Time Distribution for Application Level 2 and 3 Linking Protection

Eric E. Johnson

February 2, 1992
(revised August 6, 1992)

Introduction

The purpose of this report is to suggest operating concepts for maintaining time synchronization in HF ALE networks using Application Level 2 or 3 (AL-2 or -3) Linking Protection. With a 2 second protection interval (PI), AL-2 or -3 linking protection (LP) requires substantially tighter synchronization among stations than does AL-1 LP with its 60 second PI. The factors considered here as bearing on acquiring and maintaining that synchronization include inter-networking requirements, various time sources and distribution schemes, and the inherent time base stability of network member stations.

In some networks, every station may have continuous access to sufficiently accurate time that no over-the-air synchronization is required; other networks may provide such time standards only to selected stations and rely upon time exchange protocols to bring the entire network into synchronization. Still other networks may contain no stations with direct access to Coordinated Universal Time (UTC); such networks may simply maintain internal synchronization without regard to UTC when they do not need to interoperate with other networks, but they must synchronize with those other networks when required to interoperate.

Three time exchange protocols are discussed in the Linking Protection Implementation Guide (LPIG). The Time Service protocol provides a simple and effective way of distributing time, although the window of uncertainty of its time reports increases quickly with the number of intermediate stations between the time standard and the ultimate time user. The Time Iteration protocol overcomes this rapid increase in uncertainty by using extended handshakes in the time exchange process, with statistical techniques used to estimate the correct time arbitrarily closely, despite variations in processing and propagation times. Finally, the Precision Time protocol can be used by appropriately modified radios to transfer time with very little introduced uncertainty.

Time Service is included in the AL-1 LP standard. Time Iteration is developed but still under discussion, and Precision Time is yet to be developed. A key focus of this investigation is how well a network using AL-2 or -3 LP can operate and interoperate using only the Time Service protocol.

The Basic Timing Relationships

Time Uncertainty Windows

Certain basic timing relationships govern the operation of networks using AL-2 or -3 LP. Key among these is the requirement that the local times used for LP at any pair of stations should differ by no more than the length of the PI, or 2 seconds. (This can be relaxed if the receiving LP module examines more PIs than those specified in the LPIG, but this will reduce both P(link) and the protection afforded by LP.) Although all stations in the network may initially be precisely synchronized, their local times will drift apart due to oscillator drift in their time bases.

If we assume that time base oscillators vary by no more than ±10 ppm (the accuracy required for ALE timing), the local time at any station drifts by no more than ±864 ms per day. Thus, the difference in local times among any pair of stations increases by no more than 1728 ms per day. (In the worst case, one station drifts 864 ms ahead, while another drifts 864 ms behind, for a difference of 1728 ms). Because the stations in the network were initially synchronized, the total window of time uncertainty at any station (1728 ms in this example) is also equal to the worst-case difference between any pair of those stations.

Time Quality

The LPIG specifies that all stations must keep track of their local time uncertainty window. When reporting local time using one of the time exchange protocols, the time report is qualified by reporting this time uncertainty window using the following time quality codes. The uncertainty windows listed are intended to be relative to UTC. In stand-alone networks, an alternative time standard may be used, with some variations in operation as described later in this report (e.g., only UTC time standards should report time quality 0).

Table 1: Time Quality

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

Updates

From Table 1, it is clear that stations using AL-2 or -3 LP must maintain time quality of 4 or better to be assured of reliable protected communication within their network. When its uncertainty window approaches 2 seconds, a station should initiate attempts to reduce this window by synchronizing with a station with better time quality. When this synchronization is achieved, the local time uncertainty is set to the sum of the reported uncertainty of the new time plus any uncertainty in processing and propagation delays.

As an anti-spoofing measure, a station should only accept an updated time that makes sense in the context of its current local time and uncertainty window and the uncertainty window implied by the time quality received in the time report. Because the windows are generally relative to UTC, and each station should have drifted no more than one half of its window from standard time, the receiving station could use an acceptance range equal to (local time) ± (half of its window + half of the received window). However, the time reference in stand-alone networks may not be tied to UTC and may, in fact, drift at the same rate as other network member stations. Therefore, the acceptance range should be computed as (local time) ± (local uncertainty window), without reference to the received window.

There is one exception to this rule: when the received time quality is 0, the sending station is claiming to be a UTC time standard, and should therefore have negligible drift. Therefore, when the received time quality is 0, the acceptance window should be tightened to (local time) ± (one half of the local uncertainty window).

