one or two seconds, or more, then the overall effectiveness of the CA is not as clear.

For example, a recent experiment by the SPAWAR Systems Center San Diego, in conjunction with Fleet Battle Experiment – India, examined the performance of the collision avoidance mechanism used to mediate access to the shared acoustic channel by a collection of relay nodes. Analysis of the experiment data indicated that just over 80% of the data packets, from a single-hop perspective, were exchanged with only one RTS-CTS exchange attempt. Another 10% were successfully exchanged after the second RTS-CTS attempt [Hartfield]. Of the data messages examined, approximately 88% arrived without error and another 9% were successfully received after a single retransmission. It appears that the collision avoidance and automatic repeat request scheme was very successful in limiting the number of retransmissions. However, it may be the case that the traffic pattern could just as well have been served by an uncontrolled access method, as the lack of collisions for RTS frames indicate that the effective load was low. Unfortunately, the experiment did not address that possibility. Thus, it is worth considering when such exchanges are useful and when they are detrimental.

As an alternative to collision avoidance schemes, contention-free communications may be provided by moderated access via polling, token passing, or assigning each potential source a dedicated transmission channel prior to sending any data. Both polling and token passing introduce additional delay overhead to send either the poll or the token. Dedicated transmission capacity, or *a priori* allocations, may take the form of TDMA, FDMA, or CDMA; however, the extreme propagation delays of acoustic communications greatly complicate the use of TDMA [Hou]. An *a priori* allocation scheme must dedicate to each node a channel unique within its two hop neighborhood. This dedication ensures that a given node's transmission does not conflict with any other node's transmission, thereby eliminating the need to coordinate access before sending data. While eliminating contention, *a priori* allocation incurs a transmission penalty in that the available capacity is divided into the required discrete channels, thereby increasing the transmission time of the data. Thus, there is a trade-off between the contention resolution and the transmission delay.

Each access control method has strengths and weaknesses that should be considered before settling on a particular scheme. Several factors should be considered when selecting or designing a media access control technique. These include the propagation delay between communicating nodes; the typical frame transmission delay or frame duration, the ratio of frame size to transmission rate; network topology; and the expected traffic load and arrival pattern. Each of these factors contributes to the likelihood of a collision at the receiving node requiring retransmission of the effected frames. Each factor presents trade-offs. Understanding these trade-offs is the motivation for this paper. While there are many more schemes proposed in the literature, these two, collision avoidance and *a priori* allocation, are the focus of this study as the former is the most prolific for the underwater environment and the latter has been proposed specifically as a superior alternative for mitigating the effect of large acoustic signal propagation delays in UANs [Xie1, Xie2, Gibson1]. In order to gain insight into the issues that impact the relative performance merit of these two schemes, in terms of *end-to-end message latency*, the

authors analyzed their performance over a simple backbone type network, where all traffic originated at the host at one end of the network and terminated at the host at the other end.

Several observations were made and are examined further in this paper. First, when the propagation delay is large compared to the transmission delay over each hop, the implementation of a collision avoidance mechanism appreciably increases the total delivery delay of the traffic. Additionally, as the number of frames increase, the relative performance of the *a priori* allocation scheme improves, eventually out performing the collision avoidance scheme, for a small number of channels. As the number of *a priori* channels allocated increases, the difference to overcome also increases. Second, increasing the distance between the one-hop neighbors has more adverse impact on networks employing collision avoidance than on networks that use *a priori* allocation. Third, in general, if *a priori* channel allocation is used, sending messages as small frames results in less latency than sending them as fewer larger frames. Fourth, for a given frame size, as the number of hops traversed increases, the performance of the *a priori* allocation improves as compared to that of the collision avoidance scheme.

The remainder of the paper is organized as follows. Section II describes the topology of the network studied. Additionally, it provides the formulas used to assess the end-to-end message latency given the frame characteristics, propagation and transmission rates, individual hop length, and hop count. Section III provides results of the analysis. Finally, Section IV provides a discussion of the key observations and provides recommendations for further study and experimentation.

