IP Multicasting in HF Radio Networks

Eric E. Johnson, *Senior Member, IEEE*
New Mexico State University
Las Cruces, NM, USA

Abstract

Military applications for multicasting abound. In this paper, we propose a multicast data link protocol for third-generation (3G) high frequency (HF) radio networks, and evaluate its ability to support IP multicasting of military messages in HF radio networks.

Background

Multicasting is a one-to-many traffic delivery technique employed at the physical, link, and/or network layers. Traffic flowing from an originating node is delivered more or less simultaneously to multiple receivers. In some applications, only one node originates traffic to its multicast group, while in others any member of a multicast group can originate traffic that is sent to all members of the group.

The Internet Protocol (IP) address of a multicast group is easily distinguishable from a unicast or broadcast address. The ubiquitous version 4 (IPv4) addressing scheme reserves the range 224.0.0.0 through 239.255.255.255 for collective addresses of multicast groups. The 224.0.0.* range is reserved for use by routing and similar protocols; the remaining multicast addresses are allocated for a variety of experimental and operational purposes. Similarly, IPv6 multicast addresses are distinguished by a high-order octet of 0xFF.

IP routers that support multicasting process packets addressed to multicast addresses differently than unicast or broadcast packets. In wired networks, routers build a tree to minimize a "cost" of delivering traffic to the group. Each link in the tree carries a single copy of a multicast packet, which will be replicated downstream at a router where the tree branches. Therefore no multicast packet will appear more than once on any link.

A more robust alternative to the multicast tree is a mesh: no link carries a multicast packet more than once, but packets may follow multiple links in parallel to provide spatial path diversity. With this diversity, multicast delivery becomes more robust to noisy or lost links.

In wireless networks, the physical medium is inherently a broadcast channel, so building an efficient multicast topology consists in identifying the stations that will re-broadcast the traffic, rather than the subset of outgoing links at nodes into which each packet will be relayed [1].

Military Applications of Multicasting

Despite substantial investment in infrastructure to support network centric warfare, many communication links to and among warfighters have relatively low data bandwidths. This can result in congestion and long delays if this bandwidth is not managed optimally.

Happily, much of the necessary information flow parallels the hierarchical structure of military units, and can naturally use multicasting through the tree topology often available in military networks. Examples of such traffic include orders flowing down the chain of command, and many varieties of situational awareness updates (Blue Force Tracking, Common Operating Picture, and so on).

Military Messaging

In the 21st century, formal military messaging (i.e., record traffic) is migrating from the ACP 127 standard to an approach based on the civil standard X.400, as extended for military applications by ACP 123 and STANAG 4406 [2].

For narrowband tactical networks (lower than 20 kbps), the NATO Military Message Handling System (MMHS) working group developed Annex E to STANAG 4406. This Annex includes application level messaging protocols and related lower layer protocols, as well data compression, encoding techniques, and address mapping.

In Annex E, a Tactical Interface Agent (TIA) is defined as the interface point between high-bandwidth strategic networks and low bandwidth tactical networks. The ACP 142 P_Mul protocol [3] is specified for multicasting and for use with stations in EMCON (radio silence).

P_Mul provides reliable (acknowledged) multicast delivery under challenging tactical conditions, including low bandwidths leading to long delays, and stations in EMCON that cannot return timely acknowledgements.

Multicast group IDs in P_Mul over IP are multicast IP addresses. Both static and dynamic creation of multicast groups and selection of multicast addresses are supported.

P_Mul uses an underlying multicast network for delivering messages. When used over an IP network, P_Mul uses the connectionless transport protocol UDP for both transmission of the multicast data packets and the return of acks. Thus P_Mul does not require reliable (e.g., ARQ) service from the underlying network(s).

Because of the natural fit of military message multicasting with hierarchical military communications networks, multicasting over fixed networks, satellite networks and line-of-sight radio networks is well developed. However, it is only recently that high-frequency (HF) radio networks have become a viable alternative to these networks for formal military data messaging.

Characteristics of HF Radio

The key features of HF radio are beyond-line-of-sight communications (global reach) at low cost, and with relatively easy mobility. In HF networks, the range of each individual station is so large (thousands of kilometers) that in many cases the originator of a multicast will be able to reach all members of its group directly. In this paper, therefore, we focus on physical and link layer mechanisms for delivering traffic over challenging HF channels, and do not address network layer aspects such as building and maintaining trees or meshes for multicast traffic routing.