Update Intervals

Propagation times are bound to fall within a range of 0 to 70 ms, giving a maximum uncertainty of 70 ms. If we assume that the time of release of a transmission has an uncertainty of 8 ms (the baud width), and that the receipt and processing delay at the receiver varies by a scheduling quantum of 16 ms, the total introduced uncertainty may be approximated by 100 ms. Given this assumed uncertainty in passing time using the Time Service protocol, we can derive the figures below:

Table 2: Update Interval for 100 ms Introduced Uncertainty (±10 ppm Time Base)

	Hours to Drift Below	
Source Time Quality	Time Quality 3	Time Quality 4
0	5.6	26
1	5.2	26
2	4.1	25
3	—	19

If, on the other hand, we assume that the propagation time is known precisely, and that uncertainties in the time of release and the processing delays amount to no more than one baud (the Precision Time scenario), the above figures change as follows:

Table 3: Update Interval for 8 ms Introduced Uncertainty (±10 ppm Time Base)

	Hours to Drift Below			
Source Time Quality	Time Quality 1	Time Quality 2	Time Quality 3	Time Quality 4
0	0.167	1.2	6.8	27
1	—	1.0	6.5	27
2	—	—	5.4	26
3	—	—	—	20

Thus, a higher-precision time exchange protocol increases the interval between required updates by only one hour, which is only moderately useful in trying to maintain time quality 3, and even less so for time quality 4.

A better approach to lengthening the interval between time exchanges is to reduce time base drift below ±10 ppm, either by placing the oscillator in a controlled environment (e.g., a crystal oven), or by using a more stable oscillator. For example, improving time base stability to ±1 ppm results in a ten-fold increase in update intervals, as seen in Table 4.

Table 4: Update Interval for 8 ms Introduced Uncertainty (±1 ppm Time Base)

	Hours to Drift Below			
Source Time Quality	Time Quality 1	Time Quality 2	Time Quality 3	Time Quality 4
0	1.67	12.8	68.3	276
1	—	10.0	65.6	273
2	—	—	54.4	262
3	—	—	—	207

Note that the computations for Tables 3 and 4 assume that the actual time uncertainty at the source is equal to the maximum for the respective time quality values. The update intervals would be slightly longer if the source uncertainty is somewhat less.

Stand-Alone Networks

By a stand-alone network is meant one that *normally* does not call or accept calls from any other network. Cross net operation is possible, but is a special case, requiring special procedures.

Normal Operation

Under normal circumstances, synchronization *within* the network is all that is required; synchronization with UTC is not necessary. One station is selected to act as the network time reference (NTR). An especially stable time base would be useful at the NTR, but is not required. One or more of the other stations may be assigned to serve as alternate time references in case the primary reference becomes unreachable or is compromised.

The NTR permanently assigns itself a time quality of 1 (not 0, as discussed previously). Net members maintain their time quality at 3 or 4 via time service from the NTR or alternates. Alternate time references must maintain their time quality at 3, so that network members relying on them can achieve time qualities of 4. Thus, using the Time Service protocol figures from Table 2, alternate time references must update their time from the NTR at intervals not exceeding 5 hours. Net members require Time Service from the NTR (or an alternate) at least once per day; when time has been acquired from an alternate, time service will be required again within 19 hours. (These update intervals may be lengthened by improving time base stability beyond ±10 ppm.)

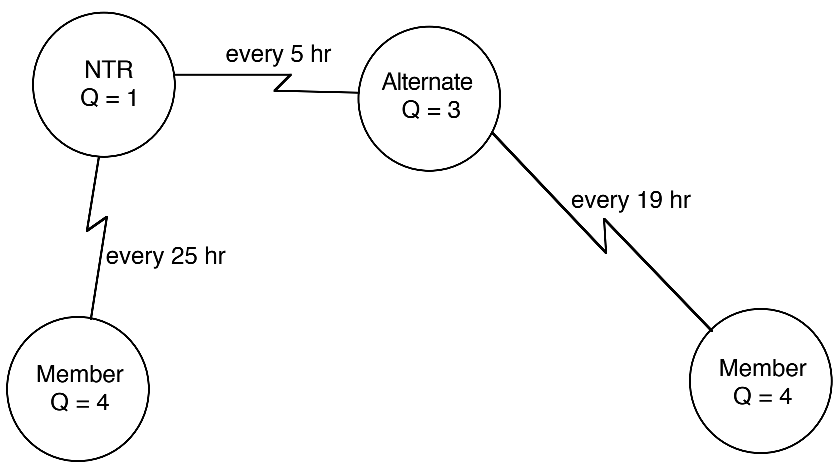

Figure 1: Stand-Alone Network

Cross-Net Operation

When a station needs to operate in more than one network, either the two networks must be synchronized, or that station will need to maintain multiple time bases, one synchronized with each network in which it operates. If the latter option is chosen, it is probably best implemented by using multiple radios at the station, with message or voice patching provided among the radios as required to permit local net members to access the other network.

