NTIA Technical Memorandum 95-166

Implementation Guide for Federal Standard 1049 Section 1, Linking Protection

Eric E. Johnson
Christopher Redding
David F. Peach
Robert T. Adair

technical memorandum series

U.S. DEPARTMENT OF COMMERCE • National Telecommunications and Information Administration

NTIA Technical Memorandum 95-166

Implementation Guide for Federal Standard 1049 Section 1, Linking Protection

Eric E. Johnson
Christopher Redding
David F. Peach
Robert T. Adair

U.S. DEPARTMENT OF COMMERCE
Ronald H. Brown, Secretary

Larry Irving, Assistant Secretary
for Communications and Information

June 1995

PREFACE

This report provides reference material for technologists who wish to understand and implement the linking protection technique developed for automated high-frequency radio networks. The technology described in this report has been incorporated in Federal Standard 1049, Section 1. This report contains data compiled from inputs from industry, academia, and the Government. Certain commercial equipment are identified in this report to adequately describe the design and conduct of the research. In no case does such identification imply recommendation or endorsement by the National Telecommunications and Information Administration (NTIA), nor does it imply that the material or equipment identified is necessarily the best available for the purpose. The views, opinions, and/or findings contained in this report are those of the authors and should not be construed as an official NTIA position unless designated by other official documentation.

CONTENTS

CONTENTS (continued)

FIGURES

TABLES

ABBREVIATIONS

ACK	Acknowledgment
AL	Application Level
ALE	Automatic Link Establishment
AMD	Automatic Message Display
ASCII	American Standard Code for Information Interchange
BER	Bit Error Ratio
bps	Bits Per Second
CRC	Cyclic Redundancy Check
CT	Cipher Text
DBM	Data Block Mode
DTM	Data Text Mode
FEC	Forward Error Correction
FED-STD	Federal Standard
FSK	Frequency Shift Keying
GPS	Global Positioning System
HF	High Frequency
ICD	Interface Control Document
ISO	International Organization for Standardization
LP	Linking Protection
LPCM	Linking Protection Control Module
LQA	Link Quality Analysis
LSB	Least Significant Bit
Mbps	Megabits Per Second
MIL-STD	Military Standard
MSB	Most Significant Bit
OSI	Open Systems Interconnection
PI	Protection Interval
PI_{opt}	Optimum Protection Interval Length

ABBREVIATIONS (continued)

ppm	Part Per Million
PT	Plain Text
p_G	Probability of Golay Success
p(link)	Probability of Linking
p_{wse}	Probability of Word Synchronization Error
SINAD	Signal-plus-noise-plus-distortion to noise-plus-distortion ratio
SNR	Signal-to-Noise Ratio
TOD	Time-of-Day
T_{lc}	Leading Call Time
T_{rw}	Redundant Word Time
T_{sc}	Scanning Call Time
T_x	Termination Section Time
t_{min}	Time of Shortest ALE Transmission
W	Window of Vulnerability
w	Word Number Field
XOR	Exclusive-or

IMPLEMENTATION GUIDE FOR FEDERAL STANDARD 1049 SECTION 1, LINKING PROTECTION

Eric E. Johnson *, Christopher Redding **, David F. Peach **, and Robert T. Adair **

Automatic Link Establishment (ALE) technology automates the selection of channels and establishment of links among high-frequency radios, but creates a vulnerability in such automated networks to hostile manipulation of network operations. The Linking Protection (LP) technique described in this report was developed to protect against such manipulation, while causing minimal degradation to network performance. This report provides a summary of ALE operation, followed by a discussion of the LP technique and suggestions for producing high-performance implementations of LP, based upon simulation studies of protected-mode performance.

Key words: authentication, automatic link establishment, ALE, high-frequency radio, HF, performance analysis, radio networks, simulation.

1. INTRODUCTION

This report is intended as an aid to implementors of Government standardized linking protection (LP) for high-frequency (HF) radio systems.

1.1 Background

The issue of protecting automated radios from hostile manipulation arose during the development of the ALE standards (i.e., FED-STD-1045 and MIL-STD-188-141A). Because ALE automates many functions of HF radio systems that were formerly closely observed by trained operators, many of the developers and users became concerned that ALE introduced an increased risk of ignorant or malicious interference to Government HF systems. The LP technique described in this report was developed specifically to prevent received ALE signalling from interacting with protected stations except when the originator of such signalling is authorized to communicate with the protected station.

---

*The author is with Johnson Research, Las Cruces, NM 88005.

**The authors are with the Institute for Telecommunication Sciences, National Telecommunications and Information Administration, U.S. Department of Commerce, Boulder, CO 80303.

1.2 Objectives

The objectives of this report are to provide reference material for technologists concerning the development of the linking protection feature, and to provide guidance for implementors. With the development of LP essentially complete, and the requirements now standardized in FED-STD-1049 and MIL-STD-188-141A, an implementor's guide is useful to promote high-performance, interoperable implementations. In this report, we summarize much of the technological developments in the course of designing and evaluating the LP techniques.

1.3 Sources

Much of this material has been published previously in other forms, but is collected here for ease of reference. The principal source is a New Mexico State University Technical Report (not copyrighted) [1]. The requirements listed in Section 2.2.1 of this report were drawn from an unpublished White Paper entitled "Linking protection requirements" by E.E. Johnson [2]. Most of Section 3 has been taken from a paper presented at the 1992 Military Communications (MILCOM) Conference sponsored by the Institute of Electrical and Electronics Engineers (IEEE) [3]. Another key paper on LP, "Linking protection for HF radio automatic link establishment," was presented at MILCOM '91 [4].

2. SYSTEM OVERVIEW

This section provides an overview of LP and its place within HF ALE radio systems.

2.1 Automatic Link Establishment

FED-STD-1045A specifies an ALE protocol for use in HF radio systems that often must link over skywave channels [3]. 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 terms of the International Organization for Standardization (ISO) Open Systems Interconnection (O-S-I) Reference Model (OSI-RM). A conceptual model of the ALE data link layer is shown in Figure 1.

The modem employs 8-ary frequency-shift keying (FSK) with 8-ms tones. Thus, 3-bit symbols are sent at a rate of 125 per second, giving a raw data rate of 375 bps. Linking is accomplished by exchanging sequences of 24-bit ALE words among ALE protocol entities.

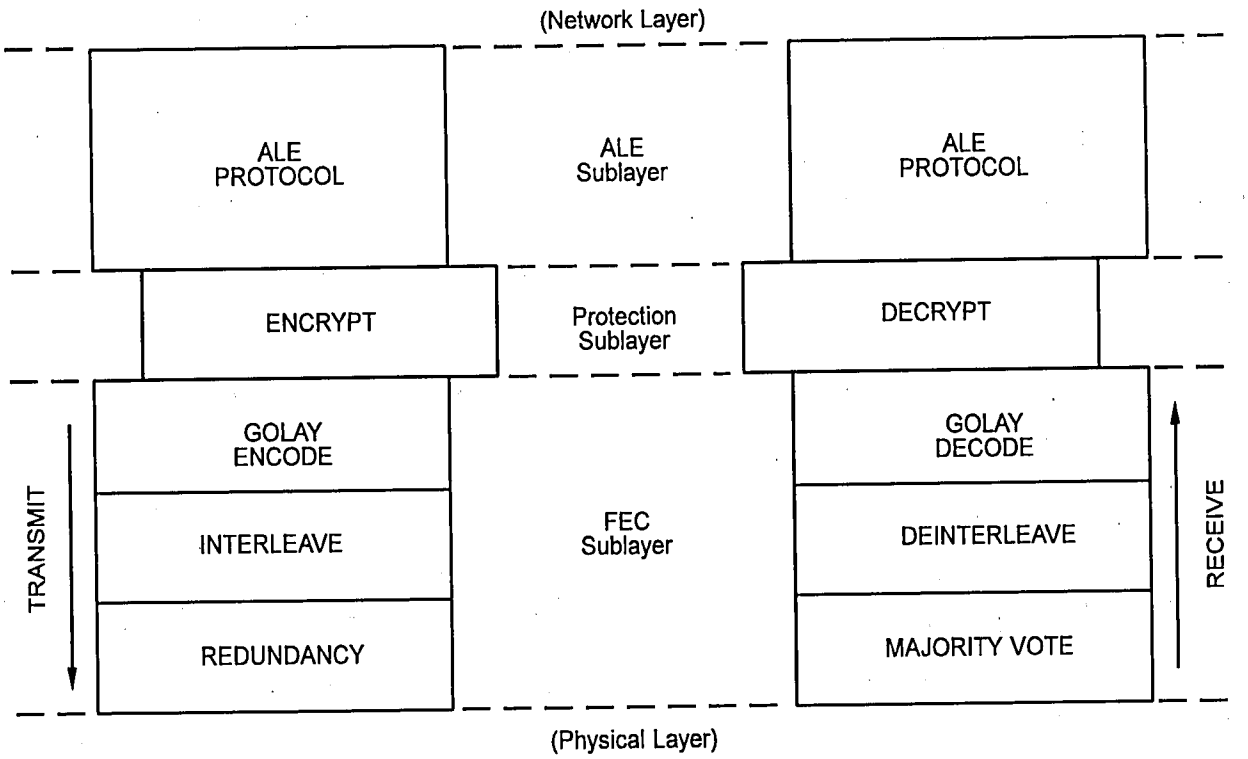

Figure 1. Conceptual model of ALE data link layer protocols.

Several means are employed within the data link layer to cope with the characteristics of HF skywave channels:

1. A (24, 12, 3) Golay code is used for forward error correction (FEC) [4], with each 12-bit half of the 24-bit ALE word encoded separately to produce two 24-bit results.

2. These two 24-bit Golay words are then interleaved bit-by-bit and a stuff bit is appended, resulting in a 49-bit word to be transmitted.

3. Finally, each 49-bit word is sent three times, which allows the receiver to correct some errors using 2 of 3 majority voting.