HF radio channels are still challenged by limited bandwidth (at most 9600 bps in 3 kHz) with widely varying error characteristics. However, with channel grouping, 64 kbps throughput has been demonstrated.

For skywave (ionospheric) HF channels, selecting a good frequency is required, since the band of usable frequencies varies throughout the day, and also depends on the locations of the participating stations, the season, the solar environment, and the noise environment at the receiver. Since the 1980s, the process of finding and tracking good frequencies has been computerized using a technology usually called Automatic Link Establishment (ALE).

Two generations of HF ALE technology are currently fielded in large numbers: second-generation (2G) equipment is widely used in large, long-haul networks. The newer 3G equipment offers special advantages in speed and power consumption that are most valuable in tactical applications, although use of 3G HF in strategic applications is contemplated as well.

Previous Work in HF Multicasting

The 2G data link protocol in STANAG 5066 offers a non-ARQ (broadcast) mode that naturally supports P_Mul multicasting. The lack of an equivalent mode in the 3G HF protocol suite has been noted by NATO communications planners, and a request to develop such a 3G HF multicast capability led to two papers published in the MILCOM 2005 proceedings:

- Zhang and Johnson [4] offered a non-ARQ proposal based on redundant transmission of the existing 3G burst waveforms.

- Koski [5] presented a more advanced proposal (patent pending) that efficiently incorporates negative ac-

knowledgements into the data link protocol in a scalable fashion, effectively providing the reliability of an ARQ protocol without time-consuming roll calls.

Either of these proposals for a Multicast Data Link (MDL) protocol could easily be incorporated into the existing 3G suite as shown in Figure 1. Because P_Mul does not require a reliable data link, we use the simpler non-ARQ approach in this paper.
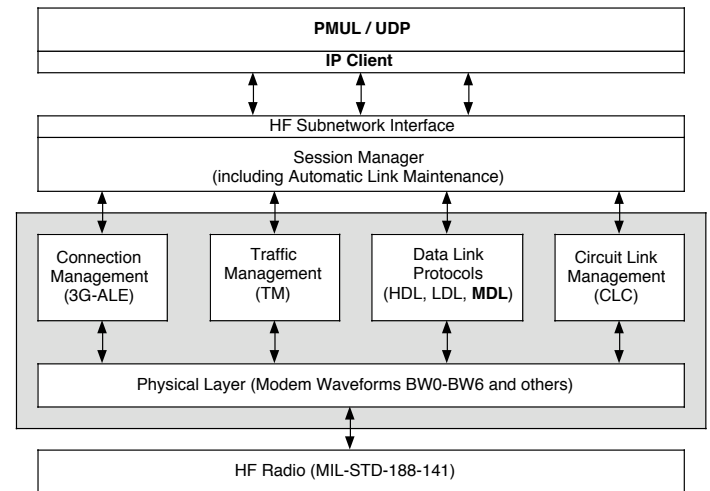
Figure 1: 3G Suite with Multicast Support

IP Multicasting in 3G HF Networks

As seen in Figure 1, support for P_Mul-style multicasting is provided by 3G HF radio networks at the two lowest layers. Robust burst waveforms in the physical layer (modem), along with code combining minimize errors in link layer frames. As discussed above, the link layer uses a non-ARQ protocol to provide best effort service. It may deliver datagrams out of order, because both ordering and reliability are relegated to upper layers.

IP Interface to HF Bearer Service

For transparency for the IP clients of HF radio subnetworks, the IP interface to 3G HF is being revised to be identical to that for 2G networks (i.e., STANAG 5066). IP datagrams are passed to and from the HF bearer service in S_UNIDATA service data units. Within the HF Session Manager, these datagrams are queued by destination and priority, and the Session Manager makes and breaks Soft Links (i.e., HF connections) to deliver them.

The interface between the Session Manager and the MDL protocol carries the IP datagrams along with metadata of two categories: universal information such as addressee(s) and priority of the traffic, and MDL-specific information such as robustness desired for the datagram (e.g., waveform and number of repetitions).

Multicast Link Establishment

The previous paper introducing 3G HF multicasting [4] described link setup for a multicast using the Robust Link Setup (RLSU) mechanism in STANAG 4538. Most currently fielded 3G HF networks use the alternative Fast Link Setup (FLSU), which will be described here.

