

Appendix I

Extended Squitter Enhanced Reception Techniques

This page intentionally left blank.

I Extended Squitter Enhanced Reception Techniques

I.1 Purpose and Scope

The purpose of this Appendix is to provide a description of improved squitter reception techniques. Elements of improved squitter reception include (1) the use of amplitude to improve bit and confidence declaration accuracy, (2) more capable error detection/correction algorithms, (3) more selective preamble detection approaches, and (4) combinations of the above.

The improved techniques presented in this Appendix represent one way of achieving the performance requirements specified in §2.2.4.4 for enhanced squitter reception. The squitter processing configuration used as the basis for these performance requirements is specified in Section I.5.

The reception techniques, as required in §2.2.4.3.4 of these MOPS, are also described in this appendix for comparison with the enhancements.

I.2 Background

Squitter reception includes the detection of the Mode S 1090 MHz waveform preamble, declaration of the bit and confidence values, error detection, and (if necessary) error correction. The current techniques for squitter reception are based upon techniques developed for use in Mode S narrow-beam interrogators and for TCAS. In both of these applications, the rate of Mode A/C fruit (see glossary) that is stronger than the Mode S waveform is relatively low, nominally less than 4,000 fruit per second.

Early applications investigated for Extended Squitter, prior to the development of this document, included long range air-ground surveillance, surface surveillance and support for TCAS. Of these three applications, the only one with the potential for operating in significantly higher fruit environments was the air-ground application. For this application, it is possible to use sectorized antennas (6 to 12 sectors) to limit the amount of fruit detected by any receiver.

Extended Squitter applications have now been defined in DO-242A, including long range (up to 90 NM) air-air surveillance in support of free flight. This type of surveillance is referred to as Cockpit Display of Traffic Information (CDTI). For this application, sectorized antennas are not an option. In high-density environments, it is possible to operate with fruit rates of 40,000 fruit per second and higher.

Even with these fruit rates, Extended Squitter reception using current techniques will provide useful performance. However, operation of Extended Squitter in very high Mode A/C fruit environments has led to the development of improved squitter reception techniques to support long range CDTI in high-density environments.

I.3 Current Squitter Reception Techniques

Receiving systems for Mode S replies and squitters are currently implemented operationally in ground based Mode S interrogator/receiver systems and in TCAS avionics. For TCAS, the current reception techniques are defined in DO-185A. These

current reception techniques are the basis for the reception requirements in these MOPS (section 2.2). The enhanced techniques that are the subject of this Appendix are not required in this version of the MOPS, but are planned to be required in a future version.

I.3.1 Overview of Current Techniques

Extended squitter uses the 112-bit Mode S waveform shown in Figure I-1. Other than bit assignments, this is the same as the long reply waveform currently used operationally for air-ground replies and for air-air coordination messages in TCAS. This waveform encodes data using pulse position modulation (PPM). A chip in the leading half of the bit position represents a binary ONE. A chip in the trailing half bit position represents a binary ZERO. Reception of the Squitter begins with the detection of the four-pulse preamble. At preamble detection a dynamic threshold is set 6 dB below the level of the preamble. Any signal received below this threshold will not be seen by the squitter reception processor. This eliminates the effect of low level Mode A/C and Mode S fruit on the reception process.

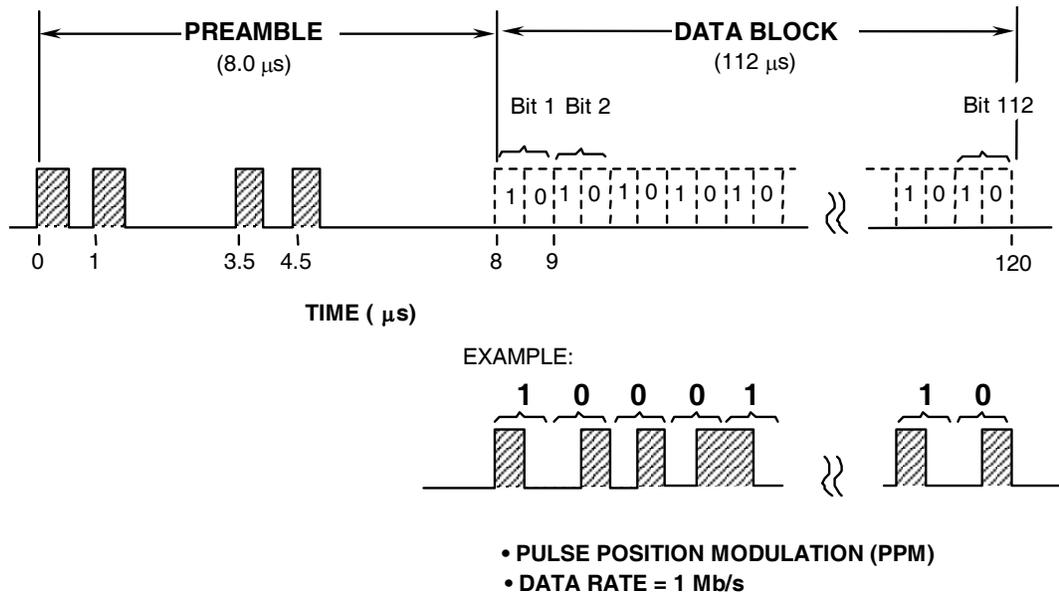


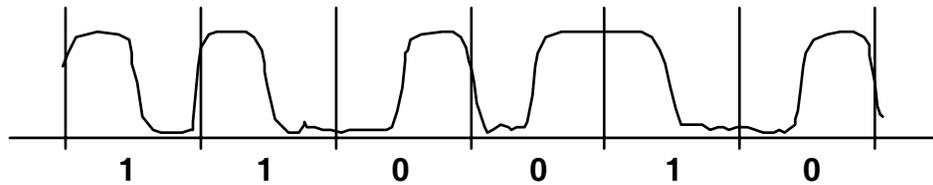
Figure I-1: Mode S Extended Squitter Waveform

Once all of the bits have been received, error detection is performed using the 24-bit CRC contained in the PI field. If no error is detected, the squitter is passed on to surveillance processing. If an error is detected (indicated by a non-zero error syndrome) an error correction technique is applied. The current error correction algorithm can correct the errors caused by one stronger overlapping Mode A/C fruit.

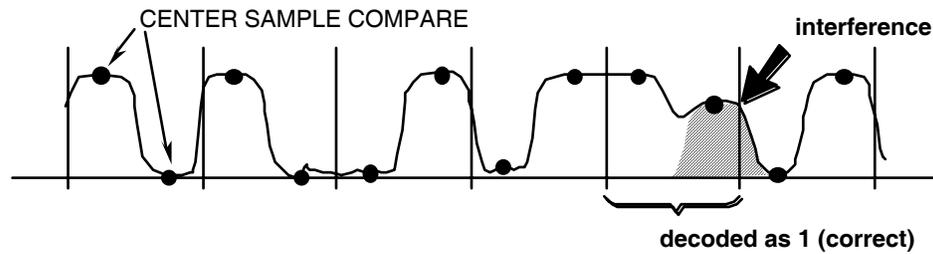
I.3.2 Current Techniques for Bit and Confidence Declaration

Currently, bit values are declared by comparing the amplitudes of the centers of the two chips; the chip with the greater amplitude is declared the bit value. This amplitude comparison technique limits bit errors caused by fruit of lower level than the squitter being received (Figure I-2, parts a and b).

