Analysis of Third-Generation HF ALE Technologies

Eric E. Johnson

New Mexico State University

ABSTRACT

Third-generation HF Automatic Link Establishment (ALE) achieves order-of-magnitude improvements over second-generation ALE in linking speed, network size, and traffic capacity. In this paper we examine the individual contributions to this performance by several of the techniques standardized in the third generation, including synchronous operation, dwell groups, and separate calling and traffic frequencies.

I. INTRODUCTION

Second-generation HF automation [1, 2, 3, 4] provided robust, reliable, and interoperable HF links that led to a resurgence of interest in HF radio for long-haul and mobile voice networks beginning in the 1980s. By the mid-1990’s, however, the growth of HF networking revealed the need for more efficient protocols so that the limited HF spectrum could support larger networks and more data traffic.

A cooperative development effort among government, industry, and academia resulted in third-generation ALE (3G-ALE) protocols [5, 6, 7, 8] that have been standardized in the U.S. MIL-STD-188-141B Appendix C and proposed NATO STANAG 4538. Some of the key 3G improvements are listed below:

- Faster link establishment
- Linking at lower SNR
- Improved channel efficiency
- ALE and data traffic use the same family of waveforms
- Higher throughput for short and long data messages
- Better support for Internet protocols and applications

The new techniques incorporated in 3G-ALE that led to this performance improvement include the following:

- Burst PSK waveforms
- Synchronous scanning of calling channels
- Partitioning of stations into dwell groups
- Trunked operation (i.e., separate calling and traffic channels)
- Multi-slot channel access using call priorities
- Carrier sense multiple access with collision avoidance (CSMA/CA) channel access procedure

This paper presents the results of an investigation of the contribution that each of these techniques makes to the improved performance of 3G HF networks.

II. OVERVIEW OF 3G TECHNOLOGY

Previous papers [5, 6, 7, 8] have discussed many of the novel aspects of 3G technology, and complete details may be found in the standards (MIL-STD-188-141B Appendix C and STANAG 4538). This section summarizes this literature, including those details important to understanding the results presented later.

A. Waveforms

Both linking and data transfer protocol data units (PDUs) are conveyed over the channel by a family of PSK waveforms that are derived from the MIL-STD-188-110A serial-tone modem. The new waveforms are optimized for bursts rather than long transmissions, which gives the system improved agility. Measurements of the 3G waveforms indicate 6–9 dB improvement in AWGN and fading channels over 2G waveforms [5].

B. Scanning

When not engaged in linking or traffic, radios scan assigned channels listening for calls, and recording occupancy (and sometimes propagation) of the channels. 3G scanning differs from the 2G case only in detail (e.g., synchronized scanning).

C. Calling

A link is established when a call and response PDU have been successfully sent over the air. In most cases, the response is sent by the called station, but may be sent by the calling station in special cases. The calling station listens on the calling channel before sending its PDU(s) to avoid interfering with other transmissions.

By comparison, 2G ALE uses a 3-way handshake. The function of the third PDU in 2G ALE is accomplished by a start-of-traffic timeout in 3G ALE.

D. Synchronous Operation

2G-ALE is an asynchronous system in that a calling station makes no assumption about when a destination station will be listening to any particular channel. 2G stations must therefore use long calls to “capture” scanning receivers. 3G-ALE includes an asynchronous mode, but it achieves its highest performance under synchronous operation.

When operating in synchronous mode, all scanning receivers in a 3G-ALE network change frequency at the same time (to within a relatively small timing uncertainty). The assigned calling frequency monitored by each 3G radio in a network can be computed at all times, so only a single, brief call PDU must be sent in a call.

Synchronous operation imposes costs, however, in infrastructure and overhead to maintain synchronization.

0-7803-6521-6/00/$10.00 (c) IEEE

E. Dwell Groups

It is not necessary that all stations monitor the same calling channel at the same time. By partitioning network members into groups that monitor different channels in each scanning dwell, calls directed to network member stations will be distributed in frequency and/or time. This greatly reduces the probability of congestion on 3G-ALE calling channels under high-traffic conditions. The set of stations that monitor the same channels at the same time is called a *dwell group*.