If, instead, a network that normally operates in stand-alone mode (as described in the previous section) needs to synchronize with another network, the approach with the minimum operational impact is for the designated NTR to synchronize with the best quality time that it can find in the new network and to continue to serve as a time server in its local network. (Note that the NTR is not serving as a *traffic* gateway into the new network, but merely as a focal point for synchronization.) If another member of the local network can reach a station in the new network with better time quality than can the NTR, the NTR function should be re-assigned accordingly.

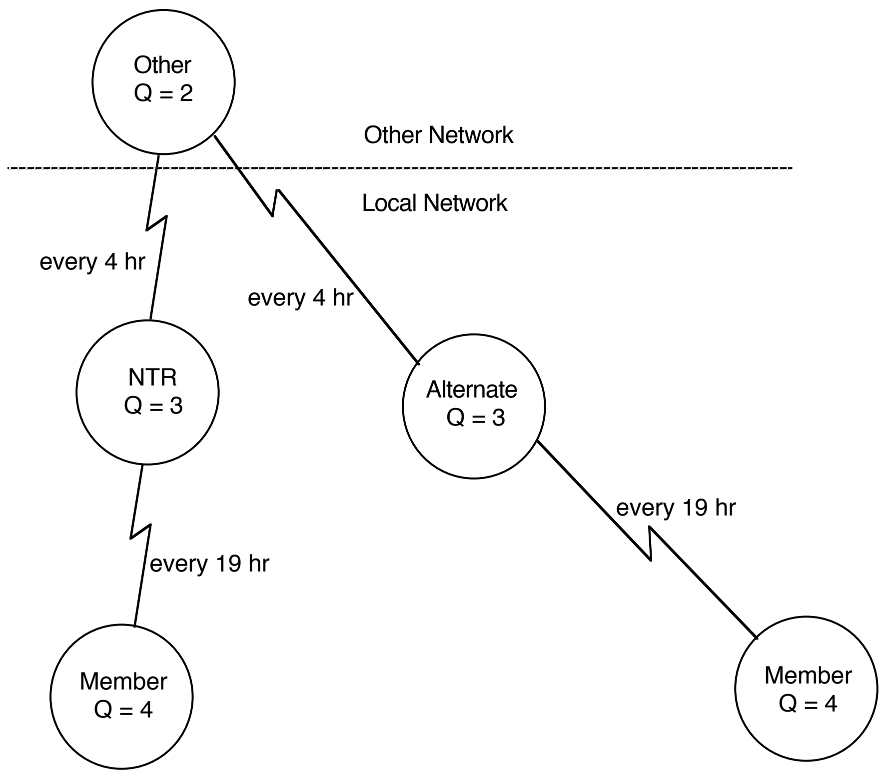

Figure 2: Cross-Net Operation

The Time Service update rates in the local network will vary depending upon the quality of time that the NTR and local alternates are able to maintain, which in turn depend upon the time quality available from the other network and the time exchange protocols available. Using only the Time Service protocol, the NTR will probably be able to achieve time quality no better than 3, assuming that quality 2 or better time is available from the other network; in this case, the NTR must serve all local network members directly. If any members cannot reach the NTR, an alternate time reference must also directly synchronize with the other network, as shown in Figure 2. If the best time quality reachable in the other network is 3, *all* local network members must synchronize directly with a time server in the other network. Finally, if only quality 4 time is reachable, we have the following options:

1. A time exchange protocol capable of bringing local network members up to quality 4 (Time Iteration or Precision Time) must be used by all local network members.
2. The network may attempt to operate with quality 5 time (which may prove unreliable).
3. The networks may drop back to AL-1 LP. In this case, it would probably be best to initially acquire time from the other network to more or less align the time bases across the boundary, but to maintain normal local synchronization for AL-2 or -3 intra-network operation, with AL-1 LP used only for cross-network links.

Note that synchronizing a stand-alone network to an external time standard will probably require a one-time intervention by operators to override local acceptance windows, because the two time standards may initially be far apart.

World-Wide Internet

By the world-wide internet is meant all networks that are synchronized to UTC. The two cases of interest here are normal operation, during which the usual coordinated time sources such as the U.S. Naval Observatory, WWV, and GPS are available, and a reconstitution scenario, in which these time sources are assumed to be unreachable by some or all networks. The following classes of networks are of interest:

1. All network members have the hardware to directly synchronize with these external time sources to within ±10 ms (time quality 1).
2. Only network time servers have such hardware.
3. No network members have such hardware.

Normal Operation

In normal operation, stations with the requisite hardware will stay at time quality 1. Class 2 network members without such hardware will easily be able to maintain time quality 4 with daily time service from their time servers.

Class 3 networks will operate similarly to stand-alone networks in cross-network mode, as described above, except that network planners will have ascertained the availability of adequate time servers in establishing the networks as part of the internet. Thus, the time servers in class 3 networks should be able to maintain time quality 3 by periodically synchronizing with stations in other networks having time quality 1. The other members of class 3 networks synchronize with their time server(s) to stay at time quality 4. Thus, class 3 networks should be able to sustain normal operations without significant difficulty.