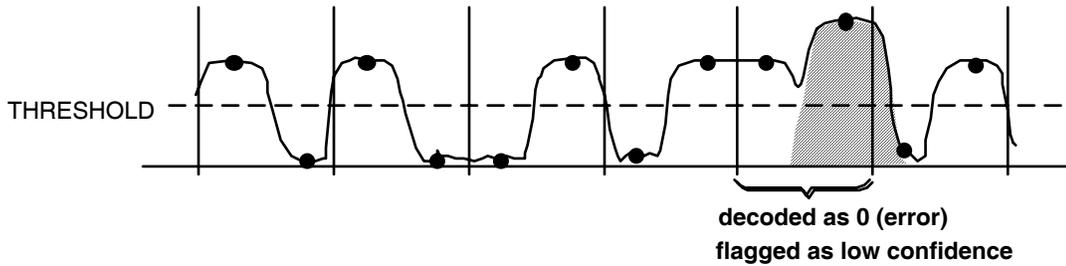
a. Pure signal



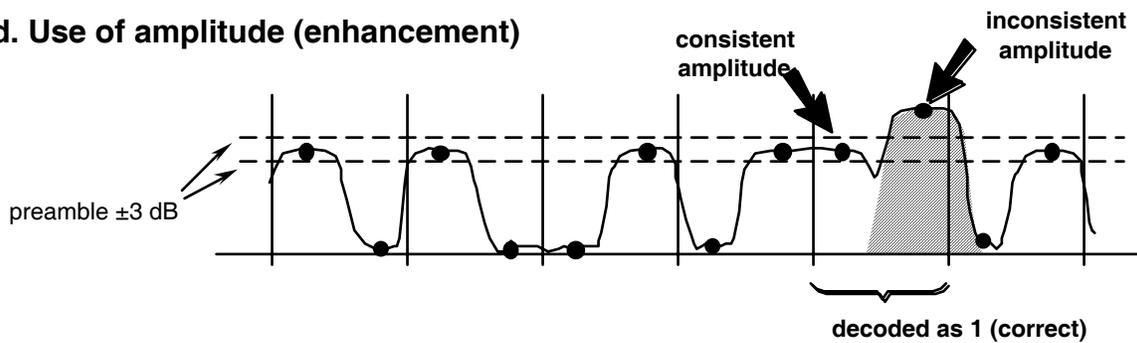
b. Signal plus weaker interference — current techniques



c. Signal plus stronger interference — current techniques



d. Use of amplitude (enhancement)



e. Multiple sample technique (enhancement)

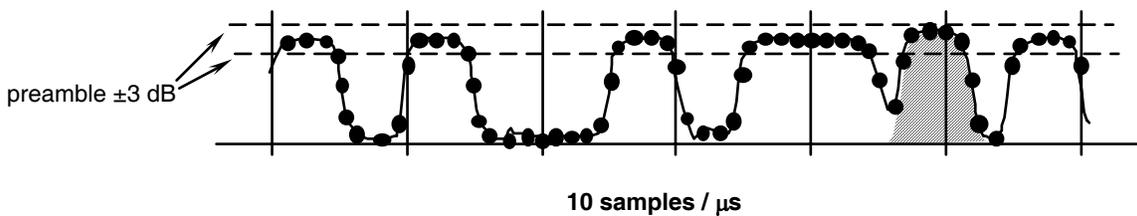


Figure I-2: Current and Enhanced Bit Demodulation Techniques

Note: The vertical scale represents log video signals.

Currently confidence for each bit is declared by observing if signals are above threshold on both chips. If only one chip is above threshold, it is declared to be high confidence since there is no evidence of overlapping fruit. If both chips have amplitude above threshold, low confidence is declared since some form of interference had to be present to cause this condition (Figure I-2, part c).

Since the nominal width of a Mode A/C pulse is 0.45 microsecond, an interfering Mode A/C pulse can affect only 1 of the 2 chips of a Mode S bit position. If the chip affected is the one with the Mode S pulse, the level of the pulse in the chip will be altered, but since it is the only chip with energy, the bit will be declared as high confidence. If instead the interfering pulse affects the blank chip, pulses will exist in both positions and the Mode S bit will be declared as low confidence. Since the bit decision is based on the larger pulse, if the Mode A/C fruit is stronger than the Mode S signal, any such bit will be declared in error, while if the Mode S signal is stronger, any such bit will be declared correctly.

I.3.3 Current Error Detection and Correction Techniques

I.3.3.1 Overview

When a Mode A/C fruit interferes with a Mode S Extended Squitter, some of the Mode S bits may be declared in error. The current Mode S error correction algorithm attempts to correct these errors by locating a string of 24 consecutive bits such that the low confidence bits in the window cover the bits in the corresponding error syndrome [reference I-1]. If such a window is found, all low confidence bits matching the syndrome are flipped, and the message is declared correct.

I.3.3.2 Sliding Window Technique

This algorithm operates by examining successive 24 bit windows, starting with bits 89-112 of the message. In order to achieve a successful error correction, each message bit in the window corresponding to a one (1) in the error syndrome must have its value complemented (i.e., a one is changed to a zero (0) and a zero is changed to a one). This complementing can only be done if each of the bits is declared to be low confidence. If so, each of the bits is complemented and the message is declared to be error corrected. If not, the window is shifted one bit downward, a transformed syndrome is computed and the process repeats. The process ends when a correctable error pattern has been found, or the sliding window reaches the beginning of the message. In order to control undetected errors, correction is not attempted if there are more than 12 low confidence bits in the window.

This technique provides error correction in cases where a Mode S message has been overlaid with one stronger Mode A/C fruit (that caused all bit errors) and one or more weaker Mode A/C fruit that were above the dynamic threshold (and caused only low confidence bits). This technique is well suited to the low levels of Mode A/C fruit observed in a narrow-beam Mode S interrogator or a TCAS. This technique is not appropriate for high fruit rate environments since it produces a high undetected error rate; for this reason, its use is prohibited in §2.2.4.4.

I.4 Enhanced Squitter Reception Techniques

Throughout the following description, a sampling rate of 8 samples per microsecond is used. This is a typical value, although other, higher, sampling rates may be used.

I.4.1 Enhanced Preamble Detection

Preamble detection identifies the beginning of an Extended Squitter reception. The process has two outputs: (1) the start time of the signal and (2) the received power level of this signal. This process includes validation that uses the receptions during the first 5 bits in the data block and several other validation tests. Following is a description of a particular enhanced preamble detection technique that has been used successfully in achieving the performance required in these MOPS.

I.4.1.1 Log Video Samples

The preamble detection process described here operates on data in the form of samples of the log video received waveform. Specifically, the sample rate is 10 samples per microsecond, although other sample rates, including 8 samples per microsecond, have been found to be effective.

I.4.1.2 Threshold

The preamble detection process includes a threshold power level used to discard very weak receptions. Typical value = -88 dBm (referred to the antenna) for an A3 receiver. The Minimum Triggering Level (MTL), which is the point of 90 percent receptions in the absence of interference, is typically about 4 dB higher than the threshold.

I.4.1.3 Valid Pulse Position

A sample that is above threshold and also is followed consecutively by **N** or more samples above threshold is defined to be a "Valid Pulse Position." If the sample rate is 10 per microsecond, then **N** = 3. This definition has the effect of defining a pulse as an event in which at least 4 consecutive samples are above threshold. For other sample rates, **N** is adjusted so that a pulse is declared when the signal is above threshold for more than 0.3 microseconds.

I.4.1.4 Leading Edge

A Leading Edge is declared for a particular sample if it is a Valid Pulse Position and also has substantial slope in the interval before this sample and less than substantial slope in the next interval. Substantial slope is defined by the power change between one sample and the next. The slope threshold is 48 dB per microsecond (applicable to receiver bandwidth of approximately 8 MHz). Therefore if the sample rate is 10 samples per microsecond, the threshold is 4.8 dB.

I.4.1.5 Initial Detection of a 4-Pulse Preamble

The preamble detection process begins when four pulses have been detected, having the spacing of the Mode S preamble. The detection criterion is:

- Finding four pulses having timing of 0 - 1.0 - 3.5 - 4.5 microsecond
- Two or more of these must be Leading Edges.
- The others can be Valid Pulse Positions
- Sample tolerance can be plus or minus 1 (but not both)

Note that the power levels in the four pulses need not agree. Note also that trailing edges are not used.

I.4.1.6 Arrival Time

The signal arrival time is initially estimated to be the leading edge of the first of these four pulses. Subsequently this is adjusted by +1 or -1 sample if two or more of the other three pulses have leading edges with that timing.

I.4.1.7 Reference Level Generation

A Reference Power Level is generated during preamble detection for use in re-triggering and during demodulation of the data block. Step 1 is to identify a set of samples to use. Among the four preamble pulses, those whose leading edges agree with the preamble timing are used; samples from the other pulses are not used. Next, for each pulse used, select the **M** samples after the leading edge sample. If the sample rate is 10 per microsecond, then **M** = 3. For other sample rates, the value of **M** is equal to the value of **N** defined in I.4.1.3.