The alternative network organization, assigning all network members to a single dwell group, permits a single call to reach all members at once (if the channel carrying the call propagates to all members). Thus, the value of dwell groups depends on the mix of point-to-point versus multipoint calls.

F. Trunked Mode

It is well known that CSMA systems saturate well before the medium is fully loaded. If traffic channels are separated from calling channels, however, the load on the traffic channels is more or less decoupled from the calling channels. Such trunked operation permits traffic channels to be heavily used while calling channels are left relatively vacant, resulting in a good balance of high traffic throughput and low link establishment latency.

Implementing trunked operation for HF radio networks involves some complications, however, compared to conventional (non-trunked) operation. First, it is necessary to estimate the usability of traffic channels rather than directly measuring the channel during link establishment. Second, even with such estimates, it is necessary to monitor the occupancy of traffic channels so that an already-active channel is not selected for a new link.

Both of these tasks can be assigned to idle radios when a station has multiple HF radios available, but for a single radio (e.g., a manpack) they require that the receiver must sometimes be used for overhead functions rather than listening for calls. This can increase the time required to link with such stations.

G. Multi-Slot Dwells

Yet another technique for reducing calling channel congestion is extending the dwell time on each channel to accommodate multiple calling slots per dwell. In such a scheme, calling stations randomly select a slot in each dwell, listen in the preceding slots for potential interference, and call if interference is not detected. (The slot following a call is used for the response.)

By spreading out the calls that would be placed during a dwell over multiple slots, and deferring when a call is detected, the rate of collisions is substantially reduced.

If deferred calls wait until the next dwell to try again, the scheme is effectively 1/N-persistent, where N is the number of slots per dwell. This improves the performance of a heavily loaded network compared to the 1-persistent case. A single-slot-per-dwell scheme must incorporate an explicit backoff after call failure to achieve a similar effect.

The optional multi-slot dwell structure specified in STANAG 4538 is shown in Figure 1. Note that the initial slot in each dwell ("slot 0") is designated for all stations to check for occupancy on a traffic channel. Since all stations do this at the same time, no calls are missed.

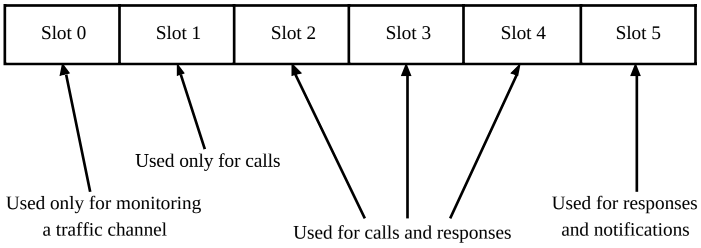

Figure 1. Dwell Structure in Multi-Slot Mode

Tuning time occurs once per dwell, and is therefore amortized over all of the slots in the dwell.

III. PROTOCOLS STUDIED

This investigation evaluated eight combinations of the 3G ALE features, as itemized in Table 1.

Protocol	Sync?	Dwell Groups?	Trunked?	Multi-Slot?
N1	●			
N1g	●	●		
N1t	●		●	
N1gt	●	●	●	
NM	●			●
NMg	●	●		●
NMt	●		●	●
NMgt	●	●	●	●

Table 1. Protocols studied

Each of the combinations studied is available in STANAG 4538.

- N1 designates the single-slot "Fast Link Setup" protocol, which was developed specifically for minimum latency in lightly loaded networks.
- N1g adds dwell groups to the N1 case.
- N1t adds trunked mode to the N1 case.
- N1gt adds both dwell groups and trunked operation to the N1 case.
- NM designates the multi-slot "Robust Link Setup" protocol, the original 3G ALE protocol. The same series of suffixes are used here as for the single-slot protocols.

IV. SIMULATION APPROACH

This investigation employed a NetSim simulator and two network scenarios described below.

A. NetSim Overview

The NetSim family of simulators [9, 10] implements a discrete-event communications network simulation architecture depicted in Fig. 2. Each of the modules shown functions independently, and implements its respective function at a level of detail appropriate for the investigation.

- Traffic sources generate voice or data messages according to specified inter-arrival time and message size distributions.
- The HF Node Controller (HFNC) at each station implements network-layer protocols and station-wide resource management.
- ALE controllers implement the ALE protocol and waveform under study. (The waveform is simulated as the probability of correct frame reception *vs.* signal-to-noise ratio.)
- Radio and antenna modules determine power and noise levels, intermodulation distortion, gain versus azimuth, etc.

