Evaluation of HF ALE Linking Protection

Dr. Eric E. Johnson, Roy S. Moore
New Mexico State University

Abstract

The resurgence of interest in high frequency (HF) radio may be largely attributed to the success of automation in making HF usable by operators with minimal HF training. Unfortunately, Automatic Link Establishment (ALE) technology opens automated HF networks to possible hostile manipulation. To counter this vulnerability, a scheme called Linking Protection (LP) has been developed and standardized that scrambles ALE transmissions so that only legitimate users can interact with protected radios. In this paper, we summarize the LP scheme, and present an evaluation of the impact of LP on linking performance. The evaluation includes both simulation results and measurements of early implementations.

1 Introduction[1]

MIL-STD-188-141A and FED-STD-1045 specify an Automatic Link Establishment (ALE) protocol for use in HF radio systems that often must link over sky-wave channels. To cope with the poor channel characteristics often encountered with such channels, the standard specifies fairly robust mechanisms at both the physical layer (modem) and the data link layer, in OSI terms. The modem employs 8-ary FSK with 8 msec tones; thus 3-bit symbols are sent at a rate of 125 per second, giving a raw data rate of 375 bps.

A conceptual model of the MIL-STD-188-141A data link layer protocols is shown in Figure 1. Linking is accomplished by exchanging 24-bit ALE words. Several means are employed to cope with the characteristics of HF sky-wave channels: a (24, 12, 3) Golay code is used for Forward Error Correction (FEC), with each (12-bit) half of the 24-bit ALE word encoded separately, producing two 24-bit results. These two 24-bit Golay words are then interleaved bit by bit, and a stuff bit is appended, to produce a 49-bit word to be transmitted. Finally, each 49-bit word is sent three times, which allows the receiver to correct some errors using 2 of 3 majority voting. The time to send this redundant word is termed T_{RW}.

At the receiver, received bits from the modem are (conceptually) shifted into a 99-bit shift register. Majority voting among the outputs of this shift register yields a 48-bit "majority word" (stuff bits are discarded), which is de-interleaved to produce two 24-bit Golay words. These are delivered to the Golay decoder, which attempts to recover a 24-bit ALE word.

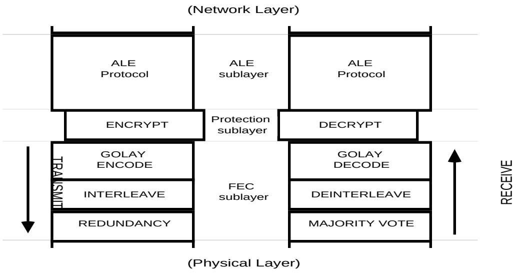

Figure 1: Conceptual Model of Data Link Layer

The format of an ALE word consists of a 3-bit preamble, which indicates the correct interpretation of the remaining bits of the word, followed by 21 data bits which are often used to hold three 7-bit ASCII characters – part of a station address, for example. The characters of an address are constrained to be chosen from a 38-character ASCII subset.

Because no bits in the ALE word are spent on synchronization, acquiring word synchronization in this system employs a series of tests on the prospective word after each received symbol (tri-bit) is shifted in. First, the number of unanimous votes in the majority voter must exceed a threshold. Next, the Golay decoder must successfully decode both halves of the 48-bit majority word. Finally, the resulting 24-bit ALE word must be acceptable to the ALE protocol module (including having an acceptable preamble, and characters from the correct ASCII subset). Once word synchronization has been achieved, it is automatically tracked for the remainder of the transmission using the same tests.

In operation, automatic link establishment is accomplished as follows:

1. The calling station transmits a call that contains the initial portion of the called station(s) address(es) for sufficient time to capture scanning receivers; the time for this "scanning call" is denoted T_{SC}.
2. Scanning receivers pause on channels containing ALE signaling. If a receiver successfully achieves word sync as described above, the ALE words are

[1] This section is derived from [1].

examined to determine whether the call is intended for that receiver.