Step 2 is an algorithm to generate the Reference Level from these samples. For each sample, compute the number of other samples that are within 2 dB. Then find the maximum of these counts. If the maximum count is unique, then the sample used to form that count is taken to be the Reference Level.

Otherwise, when there are two or more samples whose counts are maximum and equal, discard any samples whose counts are less than this maximum. For the remaining samples, find the minimum power and then discard any samples that are more than 2 dB stronger than that minimum. Compute the average of the remaining samples. This is taken to be the Reference Level for the preamble.

I.4.1.8 Overlapping Signals and Re-Triggering

The preamble detection process is capable of processing multiple overlapped preambles, but the data block processing, which is more extensive, can only accept one signal at a time. For this reason, a re-triggering function is included. This function will reject certain preamble detections when a subsequent stronger signal is received.

One step in the re-triggering process checks for overlap by later Mode S signals having certain specific timing offsets. For example, if the subsequent signal is 1 microsecond later, then two of the preamble pulses in the later signal coincide with preamble pulses in the original signal, which can cause a problem if the later signal is stronger. The

potential problem is that the stronger pulses can cause the power estimate for the first signal to be too high, and therefore prevent re-triggering. This type of problem can also occur if the timing difference is 3.5 or 4.5 microseconds. Figure I.4.1.8 illustrates the 1 microsecond overlap that motivates this test.

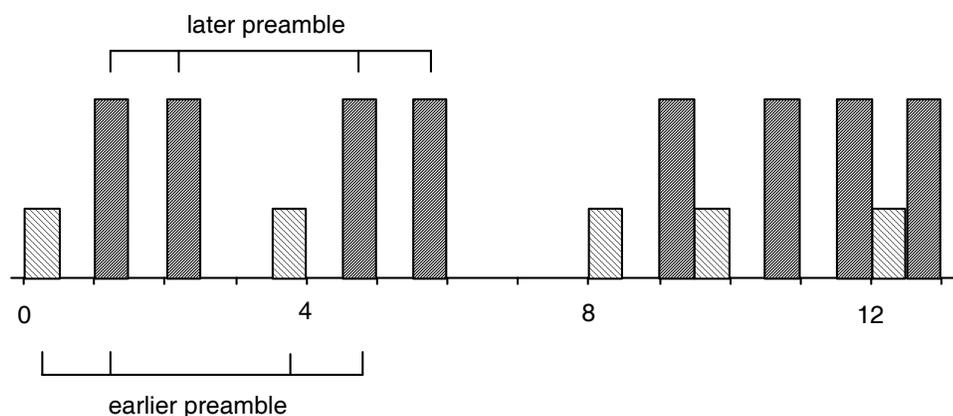


Figure I.4.1.8: Overlap of a Weak Signal by a Later and Stronger Signal

1-Microsecond Test. To counter this potential problem, the next step after preamble detection is to check for excessive power in pulse positions 1.0, 2.0, 4.5, and 5.5 microseconds after the start. For each of these pulse positions, one sample is used to estimate the power of a pulse at that time. Letting $T = 0$ denote the time one sample after the leading edge time of the first preamble pulse, then the four pulses are taken to be the samples at times $T = 1.0, 2.0, 4.5,$ and 5.5 microseconds. From these four power measurements, the minimum is used to compare against the maximum of the samples at $T = 0$ and 3.5 microseconds. If this difference indicates that the preamble under consideration is -3 dB or weaker relative to the other four samples, then this preamble is rejected.

3.5-microsecond test. A similar test is performed to protect against overlap by a stronger signal 3.5 microseconds later. The minimum power in samples at $T = 3.5, 4.5, 7.0,$ and 8.0 is compared against the maximum of the samples at $T = 0$ and 1.0 microseconds. If this difference indicates that the preamble under consideration is -3 dB or weaker relative to the other four samples, then this preamble is rejected.

4.5-microsecond test. A similar test is performed to protect against overlap by a stronger signal 4.5 microseconds later. The minimum power in samples at $T = 4.5, 5.5, 8.0,$ and 9.0 microseconds is compared against the maximum of the samples at $T = 0, 1.0,$ and 3.5 microseconds. If this difference indicates that the preamble under consideration is -3 dB or weaker relative to the other four samples, then this preamble is rejected.

I.4.1.9 Consistent Power Test

Another test is applied to validate the preamble. This test asks whether at least two of the four preamble pulses agree in power level with the Reference Level to within ± 3 dB. If not, this preamble is rejected.

I.4.1.10 DF Validation

An additional validation is made using the first five bits in the data block. Each of the five bits is considered to consist of two chips, each of duration 0.5 microseconds. Pulse detection is carried out for the 10 chips as follows. For a particular chip, a pulse is detected if a Valid Pulse Position is found at the leading edge time for this chip or within +/-1 sample of the leading edge time. The preamble is then validated if, for each of these five bits, a pulse is detected either in the first chip or in the second chip or both and the peak amplitude of the pulse is equal to 6 dB below the preamble reference level or greater. Otherwise the preamble is rejected. The peak amplitude is determined by using the highest amplitude sample of the **M** samples that comprise the pulse (where **M** = **N** + 1 for **N** as defined in I.4.1.3).

I.4.1.11 Re-Triggering

After a preamble is detected, the preamble detection process continues to be applied searching for later preambles. All of the steps described above are applied even when detected preambles are overlapped. If a particular preamble detection has survived these tests, its reference level and the amplitude of each of the five data pulses during DF validation is now compared against any earlier signal currently being processed. If the new signal is stronger by 3 dB, then the earlier signal is rejected, so that data block demodulation of the new signal can proceed. Otherwise (if an earlier signal is being processed and the new signal is not stronger by 3 dB), the new signal is rejected, so that the earlier signal processing can continue.

I.4.1.12 Preamble Detection Summary

Figure I.4.1.12 summarizes the preamble detection process.