Reconstitution

Under the reconstitution scenario, some networks will not be able to acquire time from the usual sources of UTC, and must rely upon either the ALE time exchange protocols or some other mechanism. This suggests that all stations, whether or not they are equipped with specialized time receivers, must be able to use at least the time service protocol.

In the extreme case of being unable to synchronize directly to UTC over large areas, synchronization of the internet will depend upon the ability of networks in the affected areas (the "disturbed zone") to acquire high-quality time from other stations. In such a situation, a key consideration is minimizing the height of time distribution trees, measured from surviving time standards to the most remote time users. Stations able to use their specialized time receivers will continue to maintain quality 1 time; using only the Time Service protocol, stations in direct contact with these quality 1 stations can maintain time quality 3, and can therefore serve to keep all stations that *they* can reach at time quality 4.

Thus, specialized time receivers plus the Time Service protocol support trees of up to four levels, as shown in Figure 3. Because of the heavy trafic load anticipated in the reconstitution scenario, the number of stations attempting to maintain time quality 3 should be minimized to reduce the air time consumed by Time Service handshakes as well as the overhead time consumed at the time quality 1 stations.

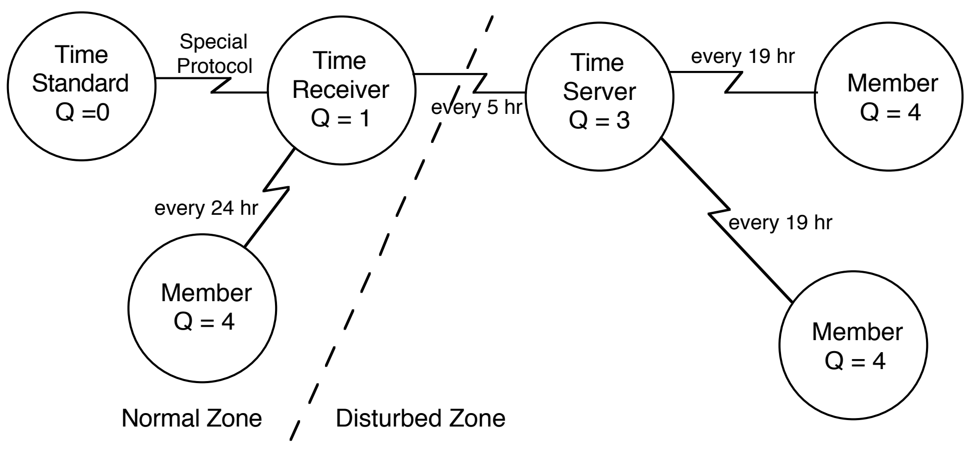

Figure 3: Reconstitution Using Time Service

If time servers in the distrubed zone can acquire time from a time quality 1 station with only 8 ms of introduced uncertainty, they can maintain time quality 2 for one to ten hours, as determined by their local time base stability (±10 to ±1 ppm). Each of these stations could then support several time quality 3 stations using any of the time exchange protocols, resulting in a five-level time distribution tree.

If all sources of UTC are lost, some station must be designated as the new time standard. If the time base at this station is quite stable (e.g., a cesium or rubidium oscillator), this new time standard will assume time quality 0; otherwise, it must use time quality 1, as discussed before. The new time standard will serve as the root of a time distribution tree whose height depends upon the protocols in common among the stations to be synchronized, and their time base stabilities. To minimize congestion at the new time standard, only those stations with highly-stable time bases (better than ±10 ppm) should synchronize directly with it, in order to maximize the update interval of these first-tier stations. The first tier then serves as time servers for all stations that they can reach, and so on. If the first-tier stations can maintain time quality 2, AL-2 or -3 LP can be sustained among all stations in a tree of height 4, using only the Time Service protocol below the first tier. If each time server provides time to an average of ten stations, this tree will contain over 1000 stations.

Conclusions

From the preceding discussion, the following conclusions may be drawn:

1. Under normal circumstances, AL-2 or -3 LP requires only the Time Service protocol for time base synchronization.
2. For normal operation, time base stability better than ±10 ppm appears to be more valuable than time exchange protocols that introduce less uncertainty than the Time Service protocol.
3. Cross-net time exchange between stand-alone networks requires only the Time Service protocol if quality 3 time is available from one network to all stations in the other, or if the time reference in a network can stay within ±500 ms of the other network's time.
4. Stations with specialized time receivers (e.g., GPS) are quite useful as time servers when the required time signals are available.
5. All LP-capable stations must support the Time Service protocol, even if equipped with GPS receivers. Stations with such time receivers may be called upon to provide high-quality time in cross-net operations, and will need some mechanism to acquire time themselves if the time standards become unreachable.