3. If a station finds the beginning of its own address in a call, it will stay on channel and read the full address(es) contained in the "leading call" that immediately follows the scanning call (this period of the call is denoted T_{lc}). If its full address is found, the receiving station reads the rest of the call, including the caller's address, and completes a 3-way handshake to establish a link.

If any of the conditions specified above is not met, a receiver will immediately return to scan.

Once a link is established, the receivers alert the operators, indicating the station ID of the distant station. At this point, operators may converse in voice mode, exchange data, or even re-program each others' automated controllers. Clearly, if an adversary could establish links with such stations, he could severely disrupt operations.

2 Linking Protection

2.1 Overview of the LP Procedure

Linking protection (LP) seeks to exclude interference and deception from the ALE function. The method chosen to shield the receiving ALE module from such distractions is time-varying encryption of ALE words. The sender encrypts ALE words using a secret key and a "seed" containing time of day and the frequency in use on a link. The receiver decrypts received words using its key, and seeds covering a narrow range of times of day (based upon the *protection interval* or PI). Unless the received word was encrypted using the same key and a time of day in that range, it will appear to be noise to the receiver and will be accepted as a legitimate call with a probability less than 10^{-7}.

In addition to the time of day and frequency in use, the seed used in the encryption algorithm also contains a "word number" that is used to sequentially number all ALE words in a transmission starting with T_{lc}. (During T_{sc}, the word number simply alternates between 0 and 1.) The state diagrams shown in Figures 2 and 3 succinctly summarize the operation of the sending and receiving LP modules, respectively, except that the procedure used to acquire word and time-of-day synchronization at the receiver is not shown. In the figures, the states are labeled using the current PI number (denoted N) and the current word number used in the seed to the encryption algorithm. The phrase "Incr. N" denotes a PI transition. Further information about this protocol may be found in [1].

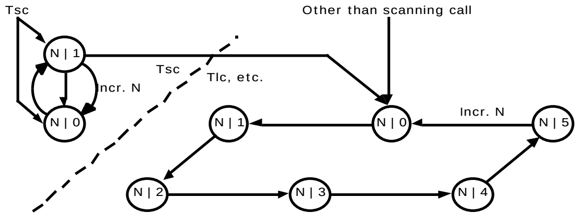

Figure 2: Transmitting State Diagram (2 second PI)

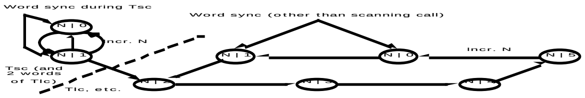

Figure 3: Receiving State Diagram (2 second PI)

The correct PI and word number to be used at the transmitting station are always determined unambiguously by the state diagram in Figure 2. However, there are several sources of ambiguity at the receiver that can cause linking attempts to fail when they result in the receiving LP module failing to follow the seed sequence used at the transmitting station:

1. Uncertainty of the time-of-day at the transmitter
2. Unknown word number during word sync acquisition
3. Unknown locations of PI transitions during T_{sc} in the received ALE word stream.
4. Unknown location of the transition from T_{sc} to T_{lc}.

The first two uncertainties require decryption of received words under a range of PI/word number combinations during word sync acquisition. In all cases, the usual word sync tests must be used to resolve ambiguity, due to the lack of synchronization codes in the ALE word stream. Techniques that may be used to resolve these ambiguities are discussed in the following sections.

2.2 Resolution of Ambiguity

Word Sync. During unprotected (non-LP) word sync acquisition, the receiver's FEC module examines the received stream of tones for bit patterns that exceed the unanimous vote threshold and produce correctable Golay words. When a candidate word is produced by the FEC module, it is checked by the ALE protocol module for acceptable preamble and ASCII subset and, if these checks pass, for compliance with the ALE protocol. When all tests concur that the received word is acceptable, word sync is assumed, and the FEC module settles into checking and returning one word every T_{rw} thereafter until otherwise notified by the ALE protocol module.