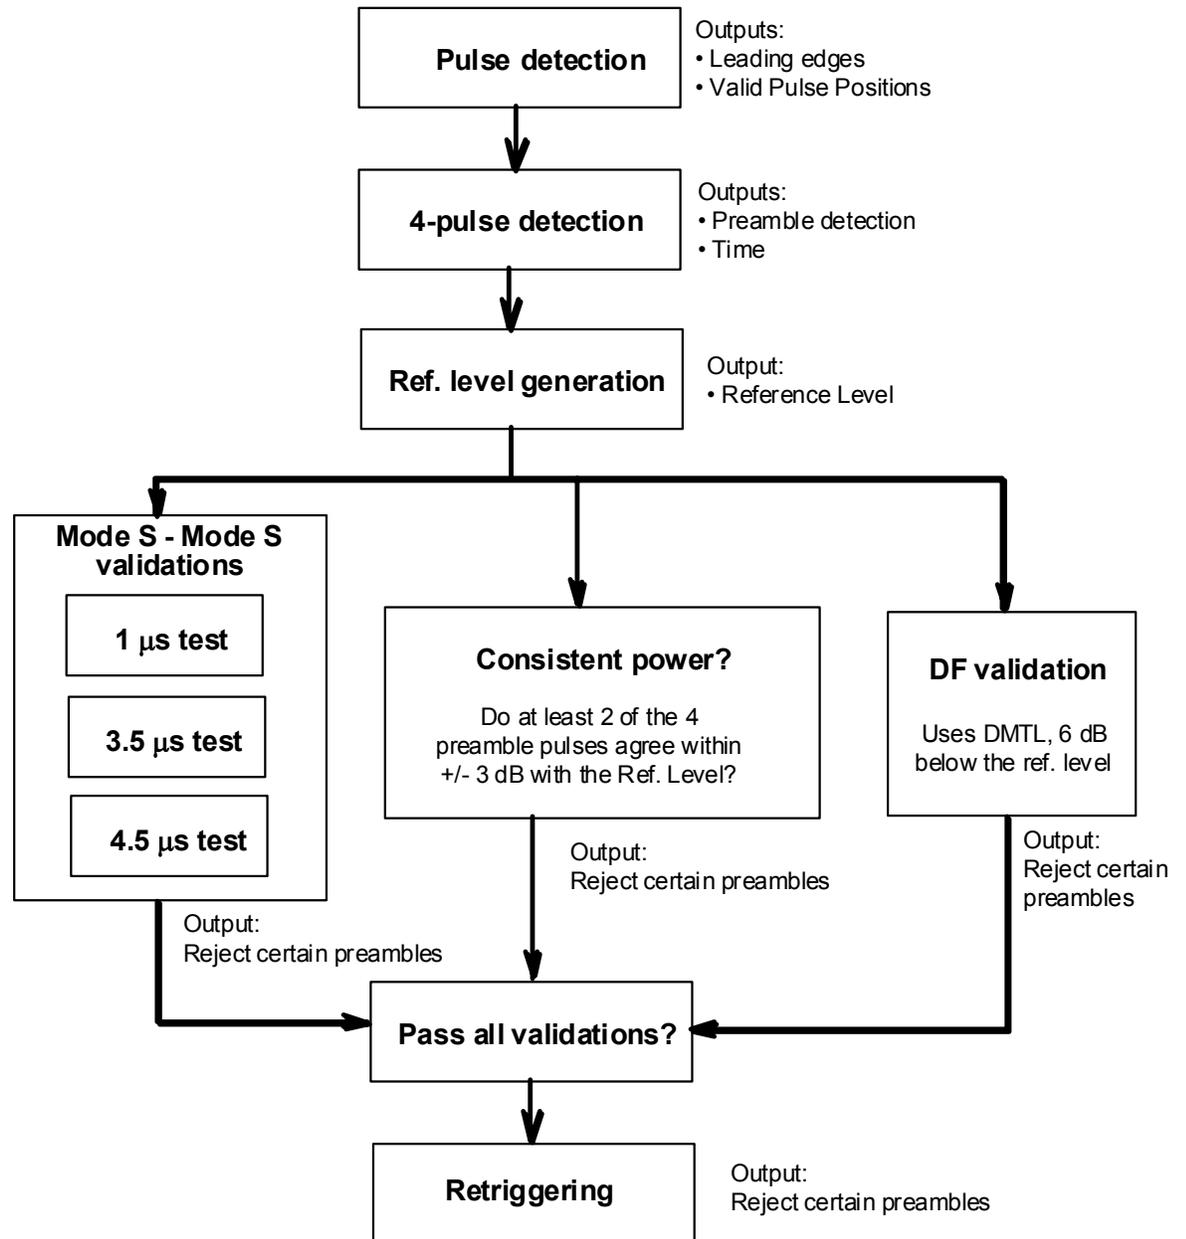


Figure I.4.1.12: Overview of Extended Squitter Preamble Detection

I.4.2 Enhanced Bit and Confidence Declaration

I.4.2.1 Overview

The current technique of declaring a bit based upon the higher of the two chips will generate bit errors in cases of higher level overlapping Mode A/C fruit (Figure I-2, part c). The use of amplitude to correlate the received pulse with the preamble pulse level will improve bit declaration accuracy. Four techniques have been investigated. One is a very simple approach that uses only the amplitude measured at the center of each chip. The remaining three use a more capable approach that takes advantage of all samples per chip that are taken to establish bit and confidence. Each of these techniques is described in the following paragraphs.

The following description of the center sample technique is intended to provide an example of processing performance needed to meet the requirements of Class A1 equipment. This description also serves as an introduction to the three more capable approaches. The center sample technique will not provide sufficient performance to meet requirements specified for Class A2 and A3 equipment in the test procedures of §2.4.4.4. The specified performance for these latter classes of equipment can only be met by an approach that performs equivalently to one of the multi-sample techniques.

I.4.2.2 Use of Center Amplitude

An improvement in the declaration of Mode S data bits can be achieved if the actual amplitudes of the center samples of the '1' and '0' chips can be measured for each data position, rather than just a comparison of which chip sample is greater. All Mode S pulses, including those of the preamble, have approximately the same level (within 1 or 2 dB). Thus if the preamble level is measured, the expected level of each data pulse will be known. Then if both center samples of a data position are above threshold, but only one is within a ± 3 dB band centered at the preamble level, it would be reasonable to assume that the corresponding chip is the correct Mode S pulse location. This is illustrated in Figure I-2, part d.

Figure I-3 illustrates the new data and confidence declaration algorithms that result when the actual sample amplitudes can be measured, and both samples are above threshold. As shown in Figure I-3, a bit is high confidence when 1 and only 1 of the two samples correlate with the preamble level. The correlating sample, rather than the larger sample, is declared to be the true data value. If both samples, or neither sample, correlates with the preamble, a low confidence bit is declared. In this case, the larger sample is selected as the bit value, as in the current technique.