II. TOPOLOGY AND ANALYSIS DERIVATION

Figure II-1 depicts a simple six node linear trip-wire topology. The nodes are evenly spaced. All traffic originates at Node A and terminates at Node B. Each intervening node forwards the data in turn. Prior to forwarding a data message each node must access the medium according to




Figure II-1: Example Backbone Topology

either the RTS-CTS exchange of the collision avoidance scheme or without delay using the *a priori* channel allocation scheme. If the available capacity is divided into four equal-capacity channels, as shown in Figure II-2, the channels can be assigned in pairs to the nodes such that each node has a channel for communicating with local sensor nodes and a separate channel for communicating with the other backbone nodes. The smaller of the concentric circles represents the range used for the local communications vice that of the backbone links. Careful allocation of the four patterns results in the topology reflected in Figure II-1, the pattern of which may be extended to model a backbone network of arbitrary length. Figure II-1 also demonstrates that no node may share the same backbone channel as any other backbone node within two hops in any direction. Thus, where a channel is used to interconnect a node as part of the backbone, that channel must be unique within that node's two hop neighborhood.




Figure II-2: Channel Pairs

Several simplifying assumptions were made for this analysis. It was assumed that no errors occur during transmission requiring retransmission. While the use of an Automatic Repeat request (ARQ) mechanism would allow for recovery of frames in error, such actions would impact both access mechanisms being considered. It was also assumed that the propagation patterns were regular, although this is highly unlikely in the real environment. Modeling the complexity of the physical environment was beyond the scope of this first-order analysis. The impact of the actual propagation patterns may lead to more dense single- and two-hop neighborhoods. This would increase the likelihood of collisions during the reservation phase for the collision avoidance schemes and require more channels in the *a priori* scheme.

While the simple topology models the backbone relay network typical of a sensor network, it does not consider the introduction of traffic at the internal nodes. Such traffic will result in increased demand on the backbone network and compete for access to the media under the collision avoidance scheme. It is expected that such competition will result in collisions of the RTS messages resulting in increased delays for the collision avoidance scheme, while for the *a priori* allocation scheme these messages would be queued for transmission and would be transmitted as soon as they reach the front of the queue. The same queuing would occur with the collision avoidance scheme as each message must wait for its turn to be forwarded. This analysis leaves the assessment of the impact of local message generation to further work where the impact of queuing on message latency will be assessed via simulation.

Three forwarding schemes for the collision avoidance method were modeled. The first, by which we refer to as Collision Avoidance with Message Switching (CA/MS), considered each message as an entity and all frames comprising a single message were forwarded similar to traditional message switching following a single RTS-CTS exchange. The second, referred to as Collision Avoidance with Frame Switching (CA/FS), forwarded each frame independently, with a separate RTS-CTS exchange prior to each frame. The third scheme, dubbed Collision Avoidance with Frame Pipelining (CA/FP), sought to take advantage of pipelining, such that after a frame had traversed three hops the next frame was sent.

For the *a priori* allocation (AA) scheme, each frame was forwarded as soon as it was received. Processing and queuing delay were assumed to be negligible as no local traffic was added at any relaying node, as noted above. Thus, the *a priori* method allowed for frame or packet switching.

r_p	Signal propagation rate: 1500 meters per second
r_t	Transmission rate
t_p	One-hop propagation time (distance/ r_p)
t_{rlc}	Transmission time of either a request-to-send or clear-to-send frame
t_f	Transmission time of a data frame (frame size/ r_t)
t_{ACK}	Transmission time of an acknowledgment frame
k	Number of frames
n	Number of nodes in message path (length of backbone in nodes)

Table II-1: Latency component variables

Following are the delay formulas for each scheme. As it was assumed that no retransmissions were necessary, both the transmission and propagation delays of the acknowledgment for the

final frame by the final node were not included as the message was available to the recipient at that point. Additionally, the propagation time for each acknowledgement is assumed to overlap the transmission time of the next frame sent by the relay node as that node sends the acknowledgement. Table II-1 provides succinct definitions for the parameters used in the formulations.