4. At the receiver, bits received 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).

5. This majority word is deinterleaved to produce two 24-bit Golay words.

6. These recovered Golay words are delivered to the Golay decoder, which attempts to recover a 24-bit ALE word.

Because no bits in the ALE protocol are spent on synchronization, acquiring word synchronization requires the receiver to employ a series of tests on the prospective word after each received symbol (tri-bit). The series of tests are described below:

1. The number of unanimous votes in the majority voter must exceed a threshold.

2. The Golay decoder must successfully decode both halves of the 48-bit majority word.

3. The resulting 24-bit ALE word must be acceptable to the ALE protocol module.

When word synchronization has been achieved, it may simply be checked once per word for the remainder of the transmission using the same tests. The receiver should nevertheless be able to acquire synchronization with a colliding signal in the midst of a transmission.

The ALE data link layer shown in Figure 1 thus comprises three sublayers: a lower sublayer concerned with error correction and detection (FEC sublayer), an upper sublayer containing the ALE protocol (ALE sublayer), and an optional sublayer in between the two called the protection sublayer. In the FEC sublayer, redundancy, majority voting, interleaving, and Golay coding is applied to the 24-bit ALE word. This constitutes the FEC-sublayer service-data-unit in terms of the OSI-RM. The ALE sublayer specifies protocols for link establishment, data communication, and rudimentary link quality analysis (LQA), based on the capability of exchanging ALE words. Data flow through these two sublayers, which are present in every ALE-equipped radio, is shown in Figure 2.

2.2 Linking Protection

Linking protection is placed in the intermediate protection sublayer so that it may make full use of the error-correcting power of the FEC sublayer while intercepting unauthorized attempts to communicate with the local ALE protocol entity. LP attempts to prevent the establishment of unauthorized links and other spurious interaction with the ALE sublayer through the following authentication process: ALE words are encrypted under a private key (which is changed at daily or longer intervals), using known randomization or "seed" information (frequency, time, date, etc.) to vary the results of this encryption on a shorter basis (a "protection interval" or PI). When received signals are decrypted by the receiving LP module, the probability that the

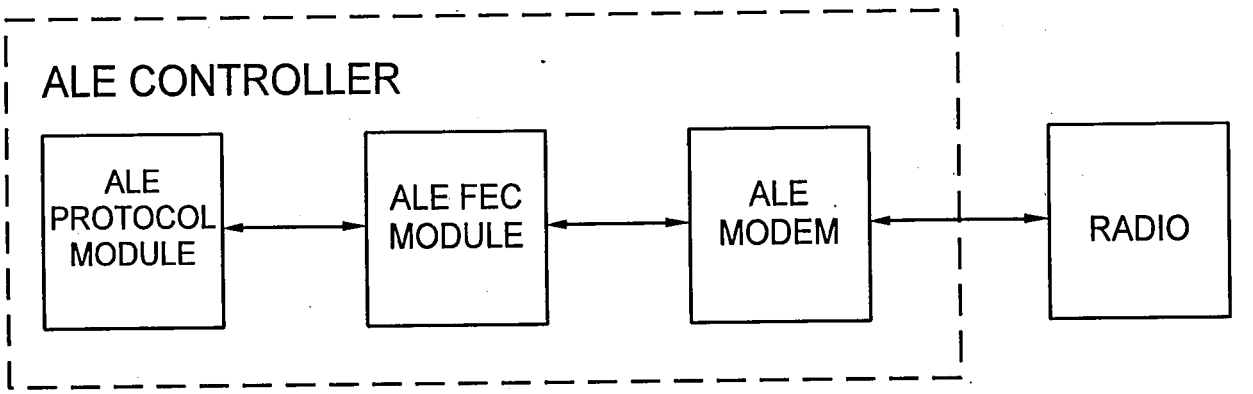

Figure 2. Data flow in a system without linking protection.

resulting ALE words will make sense to the receiving ALE protocol is very small unless the sender of the signals is encrypting ALE words using identical key and seed information.

The addition of LP to a radio involves adding the functions of a Linking Protection Control Module (LPCM), which implements the LP protocol, and a scrambler, which scrambles (or encrypts) ALE words under the control of the LPCM (Figure 3). The security of the system is based upon the inability of an adversary to "spoof" the LPCM, and relies on the difficulty of discovering the key used to scramble the ALE words. Because of the wide range of applications for LP, several different scramblers are specified in FED-STD-1049, but the LPCM is common to all LP applications, and includes a common denominator scrambler for assured interoperability of all protected radios. Note that the blocks in Figure 3 represent logical operations, not necessarily distinct hardware modules.

The standard LPCM handles unclassified ALE words only (LP for classified addresses has been defined but it requires security considerations beyond the scope of this paper). Any classified traffic must be encrypted by an appropriate cryptographic device before it reaches the ALE controller or LPCM; the resulting BLACK data may then be sent through the ALE controller or via a separate data modem. (Although some privacy is incidentally provided to the ALE words by the LP mechanism, this is not the primary purpose of LP.)

2.2.1 Transparency Requirements

A principal consideration in implementing LP is that the presence of an LP module in a radio (or its controller) should have no impact on any protocols outside of the protection sublayer in the Data Link layer. In particular, this means that achieving and maintaining synchronization of the cryptographic algorithm in scramblers (cryptographic

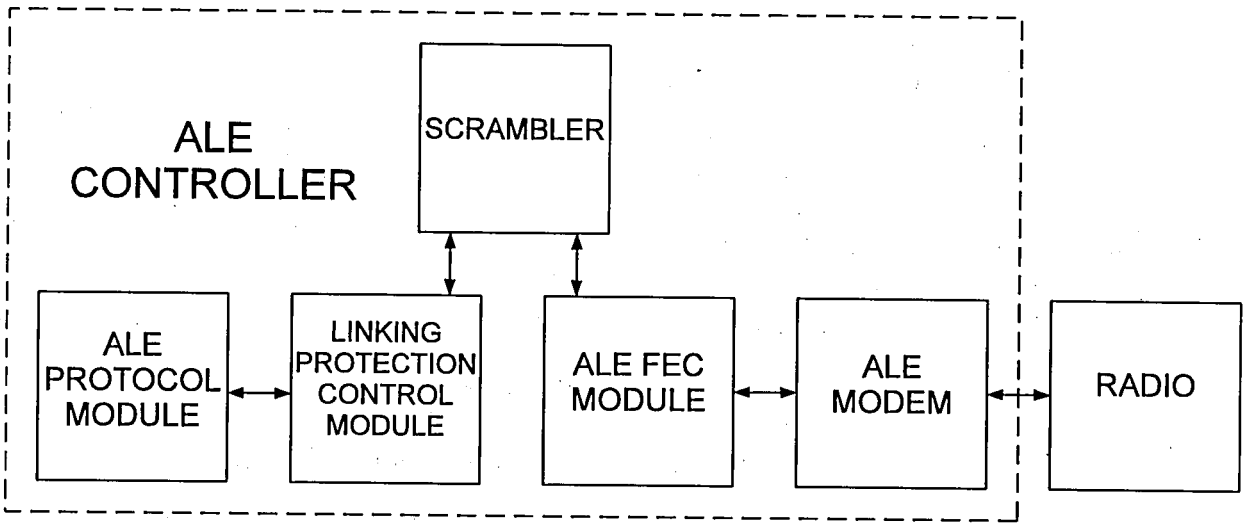

Figure 3. Data flow in a protected system.

synchronization) must occur transparently to the ALE waveform and protocols. Furthermore, scanning radios must be able to acquire cryptographic synchronization at any point in the scanning call portion of a protected transmission if this transmission was encrypted under the key in use by the receiving station. Thus, LP modules may not insert synchronization bits into the data stream, and must acquire cryptographic synchronization without the use of synchronization preambles or message indicator bits.

Transparency also requires that the decryption operations must proceed in real time so that linking time is not degraded. Also, the LP scheme must be sufficiently robust so that the effects of HF propagation do not significantly lower the linking probability for the protected mode as compared to the clear mode linking probability.

The encryption algorithm must employ a 24-bit word, so that all words are protected and no extra bits need be sent in any transmission simply to accommodate the algorithm. Many applications for LP will not accommodate physical protection of the scrambler, and some users will require an exportable algorithm for their international networks. These algorithm requirements have necessitated the development of new algorithms (or new operating modes) for each of the standardized LP application levels (see Section 4).

2.2.2 Overview of LP Operation

This section summarizes the LP procedure.

Transmit Processing

The LP module in a sending station encrypts each 24-bit ALE word to be sent using the seed data then in use (i.e., frequency and time of day) and delivers the encrypted word to the FEC module.

Receive Processing

The receiver side of an LP module is responsible for achieving cryptographic synchronization with transmitting stations, and for decrypting protected ALE words produced by the Golay decoder. In operation, when a scanning receiver arrives at a channel carrying ALE tones and timing, the FEC sublayer will process the output of the ALE modem and alert the LP receive module when an acceptable candidate word has been received. This occurs roughly once in 8 ms when the Golay decoders are correcting 3 errors per Golay word, or once in 78 ms when correcting 1 error per Golay word[1].

The receive LP module must then decipher the candidate word, and pass it to the receive ALE module which will determine if word synchronization has been achieved by checking for acceptable preamble and ASCII subset. This task is complicated by the possibility that the received word (even if properly aligned) may have been encrypted using a different time of day (TOD) than that currently at the receiver, requiring the LP module to decrypt each candidate word under seeds containing a range of TODs.

A further complication is the possibility, although small, that a word may satisfy the preamble and character set checks under multiple seeds. When this occurs, the valid successors to all seeds that produced valid words are used to decrypt the next word, and each result is evaluated in the context of the corresponding first word. It is very unlikely that after the second word is checked that a third word will be required to be decrypted under multiple protection intervals. For example, if during a scanning call (or sound) a received word decrypts "TO SAM" using seed A, and to "DATA SNV" using seed B, the next word is decrypted using the successors to those seeds, denoted A' and B'. If the result of decrypting this next word under A' is not "TO SAM," the first decrypt under seed A was spurious because the word following a TO word in a scanning call must be the same TO word. To be valid in a scanning call or sound, a word following "DATA SNV" must have three ASCII-38 characters and a THRU, REPEAT, THIS IS or THIS WAS preamble. A diagram showing all valid preamble sequences may be found in FED-STD-1045A, Figure 5 [5].