In most networks that use FLSU, all stations synchronously scan a common set of frequencies, with all stations listening for calls on the same frequency during each dwell period. Thus, once the Connection Manager has selected a frequency suitable for sending the traffic, reaching all stations in a multicast group can be accomplished in one step: a One-way Point-To-Multipoint (PTM) FLSU_Request PDU (Figure 2) is sent on that frequency, and all group members that hear the multicast call immediately stop scanning and listen for the multicast to begin.

001	Pri	Addresses	100	Chan	Traf Type	CRC
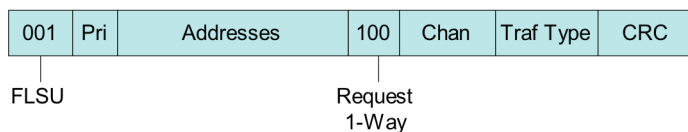
FLSU   Request 1-Way

Figure 2: One-way Point-To-Multipoint (PTM) FLSU Request PDU

- The PTM FLSU Request PDU specifies the group address as its destination and the originating station address as its source.

- The Chan (Channel) field is normally set to '111111' to indicate that the channel carrying this PDU will also be used for the traffic.

- The Traf Type (Traffic Type) field indicates which of the MDL modes described in the next section will be used to carry the traffic.

It is possible that one or more multicast group members may miss this call, either because of bad propagation or because those stations were occupied on other channels at the time of the call. Therefore, the MDL protocol provides for periodically embedding One-Way PTM FLSU Request PDUs in the multicast data stream at times that such stations would be dwelling on the traffic channel if they are scanning for calls.

Multicast PDUs

The MDL protocol uses three PDUs to offer a range of speed and robustness: the burst waveform 2 (BW2) packet from the 3G HDL protocol, and two versions of the BW3 waveform from the 3G LDL protocol, one with a 32-byte payload (most robust) and one with 512 bytes.

BW2 MDL PDU

In high-SNR channels, the BW2 PDU offers throughput approaching 4,000 bps. Each PDU carries 233 bytes of a datagram, along with Start and End flags, a Packet Byte Count, and a Packet Number that identifies the position of that 233-byte packet within a datagram of up to 14 kB. 24 such packets are sent per transmission.

Just as in the HDL protocol, the BW2 packet in MDL uses a rate 1/4 channel code, but only one of the four sets of coded bits is sent in any transmission of the packet. The Traffic Type field of the FLSU, FTM, or TM packet that precedes a BW2 MDL transmission indicates which set of BW2 bits will be sent for each packet (and datagram) during that transmission. To send additional FEC bits, a new FLSU, FTM, or TM packet will be inserted before those additional bits are sent.
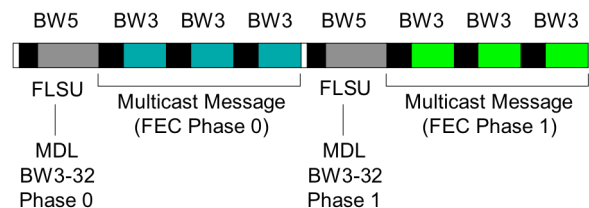
BW3 MDL PDUs

In low-SNR channels, the BW3 PDU offers robust packet delivery, with 90% reliability down to around +3 dB SNR in a 3 kHz channel. In a manner similar to the BW2 waveform, the coded BW3 burst contains one of two sets of FEC bits (rather than four sets in BW2). For the MDL protocol, the longest (512 bytes) and shortest (32 bytes) versions of the BW3 burst are used.

MDL Operation

An illustration of the integration of FLSU and MDL PDUs is shown in Figure 3, where we see

- the initial FLSU PDU that establishes the point-to-multipoint link,

- a series of three BW3-32 PDUs that carry a short multicast message,

- a second FLSU PDU that announces a repetition of the message in the alternate FEC code, and

- a second series of BW3-32 PDUs that carries the same message in that alternate FEC code.

Note that an FLSU PDU is used to switch the FEC code phase rather than FTM so that multicast group members that missed the first call might be captured at that point.