Equation (1) gives the formula for calculating the end-to-end latency when CA/MS is utilized. The factor, $2(t_{r/c}+t_p)$, is the time to perform the RTS-CTS exchange; while $(n-1)$ is the number of hops the message traverses. The total time to transmit the message is the number of frames multiplied by the frame transmission delay, which is determined by the frame size divided by the total link capacity. Only one acknowledgement is sent by each relay node in order to minimize the delay induced by the acknowledgment functionality. It is assumed that the acknowledgment provides sufficient granularity to ensure only frames received in error are retransmitted.

$$(n-1)[2(t_{r/c}+t_p) + kt_f + t_p + t_{ACK}] - t_{ACK} \tag{1}$$

In CA/FS, each frame is individually acknowledged before the next frame can be sent. This scheme provides a level of granularity over the previous scheme without adding complexity to the acknowledgment method. However, only a single RTS-CTS exchange was included for each hop. The additional propagation delay for each acknowledgment results in this scheme increasing the overall latency of the message. Equation (2) gives the formulation of this method. If an RTS-CTS exchange is required for each frame then the total latency must also include k factors of the exchange time for each hop.

$$(n - 1)((2(t_{r/c} + t_p) + k(t_f + t_{ACK} + 2t_p) - t_p) - t_{ACK}) \tag{2}$$

The third avoidance method considered, CA/FP, provides for a degree of pipelining and represents a form of packet switching. Figure II-3 depicts the effect of this pipelining scheme.

Since the entire capacity is shared by all nodes, each frame must be forwarded a sufficient distance from its source before the next frame can be transmitted. This distance can be seen to be a minimum of three hops. The reason for this is that the transmission of a frame over the third hop would result in interference with any traffic between the nodes of the first hop. The first term of Equation (3) represents the impact of the pipelining. It can be rewritten as $[n - 1 + 3(k-1)]$. Thus, it is the number of hops plus a delay of three hops for

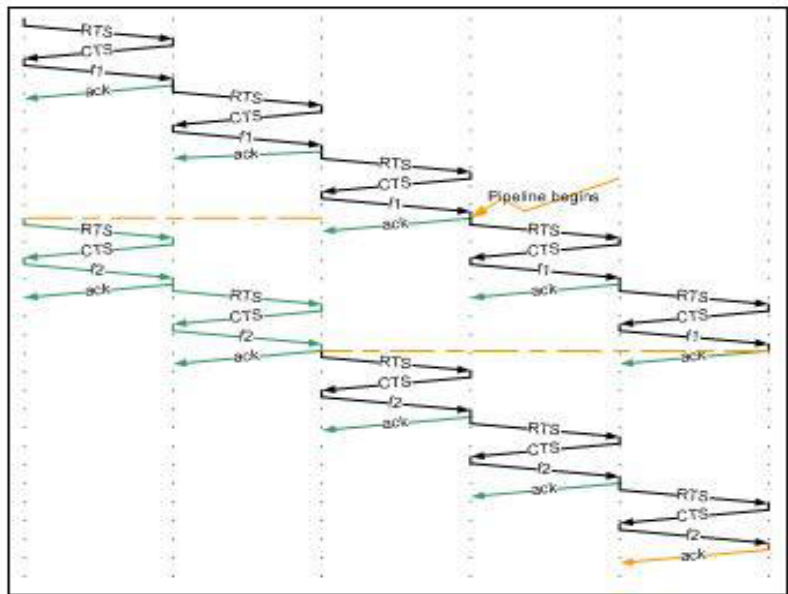


Figure II-3: CA/FP Timing Diagram

all but the first frame. The propagation time of the acknowledgment frame overlaps with the transmission time of the data frame over the next hop.

$$(n + 3k - 4)[2(t_{fc} + t_p) + t_f + t_p + t_{ACK}] - t_{ACK} \quad (3)$$

The final scheme considered, a priori allocation (AA), provides for true packet switching by assigning to each node a channel unique to its two-hop neighborhood [Xie]. This *a priori* allocation results in a reduced transmission rate for each allocated channel. Thus, the transmission time for this scheme must be increased by a factor of the number of channels into which the original capacity was divided. Therefore, t_f is scaled by the number of channels generated. It is this increase in transmission time that is in opposition to the reduction of the message latency by eliminating the need for RTS-CTS exchanges. Equation (4) provides the formula for calculating the resulting end-to-end delay using the *a priori* allocation scheme.

$$(n - 1)(t_f + t_p + t_{ACK}) - t_{ACK} + (k - 1)(t_f) \quad (4)$$

The first and second terms of Equation (4) calculate the time to send the first frame across the network, discounting the final acknowledgment. The last term represents the time necessary to forward each of the remaining frames across the final hop. Note that this method provides for pipelining the message across all hops without the need to delay transmission of a frame until the previous frame has traversed any hops. A node may forward the next frame as soon as it has received it, assuming it has finished transmitting the previous frame. Since subsequent frames may be sent as soon as the previous is transmitted, the remaining frames arrive without further propagation delay as they are in the channel as the previous frame is being received. This is the fundamental benefit of pipelining. By embedding the acknowledgements in the link layer header, such that a single header might acknowledge frames from several one-hop sources, the overhead associated with acknowledgements may be reduced, further minimizing the latency. Note that Equation (4) does not include the time for transmitting or propagating the acknowledgment for each frame across the last hop. This recognizes that the final node will acknowledge the frames using its dedicated channel, thus not delaying the transmission of any frames by its neighbor.