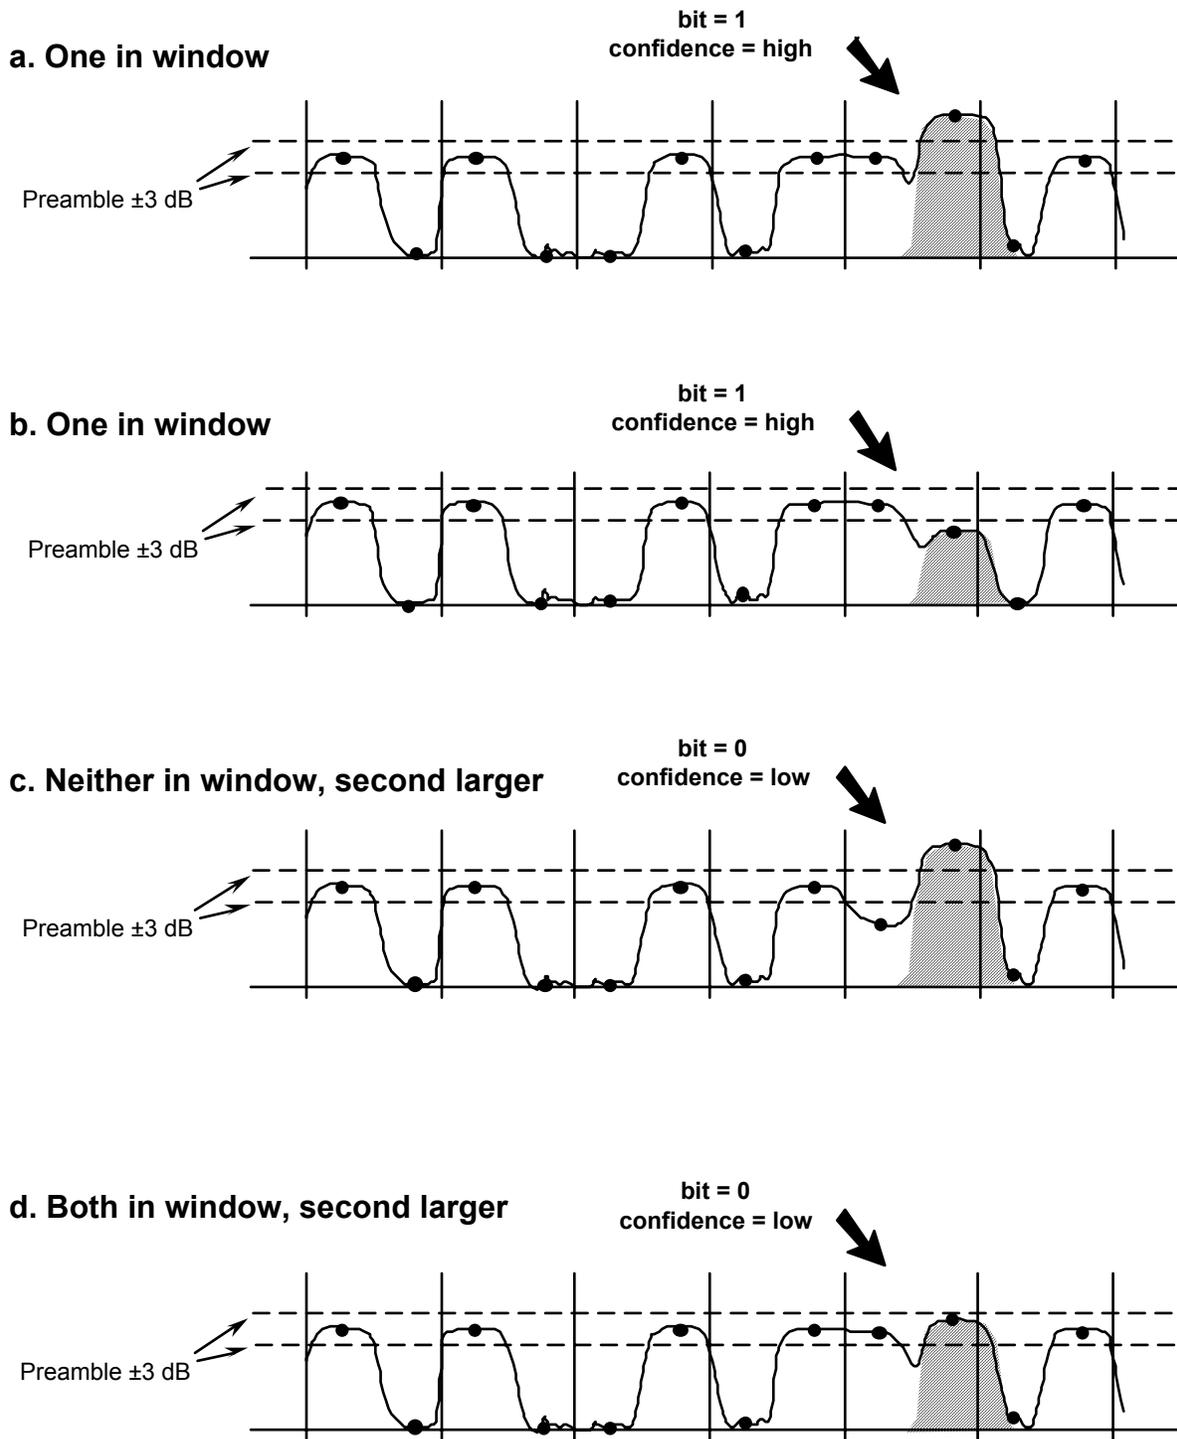


Figure I-3: Center Amplitude Bit and Confidence Declaration

The major advantage of this scheme is the significant reduction in low confidence and bit errors generated for the Mode S message. Presently, any data position with both samples above threshold is declared low confidence. This translates to a low confidence

declaration whenever an above-threshold Mode A/C fruit overlaps a Mode S empty chip position. If the Mode A/C signal is of greater power, the bit position is declared in error.

With the center amplitude technique, a low confidence bit is declared only when the Mode A/C fruit is within 3 dB of the Mode S signal. Mode A/C fruit of either lower or higher power in general cause neither errors nor low confidence bits. Thus the region of concern for Mode A/C aircraft is reduced from all the aircraft at closer range to just those at approximately co-range. This is particularly significant when the system is attempting to listen to far-range Mode S aircraft.

It is occasionally possible for this new algorithm to produce high confidence bit errors. The most likely case occurs when two Mode A/C pulses overlay the same Mode S bit position. If one overlaps the data pulse, and drives its sample out of the preamble window, then an error will result if the other Mode A/C pulse lands on the other chip and its amplitude is within the preamble window. Such occurrences will in general create a rejected message, rather than an undetected error.

I.4.2.3 Use of Multiple Amplitude Samples

The above amplitude data declaration approach can be improved if all 10 samples (5 per chip) that are taken for each Mode S bit position are utilized in the decision process. In particular, the event of both center samples within the preamble window can often be resolved. For example, consider the typical such situation depicted in Figure I-2, part e. Although both center samples are within the window, the fact that the earlier chip has all samples within the window, whereas the later chip has several samples below the window suggests that the signal is present in the earlier chip. This error situation can often be rectified when all 10 samples are examined. Also, in some cases when an interference pulse overlaps a signal pulse, the small frequency difference will produce variations in amplitude, whereas the signal by itself would be more constant, and such patterns can be useful in declaring the bit and confidence.

I.4.2.3.1 Baseline Multi-Sample Technique

The multi-sample enhanced bit and confidence declaration technique makes use of all 10 samples for each Mode S bit position to determine the bit and confidence values. Sample amplitudes in each chip are compared to the amplitude reference level established by the preamble to quantify the number of samples in each chip that:

- (a) match the preamble amplitude indicating the presence of a pulse, or
- (b) are significantly lower in amplitude indicating a lack of transmitted energy.

The first step is to establish an amplitude window that will include samples that are within +/- 3 dB of the preamble reference level and a minimum amplitude threshold set to 6 dB below the reference level. Samples that fall within the window are considered to match the preamble and samples that are below the minimum threshold are considered to indicate a lack of transmitted energy. The samples are categorized as follows:

A: within the +/- 3 dB preamble window

B: below threshold (6 dB or more below the preamble)

The second step is to count the number of samples in each chip that are of each category. Less weight is given to the samples near the transitional areas of each chip (the transitional samples are the first and last samples of each chip). To facilitate this, samples other than those at each end will count double. Therefore, with weighting factored in, and a 10 MHz sampling rate, the counts for each category will range from 0 to 8 for each chip (1 sample at each end + 3 samples in-between x 2). The four counts are summarized as follows:

1ChipTypeA = #of weighted samples in the 1 chip of type A (Match Preamble)

1ChipTypeB = #of weighted samples in the 1 chip of type B (Lack energy)

0ChipTypeA = #of weighted samples in the 0 chip of type A (Match Preamble)

0ChipTypeB = #of weighted samples in the 0 chip of type B (Lack Energy)

Next, two equations using the above counts will produce two scores that indicate how well the sample pattern matches a transmitted 0 and how well the sample pattern matches a transmitted 1. The equations are as follows:

$$1\text{Score} = 1\text{ChipTypeA} - 0\text{ChipTypeA} + 0\text{ChipTypeB} - 1\text{ChipTypeB}$$

$$0\text{Score} = 0\text{ChipTypeA} - 1\text{ChipTypeA} + 1\text{ChipTypeB} - 0\text{ChipTypeB}$$

The highest score determines the bit value. In the case of a tie, the bit defaults to zero (0). If the difference is 3 or more the bit is high confidence. The confidence threshold of 3 was determined by testing the algorithm with a 10 MHz sampling rate with thousands of iterations with a high fruit rate. If the algorithm is applied with a different sampling rate, the appropriate confidence threshold may need to be determined under similar test conditions.

I.4.2.3.2 Multi-Sample Technique With Reduced Table Lookup

To take advantage of the differences between a pure signal and a combination of signal plus interference, a technique was developed based on a lookup table, whose contents are derived from many runs of a simulation. Specifically, each of the 10 samples is quantized into four levels:

- 0: below threshold (-6 dB relative to the preamble)
- 1: above threshold but below the +/- 3 dB preamble window
- 2: within the +/- 3 dB preamble window
- 3: above the +/- 3 dB preamble window

Since there are 10 samples, with 4 possible values each, a Mode S data position can have $4^{10} = 1048576$ (1M) different sample patterns. Two 1-bit tables, each stored in a 1M x 1 ROM, are defined over the set of patterns: the first declaring the bit position to be a '1' or

'0', the second high or low confidence. Once the pattern existing for a given bit is determined, two table lookups supply the proper declaration. If higher sampling rates are used, the number of sample patterns and the table size will increase exponentially.

These tables are generated by running millions of simulations of Mode S messages in 40,000 fruit per second environments. For each bit of each trial, the pattern and the correct Mode S bit value are noted. The result could be, for example, 5876 examples of pattern 16453, of which 5477 occurred when the Mode S bit was a "1." The table values are then defined as follows, assuming an "uncertainty parameter" value of 10%:

- H1: 90% or more of the samples occurred when the bit was a '1'
- L1: 50% - 90% of the samples occurred when the bit was a '1'
- L0: 10% - 50% of the samples occurred when the bit was a '1'
- H0: 10% or fewer of the samples occurred when the bit was a '1'.

Since pulse shapes are critical to this method, live data verification of the table entries is required to establish these values.

I.4.2.3.3 Multi-Sample Technique Without Table Lookup

The above 10 sample approach requires lookup tables of size 1M, adding cost to the hardware implementation of the decoder. A variation of this method that only requires tables of size 1K, called the 5-5 approach, has been designed to reduce this expense.

The 5-5 method forms two estimates of the bit data and confidence values, one using the odd samples (1-3-5-7-9) and the other using the even samples (2-4-6-8-10); the final decision is then a combination of the individual estimates. Since each set includes samples in both chip positions, pattern matching is still possible, although the fineness of the pattern variation is cut in half. Since only 5 samples are in each set, and each sample is quantized to the same 4 levels as above, $4^5 = 1024$ patterns are possible for each set.