BW5   BW3   BW3   BW3   BW5   BW3   BW3   BW3

FLSU   Multicast Message (FEC Phase 0)   FLSU   Multicast Message (FEC Phase 1)

MDL BW3-32 Phase 0   MDL BW3-32 Phase 1

Figure 3: Example MDL Transmission

Expected MDL Performance

Although the new MDL has not yet been standardized and fielded, the proposed PDUs are sufficiently similar to current HDL and LDL PDUs that we can estimate their performance in fading HF channels. Figures 4 and 5 present estimated frame error rates versus 3 kHz SNR for the MDL PDUs based on BW2 and BW3 (the latter in both 32-byte and 512-byte sizes).
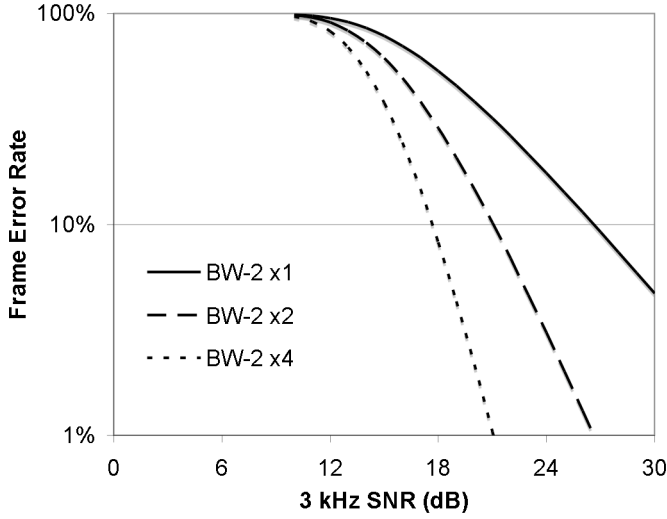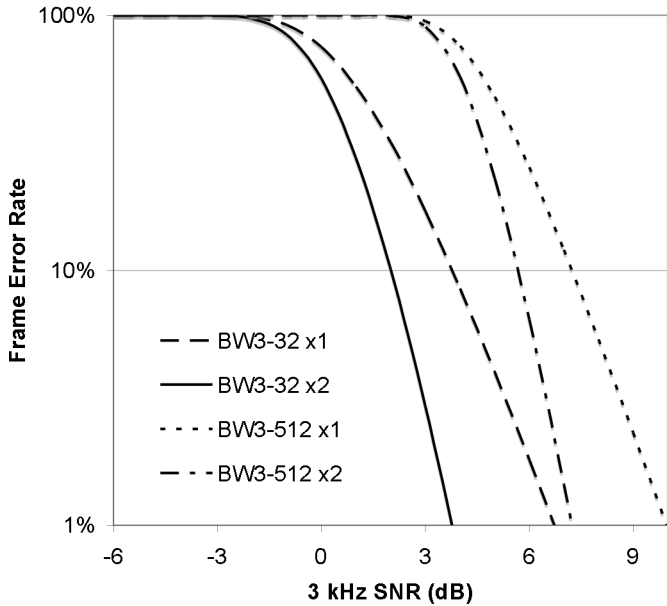

Figure 4: Estimated Frame Error Rates for BW-2 PDUs

Figure 5: Estimated Frame Error Rates for BW-3 PDUs

We can also estimate the potential throughput of the proposed waveforms, as a function of the channel SNR and message size. Table 1 lists some initial estimates of goodput for three message sizes at SNR thresholds at which

each waveform listed achieves at least 90% error-free frames.

Table 1: Goodput Estimates at 90% Reliability

3 kHz SNR	Best Waveform	32 Byte Msg	300 Byte Msg	100 kB Msg
3 dB	BW3-32 x2	68 bps	129 bps	152 bps
4	BW3-32 x1	129	244	287
8	BW3-512 x1	*	277	537
18	BW2 x4	*	**	1,002
22	BW2 x2	*	455	2,033
27	BW2 x1	*	888	3,966

* Best throughput is 141 bps using BW3-32 x1
** Best throughput is 289 bps using BW3-512 x1

Evaluation

In this section, the proposed protocol suite for IP multicasting in 3G HF networks is evaluated using the DoD-validated [6] HF network simulator, NetSim. The metrics of interest here are the fraction of multicast packets that reached their destinations, and the delay through the network.

Each multicast is preceded by FLSU (one-way, one-to-many). The frequency for each multicast in the simulations was manually pre-selected. In practice, frequency selection would be handled by an automatic channel selection (ACS) process that integrates propagation predictions with recently measured propagation on point-to-point C2 links.