Each formula was entered in the Microsoft Excel™ spreadsheet and the various parameters were varied individually in order to assess their impact on the overall message latency. The results are presented in the next section.

III. PERFORMANCE COMPARISON

The simplifying assumptions allowed for direct manipulation of the parameters affecting the latency. Once the formulas were coded, the individual variables could be varied to determine their individual effect on the total latency, thereby allowing for parametric analysis. As each formula is linear with respect to any individual variable, the resulting graphs are linear if only one independent variable is allowed to vary. The respective slope of each line indicates the degree to which that variable impacts the latency.

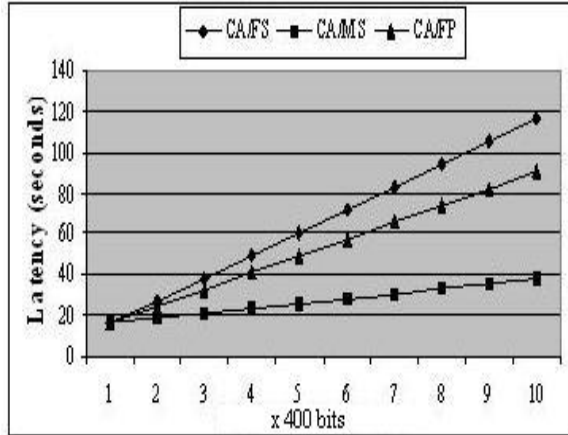


Figure III-1: Relative Performance of the CA Variants

Three variants of the Collision Avoidance scheme were considered as derived by Equations (1)-(3) and are presented in Figure III-1. The figure shows the end-to-end latency for messages ranging from 400 to 4000 bits, divided into 400 bit frames, over a six hop network composed of one kilometer hops. When the message is only one frame, the performance of the three schemes is identical, as would be expected. The best performing of the three was observed to be message switching, CA/MS, while the worst performing was transmission of independent frames, CA/FS. The performance of Frame Pipelining, CA/FP, fell between the other two. As the relative performance of CA to AA is the focus of this paper, the remainder of the analysis considered only CA/MS and AA.

Impact of Hop Distance and Propagation Rate

One of the fundamental differences between acoustic communications and aerial (radio frequency) communications is the signal propagation rate. It is this difference that significantly impacts the performance of collision avoidance schemes. Figure III-2 provides insight into the relative performance of CA and AA over both aerial and acoustic links. The network configuration ranged from 3 to 15 hops,