From Figure 1, the place of the LP function in this chain of events is clear: it is interposed between the FEC sublayer and the ALE protocol. Thus, when the FEC module returns a candidate word, the LP sublayer must decrypt it using seeds containing all valid PI/word number combinations, and deliver the results to the ALE protocol module where the final series of tests is applied. In most cases, at most one seed will produce a word that is acceptable to the ALE module, and time-of-day synchronization between the transmitter and the receiver will be achieved simultaneously with word sync. However, on rare occasions a candidate word from FEC will produce acceptable ALE words under two or more seeds, resulting in an ambiguity that must be resolved before the LP function can properly decrypt subsequent words.

The word number sequencing in the LP protocol was designed specifically to assist in the resolution of this ambiguity. For individual and net calls with single-word addresses (a common case), the next word received following word sync acquisition must be identical to the first word when decrypted under word numbers alternating between 0 and 1, possibly with a PI change, whether word sync is achieved during T_{sc} or T_{lc}. In this case the word sync function can resolve seed ambiguity by simply waiting for the next word and decrypting it under the appropriate combinations to determine which PI/word number combination was the correct one for the first (word sync) word. This next word is then retained to be returned when requested by the receiving ALE protocol software.

Transitions During T_{sc}. Note that in Figure 3, three arrows emerge from the N/1 state during T_{sc}, corresponding to simple alternation to word number 0; transition to the next PI and word 0; and a transition to N/2, indicating that T_{lc} began two words ago. These transitions are identified by evaluating incoming words under all possibilities, and selecting the most probable case when more than one case passes the tests.

3 Performance Evaluation

In this section, results obtained from a detailed simulator are compared to results of non-protected simulations, and to measurements of early implementations of linking protection. The metric used in these comparisons is the probability of successful link establishment as a function of channel conditions. The conditions include a range of signal-to-noise ratio for three standard channels: a pure Gaussian noise channel (no fading or multipath), a “good” channel (0.1 Hz fading bandwidth with 0.5 ms multipath delay) and a “poor” channel (1 Hz fading bandwidth with 2 ms multipath delay).

3.1 The Simulator

The simulator used to evaluate Linking Protection performance is written in C, and structured along the lines of the protocols simulated, as shown in Figure 4. Each simulator shown was validated against measured results. The HF channel simulator implements the well-known Watterson model [2]; the remaining simulators implement the protocols specified in the standards, with a few minor limitations in the ALE protocols simulated (i.e., only single-word addresses and individual or net calls are simulated; this results in simpler protocol processing in the simulator than in fielded ALE controllers).

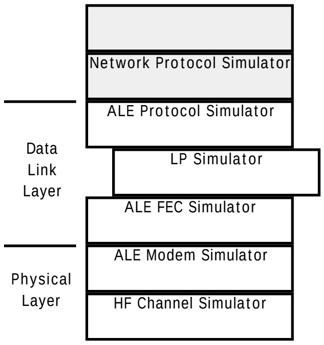

Figure 4: Simulator Stack

3.2 Protected *vs* Non-Protected Simulations

The first group of graphs compares the predicted linking probabilities for three cases:

- ALE with no linking protection (sometimes called AL-0 LP)
- AL-1 LP (60 second protection interval)
- AL-2 LP (2 second protection interval)

Other experimental conditions were as follows:

- individual calls with no embedded messages

- unanimous vote threshold in FEC module set to 0
- 2 ch/s scanning rate (500 ms dwell time)
- channel fading gains recomputed at 8 ms intervals

From these results, it is apparent that linking protection, as implemented in these simulations, produces a small but measurable degradation in linking probability. In nearly error-free channels, the linking probability achieved when using LP (especially AL-2) falls short of 100%, due to increased sensitivity to a false word sync problem (described below) that is also present in the non-protected ALE implementation.