All three scenarios are run for 24 hours to assess the effects of diurnal propagation variations on multicast delivery. Each scenario is simulated for three combinations of month and smoothed sunspot number (SSN): July with SSN = 10 (most challenging), October with SSN = 130 (best propagation), and April with SSN = 70 (moderate conditions). The propagation model used the lowest-decile SNR values in every case (i.e., exceeded 90% of the time).

Simulation Scenarios

Three scenarios were developed for evaluating this multicast stack: a regional multicast within a theater of operations, a strategic (continental) dissemination of Emergency Action Messages, and long-haul multicast of Air Tasking Orders with acknowledgements.

Regional, no acknowledgements

The regional multicast scenario uses NVIS propagation to deliver situational awareness updates every 5 minutes from a ship standing offshore to a Marine regiment ashore.

Each update is a compressed packet, 10 kB in size. Six HF radios receive the multicasts, and forward the data into the tactical line-of-sight radio networks ashore. No acknowledgements are sent by the recipients.

This message is of sufficient size to make good use of the highest-throughput MDL waveform, BW2, but we have seen that SNR of at least 18 dB is required for effective use of BW2. Fortunately, a 1 kW shipboard transmitter is able to provide enough signal strength at the receivers for SNRs ranging from 20 to 34 dB in the conditions considered. As seen in Figures 6-8, this is sufficient to provide reliable multicast message delivery under most (but not all) conditions.

Note that because this is a compressed message, all 44 packets must be received error-free before decompression can be successful. Partial delivery is not possible.
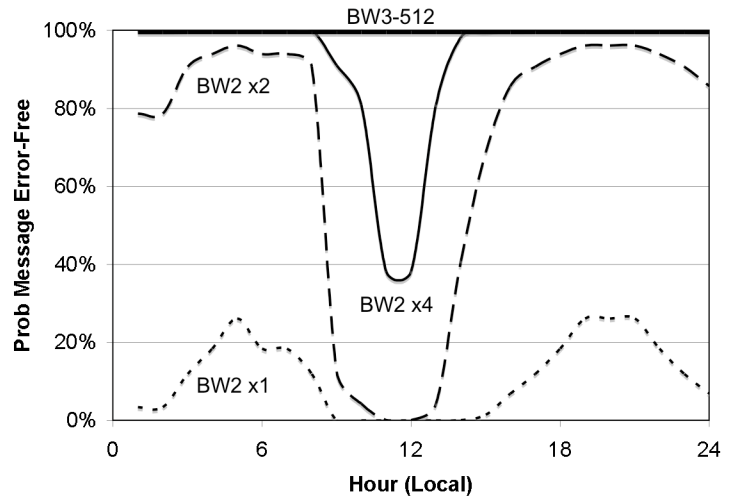
When SNR becomes marginal (e.g., noon in the summer with low solar activity, Figure 8), BW3-512 can be used instead of BW2. This offers 100% reliability all day.
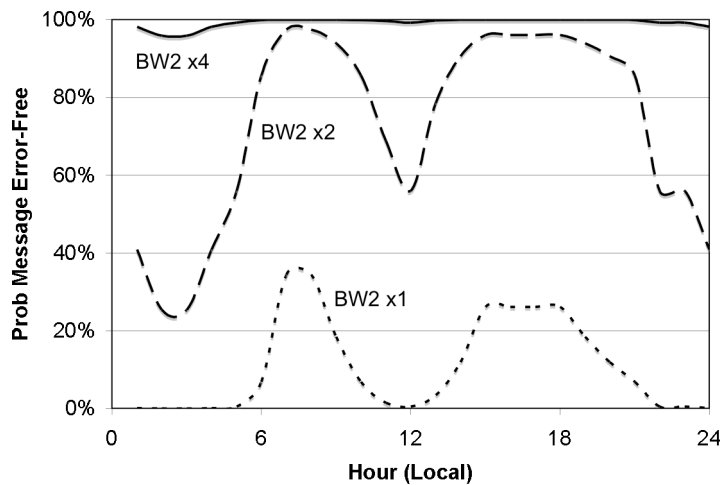
Figure 6: NVIS Message Delivery in October, SSN 130
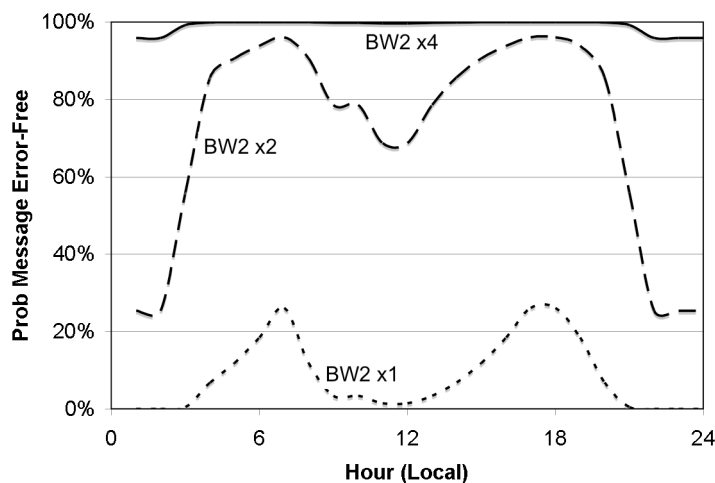
Figure 7: NVIS Message Delivery in April, SSN 70