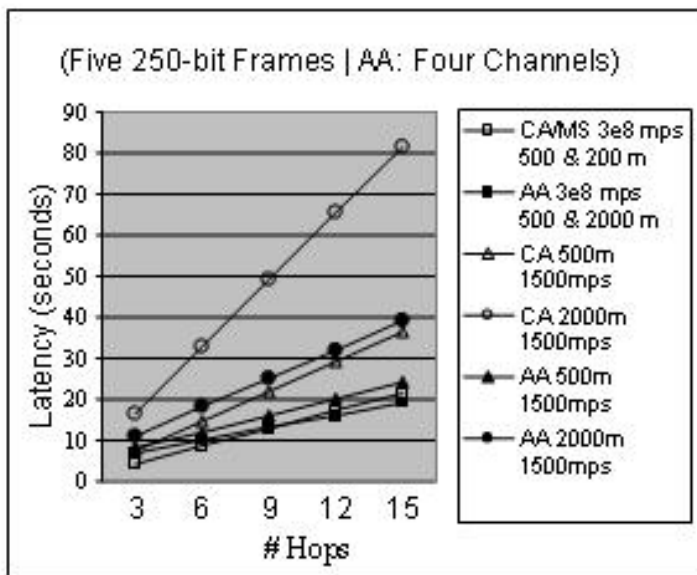


Figure III-2 Impact of Propagation Rate, Hop Distance, and Hop Count

Relative Performance of the CA Schemes

The network configuration ranged from 3 to 15 hops, where each hop was either 500 meters or 2000 meters, as indicated by the legend. Four channels were provided for the AA scheme and the message size was 1250 bits, divided for the AA scheme into 5 frames. As the total capacity modeled was 1000 bits per second, each of the four channels implementing the AA scheme were modeled at 250 bps. The effect of single hop length on latency over aerial links is negligible for both CA and AA schemes, where the propagation rate is 300,000,000 meters per second. Indeed, the difference in latency for hops of 500 meters or hops of 2000 meters was less than 1 percent and, therefore, Figure III-2 has a single line for each of the aerial access methods.

The CA/MS scheme is impacted to a larger degree with increasing hop length for the acoustic communications links. The CA scheme performs better for the 500 m hops than does AA, but

only for two-hop networks. For longer hops distances and for hops counts greater than two, the AA scheme performs better. As hop counts increase beyond nine, AA performs better even for aerial links, given a channelization degree of four.

This suggests that for longer backbone networks an AA scheme is superior to a CA scheme.

Impact of Hop Count & Distance, and Message Size

It was observed that both the number of hops traversed and the single-hop distance impacted the relative performance of the two schemes considered. Figure III-2 also shows that as the number of hops traversed increases the AA scheme's relative performance to that of the CA/MS scheme continues to improve. This impact is more apparent in Figure III-3. Here a message composed of 200-bit frames is sent across the network. As the number of hops is increased the performance of the AA scheme, composed of four channels, continues to improve over that of the CA/MS scheme. Note that the total message size, depicted as the number of frames is also an important factor. Where the message is 400 bits, the CA/MS scheme performs worse than the AA scheme for all hop counts and 1-hop distances considered. However, when the message is 1400 bits, the CA/MS scheme performs better over smaller hop counts, but for greater than 3 hops, the AA scheme is superior.

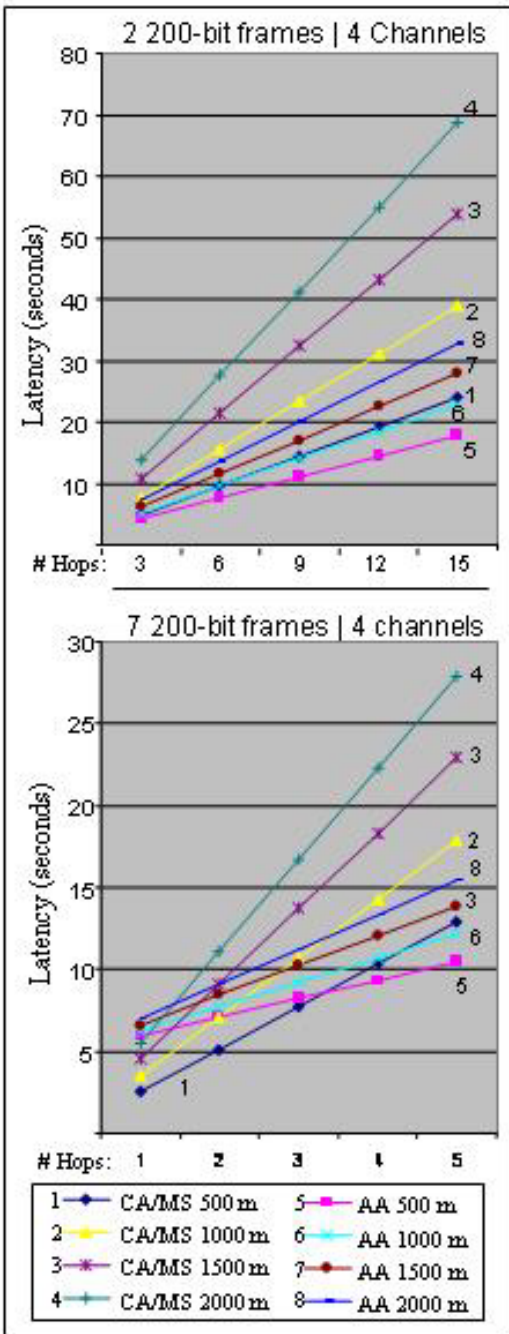


Figure III-3: Impact of Hop & Frame Count

If the 1400 bit message is reduced to a single frame, then the CA/MS scheme out-performs the AA scheme over all hop counts and 1-hop distances considered