When the word sync algorithm is continuously reading symbols from the modem, there is a non-negligible probability that both Golay decodes will succeed at an erroneous word phase; in Detect-6 / Correct-1 mode, this probability is approximately 3% for each candidate word checked. r

When the Golay check accepts a mis-aligned word, the preamble and ASCII subset checks will usually reject the word: of all 24-bit words, fewer than 1% contain three ASCII-38 characters and one of the three preambles accepted by the word sync algorithm in the simulator. However, if Golay determines that a mis-aligned word is "correctable" and the resulting ALE word passes the preamble and ASCII checks, the word sync algorithm will accept the erroneous word phase and the linking attempt will almost certainly fail.

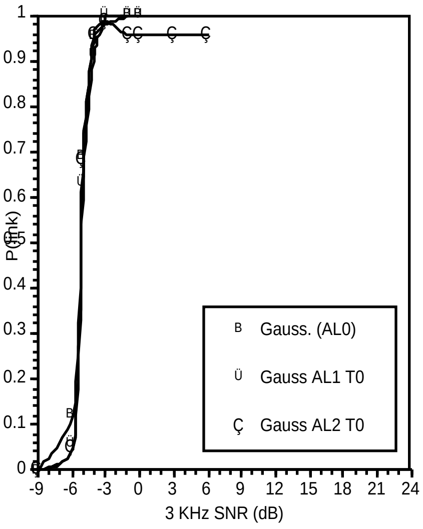

Figure 5: Effect of LP on P(link) — Gaussian Channel

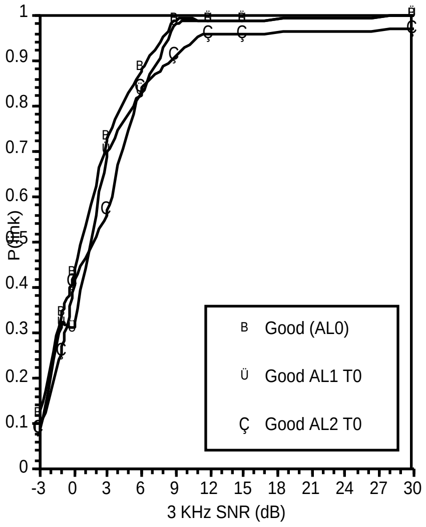

Figure 6: Effect of LP on P(link) — Good Channel

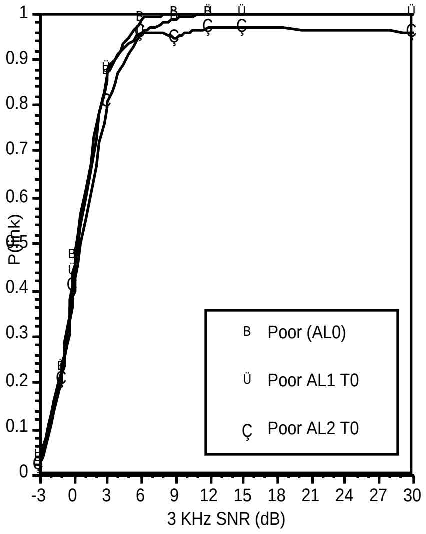

Figure 7: Effect of LP on P(link) — Poor Channel

If we denote the probability of Golay success on random inputs as p_G and the probability that a random word passes the preamble and ASCII checks as p_{pac}, then we can estimate the probability of a word sync error in non-protected mode, p_{wse}, as follows (assuming a uniform relative word phase distribution, and that the FEC process is invoked for each arriving symbol):

$$p_{wse} = \left[\, 1 - P\left(\begin{array}{c} \text{no false Golay before} \\ \text{correct word phase} \end{array} \right) \right] p_{pac} \;= \left[\, 1 - \sum_{i=0}^{48} \begin{array}{c} P\left(\text{i symbols before correct word phase}\right) \bullet \\ P\left(\text{no false Golay in first i symbols}\right) \end{array} \right] p_{pac}$$

$$= \left[1 - \frac{1}{49} \sum_{i=0}^{48} \left(1 - p_G\right)^i \right] p_{pac} = \left(1 - \frac{1}{49} \left[\frac{1 - \left(1 - p_G\right)^{49}}{p_G} \right] \right) p_{pac} \; \approx \; 0.005$$

Thus, in the long run, we would expect that a non-protected ALE station would experience a linking failure on ideal channels about once in 200 attempts due to false word sync, unless the word sync algorithm is improved over that described above. Use of a higher unanimous vote threshold should reduce p_{WSe} at the possible expense of reduced linking probability over marginal channels. (Another possibility is “soft” word sync detection.)

