NOT USING ALE FOR FREQUENCY SELECTION

Steve Kille
Isode Ltd.
14 Castle Mews, Hampton
TW12 2NP
United Kingdom
steve.kille@isode.com

SUMMARY

The core function of ALE (Automatic Link Establishment) has historically been to select a frequency to use prior to exchanging data or voice. An application will request to link to a peer identified by an ALE address and the ALE layer will select the frequency and establish a link. This paper explores an architecture where the choice of frequency is made above the ALE level. The ALE layer is still used to establish links using a frequency explicitly requested from the higher layer.

The paper starts by considering a number of downsides of ALE frequency selection, including:

- ALE waveforms are different to the ones used for data transfer
- For Wideband HF, ALE measurements use only 3kHz and not the full bandwidth
- Sounding times are too short to allow for HF Intermediate Term Variation
- Sounding overhead reduces time channel can be used and increases latency
- ALE units cannot generally take into account peer location
- ALE units often pick a poor frequency after an outage that affects many frequencies
- Poor handling of asymmetric links

The paper then looks at how all of these problems can be addressed by moving frequency selection above the ALE layer and in particular integrating the choice with the STANAG 5066 layer and using frequency pools to minimize frequency count. It shows how frequency predication models and knowledge of mobile unit (peer) location can be used to more intelligently select the best frequency and improve overall performance.

It then considers the related problem of terminating ALE links and making full use of the ALE link management procedures specified in STANAG 5066 Edition 4[1]. There are several trade-offs, making this a harder problem than initial link establishment.

Finally, the paper looks at experience with current systems implementation considerations for the approaches described in the paper.

1 ALE AND STANAG 5066

The specification of Automatic Link Establishment (ALE) referenced in this paper are specified in US MIL-STD-188-141D[2]. There are three distinct ALE variants specified that are referenced in this paper:

1. 2G ALE. This is the oldest ALE variant still in widespread use. It is simple, robust and uses an asynchronous approach.

2. 3G ALE. This adds a number of improvements. The key difference with 2G is that it is synchronous and relies on all systems having synchronized clocks.
3. 4G ALE. The newest variant can operate synchronous or asynchronous. It also adds a third "staring" mode which uses sampling radios to listen on all frequencies in order to support much faster linking. 2G and 3G are for used with 3kHz narrowband HF. 4G supports Wideband HF with bandwidth up to 48 kHz by negotiating the bandwidth to be used in each direction. This function makes 4G ALE a core component of any Wideband HF deployment.

This paper is concerned with data applications operating over HF and it considers the STANAG 5066 data link layer. Edition 4 of STANAG 5066 introduces a specification of using ALE; some detailed aspects of this are discussed below.

Although the details of ALE are complex, the basic model is very simple. There is a list of HF frequencies agreed by all nodes. A receiving system will listen on each frequency in turn or an all frequencies together in 4G staring mode. Then the sending system will test frequencies by sending a short message. When the receiving system hears a transmission from a sending system, it will respond and a link is established.

The standard ALE usage model is that an application which desires to send data to a destination will request to send data to an identified destination and the ALE layer will try frequencies in turn to establish a link. The choice of frequency is completely transparent to the application. This is ideal for a voice call. For data, it provides a clean separation between application and ALE.

Although this standard model works for data, it has a number of draw backs discussed below. This leads to the superficially strange new approach that ALE, which originally had a primary goal of selecting a frequency is not actually used for the frequency selection.

2 DOWNSIDES OF ALE LAYER SELECTING FREQUENCY

We now consider various consequences of the standard model and why this leads to sub-optimal performance.

2.1 DIFFERENT WAVEFORMS

2G ALE used a very different waveform than associated data waveforms. Modern ALE protocols make use of specially optimized waveforms which are similar to the waveforms used for data transfer. The ALE layer will make choices of frequency based on characteristics of the ALE waveforms. If there are differences in propagation characteristics of the data waveforms that will be used to transfer data, then there is no possibility for the ALE layer to take this into account. It is impractical to measure this and it is likely that this is not a major issue.

2.2 VARIATION ACROSS THE BANDWIDTH

For Wideband HF, 4G ALE can negotiate channels of up to 48 kHz. Each of the ALE frequencies will have a maximum bandwidth available. However the ALE protocol only operates over a 3kHz channel within the target frequency range. This means that ALE frequency choice is made only on the characteristics of this one part of the channel. If this part of the channel does not have characteristics equivalent to the whole channel, this has potential to lead to sub-optimal frequency choice.