3. PROTECTION INTERVAL ANALYSIS

Because LP employs protection intervals (which are time-based), all protected stations must maintain accurate clocks. The seed data contains the time of day element, referred

to as the TOD-portion of the seed. Practical considerations suggest that local times of stations may differ by significant fractions of a minute unless some means is employed to maintain tighter synchronization. Because the effectiveness of LP increases as the length of the PI decreases, there is a trade-off between protection and the cost of implementing and using a time synchronization protocol. The following section is based on the MILCOM 92 paper, "Analysis of high-frequency radio linking protection."

One of the fields in the seed used in LP contains a count of protection intervals since midnight each day, and is a principal source of the time variation in the protection process. This field is updated at the beginning of each PI. Because the receiving LPCM must use the same seed to recover each plaintext ALE word as that used in the word's encryption, all stations must be synchronized to some degree. Receiving LPCMs will usually attempt to decrypt received ciphertext under several seeds with adjacent PI numbers to discover the one currently in use at the transmitting LPCM.

To keep the number of valid seeds manageable, the range of local times among LPCMs in a network should vary by no more than one PI. As shown below, the effectiveness of linking protection increases as the length of the PI decreases, making a tightly synchronized system desirable from a security standpoint, as well. However, the cost of implementing and operating the protected system increases with the degree of synchronization required, so the choice of a protection interval length must balance security with the costs to produce and to use the system.

3.1 Vulnerability Analysis

The encryption of ALE words forces an adversary who wishes to interact with protected stations to transmit words that will be accepted by those stations when decrypted. Two possible techniques for such attacks upon the system are as follows: 1) the adversary generates such ALE words locally ("guessing attack"), and 2) the adversary plays back properly encrypted ALE words intercepted from other stations.

A sufficient counter to a guessing attack is to force the adversary to correctly guess the encrypted versions of two ALE words, each under two different TODs. This may be achieved either by ignoring transmissions that reuse a TOD, or by choosing a PI length shorter than the minimum interval between the start of the "leading call" in the call and the start of the acknowledgment (about 3 s) [5].

The vulnerability of a network to playback attacks may be measured by the maximum time that a recorded transmission will be accepted by some network member. If each LPCM in a network is prepared to accept transmissions encrypted in any of δPIs centered on its own local time, this window of vulnerability (W) is given by:

$$W = \max(0, \delta PI - t_{min})$$

where t_{min} is the duration of the shortest transmission and PI is the length of the protection interval.

If $\Delta T = d \bullet PI$ (where ΔT denotes the range of LPCM times among the network members), then $\delta = 2d + 1$, and the optimum PI length (i.e., the maximum PI length that produces W = 0) is:

$$PI_{opt} = \frac{t_{min}}{2d + 1}$$

Use of this PI length ensures that successive transmissions must use different TODs, regardless of the protocol used. For PI = ΔT, d = 1 and $PI_{opt} = t_{min}/3$, the result is 392 ms for linking transmissions.

On the other hand, if only one transmission is accepted using any given TOD, the upper limit on the length of the PI comes not from security but from communication considerations: throughput suffers if the PI becomes longer than the minimum time between the receipt of valid transmissions, due to the rejection of some valid transmissions. When "hand-shaking" with another station, this interval is 3.14 s; when receiving slotted responses, however, the rate may climb to one response every 1.83 s.

If PI = ΔT, eliminating the possibility of multiple legitimate slotted responses using the same TOD requires $PI = t_{slot}/2 = 915$ ms. A longer PI may be used if ΔT is reduced (or t_{slot} is increased) commensurably, or if the occasional rejection of valid slotted responses is tolerable.

To summarize the analysis of vulnerability versus PI length, the guessing attack may be foiled either by using a PI of 3 s or less, or by ignoring transmissions that reuse a TOD. To defeat the playback attack, it is sufficient to employ a PI equal to one T_{rw} (392 ms) when PI = ΔT; the PI may be extended to up to 1 s with tighter synchronization. If stations reject reuse of a TOD, the system is vulnerable only to transmissions recorded by an adversary that were not heard by an addressee, and then only for the remaining period of validity of such a recorded transmission. The length of the PI then determines the probability of unwarranted rejection of transmissions: a PI of 3 s or less will only affect slotted responses, and a 1-s PI will virtually preclude the rejection of these.

3.2 Feasibility Analysis

The requirement for synchronization among stations using LP brings with it a variety of costs, both in implementation and in operation, which increase with the degree to which stations must be synchronized. Implementation costs are incurred in both the hardware and the software of the controller for a protected radio: the hardware must include an accurate time base, and the software must include mechanisms for accurately

setting this time base and for synchronizing it with other network members. Operational costs include the necessity to accurately set the time base before communication may take place, and propagation delay for over-the-air synchronization.

The accuracy of each local time base depends upon how accurately the time was last set, and how far the time base has drifted since then. Operators can usually set time to within 100 ms, given that they have accurate time. Stations may also decode standard time broadcasts, when available, and automatically compensate for propagation delays to achieve time settings with millisecond accuracy. To these timing uncertainties must be added any variation in internal processing time to set the time base; such variations range up to hundreds of milliseconds in existing ALE controller designs.

Time base drift is dependent upon the stability of the oscillator used: a 10-ppm source may be available from the synthesizer in the RF section of the radio, which will hold this drift to less than 1 s per day; if a 200-ppm watch crystal is used instead, this drift may be up to 17 s per day.

Once network members are in synchronization, they may periodically exchange time broadcasts to stay synchronized. The accuracy of time obtained in this fashion includes uncertainties in the time at the sending station, the release accuracy of the time broadcast, variations in propagation delay among stations, and processing variations at the receiving stations, many of which may be minimized through careful design of the controller software. If we assume a conservative figure of ±200 ms for the accuracy of time setting at each station in a network, controllers with 10-ppm time bases will be able to hold time to within ±500 ms (for a 1-s PI) for only 8 hours before time must be reset to within ±200 ms; a 2-s PI would permit these stations to go 22 hours between updates. Controllers with 200-ppm time bases would require updates every 25 min for a 1-s PI, every hour for a 2-s PI, or every 40 hours for a 1-min PI.

HF systems are often intended for emergency use. Therefore, to determine a feasible PI length, the abnormal situations in which protected systems may have to be brought into synchronization must be considered. For example, operators may be under stress and standard time broadcasts may be unavailable. In such cases, the accuracy of initial time setting may depend upon the accuracy of time obtained from someone's wrist watch, which may be accurate only to within ±15 s. Thus, feasible PI lengths may be on the order of a minute, rather than a second.

The approach selected for the standard system is to rely upon operators to get station times synchronized to within 1 min (±30 s), and then to employ a protocol to synchronize stations to within 1 or 2 s (fine sync) for full linking protection. While it is possible to operate networks with only coarse (1 min) time synchronization, this reduces the protection offered by this system against playback (tape recorder) attacks.

4. LINKING PROTECTION APPLICATION LEVELS

The six LP application levels shown in Table 1 are defined in FED-STD-1049 [7], with Application Level 5 (AL-5) providing the highest level of protection. Application Levels 3 through 5 require distinct hardware scramblers, while Application Level 1 and 2 scramblers may alternatively be implemented in software or firmware.

All protected radios must be capable of operating at Application Level 1 for interoperability. However, a means must also be provided to disable automatic link establishment at application levels less secure than those in use by a station. For example, a station that is operating at Application Level 3 should be able to disable the receiver from listening for linking attempts at Application Levels 0, 1, and 2. This mechanism must not preclude the operator from manually initiating a link using a disabled application level, however; this manual override capability is required for interoperability. When a valid call is received that uses a disabled application level of LP, one option is to alert the operator but not link.

5. LINKING PROTECTION SPECIFICATION

The LP procedure specified below should be implemented as distinct functional entities for control functions and cryptographic functions. (Distinct *hardware* for each function is not required unless otherwise indicated). The linking protection control module (LPCM) performs the control functions specified below, and interfaces to the ALE controller (Figure 2). Scrambler(s) perform the cryptographic operations on ALE words, under the control of the LPCM. A means should be provided to disable the LP functions and operate the radio in clear (unprotected) mode, and hardware scramblers should be removable without impairment of the clear-mode functionality of a radio.

Use of linking protection should neither increase the time to establish a link compared to a nonprotected radio, nor degrade the probability of linking below the requirements for nonprotected linking specified in [5].

5.1 Linking Protection Control Module

The LPCM controls the attached scrambler(s) as specified in this section.

5.1.1 Scrambler Interfaces

The LPCM interacts with hardware scrambler(s) according to the circuits and protocols specified in the Interface Control Document (ICD) for each scrambler. Interaction with

Table 1. Linking Protection Application Levels

Application Level	Definition
AL-0	Application Level 0 (AL-0) refers to nonprotected, or clear mode, operation. No scrambler or LPCM is used.
AL-1	The AL-1 scrambler uses the Lattice Algorithm specified in USAISEC Technical Report ASQB-OSI-S-TR-92-04 [8] and may be implemented in hardware, firmware, or software, with manufacturer-specified interfaces. AL-1 is mandatory for all protected stations, and therefore provides protected interoperability among them. The AL-1 protection interval is 60 s, which provides slightly lower protection than that available using Application Level 2, but with relaxed synchronization requirements.
AL-2	The AL-2 scrambler uses the same algorithm as specified for AL-1. This application level is for general U.S. Government and commercial use. The AL-2 protection interval is 2 s.
AL-3	The AL-3 scrambler uses the DES algorithm. This application level is for general U.S. Government and commercial use. The AL-3 protection interval is 2 s.
AL-4	The AL-4 scrambler is for U.S. Government use only and employs a hardware scrambler and Interface Control Document (ICD) developed by NSA. The LPCM interface to this module must be certified as secure. The AL-4 protection interval is 2 s.
AL-5	The AL-5 scrambler is for U.S. Government use only and employs a hardware scrambler and ICD developed by NSA. An AL-5 scrambler may be used to protect classified traffic, and is an accountable Controlled Cryptographic Item. Systems employing AL-5 LP must meet security requirements beyond those for AL-4. The AL-5 protection interval is 1 s, maximum.

software implementations of scramblers need only comply with the applicable function call ICD, if one is specified.

