DADF

User Station Design Justification - Baseband Processing

Document Signature Table

	Name	Function	Signature	Date
Author	Dr. Paul Crawford			20.08.1999
Approval				
Approval				
Release				
Eumetsat Approval				

Document Change Record

Issue/Revision	Date	DCN No.	Changed Pages/Paragraphs
1.0	13.01.98		Initial Release
4.0	22.06.98		1.3.2 p. 11, 1.4.1 p. 18, Table 4 p. 32, VCS Ref. No.
			1.2 eqn 13, Figures 3-5, 1.2.3 text, 1.73 eqn 48-51, 1.8 table 5 ISI values.
	13.01.99		Minor on review of RIDs.
5.0	20.08.1999		No changes.

Table of Contents

Table of Figures

Table of Tables

1 Filter Design and Bit Synchronisation

1.1 Objectives

To design the optimum reception system we use the matched filter theorem and the requirement for identical filters on both the TX and RX side of the communication system. The MSG system is specified as having an overall response which is from the Root Raised Cosine family, with a=0.7 for HRIT and a=1.0 for LRIT (US.390 via the spacecraft ICD 1.0, p62 and p70).

The requirement for low bandwidth and zero Inter Symbol Interference (ISI) can be achieved with an overall system impulse response *s(t)* with zeros at all multiples of the bit period. The most common are based on the:

$$\mathrm{sinc}(x) = \frac{\sin(x)}{x} \qquad 1$$

shape (for minimum bandwidth) but with slight modification to reduce their sensitivity to filter tolerance, timing error, etc. These give rise to an overall power spectrum *S(f)*. For the matched filter system, the transmit filter $H_{TX}(f)$ and the receive filter $H_{RX}(f)$ are related by:

$$H_{TX}(f) = H_{RX}^{*}(f) \qquad 2$$

Where * implies the complex conjugate (time reversal of the impulse response for real valued signals), and:

$$H_{TX}(f)H_{RX}(f) = S(f) \qquad 3$$

Since the two filters have the same amplitude response (equal apportioning) for best performance in an additive Gaussian noise channel, we have:

$$\left|H_{TX}(f)\right| = \left|H_{RX}(f)\right| = \sqrt{S(f)} \qquad 4$$

It is normal to allocate identical linear phase $\sqrt{S(f)}$ filters to both transmit and receive since true linear phase systems are the conjugate of themselves and hence of each other, pure time delay ignored.

The “Raised Cosine” family have a frequency response of the form (Cooper, 1986, p176, eqn 4-143):

$$S(f) = \begin{cases} 1 & |f| < \dfrac{1-a}{2T_B} \\ \dfrac{1}{2}\left[1+\cos\left(\dfrac{\pi\left(2fT_B - 1 + a\right)}{2a}\right)\right] & \dfrac{1+a}{2T_B} \le |f| \le \dfrac{1+a}{2T_B} \\ 0 & |f| > \dfrac{1+a}{2T_B} \end{cases} \qquad 5$$

Where *a* is the roll-off factor and T_B is the bit period. The required “Root Raised Cosine” TX and RX filters, assuming equal apportioning and linear phase are given by:

$$H(f) = \sqrt{S(f)} \qquad 6$$

For the HRIT system (a=0.7), the overall responses *S(f)* and RRC filter *H(f)* look like:

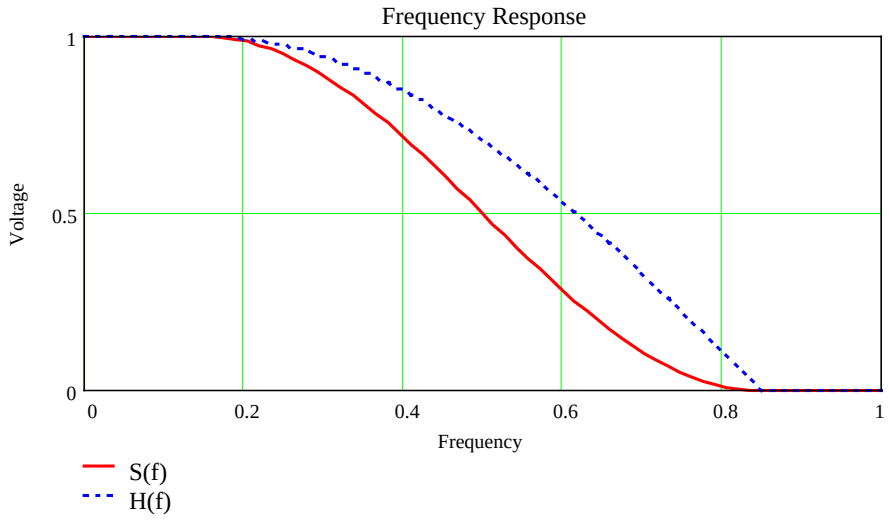

Figure 1: Frequency Response of Raised Cosine Filters

The overall time domain response of the system, *s(t)*, can be computed from the Fourier transform of the system transfer function *S(f)*. This results in (Cooper, 1986, p177, eqn 4-145):

$$s(t) = \frac{\cos\left(\frac{\pi a t}{T_B}\right)}{1-\left(\frac{2 a t}{T_B}\right)^2} \operatorname{sinc}\left(\frac{\pi t}{T_B}\right) \frac{1}{T_B} \qquad 7$$