2.3 LQA AND SOUNDING CONSIDERATIONS

ALE systems will use Link Quality Analysis (LQA) information in order to determine the order in which frequencies are tried. LQA information is obtained when the ALE protocol operates, but this is generally insufficient to support making good frequency choices. Sounding is an ALE capability used to provide better LQA data for frequency choice. There are two primary approaches:

1. Standard Sounding. Each node will, at intervals, send out data on each frequency. Other nodes will receive the sounding data and record information on channel quality. This is a simple approach, which works well for channels that are not used for a large percentage of the time.
2. Sounding Exchange. This is a pro-active approach where a node will attempt to communicate with a single peer on all frequencies and the peer will respond on each frequency that it hears.

Isode's Icon-5066 product supports both mechanism. Standard sounding is handled by the underlying ALE Unit. Focus has been on Sounding Exchange, as target environments will typically have higher loads that are not appropriate for standard sounding. Prior to establishing a link, Icon-5066 will request Sounding Exchange if link quality information for the peer is older than a configurable period (typically 30-60 minutes). This has an overhead of a minute or two for each peer, once or twice per hour. It would be desirable to eliminate this overhead.

A second issue is HF Intermediate Term Variation. The paper "Optimizing applications and data links for HF radio intermediate term variations: Can you ride the wave?"[3] considers the sampling lengths desirable to make good assessment of channel quality. It concludes that long sampling leads to the best choices. ALE sounding uses much shorter transmission times, in order to keep the sounding overhead reasonable. A consequence of this is that sample time is significantly shorter than optimal.

2.4 ASYMMETRIC LINKS

ALE is commonly used to establish links between a pair of nodes where data is transferred in both directions. ALE systems using standard sounding will make choices based on receive measurements only. This can lead to poor channel choice when channels have asymmetric characteristics.

2.5 PEER LOCATION

When the location of a peer is known, this could be used to significantly optimize frequency choice base on HF Prediction models such as ITU P.533[4]. This is not something that can be done using the standard ALE model, as an ALE Unit is not aware of peer location.

2.6 BEHAVIOUR AFTER OUTAGES

An ALE Unit will generally try first the frequency it believes will be best and then go through other frequencies in turn. If none connect, it will then cycle through all of the frequencies.

This works well when some frequencies are propagating and others are not. A common HF scenario is that none of the available frequencies will propagate for a period. When this happens, the frequency chosen after the outage will be driven by which one happens to be tried next; it will often be a frequency that propagates, but is not the best choice.

3 MOVING FREQUENCY CHOICE UP A LAYER

The key change to address these issues is to move the frequency choice up a layer, which for standardized data applications means using STANAG 5066. The change is that the application directive to ALE changes from "link to this ALE address" to "link to this ALE address on this frequency". Isode investigation of ALE products suggest that many, but not all, provide interfaces to support this new mode of operation.

The ALE protocol and STANAG 5066 protocol are unchanged; this is just a shift to making frequency choice at a higher layer. This can happen because the relevant NATO protocols follow the OSI model of standardizing protocols with service interfaces as descriptive; compliance and interoperability is to the protocols. An implementation is free to modify the service interfaces, as is being done here, provided that protocols are maintained.

3.1 WHY NOT DOWN A LAYER

The same overall effect could be achieve by the STANAG 5066 Layer providing information downwards to ALE, including Quality of Service requirements, information on 5066 level LQA, peer location and information on peer movement. This is an implementation choice.

This approach as been taken by the US Navy in the PAALE (Propagation-Aware ALE) system. Details of this system are not know.

Isode provides a STANAG 5066 server, which it wishes to work with many ALE units. There is a clear commercial rationale to put the new functionality in the STANAG 5066 layer. Moving the functionality down to ALE would also significantly increase the complexity of the STANAG 5066 / ALE interface.

3.1.1 Red/Black Considerations

Military applications deploying STANAG 5066 generally use a red/black architecture where applications sit on the red side, with modems and radios sitting on black side with all data flowing through a crypto. There are a number of ways to split functionality across the boundary, noting that traditionally STANAG 5066 sits red side and ALE on black side. However the split is made, there is a need for control information to flow red to black and status information to flow black to red, with separation provided by crypto bypass.

From a security perspective it seems preferable to keep as much information as possible on red side. If mobile unit location and direction/speed of travel is being used to determine which frequency is used, it seems desirable that this information is managed red side. This consideration suggests it is preferable to move the decision functionality upwards rather than passing additional information downwards.