When LP is added to the receiver, several words are presented to the preamble and ASCII checks for each candidate word that passes the FEC-sublayer checks, because each such candidate word is decrypted under several PI/word number combinations (either 6 or 8). P_{WSe} is therefore increased by this factor, resulting in a drop in linking probability of a few percent. In the simulations shown in Figures 5 through 7, P(link) reached a plateau of 96 to 97%. When diagnostic simulations were made at 30 dB SNR, the failures in every case were due to false word sync acquisition.

Subsequent simulations have shown that adjustment of the unanimous vote threshold to 25 provides 99 to 100% linking performance in high SNR channels, with minimal performance degradation at low SNR.

3.3 Simulation versus Measured Results

Implementations of linking protection have been successfully tested recently by NTIA. Figure 8 compares measurements made at the SNR values for which performance standards are specified in FED-STD-1045, compared to the results of the simulations described before (with the unanimous vote threshold set to 25).

The measurements are qualitatively similar to the simulation results, with somewhat lower performance in general. This is not unexpected, because the simulator supports only a subset of the protocols supported by the tested equipment. In particular, because the simulator does not support group calls or multi-word addresses, the number of possible protocol branches that it must examine during word sync acquisition and during T_{SC} is reduced to exactly one: every word must be identical to the last one until T_{lc} is reached. Furthermore, there is only one format of T_{lc} that must be accommodated by the simulator. Examination of the simulator output suggests that the elimination of these possibilities for going astray in following the protocol may be at least partially responsible for the better results produced by the simulator. r

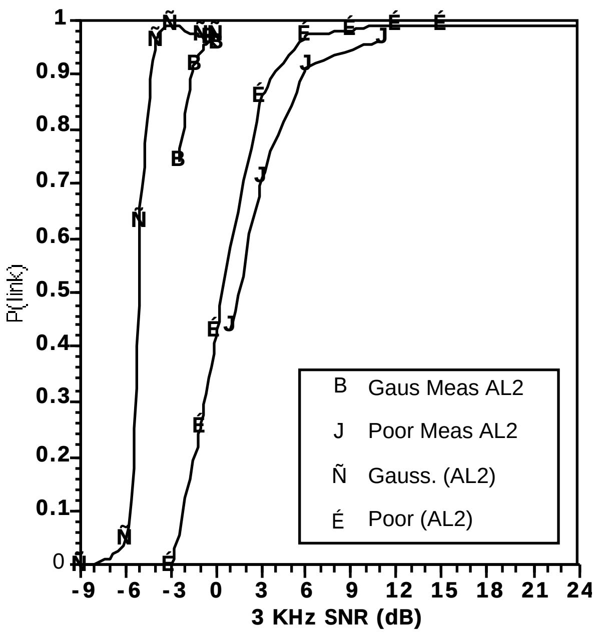

Figure 8: Simulation Vs Measurement (AL–2)

4 Conclusion

Simulation of the recently-standardized HF linking protection technology has shown that LP can provide its protection from imitative ALE calls with minimal degradation in linking performance. A unanimous vote threshold near 25 provides a good balance of performance over a broad range of channel conditions.

References

1. Johnson, E.E. *Linking Protection Implementation Guide*. Technical Report NMSU-ECE-91-004A, New Mexico State University, March 1992.

2. Watterson C., *et al.* "Experimental Confirmation of an HF Channel Model" *IEEE Transactions on Communication Technology* **(COM-18)** 6: 792-803, 1970.