To counteract the loss of resolution, and to aid in the combining operation, 3 levels of confidence (high, medium, and low) are defined for each pattern. Following the simulation generation scheme described above, the table values are defined as follows:

- H1: 90% or more of the samples occurred when the bit was a '1'
- M1: 70% - 90% of the samples occurred when the bit was a '1'
- L1: 50% - 70% of the samples occurred when the bit was a '1'
- L0: 30% - 50% of the samples occurred when the bit was a '1'
- M0: 10% - 30% of the samples occurred when the bit was a '1'
- H0: 10% or fewer of the samples occurred when the bit was a '1'.

The values of 10% and 30% are parameters; the value of 30% was selected to provide the performance discussed below. Since 3 confidence levels now exist, the confidence lookup table for each sample set is sized 1024x2. The total table requirement for data and confidence, for the two sample sets, is thus 1024x6.

Once the values and confidences are determined for each set of samples, odd and even, the composite values actually declared for the bit are found according to Table I-1.

Table I-1: Combining Odd and Even Outputs

Odd	Even					
	H1	M1	L1	H0	M0	L0
H1	H1	H1	H1	L0	H1	H1
M1	H1	H1	L1	H0	L0	L1
L1	H1	L1	L1	H0	L0	L0
H0	L0	H0	H0	H0	H0	H0
M0	H1	L0	L0	H0	H0	L0
L0	H1	L1	L0	H0	L0	L0

Note that if either sample set is high confidence, that set's data value rules unless the samples conflict and are both high confidence. Also notice that two agreeing medium confidence samples produce a high confidence result. With the parameter for medium confidence set at 30%, the probability of error for medium confidence agreement is $0.3 \times 0.3 = 0.1$, which matches the 10 sample probability of error for a high confidence decision. Finally, when the sample decisions conflict at the same confidence level, a low confidence 0 is declared for lack of a better estimate.

If a sampling rate other than 10 MHz is used, the number of samples from the odd and even samples and the resulting table sizes will be adjusted accordingly. For example, if a 8 MHz sampling rate is used, there are 4 samples in each set, and each sample is quantized to the same 4 levels as above, $4^4 = 256$ patterns are possible for each set. The total table requirement for data and confidence, for the two sample sets, is thus 256x6. If a 16 MHz sampling rate is used, there are 8 samples in each set, and each sample is quantized to the same 4 levels as above, $4^8 = 65,536$ (64K) patterns are possible for each set. The total table requirement for data and confidence, for the two sample sets, is thus 64Kx6. Once the values and confidences are determined for each set of samples, odd and even, the composite values are declared according to Table I-1 independent of the sampling rate.

I.4.3 Enhanced Error Detection and Correction Techniques

I.4.3.1 Overview

Three new error detection/correction techniques have been developed. The first, termed the "Conservative" technique is a variation on the current Sliding Window technique, intended to reduce the undetected error rate. The second, termed the "Whole Message" technique, models the effect of whole Mode A/C fruit on the bit and confidence declarations. The third, termed the "Brute Force" technique, performs a bounded exhaustive search of all combinations of bit reversals for low confidence bits. The following sections describe these techniques in more detail.

I.4.3.2 Conservative Technique

The sliding window technique is suitable for the low fruit environments of a rotating beam antenna (long range, narrow beam) or a TCAS (omni-directional, short range). However, the fruit environment for the long-range air-to-air application (which uses omni-directional antennas) may be very severe and therefore, the sliding window technique cannot be used because of undetected error considerations.

For the very severe fruit environments, a simpler approach, known as the conservative technique is used. Using this technique, error correction is only attempted if all of the low confidence bits in the message are within a 24-bit window, and there are no more than 12 low confidence bits. This constraint limits the application of error correction to signals that nominally had only a single overlapping stronger Mode A/C fruit. This is a conservative approach in that the conditions for attempting error correction are much more restrictive than with sliding window. It produces a lower level of successful error correction since it does not attempt to correct messages with multiple Mode A/C overlaps. However, it produces a very low undetected error rate, as intended.

If the conditions for applying the conservative technique are met, the error syndrome is generated for the window position and (as for the sliding window technique) a check is made to see if the ones (1) in the error syndrome correspond to low confidence bits in the window. If so, error correction is accomplished. If not, the process is terminated.

If the low confidence bits span less than 24-bits, more than one window could be defined to span them. This will not effect the error correction action, since regardless of the 24-bit window selected to span the low confidence bits, the ones (1) in the error syndrome will identify the same message bits. That is, if the window is moved one bit, the error syndrome will shift by one bit.

Note that in the above description, there is only one possible successful error correction possibility. Regardless of the specific 24-bit window position, the same message bits are identified. All of the bits corresponding to a one (1) in the error syndrome must be complemented. This can only happen if they are all low confidence. Therefore, there is at most, one correctable error pattern that can be achieved with the conservative error correction technique.

I.4.3.3 Whole Message Error Detection and Correction Technique

The limitation of the current error correction technique described in I.3.3 of being able to handle only a single overlapping Mode A/C fruit has led to the development of the Whole Message error correction technique, which can be applied when the Conservative technique has failed. This technique is designed to handle up to five overlapping fruit, provided they do not overlap each other, hence creating "individual interference regions." This technique is applicable only when the center sampling technique is used (§I.4.2.2).

As described in I.3.2, the correctness of low confidence bits depend upon the amplitude of the Mode A/C interference relative to the Mode S squitter: if the Mode A/C fruit is stronger than the Mode S signal, any low confidence bit will usually be declared in error, while if the Mode S signal is stronger, any low confidence bit will usually be declared correctly. This observation leads to the following hypothesis, which serves as the basis of the Whole Message error correction routine:

For a given Mode A/C interference region, either all low confidence bits will be correct, or all low confidence bits will be wrong.

Although this hypothesis is generally correct, it should be noted that three effects can lead to its violation:

1. If the interfering Mode A/C fruit is at approximately the same power level as the Mode S signal, or
2. If the interfering Mode A/C fruit pulses are wider than 0.50 microseconds (the spec allows up to 0.55 microseconds, and some out-of-spec transponders even go beyond this value), or
3. If the Mode S signal is near enough to the noise level for noise to corrupt some of the samples.

The Whole Message algorithm attempts to divide the Mode S message into disjoint 24-bit regions, one for each presumed Mode A/C interfering fruit. Then the syndrome of each region, assuming all bits to be in error, is generated (see below in §1.4.3.4 for an explanation of syndrome generation for bit sets). Finally, all possible combinations of syndromes are considered. If one and only one combination matches the Mode S syndrome, all low confidence bits within the interference regions corresponding to the winning combination are reversed. Flowcharts for the Whole Message error correction technique are presented in Figure I-4a and Figure I-4b. Figure I-5 presents an example of the application of this method.