3.2 OPTIMIZING BASED ON PEER LOCATION

Support for HF Prediction can be added cleanly with this new model. This requires knowledge of peer location. This might be achieved over STANAG 5066 using a mechanism such as the HF-LISP[5] location protocol proposed by Isode. This can be used as input to selecting frequency. This will obtain for each frequency a prediction of probability of propagation and expected link quality. Frequencies that will not propagate can be eliminated immediately. Other frequencies can be tried in order giving preference to links predicted to be better. A secondary consideration is to look at prediction of how a link will change. It is sensible to give preference to an improving frequency when the link is expected to be long lived.

When all frequencies have been tried, it is recommended to retry frequencies that are predicted to be better more often than those predicted to be poorer. This will reduce the chance of the poor behaviour after outages described above.

3.3 STANAG 5066 DOING SOUNDING

STANAG 5066 Edition 4 communicates link quality to peers be sending a recommended transmission speed for bulk data. This means that STANAG 5066 layer will have a clear view of the quality of any link in both directions. Measurements can be made over an extended period, which allows for Intermediate Term Variation and avoids the issues caused by short ALE sounding measurements. This STANAG 5066 level LQA can be used in conjunction with ALE level LQA provided by ALE Unit to STANAG 5066.

It is to be hoped that use of HF Prediction will often lead to a good frequency choice. The HF Prediction will also provide enough information to give a baseline expectation of performance for a link over a given frequency. If a given link is exceeding expectations, it will generally make sense to continue. If not, another frequency can be tried. STANAG 5066 performance information will be obtained for the new link, which will enable a choice of keep going vs revert to previous vs try another frequency. Similarly a choice can be made for a new link to either use a new frequency or to use a recent one for which there is 5066 sounding data.

Because sounding is done by STANAG 5066 layer, data is being sent and the time for sounding is not being wasted. It also means that sounding is being done using the “correct” waveforms and bandwidth, addressing the first concerns.

4 FREQUENCY POOLS

ALE linking time, apart from 4G ALE staring, is heavily dependent on the number of frequencies. HF ALE Networks are generally designed to provide coverage over a target geographical area. HF frequency propagation varies significantly over the day for a given frequency, so to provide 24 hour coverage a quite large number of frequencies is generally needed. Using all of these frequencies all of the time, as many ALE systems do, will slow down ALE linking significantly. At any given time, it is likely that a number of frequencies in the pool will not propagate. While the HF Prediction described above will avoid wasting time trying to connect over such frequencies, the ALE searching overhead remains. It is desirable to optimize here.

The solution to this is to vary the frequency pool, as originally proposed in HF2000[6]. HF Prediction can be used to determine, for each of a set of periods during the day, which are the useful frequencies for the target ALE network area. Then each node on the network can change ALE configuration at the allotted time. This will ensure that ALE keeps working without break and performance will be optimized by removing poor options from the ALE frequency pool in each period.

5 CLOSING ALE LINKS

A problem closely related to frequency selection is determining when to close an ALE links, which for a number of reasons is a harder problem. The first consideration is STANAG 5066 link management

5.1 CLOSING ALE LINKS IN STANAG 5066

STANAG 5066 Ed4 added a specification of using ALE and two mechanisms to close links.

STANAG 5066 in ALE 1:1 mode maintains a logical soft link (CAS-1) over the ALE link. In some situations a link is run until there is no more data and both links are then closed. Two situations require the link to be closed while there is more data to send:

1. One of the nodes has higher priority data to send to another node and needs to close the current ALE link in order to do this.
2. There is a "fairness" policy and a node desires to close a link to send data of the same priority to another node or to allow another node to connect.

To support this, there is a negotiated "clean close" mechanism, where both ends stop sending new data and ensure that all PDUs where transmission has started are fully delivered and acknowledged. This gives a mechanism to ensure fairness in a system with many nodes.

The second mechanism provides ALE re-linking. The CAS-1 link is retained and partially transferred PDUs will continue after the re-link. The goal of this is to allow the STANAG 5066 server to decide that it wants to try a different ALE link to the current one.