Figure 8: NVIS Message Delivery in July, SSN 10

Spatial diversity is not reflected in these graphs. It is possible that packet losses will not occur at all receivers, perhaps due to differing local noise environments. In such a case, it may be possible to achieve 100% reliability under even the most hostile conditions (e.g., noon in July with low solar activity) by delivering the message throughout the unit as soon as one receiver decodes it.

The cycle time using 24-packet BW2, including an FLSU PDU each cycle, is 11.25 s. Two such cycles are needed to deliver a 10 kB message. When all four code phases are required to decode the message error-free, the delivery time will be 90 s, but some receivers could decode and deliver an error-free copy as soon as 22.5 s.

When BW3-512 is used, each transmission of the compressed message requires 145 s, but BW3 achieves 100% reliability throughout the day in this scenario.

Strategic, no acknowledgements

The first strategic scenario involves dissemination of Emergency Action Messages (EAMs) via skywave channels to 24 aircraft dispersed over North America (Figure 9). The EAM is assumed to fit in a 32-byte packet; therefore it can be transmitted completely by the smallest multicast PDU (BW3-32). An EAM is sent twice an hour simultaneously on two frequencies each by two high-power stations. The aircraft are in EMCON (not permitted to transmit), so this scenario uses no acknowledgements.

For this simulation, only the two stations shown in Figure 9 as filled circles (Nebraska and Alaska) were used. Two additional hypothetical station locations are shown as open circles. Transmitters are 4 kW, and omnidirectional antennas with 8 dBi gain were simulated. Aircraft noise varied with frequency as usual, with the noise level at 3 MHz set to -106 dBW. Each BW3 packet is sent 4 times per transmission: twice in each code phase in alternation.
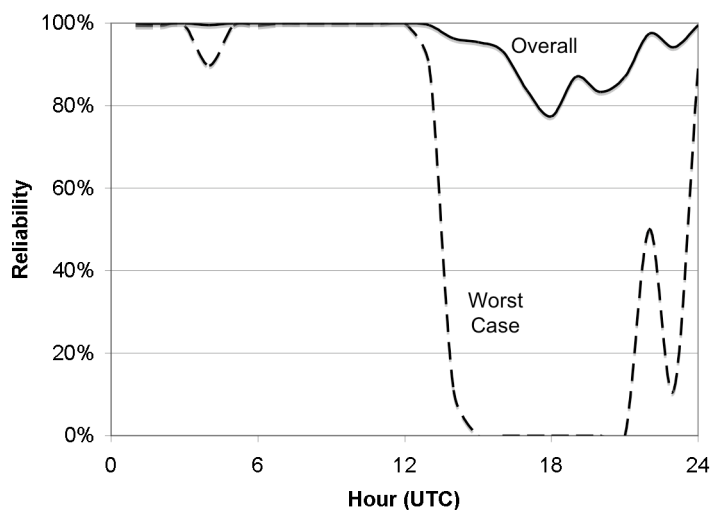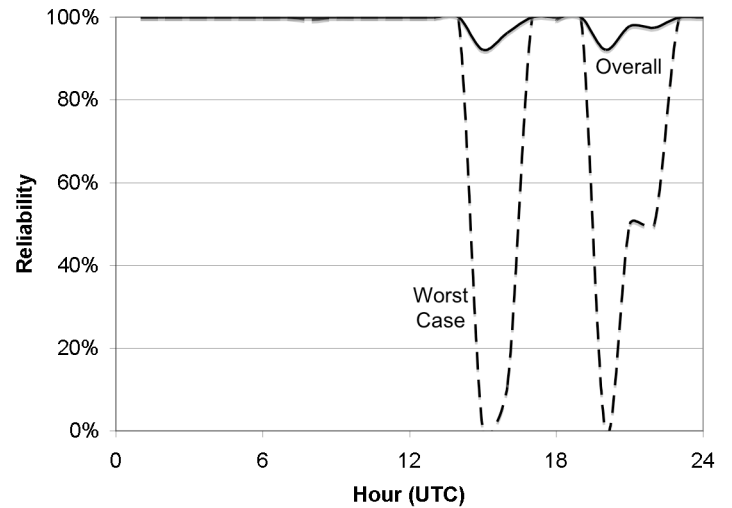
Figure 9: EAM Scenario