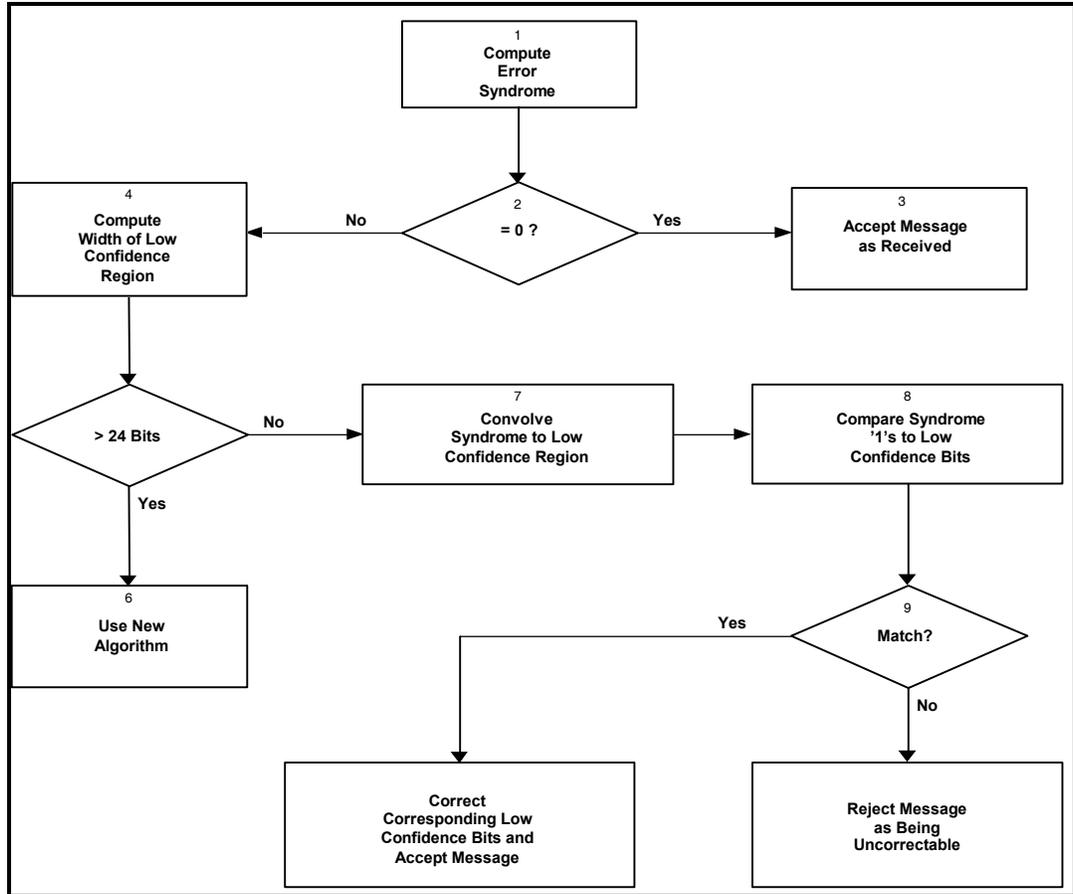


Figure I-4a: Whole Message Top Level Error Correction Algorithm

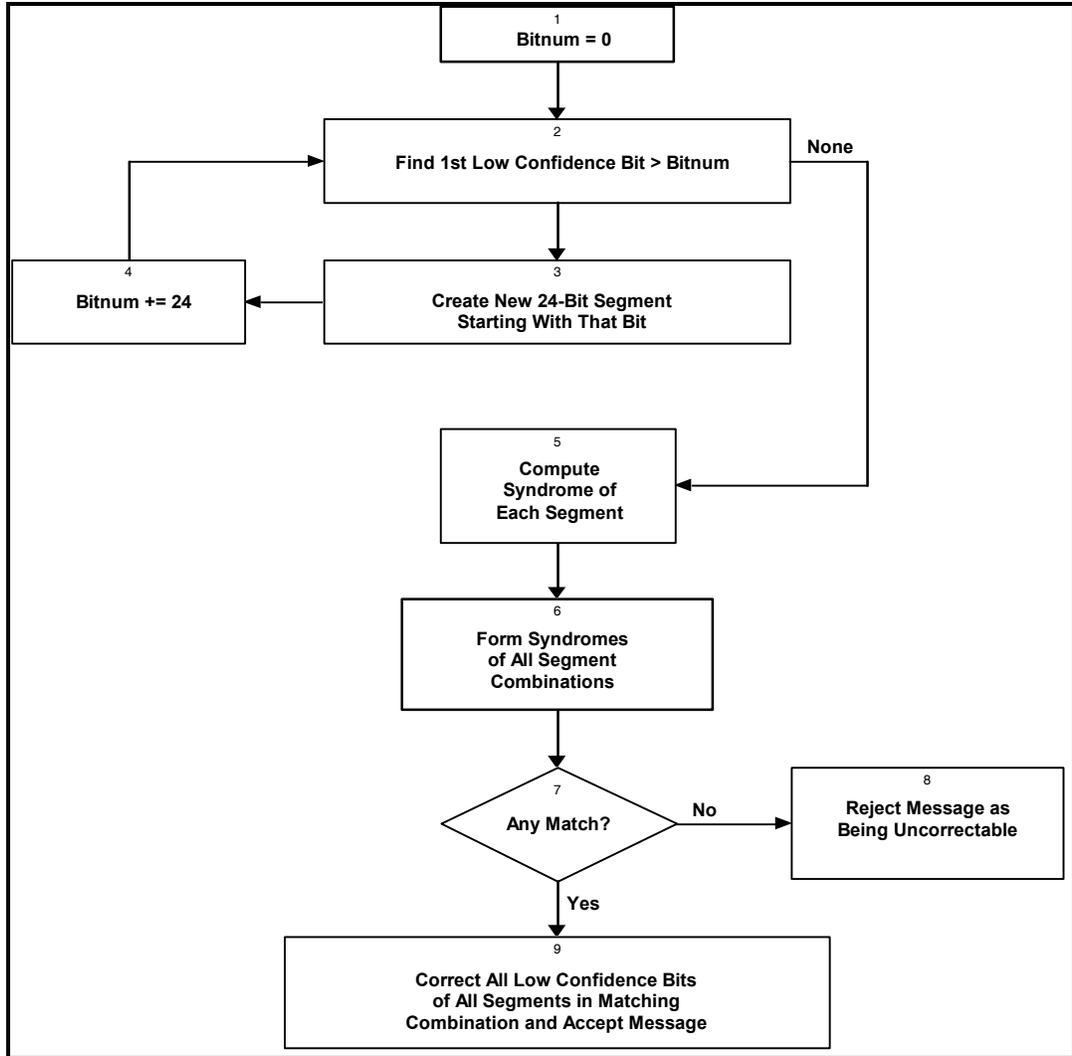


Figure I-4b: Whole Message Detailed Error Correction Algorithm

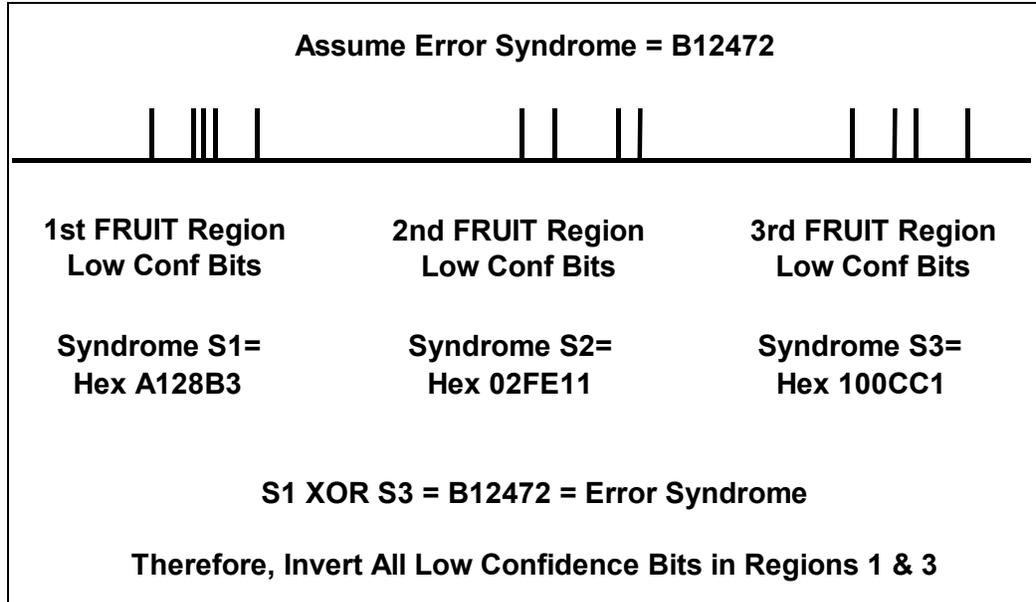


Figure I-5: Whole Message Error Correction Example

I.4.3.4 Brute Force Error Correction Technique

If the bit declaration algorithm has performed its function properly, all errors in Mode S data values will reside in bits declared low confidence. If this is true, a simple approach to error correction is to try all possible combination of low confidence bits, and accept the set that matches the error syndrome (provided only one success is discovered). For obvious reasons, this method has been named the Brute Force Technique. It is applicable to any method of data and confidence declaration, with or without amplitude. As illustrated in Figure I-6, Brute Force error correction is applied after the other techniques have failed.