5.1.2 Time of Day

The LPCM requires accurate time and date information for use in the LP procedure. Therefore, the local time base should not drift more than ±1 s per day when the station is in operation.

A means must be provided for entry of date and time via either an operator interface or an electronic fill port such as an interface to a Global Positioning System (GPS) receiver. Ideally, both will be provided. This interface may also provide for the entry of uncertainty of the time entered (see Section A-1 in Appendix A). If time uncertainty is not provided, a default time uncertainty must be used. Defaults for the various time fill ports may be separately programmable; unless otherwise programmed, the default uncertainty should be ± 15 s.

After initialization of time of day, the LPCM employs the time protocols of Appendix A. The time protocol selected should maintain total time uncertainty less than that of the protection interval length of the most secure LP application level it is using. LPCMs respond to time requests (see Appendix A) unless this function is disabled by the operator.

5.2 Seed Format

The LPCM maintains randomization information for use by scrambler(s), and provides this "seed" information to each scrambler according to the applicable ICD.

The 64-bit seed contains the frequency carrying the protected transmission, the current PI number, the date, and a word number in the format shown in Figure 4. The most significant bit of the seed, and of each field, is on the left. The TOD portion of the seed must be monotone nondecreasing; the remaining bits are not so constrained.

The Date field is formatted as shown in Figure 5. The Month field contains a 4-bit integer for the current month (1 for January through 12 for December). The Day field contains a 5-bit integer for the current day of the month (1 through 31). Some mechanism should be provided to accommodate leap years.

The PI field is formatted as shown in Figure 6. The Coarse Time field contains an 11-bit integer that counts minutes since midnight. The 6-bit Fine Time field is set to all 1's when using AL-1 LP; when a time synchronization protocol (see Appendix A) is employed to obtain more accurate time, the Fine Time field is set to the time obtained

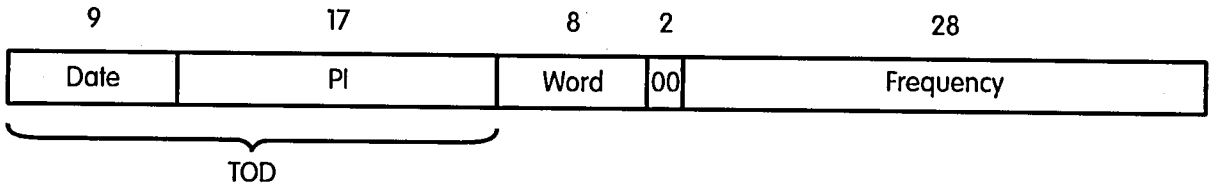

Figure 4. Seed format [1].

using this protocol. The Fine Time field will always be a multiple of the PI length, aligned to PI boundaries (e.g., with a 2-s PI, Fine Time is always even). The word number field is used to count words sent during each PI.

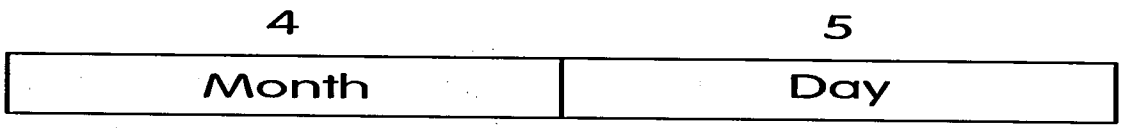

Figure 5. Date format [1].

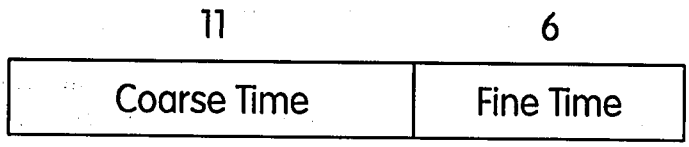

Figure 6. PI number format [1].

The Frequency field is formatted as shown in Figure 7. Each 4-bit field contains one binary-coded decimal digit of the frequency of the current protected transmission.

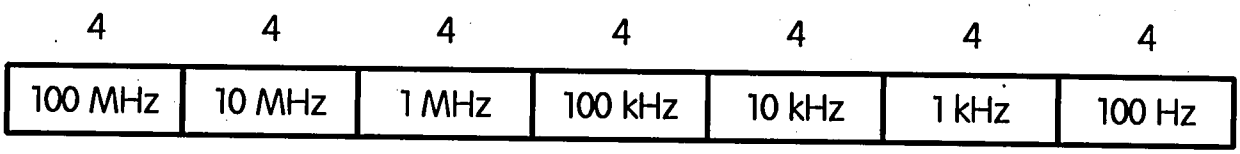

Figure 7. Frequency format [1].

5.3 Procedure

The procedure used to protect transmissions consisting entirely of 24-bit ALE words is presented in Sections 5.3.1 and 5.3.2, followed by the procedure for the data block portion of Data Block Message transmissions in Section 5.3.3. The procedure for orderwire messages using AL-1 and AL-2 is in 5.3.4.

When a radio is neither transmitting nor receiving, the PI number is incremented as follows: 1) when using Application Levels 2 through 5 LP, the Fine Time field is incremented at the end of each PI by the length of the PI; when the Fine Time field rolls over to 0, the Coarse Time Field is incremented, and 2) at midnight, the Coarse and Fine Time fields are set to 0, and the Date fields are updated.

When using Application Level 1 LP, the Fine Time field contains all 1's, and the Coarse Time field is incremented once per minute. At midnight, the Coarse Time field is set to 0, and the Date fields are updated.

5.3.1 Transmitting Station

Each word to be transmitted is encrypted by the scrambler using the current seed information. In the course of a transmission, the protocol described below may cause a discrepancy between the TOD fields in the seed and the real time; such a discrepancy is a normal consequence of the LP procedure, and will persist until the conclusion of each transmission, whereupon the TOD fields of the seed are corrected. The TOD and word number field (w) combination is collectively referred to as TOD/w. The word number field is used as follows:

1. During the scanning call phase (T_{sc}) of a call, the calling station alternates transmission of words encrypted using w = 0 and w = 1. The first word of T_{sc} uses whichever value of w will result in w = 1 for the last word of T_{sc}. The TOD used during T_{sc} will change as required to keep pace with real time, as long as w = 0; words encrypted with w = 1 must use the same TOD as the preceding word.

2. At the beginning of the leading call phase (T_{lc}) of a call (which is the beginning of a single-channel call), the first word must be encrypted using w = 0 and the correct TOD for the time of transmission of that word.

3. All succeeding words of the call use succeeding word numbers up to and including $w = w_{max}$; for the word following a word encrypted with $w = w_{max}$, the TOD is incremented and w is reset to

0. w_{max} = 2 for a 1-s PI, w_{max} = 5 for a 2-s PI, and w_{max} = 153 for a 60-s PI.

4. Responses and all succeeding transmissions start with w = 0 and the current (corrected) TOD, with these fields incremented as described in 3 above for each succeeding word.

Figure 8 illustrates the permissible TOD/w combinations (LPCM states) for a transmitting station using a 2-s PI (w_{max} = 5), and the permissible sequences of these combinations. In this figure, T represents the TOD in use.

Sounds are protected in the same fashion, with the redundant sound phase (T_{rs}) in the place of T_{lc}. A single-channel sound is analogous to a single-channel call, and begins the above procedure at step 2. A multi-channel sound is analogous to a scanning call, and begins with step 1.

5.3.2 Receiving Station

Because of the possibility of acceptable decodes under multiple TOD/w combinations, receivers must attempt to decode received words under all allowed combinations (the current and adjacent PIs, future and past, and both w = 0 and w = 1) when attempting to achieve word synchronization with a calling station (six combinations). Stations prepared to accept time requests (see Appendix A) will also need to decode received words using coarse TOD (Fine Time = all 1's, with correct Coarse Time) with both w = 0 and w = 1 (eight combinations total). All valid combinations must be checked while seeking word synchronization; after achieving word synchronization, the number of valid combinations is greatly reduced by the LP protocol.

Figure 9 illustrates the permissible TOD/w sequences for a receiving station using a 2-s PI after word synchronization is achieved. Note that, unlike the transmitter, the receiving station LPCM state machine may be nondeterministic. For example, when in T_{sc} and in state T/1, a received word may yield valid preambles and acceptable characters when decrypted using all of the valid combinations, i.e., T/0, (T+1)/0, and T/2 (the latter implying that T_{lc} started two words previously). Therefore, the receiver will be in three states at once until the ambiguity is resolved by evaluating the decrypted words for compliance with the LP and ALE protocols.

Receivers using a protection interval of 2 s or less are not permitted to accept more than one transmission encrypted using a given TOD, and need not check combinations using that TOD. For example, if a call is decrypted using TOD = X, no TOD before X + 1 is valid for the acknowledgment (the *response* travels in the opposite direction, and may use TOD = X).