Impact of Channelization Degree




Figure III-4: Impact of Channelization Degree

Increasing the number of channels into which the network's transmission capacity is divided, so as to support increased node densities, adversely impacts the performance of the AA scheme. As the degree of channelization increases it becomes more and more difficult for an AA scheme to match the performance of the CA scheme. Figure III-4 highlights this impact over both a five and seven hop backbone network. Three message lengths were considered: 500 bits, 3000 bits and 15000 bits. The graph clusters the respective curves accordingly. Each curve represents the latency for the given message size where the first data point of a curve is the latency if CA/MS is used and each of the other data points reflect a particular frame size for the AA scheme. Each of the four curves radiating from the CA point represents a different degree of channelization, ranging from four to seven channels. Note that the transmission rate for the CA/MS scheme is 1000 bps, while the transmission rate for each of the AA curves is 1000 bps divided by the number of respective channels.

As the degree of channelization increases, it is more difficult to find a frame size for which AA can out perform CA/MS. Further, as the message size increases the percentage of total time the CA/MS scheme is stalled waiting for handshake messages to propagate decreases. At the same time, the amount of time spent transmitting the first frame for the AA scheme increases as the channelization degree increases. With five 1-km hops, if more than

five channels are allocated the AA scheme is unable to improve over the latency of the CA/MS scheme if the message size is at least 3000 bits. However, for hop counts of seven or more and channelization of up to seven channels, the AA scheme matches or exceeds the performance of CA/MS for even extremely large (more than 200,000 bits. This underscores the advantage of data pipelining that is enabled by the a priori allocation scheme.

Summary of Results

While each of the parameters impacted the latency to some amount, they affect the two general schemes in varying degrees. For example, while the degree of channelization is not an issue with CA/MS, as no channelization is required, it only requires subdividing the total capacity into a few channels before the impact of the decreased transmission rate is apparent for the AA scheme. Channelization should not be considered in absence of the other factors, as the hop count, single hop distance, and message size may compensate for the reduced transmission rate. Increasing the hop count and decreasing the 1-hop distance so as to maintain the same backbone length adversely impacts both schemes, however CA/MS suffers to a greater degree than does AA. Thus, if the CA/MS scheme is used the total hop count must be minimized by increasing the distance between relay nodes. This does not consider the impact to power consumption of transmission range, only the impact to net message latency. If shorter hop distances are used to conserve power in spite of the adverse impact to latency, then AA provides benefit over CA/MS for the increased hop count. Overall it appears that the CA/MS scheme is much more sensitive to hop count and message size than is AA.

IV. CONCLUSIONS AND RECOMMENDATIONS FOR FURTHER STUDY

It is clear from the analysis results that the selection of the access methodology can have a significant impact on end-to-end message latency. While the network analyzed had many simplifying it is nonetheless a reasonable approximation to some UAN implementations. However, the simplifying assumptions allowed for the impact of frame size, transmission rate, propagation distance, total frame size, and number of hops traversed on total message latency to be examined in light of the media access methodology. Following are several observations from the analysis:

- Collision avoidance (CA), implementing a form of message switching, improves latency over *a priori* allocation (AA) as message size increases relative to distance for very limited hop count networks.
- Attempting to pipeline CA-based schemes results in increased latency as the ratio of frame delay to propagation delay decreases. Further, in order to prevent interference between frames they must be spaced at least four hops apart. That is, before a new frame can be sent the previous must have been able to complete three hops.
- The relative performance of AA compared to CA can be improved by dividing messages into smaller blocks to allow more pipelining over the AA links.
- The degree of channelization has a significant impact on the utility of the AA scheme.
 - Given a channelization of at least four and a hop distance of one kilometer, for hop counts less than three, CA/MS provides better latency than does AA.
 - For a channelization of four, AA will outperform CA/MS over at least four hops for any message size considered. Thus, large files such as imagery will be better served over backbones of four or more hops using an AA scheme.
 - For hop counts of seven or more, with a channelization of up to seven, AA outperforms CA/MS.
 - Further, as the degree of channelization increases, in order for AA to surpass the latency performance of CA the hop count must increase. Thus, the longer the

backbone network, in terms of hop count, the greater the degree of channelization can be sustained without resulting in decreased performance relative to CA/MS. However, decreasing the hop distance in order to increase the hop count for a given backbone length, in terms of kilometers spanned, results in an overall increase in latency for either scheme.