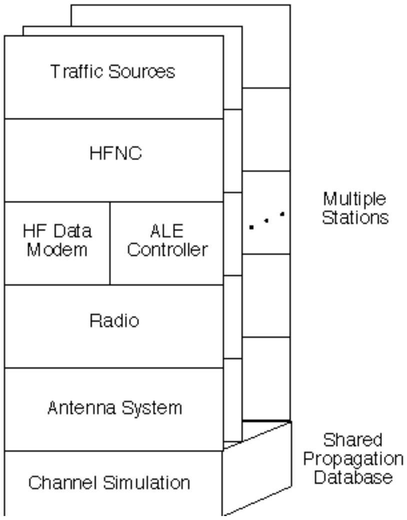

Figure 2. NetSim Simulator Structure

When a radio is tuned to a particular frequency, it receives a composite signal that includes the effects of all transmissions worldwide that are in progress on that frequency. To compute the effective SNR of the composite signal (for receivers that lock onto the strongest arriving signal), it is sufficient to compute the ratio of the strength of the strongest arriving signal to the noncoherent sum of natural noise plus all other signals plus distortion.

The channel model used in most NetSim simulations is the "Walnut Street" model of ionospheric propagation (validated by the Joint Interoperability Test Center), along with a direct wave model for aircraft within line of sight at altitude. Estimates of median SNR and first and ninth deciles of SNR for a link are computed using VOACAP. These are then used to generate representative random processes for path loss on the links of interest, using linear interpolation between hourly predicted values plus lognormal variation about this line for intermediate-term variation plus Rayleigh fading. For comparison with other simulations, however, the Walnut Street model can be replaced with a fixed-SNR model.

B. Air-to-Ground Scenario

The first of the two scenarios, an air-to-ground scenario, is representative of a large-scale mostly-voice network. The traffic is unevenly distributed over a large region and over time, resulting in less congestion than might otherwise be expected.

A fleet of 115 aircraft fly at various times during daylight hours (local time) among bases on the East Coast of the US, Europe, and the Central Command theater, as shown in Figure 3. They communicate with 14 base stations situated throughout the northern hemisphere and just south of the equator.

The simulation covers 24 hours in June with a sunspot number of 100.

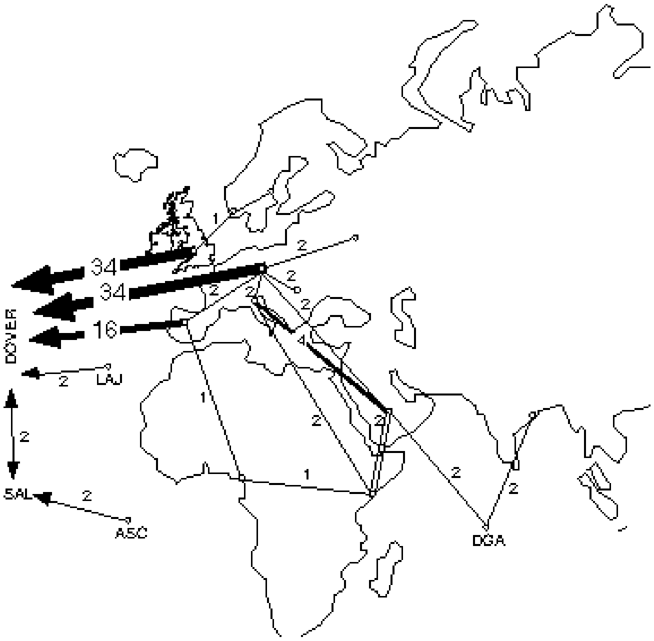

Fig. 3. Air-to-ground scenario

- The aircraft take off at intervals throughout the day.
- Each aircraft places (on average) one 5-minute voice call per hour while en route. (Intervals and durations are exponentially distributed.) Ground stations are selected adaptively.
- Each aircraft carries a single 400W ALE radio.
- Each of the 14 ground stations spread around the globe has nine or ten identical 3G-ALE-equipped radios, each with 4 kW output power.
- All antennas are omnidirectional.
- Eighteen channels are available. For trunked operation, five are used for calling channels and thirteen for traffic. Otherwise, all eighteen are used for both calling and traffic.
- Ground stations sound on each calling channel every 45 minutes so that aircraft can adaptively select ground stations for each call. Aircraft do not sound.