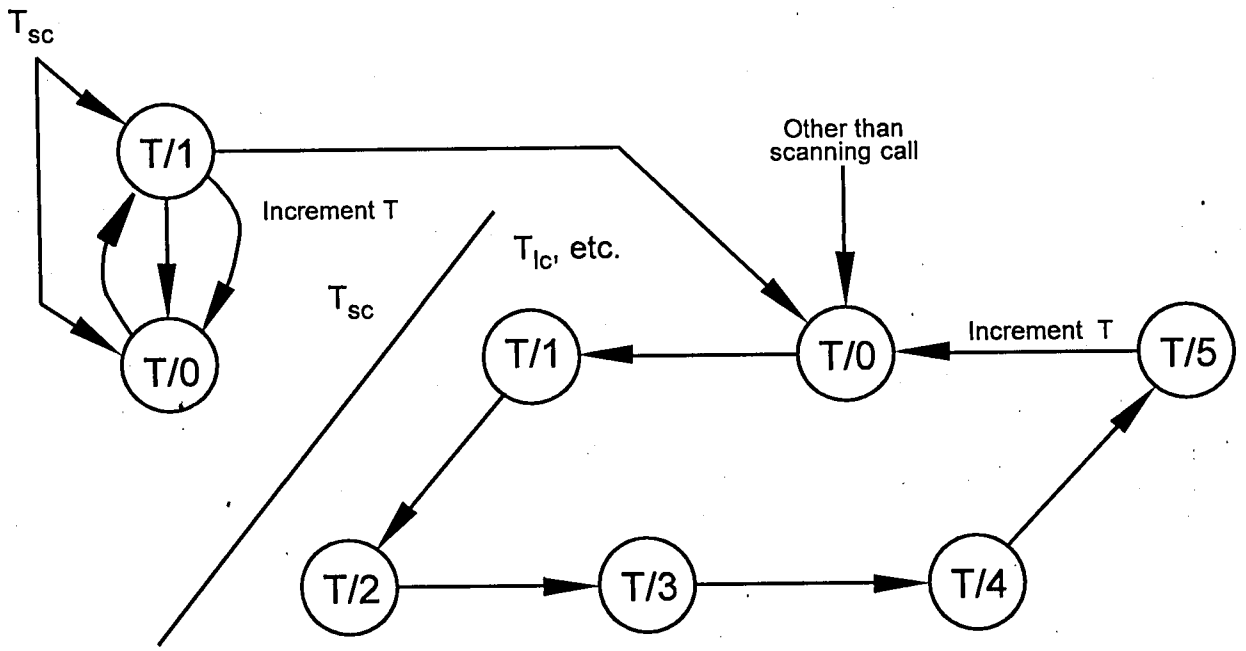

Figure 8. Transmitting station state diagram for a 2-s PI [1].

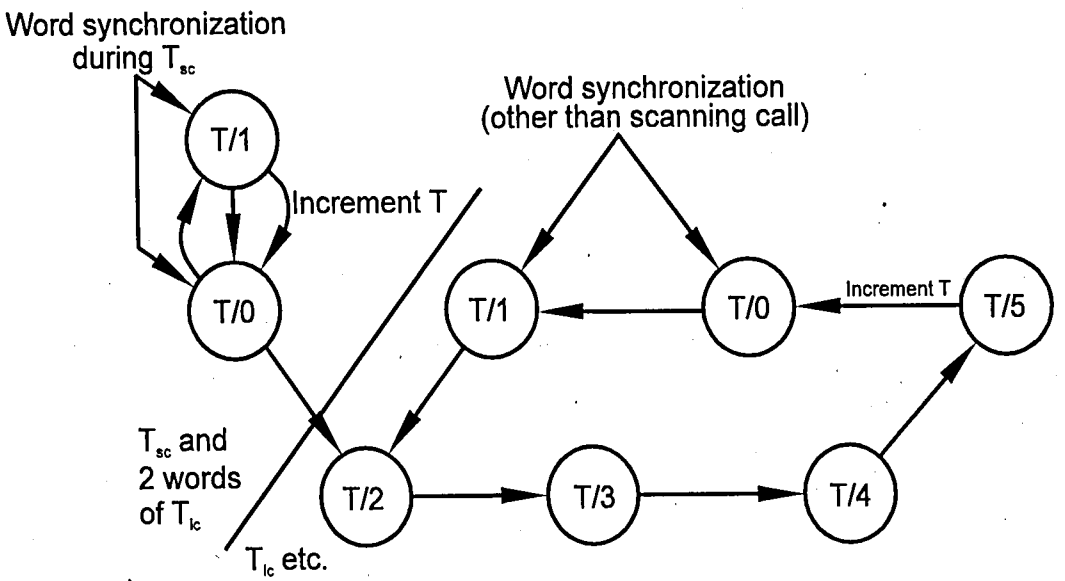

Figure 9. Receiver state diagram for a 2-s PI [1].

5.3.3 Data Block Messages

A Data Block Message (DBM) data block contains an integral number of 12-bit words, the last of which comprises the least-significant 12 bits of a cyclic redundancy check (CRC). These 12-bit words are encrypted in pairs, with the first 12-bit word presented to the LPCM by the ALE protocol as the more significant of the two. When a data block contains an odd number of 12-bit words (i.e., all basic DBM data blocks and those extended DBM data blocks with odd N), the final 12-bit word is not encrypted, but is passed directly to the FEC sublayer.

The word number field of the seed is incremented only after three pairs of 12-bit words have been encrypted, rather than after every 24-bit word as in normal operation. However, the word number field is incremented exactly once after the last pair of 12-bit words in a DBM data block is encrypted, whether or not it was the third pair to use that word number.

5.3.4 Orderwire Messages Using AL-1 and AL-2

Orderwire messages [3] are treated differently under AL-1 and AL-2 than under AL-3 and higher. All command (CMD preamble) words are encrypted. However, the data (DATA preamble) and repeat (REP preamble) words of the Automatic Message Display (AMD) and the Data Text Message (DTM) commands, and the Data Block Message (DBM) data block are not encrypted. Note that the first three characters of an AMD message (contained in the CMD word that begins the AMD message) are encrypted, but the remaining characters of the AMD message are not encrypted. While sending a nonencrypted message, the seed is still sequenced as described above, so that the word following the message (CMD or frame termination) is encrypted using the same seed whether or not the message was protected.

The encryption of all CMD words is necessary so that the receiver can always anticipate the transition to clear mode. The transition back to protected mode is predicted in DTM and DBM CMD words, but must be identified during an unencrypted AMD message by evaluating each received word for protocol compliance in both the clear mode and protected mode.

6. LINKING PROTECTION IMPLEMENTATION

This section presents some suggestions for implementing LP, based on implementation of LP in an ALE system simulator.

6.1 Timekeeping

Synchronization of local times for LP requires some cooperation between the protocol entity and the LP time base. One concept of how the coordination across the ALE-LP sublayer boundary may be affected in this case is as follows:

1. TOD is maintained by the ALE entity, and is provided to the LP entity as required.

2. The transmit LP entity uses the TOD provided by the transmit ALE entity to form seeds during T_{sc} and for the initial time setting for T_{lc}. Thereafter, the TOD from ALE is ignored and the transmit LP entity sequences seeds as shown in the state diagram in Figure 8.

3. On the receive side, seed sequencing is performed by the functions responsible for achieving and maintaining word synchronization. These functions may be implemented within either the LP or the ALE module, but must know the current phase of the ALE protocol (e.g., T_{sc}, T_{lc}, etc.).

4. For authentication of clear mode time exchanges (Appendix A), the ALE module must be able to call upon the LP module to encrypt and decrypt individual ALE words "off-line."

6.2 Ambiguity Resolution

The correct PI and word number to be used at the transmitting station are always determined unambiguously by the state diagram in Figure 8. However, there are several sources of ambiguity at the receiver (see Figure 9):

1. Uncertainty of the time of day at the transmitter.

2. Unknown word number during word synchronization acquisition during T_{sc}.

3. Unknown locations of PI transitions in the received stream of ALE words.

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 synchronization acquisition. In all cases, the usual word synchronization tests (see Section 2.1) must be used to resolve ambiguity,

due to the lack of special synchronization information in the ALE word stream. The following approaches to resolving ambiguity worked satisfactorily in the simulator.

6.2.1 Word Synchronization

During unprotected (non-LP) word synchronization acquisition, the receiver's FEC module examines the received bit stream for 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 synchronization is assumed. The FEC module then checks and returns one word every redundant word time until otherwise notified by the ALE protocol module.

From Figure 1, the place of the LP function in this chain of events 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 under all valid TOD/w number combinations and deliver the results to the ALE protocol module where the final series of tests is applied. In most cases, no more than one combination will produce a word that is acceptable to the ALE module, and TOD synchronization between the transmitter and the receiver is achieved simultaneously with word synchronization. However, on rare occasions a candidate word from FEC will produce acceptable ALE words under two or more combinations, 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. The word that is received after word synchronization acquisition must satisfy the ALE protocol in the context of the previous word when decrypted under word numbers alternating between 0 and 1 (and with a potential PI change). Thus, the word synchronization function can simply wait for the next word following word synchronization and decrypt it under the appropriate combinations to determine which PI/word number combination was the correct one for the previous synchronized word.

When the following word returned by the FEC module contains uncorrectable errors, however, it cannot be used to resolve the ambiguity. One approach to this situation is for the word synchronization function to simply select the most likely successful decrypted word as a correct guess. A more sophisticated implementation could look farther ahead in the data stream for correct words, but the added benefit would probably be small, as the frequency of ambiguities followed by errors should be insignificant for usable channels.

6.2.2 Transitions During the Scanning Call

Note that in Figure 9, three arrows emerge from the T/1 state during T_{sc}. These correspond to simple alteration to word number 0, transition to the next PI and word number 0, and a transition to T/2, indicating that T_{lc} began two words ago.

PI transitions may be detected by looking either forward or backward in the ALE word stream. However, T_{lc} transitions should not be checked by examining future words, for two reasons:

1. The third word following the T_{lc} transition (the T/2 word) may be the first word of a message section or of T_x (the conclusion or the transmission). This results in so many valid possibilities that little is learned by examination.

2. If the third word following the T_{lc} transition is the first word of T_x, there may not *be* any following words to examine.