Reliability results are shown in Figures 10-12. The solid line plots the overall reliability for each hour (fraction of all packets successfully received over all destinations). The dashed line shows the worst case at each hour.
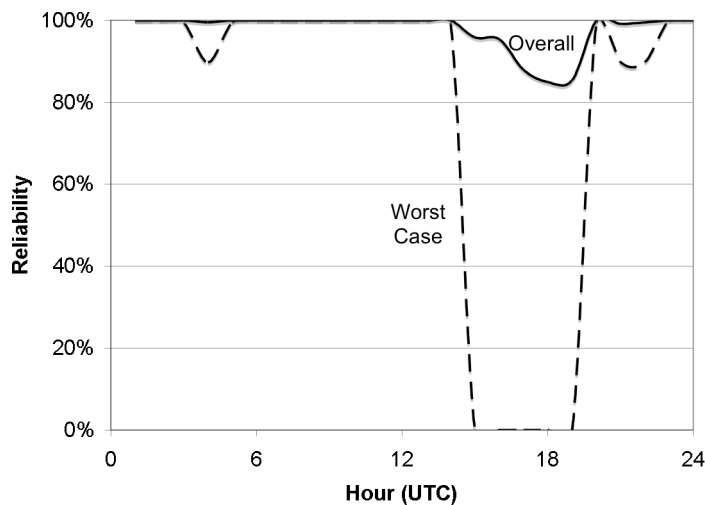
Figure 10: EAM Reliability in October, SSN 130

Figure 11: EAM Reliability in April, SSN 70

Figure 12: EAM Reliability in July, SSN 10

The worst case at each hour is not always the same aircraft, but is usually one of the aircraft that is roughly equidistant from the two transmitting sites. (The exception is the two Arctic locations at dawn with moderate to high solar activity.) Thus, we might expect that using the two additional sites shown in Figure 9 would be helpful. This was confirmed by additional simulations: when those two sites multicast the EAM along with the two original sites, *reliability at all of the aircraft locations reaches 100% under nearly all conditions*.

The time required to send the BW3-32 packet four times, including a FLSU PDU before each, is 7.25 s.

Strategic, with acknowledgements

The second strategic scenario involves acknowledged delivery of Air Tasking Orders to 8 aircraft in the CENTCOM AOR. A new multicast begins every six hours, and it is repeated twice an hour until all recipients have acknowledged receiving it.

The multicast message size is 100 kB. A hypothetical transmitting station at Doha, Qatar is equipped with a 4 kW transmitter and 8 dBi omni antenna. Aircraft radio parameters in this scenario are the same as above.

Message delivery using BW2 with all four code phases ranges from 98% to 100% through all hours of the day in the moderate and good conditions (April with SSN 70 and October with SSN 130). However, in the summertime when solar activity is low, signal strength falls off sufficiently at midday to noticeably impact reliability (Figure 14). As we saw in the NVIS scenario, BW3-512 provides full reliability under such conditions, although at about half the speed: 24 minutes to send the air tasking order once using BW3-512, versus 14 minutes using BW2 x4.
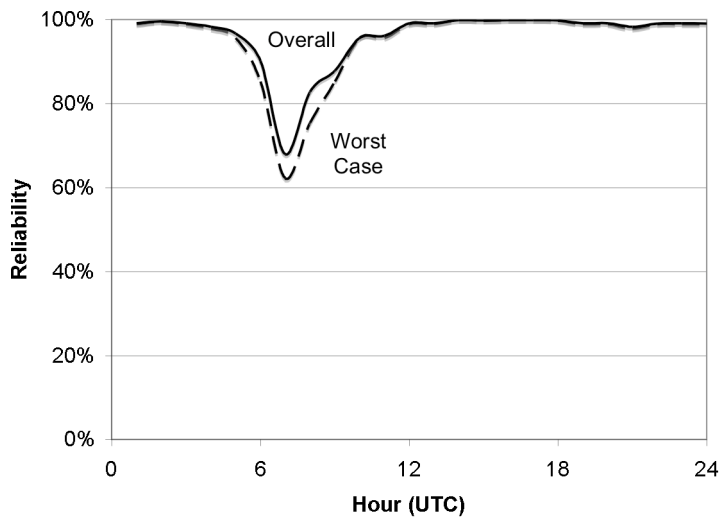
Figure 14: Air Tasking Order Reliability in July, SSN 10