Two variations of this scenario were simulated:

0-7803-6521-6/00/$10.00 (c) IEEE

AG Full	The full air fleet, to investigate a moderately high load with mobile, dispersed units. (The size of this network requires dwell groups.)
AG 2	A single flight of two aircraft, to investigate a very light load.

C. Ground-to-Ground Scenario

A network of ten fixed stations is shown in Figure 4, along with the number of messages per hour (total) sent on each link. The stations, each with several radios, are positioned geographically as shown, in temperate latitudes, with a span of about 150 km from the leftmost station to the rightmost. Every station in this network can usually hear transmissions from every other station.

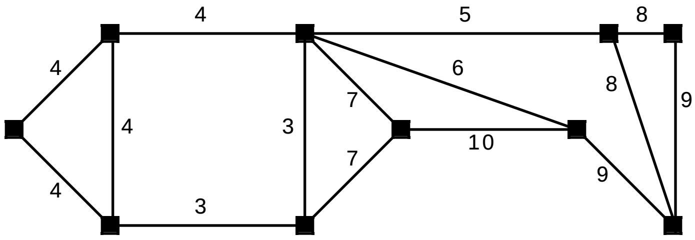

Figure 4. Ground-to-ground scenario

Eighteen channels were allocated (5 calling, 13 traffic for trunked operation). As the paths used in this scenario were expected to be NVIS, the frequencies were in the 3 - 11 MHz range. Sounding was not needed, and was not used.

Two scenarios were simulated for this network:

GG Voice	The traffic is all voice calls, each lasting 2 minutes on average. The indicated number of calls is sent on each link.
GG Data	The traffic is all packet data. Half are low-priority, half are high-priority. The number of calls at each priority is as shown in the figure (a total of twice the numbers shown).

Note that data traffic can use channels with negative SNR, while voice traffic requires at least 12 dB SNR. When a link is established that is found to have inadequate SNR for the traffic to be sent, the link is rejected and a new link is requested.

V. RESULTS

There are many metrics that could be used to compare the alternative protocols. This investigation used a metric that has been cited in procurements for HF networks such as SCOPE Command and MHFCS: the fraction of all messages for which a usable link was established within a specified time (30 seconds in this case). This metric reportedly correlates well with user satisfaction with the network.

Simulation results for the four scenarios and the eight protocols are presented in Table 2. As noted previously, only the protocols that use dwell groups support the AG Full scenario.

Factorial analyses [11] of the results are summarized in Table 3. For each scenario, the table indicates the

				Simulated Scenario			
Protocol	Dwell Groups?	Trunked?	Multi-Slot?	AG 2	GG Voice	GG Data	AG Full
N1				100	76	81	
N1g	●			86	65	68	69
N1t		●		86	67	71	
N1gt	●	●		79	62	78	63
NM			●	93	79	79	
NMg	●		●	100	88	90	79
NMt		●	●	100	80	85	
NMgt	●	●	●	93	84	89	82

Table 2. Percentage of usable links completed within 30 seconds

	Simulated Scenario			
Factor or interaction	AG 2	GG Voice	GG Data	AG Full
Multi-slot dwells	35% +	74% +	58% +	90% +
Dwell groups	13% –	0%	2% +	
Trunked operation	13% –	5% –	1% +	1% –
Multi-slot and groups	13% +	17% +	13% +	
Multi-slot and trunking	13% +	2% +	1% +	9% +
Groups and trunking	1% –	0%	5% +	
All three features	13% –	2% –	21% –	

Table 3. Effects of 3G features: Fractions of overall performance variation and polarity of effect

contribution of each principal factor (the 3G protocol features) and each interaction to the total performance variation observed, along with the direction that use of that feature moved the performance in that scenario.

For example, in the GG Voice scenario, use of multi-slot dwells made, on average, a positive contribution to performance that amounted to 74% of the total variation in linking performance among the eight protocols studied. The average impact of either dwell groups or trunked operation alone was minimal, but dwell groups interact strongly with multi-slot dwells. This interaction was strongly positive, accounting for most of the remaining performance variation (17%) among the eight protocols.