Thus the following algorithm (which attempts to maximize *a priori* probability, given the previous word) is preferred for detecting PI and T_{lc} transitions during T_{sc}:

1. If the T/1 (previous) word had uncorrectable errors, assume no transition (i.e., T/0).

2. Otherwise, if the current word decrypted under T/0 is valid following the decrypted T/1 word, assume no transition.

3. Otherwise, try T+1/0, and assume a PI transition if the result satisfies a valid T_{sc} word sequence.

4. Otherwise assume the T_{lc} transition (T/2).

6.3 Uniform Vote Threshold

From simulation results, it is apparent that linking protection can produce small but measurable degradations in linking probability unless the unanimous vote threshold (number of unanimous votes output from the majority voter) is employed to reduce the probability of false word synchronization. When the word synchronization 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. With the Golay decoder operating in (6, 1) mode (correcting single errors and detecting up to 6 errors per Golay word), this probability is approximately 3% for each candidate word checked. When this occurs, the preamble and ASCII subset checks will usually reject the word.

Of all 24-bit words, roughly 1% contain three ASCII-38 characters and a valid preamble. However, if the Golay decoders determine that a misaligned word is "correctable" and the resulting ALE word passes the preamble and ASCII checks, the word synchronization function will accept the erroneous word phase and the linking attempt will almost certainly fail.

To estimate the probability of a word synchronization error in nonprotected mode, the equation on the following page should be used. This derivation assumes that the FEC checks a new word for each symbol received (versus each bit), and that correct alignments are uniformly distributed over 0 to 48 symbols, inclusive, after the word synchronization algorithm begins. The following variables are used in the word synchronization error equation on the following page and are defined below.

p_{wse}: The probability of a word synchronization error in nonprotected mode.
p_G: The probability of Golay success on random inputs.
p_{pac}: The probability that a random word passes the preamble and ASCII checks.
p_F: The probability that no false Golay outputs occur before correct word phase.
p_i: The probability that i number of symbols are returned before correct word phase.
p_{FI}: The probability that no false Golay outputs occur in the first i number of symbols.

Thus, in the long run, it would be expected that a nonprotected ALE station would experience a linking failure on ideal channels about once in 200 attempts due to false word synchronization, unless a high unanimous vote threshold is used to reduce p_{wse} (at the possible expense of reduced linking probability over marginal channels). Another possibility for reducing p_{wse} is "soft word synchronization detection." In this case, the receiver stays in the word synchronization seeking mode for at least one additional ALE redundant word time (T_{rw}) after finding an acceptable word, storing and evaluating all potential words until the ALE protocol can unambiguously identify the correct word phase by checking subsequent words.

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 TOD/w combinations (either 6 or 8). The p_{wse} is therefore increased by this factor, resulting in a drop in linking probability of a few percent.

In simulations with a unanimous vote threshold of 0 (e.g., 48 unanimous votes output from the majority voter), p(link) reached a plateau of 98 to 99%; the failures in every case at high signal-to-noise ratio (SNR) were due to false word synchronization acquisition. However, adjustment of the unanimous vote threshold to 25 (23 unanimous votes output from the majority voter) provided full performance in high SNR channels, while minimizing any performance degradation at low SNR.

$$p_{wse} = [1 - p_F]p_{pac}$$

$$= \left[1 - \sum_{i=0}^{48} p_I \, p_{FI}\right] P_{pac}$$

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

$$= \left(1 - \frac{1}{49}\left[\frac{1 - (1 - p_G)^{49}}{p_G}\right]\right) p_{pac}$$

$$\approx 0.005$$

6.4 Orderwire Messages Using AL-1 and AL-2

The requirement to send human-entered unencrypted messages under AL-1 and AL-2 presents another hazard to LP implementors by introducing mode changes in the middle of ALE words. Failure to precisely follow these mode changes at the receiver can result in the loss of otherwise good words. Fortunately, the location in the ALE word stream of all but one of these mode changes can be precisely predicted, minimizing the opportunities for difficulty. There are four cases of AL-1 and AL-2 orderwire messages:

1. Data Text Messages (DTM)
2. Data Block Messages (DBM)
3. Automatic Message Display (AMD)
4. All other messages (LQA, BER, SINAD, time, etc.)

These cases are discussed in the following paragraphs. A key point in preserving ALE performance despite the mode change requirement is that all CMD words are encrypted, regardless of the message type. This ensures that the first word following T_{lc} is always encrypted. If this were not the case, the receiver would always have to allow for a transition to clear mode coincident with the transition from T_{lc}, which introduces a significant loss in linking probability even when no messages are sent.

6.4.1 Data Text Messages

As noted above, the CMD word that begins a DTM is always encrypted. This CMD contains a count of the words in the message proper. The LP function should switch to clear mode for exactly this number of words, starting with the word that immediately follows the CMD word. Any word that follows the DTM will be encrypted. If another orderwire message follows the DTM, it begins with a CMD word (always encrypted); otherwise the next word begins the frame termination, which is likewise always encrypted. During the clear mode message, the TOD fields of the seed are sequenced exactly as if the message were encrypted.

6.4.2 Data Block Messages

The DBM case is identical to the DTM case, except that the seed sequencing follows the usual procedure for DBM mode: the TOD is incremented once for every three words received, and exactly once at the end of the message block.

6.4.3 Automatic Message Display

An AMD message begins with a CMD word, which is always encrypted and indicates that a change to clear mode is required if the message is longer than one ALE word. Having received an AMD CMD word, the receiving LPCM must consider two cases (listed in order of decreasing *a priori* probability):

1. The next word probably continues the AMD message if the word delivered by the receiving FEC module (not decrypted) carries a DATA preamble. The LPCM should switch to clear mode.

2. Otherwise, if the decrypted version of the word carries a REPEAT, THIS IS, THIS WAS, or FROM preamble, the AMD message was probably only one word long. The LPCM should continue in protected mode.

If case 1 is chosen, the LPCM must continually evaluate each arriving word under two hypotheses:

1. If the word continues the AMD message, the nondecrypted word must alternate DATA and REPEAT preambles and contain only ASCII-64 characters.

2. If the word follows the AMD message, it must decrypt to a word acceptable to the ALE protocol.

Because there is a non-negligible probability that a word will satisfy both possibilities, the LPCM must have a rule for choosing whether to stay in clear mode or switch back to protected mode. In such ambiguous cases, it again seems best to select the result having the higher *a priori* probability. For example, if staying in clear mode would allow an AMD message to be longer than allowed, switch back to protected mode. Otherwise, if the decrypted word contains a THIS IS or THIS WAS preamble but an unfamiliar address, assume that the AMD message is continuing.

6.4.4 All Other Messages

All other orderwire messages, such as LQA, bit error ratio (BER), SINAD, and multipath information are internally generated, and will not require a switch to clear mode.

7. CONCLUSION

The use of Linking Protection in HF ALE networks provides a useful degree of security from imitative deception. This protection may come at the cost of a slight decrease in linking probability unless both the LP and ALE implementations are carefully designed to resolve ambiguity in the direction of maximum *a priori* probability.

8. ACKNOWLEDGMENTS

Work on the techniques that eventually became Linking Protection began in earnest in 1988 via a series of meetings among Gene Harrison (MITRE), Bill Beamish (Harris Corporation RF Communications Division), and Eric Johnson (New Mexico State University). As the scheme took shape, Harris Corporation began a proof-of-concept implementation. Chuck Linn, project engineer at Harris, contributed significantly to the effort by finding problems in the LP techniques and suggesting improvements in both the concepts and the documentation.

Chris Redding (then of NSA) and his colleagues contributed to the cryptographic aspects of LP, and helped to clarify that LP is an authentication rather than a privacy scheme. This prompted Dr. Johnson's change in terminology from "link protection" to "linking protection" to emphasize that the technique was intended to protect the linking function, rather than the ALE traffic.

The implementation of LP by Lars Stromberg and his colleagues at Sunair Electronics demonstrated that a protected radio can indeed meet the linking performance requirements of the ALE standards. The members of the HF Radio Standards Development Working Group Technical Advisory Committee all contributed to refining the concepts and improving the wording of the LP standards.

9. REFERENCES

[1] Johnson, E.E., "Linking protection implementation guide," New Mexico State University Technical Report NMSU-ECE-92-005, March 1992.

[2] Johnson, E.E., "Linking protection requirements," New Mexico State University, May 1991.

[3] Johnson, E.E., "Analysis of high-frequency radio linking protection," in Proc. *IEEE Military Communications Conference (MILCOM 92)*, San Diego, CA, 1992, pp. C20.2.1 - C20.2.5.

[4] Redding, C., and E.E. Johnson, "Linking protection for HF radio automatic link establishment," in Proc. *IEEE Military Communications Conference (MILCOM 91)*, McLean, VA, 1991, pp. 49.1.1 - 49.1.5.

[5] National Communications System, *Federal Standard 1045A - Telecommunications: HF Radio Automatic Link Establishment*, 1993.

[6] Lin, S., *An Introduction to Error-correcting Codes*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1970, pp. 32-41.

[7] National Communications System, *Federal Standard 1049 - Telecommunications: HF Radio Automatic Operation in Stressed Environments, Section 1: Linking Protection*, 1993.

[8] U.S. Army Information Systems Engineering Command, USAISEC Technical Report ASQB-OSI-S-TR-92-04, Lattice Algorithm, March 1992.

APPENDIX A: TIME EXCHANGE PROTOCOLS

The following protocols are employed to synchronize LP time bases. The time service protocols for active time acquisition, both protected and nonprotected, are mandatory for all implementations of LP.

A.1 Time Quality

Every time exchange command word transmitted reports the current uncertainty in time of day at the sending station, and whether or not time is transmitted in the command word. The codes listed in Table A-1 are employed for this purpose. The time uncertainty windows in the table are upper bounds on total uncertainty (with respect to coordinated universal time); for example, uncertainty of ±6 seconds is 12 seconds total, and requires a transmitted time quality value of 6.

Table A-1. Time Quality and Corresponding Time Uncertainties

Time Quality Code	Time Uncertainty Window
0	none
1	20 ms
2	100 ms
3	500 ms
4	2 s
5	10 s
6	60 s
7	unbounded

Stations power up from a cold start with a time quality of 7. Time uncertainty is initialized when time is entered, and is maintained thereafter as follows: the uncertainty increases at a rate set by oscillator stability (e.g., 72 ms per hour with a ±10 ppm time base), until the uncertainty is reduced upon the acceptance of time with less uncertainty from an external source, after which the uncertainty continues to increase at the above rate.