- 3 Channels allows for AA of simple backbone type configuration.
- 4 Channels allows for AA along the backbone with concurrent CA local networks.

From these observations the authors conclude that the medium access methodology must consider the traffic being supported and the topology of the network. These factors impact the number and size of frames generated, as well as the number of hops traversed, the length of each hop, and the available transmission rate. The available transmission rate is constrained by the number of channels that must be allocated in order to dedicate to each node a unique channel within a two-hop neighborhood.

This analysis should be extended in the following ways:

- More extensive topologies should be addressed. While the trip-wire or backbone topology mimics some of the current UAN implementations, denser node neighborhoods should also be considered to measure the full impact of the access schemes.
- More realistic traffic characteristics should be included. The effect of frame loss and the insertion of traffic at internal nodes may have a large impact on the performance of a given access protocol. These changes will impact the overall traffic load of the network. The simulation should also allow for an analysis of the relative performance impact of the collision avoidance scheme under various intensities of traffic load. It is anticipated that for very low loads a very simple Aloha-like access mechanism may be appropriate.
- The effect of queuing should be modeled. Since only the transmission of a single message was considered in this analysis, it can be expected that the recommended changes will impact queuing at each hop.
- This analysis assumed that the physical characteristics of all links were constant. However, in practice, the physical parameters of each link may vary. The available transmission rate may vary from hop to hop due to the non constant nature of the water channel. Such things as ambient noise, temperature, salinity, or turbidity can impact the available signal to noise ratio thus limiting the theoretic maximum data capacity. Further, a non zero bit error rate will either require frame retransmission or sufficient redundancy in the data the support forward error correction. Finally, the individual hop distances will likely vary, changing the relative propagation delay for each hop. The composite effect of these and other factors can have an impact on the appropriate access scheme and frame size. A more thorough model should consider such differences in the link-by-link performance.
- Further ideas should be explored, such as Aloha-like protocols for low traffic loads or dynamic capacity allocation, such as bandwidth-on-demand to mitigate channelization induced transmission delays, should be considered.

Finally, the result should be validated by experimentation in situ to validate the model.

REFERENCES

- K. Fall, "A Delay Tolerant Network Architecture for Challenged Internets," Intel Berkeley Research Technical Report 03-003, Feb 2003
- L. Freitag, M. Stojanovic, M. Grund, S. Singh, "Acoustic Communications for Regional Undersea Observations," Proceedings of Oceanology International 2002, London, March 2002
- J. Gibson, A. Larraza, J. Rice, K. Smith, and G. Xie, "On the Impacts and Benefits of Implementing Full-Duplex Communications Links in an Underwater Acoustic Network," Proceedings of the 5th Mine Symposium, April 2002
- J. Gibson, G. Xie, A. Kaminski, "Demand Assigned Channel Allocation Applied to Full Duplex Underwater Acoustic Networking," pending publication by PACON International (2005)
- B. Hou, O.R. Hinton, A.E. Adams, and B. S. Sharif, "An Time-Domain-Oriented Multiple Access Protocol for Underwater Acoustic Network Communications," Proceedings of the IEEE Oceans '99 Conference, 1999
- G. Hartfiled, "Link Layer and Network Layer Performance of an Undersea Acoustic Network at Fleet Battle Experiment-India," Master's Thesis (MS-ISO), Naval Postgraduate School, June 2003
- J. Kurose, K. Ross, Computer Networking A Top-Down Approach Featuring the Internet (2nd Ed.), Addison Wesley, 2003
- W. Stallings, Data and Computer Communications (7th Ed.), Pearson-Prentice Hall, 2004
- G. Xie, J. Gibson, "A Network Layer Protocol for UANs to Address Propagation Delay Induced Performance Limitations," Proceedings of the MTS/IEEE Oceans 2001 Conference, November 2001
- G. Xie, J. Gibson, K. Bektas, "Evaluating the Feasibility of Establishing Full-Duplex Underwater Acoustic Channels," Proceedings of the Third Annual Mediterranean Ad Hoc Networking Workshop, June 2004