VI. DISCUSSION

Before we proceed to analyze the simulation results, it should be noted that the performance in Table 2 is reduced by "cold start" effects that might be eliminated by pre-programmed link data in real networks.

A. AG 2 Scenario

The light loading in this scenario produces only fourteen calls in the simulation. Each call thus accounts for 7% of the total, producing the quantization seen in Table 2. From the analysis of the results we note that even in this light loading, the use of multi-slot dwells improves linking performance, and was the largest effect observed.

B. GG Voice Scenario

The heavier loading in this scenario (up to 89 messages per hour) increased the importance of multi-slot dwells. Multi-slot dwells, especially in combination with dwell groups, reduced calling channel congestion by an order of magnitude, resulting in fewer deferred calls.

In this, as in the other scenarios, addition of either dwell groups or trunking to the single-slot scheme reduces its performance. Use of dwell groups reduced the number of third-party calls overheard by the single-slot stations; this eliminated much incidental link quality data and reduced the likelihood of first-call success. The backoff mechanism used in the N1 protocols then delayed subsequent re-call attempts. Lack of a dedicated slot for monitoring traffic channels reduced the performance of trunked operation in single-slot mode.

C. GG Data Scenario

This scenario places twice as many calls as the GG Voice scenario, but with roughly the same total traffic due to shorter message durations. The shorter messages and robust modem improved overall network performance.

D. AG Full Scenario

The full-scale air-to-ground scenario includes the heaviest loading and the most dynamic propagation of the four scenarios studied. Use of dwell groups is required, since the size of the network is greater than 60 stations (the maximum size of a single dwell group).

For this traffic level (as well as for heavier levels studied previously), the full range of congestion control techniques is beneficial.

VII. CONCLUSIONS

In each of the scenarios studied, the most important factor for maximizing the number of links established within 30 seconds was the use of multi-slot dwells.

The effects of the other two factors studied were mixed. Adding dwell groups and/or trunking to the single-slot protocol had uniformly negative effects on performance. For the multi-slot protocol, the use of dwell groups was uniformly beneficial. Addition of trunking to the multi-slot protocol was important for supporting the heaviest traffic load (5 minute voice calls in the full air-ground scenario), but less so under lighter loading.

ACKNOWLEDGMENTS

The third generation protocols described here include contributions from a wide cross-section of the global HF community. Particular thanks go to the NATO ARCS working group, including Paul Arthur and Andrew Gillespie from DERA (UK), Patrick Bruas from Thomson (France), Hans Denk of DASA (Germany), and Tom Kenney of Harris (US). Vive le CoTS!

REFERENCES

1. MIL-STD-188-141A, *Interoperability and performance standards for medium and high frequency radio systems, 1988.*
2. FED-STD-1045, *Telecommunications: HF Radio Automatic Link Establishment*, 1990.
3. FED-STD-1052, *Telecommunications: HF Radio Modems*, 1990.
4. STANAG 5066, *NATO Standardization Agreement: Profile for High Frequency (HF) Radio Data Communications*, 1999.
5. M. Chamberlain, W. Furman, and E. Leiby. "A scalable burst HF modem." *Proceedings of HF98, The Nordic Shortwave Conference*, 1998.
6. E. Johnson, T. Kenney, M. Chamberlain, W. Furman, E. Koski, E. Leiby, and M. Wadsworth, "U.S. MIL-STD-188-141B appendix C - a unified 3rd generation HF messaging protocol," *Proceedings of HF98, The Nordic Shortwave Conference*, 1998.
7. E. Johnson, "Third-generation technologies for HF radio networking," *Proceedings of MILCOM '98,* 1998.
8. E. Johnson, "Simulation results for third-generation HF automatic link establishment," *Proceedings of MILCOM '99*, 1999.
9. E. Johnson, R. Desourdis, and M. Rager, "Simulation of MIL-STD-187-721C automated HF networking," *Proceedings of the Ionospheric Effects Conference IES '96*, 1996.
10. E. Johnson, "Fast propagation predictions for HF network simulations," *Proceedings of MILCOM '97*, 1997.
11. R. Jain, *The art of computer system performance analysis*, Wiley, 1991, chapter 17.