5.2 LINK QUALITY CONSIDERATIONS IN LINK CLOSE

When opening an ALE link, essentially there is a need to pick a frequency and the requirement is to pick the best one. Closing and re-opening an active link is a similar decision, but there needs to be qualitative view that there is a good probability of being able to open a link that is better than the current one. Considerations include:

1. What is the expected performance of the current frequency and alternative frequencies. HF Prediction can help with this to provide an appropriate benchmark for the current frequency and relative performance of other frequencies.
2. When to change. HF performance varies over time and it may be preferable to keep going rather than switch. It is anticipated that it will generally be sensible to keep going with a link for multiple minutes before switching. This will also enable confidence in recording frequency performance.
3. Recent performance of other frequencies. Recent switches can provide information on the benefits of switching and potentially giving indications of performance being affected on a range of frequencies.
4. Bandwidth too low. It may be that re-linking on the same frequency with a higher bandwidth will enable better performance to be obtained from the frequency under improving conditions.

5. Bandwidth too high. It may be that re-linking on the same frequency with a lower bandwidth will enable better performance when conditions are deteriorating.

It can be seen that there are a number of factors influencing choice and that it will be a complex decision to optimize.

6 ISODE EXPERIENCE AND PLANS

Some of the issues noted with letting ALE module choose frequency are theoretical, but most are arising from Isode observations which are described here. Isode's initial approach to ALE implementation in Icon-5066 was to treat ALE as a black box. This did not facilitate effective analysis, and the need to audit Link Quality Analysis (LQA) became clear. This audit proved that standard sounding was inadequate, as it did not provide sufficient information for the ALE to make good frequency choices. This was addressed by use of Sounding Exchange when LQA information is stale. This improved ALE performance, but led to delays of around two minutes in establishing some ALE links. This is seen as acceptable, but sub-optimal.

Experience suggests that ALE often picks a good frequency with appropriate bandwidth, but not always. Icon-5066 currently will use an ALE link until there is no more traffic or the link fails. There are regularly situations where a link is used for an extended time where it is clear from historical analysis that it would have been preferable to drop the link and either move to a new frequency or change bandwidth. Specific scenarios observed:

1. A frequency is selected which performs poorly and knowledge of expected propagation suggests that an alternative frequency would have been better. This typically occurs after a long ALE linking time, which suggests there has been a period when no frequencies were propagating. It seems likely that the link is made on the first frequency tried as conditions generally improve.
2. Bandwidth too low. A link will operate well, but analysis suggests that better performance would have been obtained at a higher bandwidth which the good SNR suggests could have been supported.
3. Less frequently, a link struggles on with bandwidth that is too high. Operation will be at the lowest speed for the bandwidth and a high percentage of data is lost. A narrower bandwidth would allow more average power and lower speeds that would give better performance.

There was also a general concern that it was not clear why the ALE picked a given frequency. By moving control of frequency to the STANAG 5066 layer, it would be possible to better understand the whole system.

The analysis of addressing the issues presented in this paper reflects the intended direction of Isode product evolution. Given the detailed ALE handling currently provided, this will be an incremental shift of capability.

7 CONCLUSIONS

This paper has described how the standard ALE approach of having the ALE module select frequency is sub-optimal. This feels initially strange, as the initial goal of ALE was frequency selection. By moving frequency selection into the STANAG 5066 layer, performance can be enhanced, particularly by removing the need for independent sounding and by making use of HF Prediction based on location.

8 ACKNOWLEDGEMENTS

Thanks to Dr Eric Johnson for an insightful review of a draft of this paper.

REFERENCES

[1] *STANAG 5066 - Edition 4 " TECHNICAL STANDARDS FOR HF RADIO LINK LAYER AND APPLICATION SUPPORT PROTOCOLS FOR SINGLE CHANNEL WAVEFORMS"*, NATO Specification, November 2021

[2] *MIL-STD-188/141D – "INTEROPERABILITY AND PERFORMANCE STANDARDS FOR MEDIUM AND HIGH FREQUENCY RADIO SYSTEMS"*, DEPARTMENT OF DEFENSE INTERFACE STANDARD, December 2017

[3] Steve Kille, *Optimizing applications and data links for HF radio intermediate term variations: Can you ride the wave?*, The Nordic Shortwave Conference HF 16, ISBN 978-91-639-1642-7, 2016

[4] *P.533 – "Method for the prediction of the performance of HF circuits"*, ITU, August 2019

[5] *HF Location & Information Sharing Protocol – HF-LISP*, published by Isode as S5066-APP10, https://www.isode.com/whitepapers/S5066-APP10.html (April 2022).

[6] Rickard Berg & Håkan Bergzén, *Some frequency management results from over-air testing with HF 2000*, Presentation to HF Industries Association, San Diego, February 2014.