A station accepting time from another station must add uncertainty in propagation delays and variations in its processing time to determine its own internal time uncertainty. For example, if a station receives time of quality 2, it adds to the received uncertainty of 100 ms (±50 ms) its own processing delay uncertainty of, ±100 ms for example, and a propagation delay bound of ±35 ms. Therefore, a new time

uncertainty of ± 185 ms, or 370 ms total will be seen at the station. With a ±10-ppm time source, this uncertainty window would grow by 72 ms per hour, so after two hours the uncertainty becomes 514 ms, and the time quality has dropped to 4. In another 20 hours, the uncertainty approaches 2 seconds, and the station (if using a 2-second PI) should request the correct time before it drops to time quality 5 and is presumed to have lost fine time synchronization.

If a low-power clock is used to maintain time while the rest of the unit is "powered off," the quality of this clock is used to assign time quality upon resumption of normal operation. For example, if the backup clock maintains an accuracy of ±100 ppm under the conditions expected while the station is "powered off," the time uncertainty window is increased by 17 seconds per day; such a radio that has been "powered-off" for much over three days may not be presumed to retain even coarse sync, despite its backup clock, and may require manual entry of time.

A.2 Time Service Protocol

The Time Service Protocol is used to efficiently obtain time from another station.

A.2.1 Word Formats

The mandatory time protocols employ the following three types of ALE words: command words, coarse time words, and authentication words, in the formats listed below.

A.2.1.1 Command Word

Time exchange command words, used to request and to provide TOD data, are formatted as shown in Figure A-1. The three most significant bits (W_{1-3}) contain the standard CMD preamble 110. The next seven bits (W_{4-10}) contain the ASCII character '~' (1111110), indicating a time exchange command word.

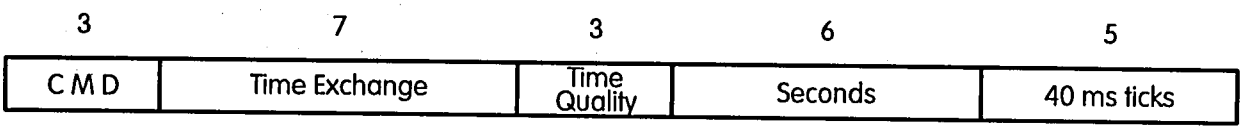

Figure A-1. Time exchange command word.

Time Is Command

The Time Is command word carries the fine time current at the sending station as of the time of transmission of the end of the Time Is command word, and is used in protected time requests and all responses. In a Time Is command word, the Seconds field is set to the current number of seconds elapsed in the current minute (0 - 59), and the Ticks field is set (or rounded) to the number of 40-ms intervals that have elapsed in the current second (0 - 24). The Time Quality command reflects the sum of the uncertainty of the local time and the uncertainty of the time of transmission of the Time Is command, encoded as shown in Table A-1.

When a protocol requires transmission of a Time Is command word, but no time value is available, a NULL Time Is command word is sent, containing a time quality of 7 and the Seconds and Ticks fields both set to all 1's.

Time Request Command

The Time Request command word is used to request time when no local time value is available, and is only used in nonprotected transmissions. In a time request command word, time quality is set to 7, the Seconds field to all 1's, and the Ticks field to 30 (11110).

Other Encodings

All encodings of the Seconds and Ticks fields not specified here are reserved, and should not be used until standardized.

A.2.1.2 Coarse Time Word

Coarse time words are formatted as shown in Figure A-2., and contain the coarse time current as of the transmission of the beginning of that word.

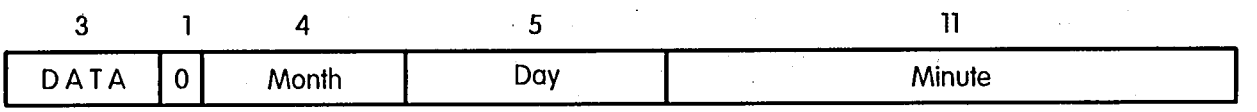

Figure A-2. Coarse time word.

A.2.1.3 Authentication Word

Authentication words, formatted as shown in Figure A-3., are used to authenticate the times exchanged using the time protocols. The 21-bit authenticator is generated by the sender as follows:

1. All words in the time command message preceding the authentication word (starting with the Time Is or Time Request command word that begins the message) are exclusive-ored (a bitwise logical operation that produces a 1 result in each bit position that corresponds to an odd number of 1 bits in the corresponding input bits).

2. If the message to be authenticated is in response to a preceding time command message, the authenticator from that message is exclusive-ored (XOR) with the result.

3. The 21 least significant bits of the final result are used as the authenticator.

(NOTE: The nonlinearity necessary to produce reliable authentication is provided by the normal LP process when an authentication word is so protected; when an authenticator is to be sent in the clear, it can only provide authentication if it is encrypted off line.)

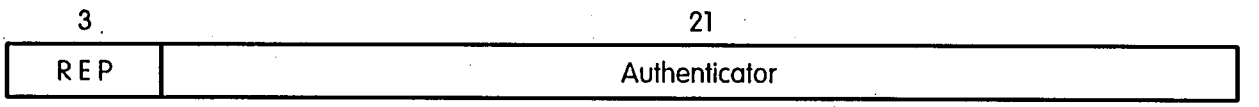

Figure A-3. Authentication word.

A.2.2 Active Time Acquisition (Protected)

A station that knows the correct date and time to within one minute may attempt to actively acquire time from any station with which it can communicate by employing the protocol in the following paragraphs. The quality of time so acquired is necessarily at least one grade more uncertain than that of the selected time server, due to the nature of the Time Service Protocol. A station that does not know the correct date and time to within one minute may nevertheless employ this protected protocol by repeatedly guessing the time until it successfully communicates with a time server.

Time Request Call

A station requiring fine time requests the current value of the network time by transmitting a Time Request Call, formatted as follows:

TO <time server> **CMD** Time Is <Time> **DATA** <Coarse Time> **REP** <authenticator> **THIS IS** <requester>

In principle, any station may be asked for the time, but some stations may not be programmed to respond, and others may have poor time quality; thus, multiple servers may need to be tried before sufficient time quality is achieved.

The Time Is command is immediately followed by a coarse time word and an authentication word. The authenticator is generated by the XOR of the command word and the coarse time word, as specified above.

The Time Request Call transmission is protected using the usual LP procedure. When acquiring time synchronization, the coarse seed (Fine Time field in the seed set to all 1's) current at the requesting station is used; when used to reduce the time uncertainty of a station already in time synchronization, the current fine seed is used.

Time Service Response

A station that receives and accepts a Time Request Call responds with a Time Service Response formatted as follows:

TO <requester> **CMD** Time Is <Time> **DATA** <Coarse Time> **REP** <authenticator> **THIS IS/WAS** <time server>

The Time Is command is immediately followed by a coarse time word and an authentication word. The authenticator is generated by the 3-way XOR of the command word and the coarse time word from this transmission and the authentication word (including the REP preamble) from the requester, as specified above. The entire Time Service Response is protected as usual, using the time server's current coarse seed if the request used a coarse seed, or the current fine seed otherwise. Note that the seed used in protecting a Time Service Response may differ from that used in the request that caused that response.

A time server should only respond to the first Time Request Call using each fine or coarse seed; i.e., one coarse request per minute, and one fine request per fine PI. Acceptance of each class of time requests (coarse and fine) may be disabled by the operator. Stations that are prepared to accept coarse Time Request commands must decrypt the initial words of incoming calls under eight (versus six) possible seeds: w = 0 and w = 1 with the current coarse TOD, and with the current fine TOD ± 1 PI. Note that only one coarse TOD is checked versus checking three fine TODs.

Time Server Request

Normally, the time server concludes the time service protocol by terminating its response with the THIS WAS preamble. A time server may instead request authenticated time from the original requester by returning a Time Server Request, which is identical to the Time Service Response as discussed above except that the THIS WAS termination is replaced by THIS IS preamble. The original requester must then respond with a Time Service Response with an authenticator generated by the 3-way XOR of the command word and the coarse time word from its Time Service Response and the authentication word (including the REP preamble) from the Time Server Request.

Authentication and Adjustment

A station awaiting a Time Service Response attempts to decrypt received words under the appropriate seeds:

1. If the request used a coarse seed, the waiting station tries the coarse seeds used to encrypt its request with w = 0 and w = 1, and those corresponding to one minute later.

2. If the request used a fine seed, the waiting station tries the usual six seeds: w = 0 and w = 1 with the current fine TOD ± 1 PI.

Upon successful decryption of a Time Service Response, the requesting station exclusive-ors the received command and coarse time words with the authentication word it sent in its request. If the 21 least-significant bits of the result match the corresponding 21 bits of the received authentication word, the internal time may be adjusted using the time received in the Time Is command and Coarse Time word, and the time uncertainty updated. (Note that the computed time uncertainty may be greater than the current local time uncertainty, even if the received time uncertainty was less.)

A.2.3 Active Time Acquisition (Nonprotected)

A station that does not know the correct date and time to within one minute may attempt to actively acquire time from any station with which it can communicate in nonprotected mode by employing the protocol in the following paragraphs. Because time is not known in this case with sufficient accuracy to employ LP, the entire exchange takes place in the clear, with the authentication procedure as the only barrier against deception.

Time Request Call (Nonprotected)

A station requiring time requests the current value of the network time by transmitting a Nonprotected Time Request Call, formatted as follows:

TO <time server> **CMD** Time Request **DATA** <Coarse Time> **REP** <"random" #> **THIS IS** <requester>

The Time Request command is immediately followed by a coarse time word, followed by an authentication word containing a 21-bit number, generated in such a fashion that future numbers are not predictable from recently used numbers from any net member. Encrypting a function of a radio-unique quantity and a sequence number that is incremented with each use (and which is retained while the radio is powered off) may meet this requirement.