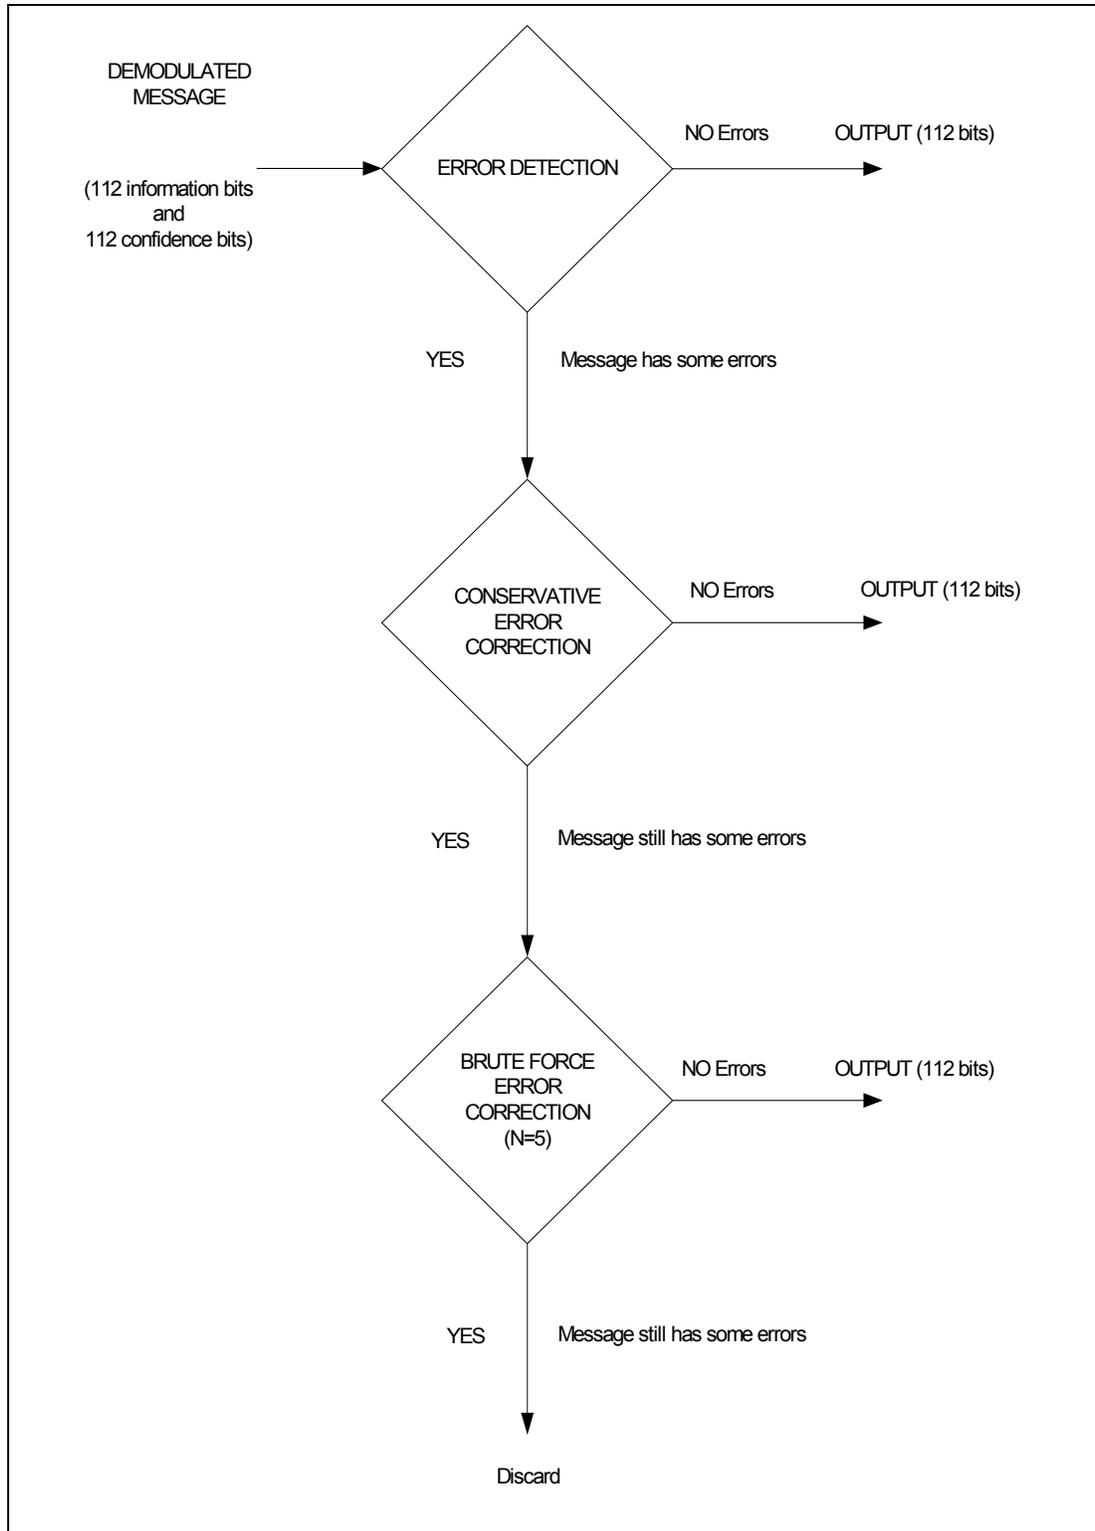


Figure I-6: Error Detection and Correction

Implementation of this technique depends upon the fact that each Mode S bit position corresponds to a unique syndrome, and that sets of bits produce a syndrome that is the exclusive OR of all individual bit syndromes. For example, if bit 1 is the only bit declared in error, the error syndrome at the receiver will be hex 3935EA, while bit 31 produces hex FDB444, and bit 111 has syndrome hex 000002. Thus if those three bits are all declared in error, the error syndrome will be calculated to be hex C481AC. The table of individual bit syndromes is pre-calculated and stored in the receiver.

It is possible for two or more subsets of the low confidence bits to match the syndrome. In such cases, the message is rejected, and no harm is done. However, if a high confidence bit has been declared in error, and a single subset of the low confidence bits matches the syndrome, the message will be "corrected" to the wrong message, producing an undetected error. (If no subset matches the syndrome, it must be true that a high confidence bit error has been made, and the message is rejected.)

Clearly, for processing time and error bounding reasons, the maximum number of low confidence bits to process must be limited. The number of cases to consider is given by 2^n . If n low confidence bits exist for a message; this grows exponentially with n (32 at $n = 5$, 4096 at $n = 12$). The undetected error rate is proportional to the number of cases, and thus also grows exponentially with n . Fortunately, the Hamming distance of 6 for the Mode S parity code implies that undetected errors are essentially zero if $n \leq 5$ is enforced. For this reason, a value of $n = 5$ has been used in the development of the brute force algorithm.

I.5 Improved Reception Performance in a High Fruit Environment

For improved reception performance in a high fruit environment, the optimum configuration has been found to be:

1. Enhanced preamble detection (§I.4.1)
2. Bit and confidence declaration based on the Baseline Multi-sample Technique (§I.4.2.3.1)
3. Error detection using the Mode S 24-bit CRC technique (Ref. RTCA DO-185A)
4. First pass error correction using the conservative technique (§I.4.3.2)
5. Second pass error correction using the brute force technique with $n=5$ (§I.4.3.4)

The process proceeds as follows. Preamble detection and bit and confidence declaration are performed. Next, Mode S error detection is applied. If the message passes, the process ends, and the message is delivered. If an error is detected, then conservative error correction is applied, and if a correction results, then the process ends, and the message is delivered. If the constraint for conservative correction is not satisfied, then the brute force technique is applied.

The above configuration was used as the basis for the performance required for enhanced squitter reception as specified in §2.2.4.4.

I.6 Summary

New techniques have been developed for enhancing reception of Extended Squitters in environments of high interference. The new techniques include improvements in preamble detection, improvements in declaration of information bits and confidence bits within the squitter message, and improvements in error detection/correction.

These developments were originally carried out using pulse-level simulation to assess the resulting performance. Subsequently flight tests have been conducted using these techniques, making comparisons relative to the current techniques. Both the simulation and flight test results indicate that substantial improvements in performance are achievable using these techniques when operating in an environment of high interference.

APPENDIX I REFERENCES

- I-1. J. L. Gertz, "Fundamentals of Mode S Parity Coding," M I T Lincoln Laboratory Project Report ATC-117, April 1984.

This page intentionally left blank.