In this scenario, we are also interested in the time required for the aircraft to return acknowledgments, using FLSU to contend for the channel and BW3-32 for the P_Mul ACK. The results of simulations are shown in Figure 15 for the best and worst propagation conditions, and are seen to be nearly identical. Even when all 8 aircraft begin contending for the channel at the same time (the worst case), the median time to acquire the channel and return the ACK is approximately 20 seconds.
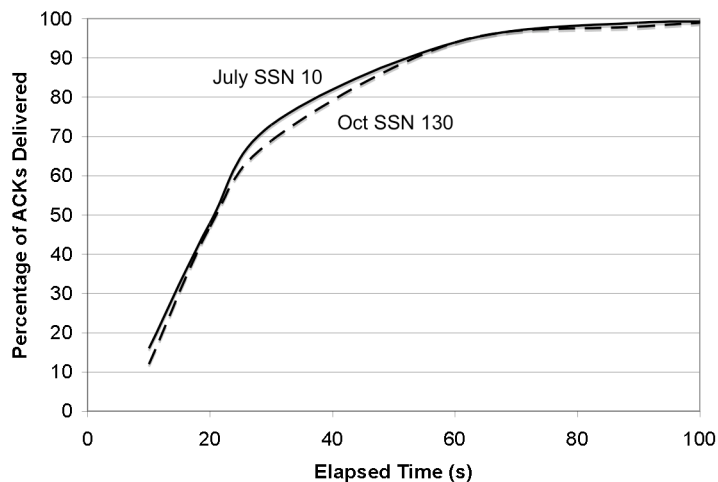
Figure 15: Cumulative Probability of Delivering P_Mul ACKs

Conclusions And Future Work

This paper has presented and evaluated a proposed Multicast Data Link (MDL) protocol that will augment the existing 3G HF protocol suite in NATO STANAG 4538. The proposed MDL protocol provides a 3G alternative to the popular 2G HF data link of STANAG 5066, and should be interchangeable with STANAG 5066 in supporting military messaging (STANAG 4406).

MDL integrates completely with the link establishment and traffic management functions of both "robust" and "fast" 3G HF protocols, and supports trading off robustness versus speed over a decade range.

In this study, we examined only single-hop multicasting, so it was not necessary to consider multicast routing, nor the construction and maintenance of multicast trees or meshes. However, not all HF networks are completely connected (i.e., have a diameter of one hop), so future work should evaluate the suitability and especially the efficiency of state-of-the-art multicast routing protocols for use in HF networks.

It would also be interesting to compare the MDL described here (which uses the 3G burst waveforms) with use of a standard HF data modem after link setup.

References

1. J. Bani-Salameh, *Military Multipath Multicast Routing Protocol*, PhD Dissertation, New Mexico State University, 2005.

2. NATO Standardization Agreement (STANAG) 4406, "Military Message Handling System," Ed. 1, Military Agency for Standardisation, 1999.

3. ACP 142, P_MUL – *A Protocol for the Reliable Multicast Messaging in Bandwidth Constrained and Delayed Acknowledgement Environments*, Allied Communications Publication, August 2001.

4. H. Zhang and E.E. Johnson, "A Third-Generation Multicast Protocol for HF Wireless Networks," *Proceedings of MILCOM 2005*, IEEE, 2005.

5. E. Koski "Concepts for a Reliable Multicast Data Link Protocol for HF Radio Communications," *Proceedings of MILCOM 2005*, IEEE, 2005.

6. R.K. Bullock, "SCOPE Command High Frequency (HF) Network Simulation Model (NetSim-SC) Verification and Validation Report," Defense Information Systems Agency, Joint Interoperability Test Command, Ft Huachuca, AZ, January 1998.