Time Service Response (Nonprotected)

A station that receives and accepts a Nonprotected Time Request Call responds with a Nonprotected Time Service Response formatted as follows:

TO <requester> **CMD** Time Is <Time> **DATA** <Coarse Time> **REP** <authenticator> **THIS WAS** <time server>

The Time Is command is immediately followed by a coarse time word and an authentication word. The 21-bit authenticator is generated by encrypting the 24-bit result of the 3-way XOR of the command word and the coarse time word from this transmission and the entire random number word (including the REP preamble) from the requester. The encryption uses the AL-1 algorithm, and a seed containing the time sent with w = all 1's.

A time server should only respond to the first error-free Nonprotected Time Request Call received each minute (according to its internal time). Acceptance of nonprotected time requests may be disabled by the operator.

Authentication and Adjustment (Nonprotected Mode)

Upon receipt of a Nonprotected Time Service Response, the requesting station XORs the received coarse time word with the received Time Is command word, XORs the result with the entire random number word it sent in its Time Request Call, and encrypts this result using w = all 1's and the coarse time contained in the Time Service Response. If the 21 least-significant bits of the result match the corresponding 21 bits of the received authentication word, the internal time may be adjusted using the received coarse and fine time, and the time uncertainty updated.

A.2.4 Passive Time Acquisition

As an alternative to the active time acquisition protocols specified above, stations may attempt to determine the correct network time passively by monitoring protected transmissions. Regardless of the technique used to otherwise accept or reject time so acquired, passive time acquisition must include the following constraints to reduce the probability of manipulative deception:

1. Local time may only be adjusted to times within the local window of uncertainty. Received transmissions using times outside of the local uncertainty window must be ignored.

2. Local time uncertainty may be adjusted only after receipt of transmissions from at least two stations, both of which include time quality values, and whose times are consistent with each other within the windows implied by those time qualities.

A passive time acquisition mechanism may also be used to maintain network synchronization once achieved. Passive time acquisition is optional; if provided, the operator should be able to disable it.

A.3 Time Broadcast

To maintain network synchronization, stations should be capable of broadcasting unsolicited Time Is commands to the network, periodically or upon request by the operator:

TO <NET> **CMD** Time Is <Time> **DATA** <Coarse Time> **REP** <authenticator> **THIS WAS** <time server>

The Time Is command is immediately followed by a coarse time word and an authentication word. The authenticator is generated by the XOR of the command word and the coarse time word from this transmission. If the broadcast is made without LP (i.e., in the clear), the authenticator must be encrypted as described in Section 2.3.2 to provide any authentication.

Note that the use of an authenticator that does not depend upon a challenge from a requesting station provides no protection against playback of such broadcasts. A station receiving such broadcasts must verify that the time and the time uncertainty that they contain are consistent with the local time and uncertainty before such received time is used for any purpose.

APPENDIX B: EXAMPLES OF SEED ENCODING AND TIME SERVICE PROTOCOL

This appendix contains examples of seed encoding (Figure B-1.) and of Time Service words (Figure B-2.).

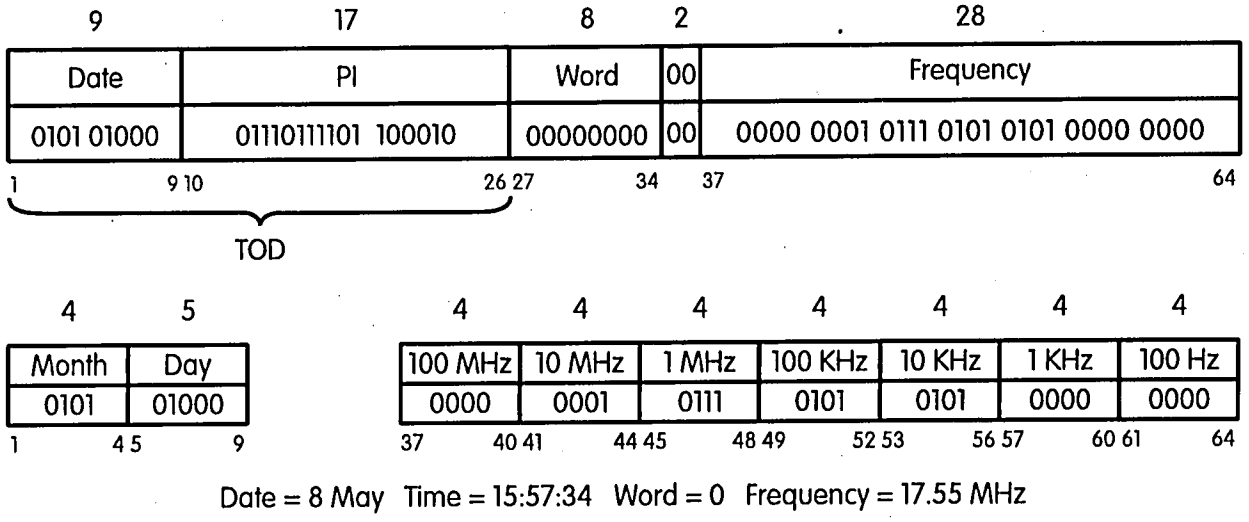

Figure B-1. Example of seed encoding.

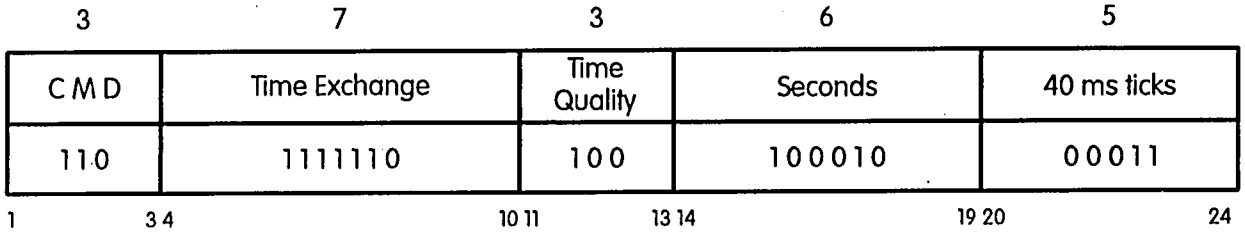

Figure B-2. Example of Time Service protocol.

FORM NTIA-29
(4-80)

U.S. DEPARTMENT OF COMMERCE
NAT'L. TELECOMMUNICATIONS AND INFORMATION ADMINISTRATION

BIBLIOGRAPHIC DATA SHEET

	1. PUBLICATION NO.	2. Gov't Accession No.	3. Recipient's Accession No.
	95-166		

4. TITLE AND SUBTITLE
Implementation Guide for Federal Standard 1049
Section 1, Linking Protection

5. Publication Date

6. Performing Organization Code
NTIA/ITS.N1

7. AUTHOR(S)
E.E. Johnson, C. Redding, D.F. Peach, & R.T. Adair

9. Project/Task/Work Unit No.

8. PERFORMING ORGANIZATION NAME AND ADDRESS
National Telecommunications & Information Administration
Institute for Telecommunication Sciences
325 Broadway
Boulder, CO 80303

10. Contract/Grant No.

11. Sponsoring Organization Name and Address
NTIA, Herbert Hoover Bldg.
14th & Constitution Ave., NW
Washington, DC 20230

12. Type of Report and Period Covered

13.

14. SUPPLEMENTARY NOTES

15. ABSTRACT *(A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here.)*

Automatic Link Establishment (ALE) technology automates the selection of channels and establishment of links among high frequency radios, but creates a vulnerability in such automated networks to hostile manipulation of network operations. The Linking Protection (LP) technique described in this report was developed to protect against such manipulation, while causing minimal degradation to network performance. This report provides a summary of ALE operation, followed by a discussion of the LP technique and suggestions for producing high-performance implementations of LP, based upon simulation studies of protected-mode performance.

16. Key Words *(Alphabetical order, separated by semicolons)*

authentication, automatic link establishment, ALE, high frequency radio, HF, performance analysis, radio networks, simulation

17. AVAILABILITY STATEMENT

☐ UNLIMITED.

☒ FOR OFFICIAL DISTRIBUTION.

18. Security Class. *(This report)*
unclassified

20. Number of pages
44

19. Security Class. *(This page)*
unclassified

21. Price:

*U.S. GOVERNMENT PRINTING OFFICE:1995-673-018/00109

NTIA FORMAL PUBLICATION SERIES

NTIA CONTRACTOR REPORT

Information generated under an NTIA contract or grant and considered an important contribution to existing knowledge.

NTIA HANDBOOK

Information pertaining to technical procedures, reference and data guides, and formal user's manuals that are expected to be pertinent for a long time.

NTIA MONOGRAPH

A scholarly, professionally oriented publication dealing with state-of-the-art research or an authoritative treatment of a broad area. A monograph is expected to have a long lifespan.

NTIA REPORT

Important contributions to existing knowledge but of less breadth than a monograph, such as results of completed projects and major activities, specific major accomplishments, or NTIA-coordinated activities.

NTIA RESTRICTED REPORT

Contributions that fit the NTIA Report classification but that are limited in distribution because of national security classification or Departmental constraints. This material receives full review and quality control equivalent to the open-literature report series.

NTIA SPECIAL PUBLICATION

Information derived from or of value to NTIA activities such as conference proceedings, bibliographies, selected speeches, course and instructional materials, and directories.

SPONSOR-ISSUED REPORTS

NTIA authors occasionally produce reports issued under an other-agency sponsor's cover. These reports generally embody the criteria of the NTIA Report series.

For information about NTIA publications, contact the ITS Technical Publications Office at 325 Broadway, Boulder, CO, 80303 Tel. (303) 497-3572 or e-mail sexton@its.bldrdoc.gov.

This report is for sale by the National Technical Information Service, 5285 Port Royal Road, Springfield, VA 22161, Tel. (703) 487-4650.