This is symmetric about *t*=0 (since we are simplifying the equations by assuming non-causal responses) and has the shape:

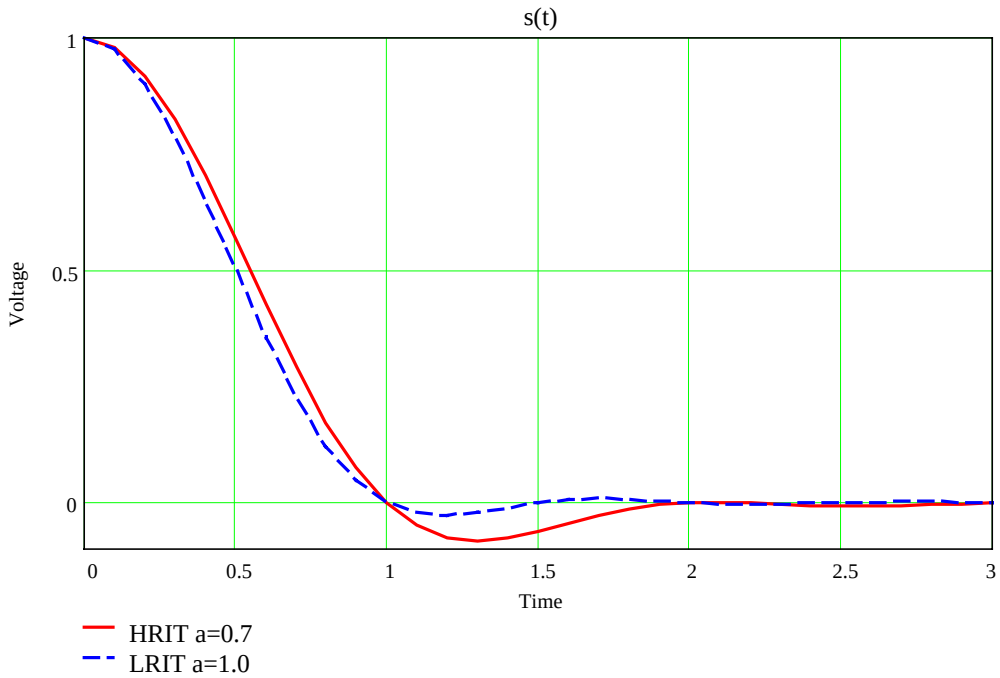

Figure 2: Overall System Impulse Response

Unfortunately, the implementation in a digital asynchronous system is not quite as straight forward. It is desirable to have as low a sample rate as possible in to the DSP processor to minimise the computational burden. From the Nyquist limit we need a sampling frequency of greater than twice the highest frequency present, in the above example, this is given by:

$$f_S > 2\left(\frac{1+a}{2T_B}\right) = \frac{1+a}{T_B} \qquad 8$$

This implies we can accept less then 2 samples per bit if $a < 1$, however, the other requirement of the DSP algorithm is accurate timing recovery for bit synchronisation. If we simply make f_S very high we can chose the nearest sample value to the “true” optimum time and this would be satisfactory. Unfortunately, this results in a very high sample rate if an acceptably low technology loss is to be achieved, hence the normal use of some interpolation method.

1.2 Calculation of Technology Loss Due to Bit Timing Errors

The calculation of the effects on bit timing errors is based on the same basic principal of the PLL design justification. The procedure consists of evaluating the probability distribution function of the timing errors then evaluating the impact on the error rate. The overall loss is the extra Eb/No required to overcome the degradation in effective Eb/No.

There are two effects which increase the error rate in the presence of sample point timing errors:

- The reduction in signal voltage as the sample point moves from the optimum time t=0.
- The effect of Inter Symbol Interference (ISI) due to the non-zero values of adjacent bits whenever the sample point is not exactly one bit away from their peak value.

As for the inter-phasor crosstalk in the QPSK error rate calculations, it is the 2nd effect, the ISI from adjacent bits, that dominates the loss in effective Eb/No. The calculation method for evaluating the effective Eb/No at any given time offset uses the formula for the signal voltage at that time:

$$v(t) = \sum_{i=-\infty}^{\infty} d_i \cdot s(t + i \cdot T_B) \qquad 9$$

Where the d_i represents the data value for the ith bit (equiprobable data of value ±1). In the calculations we assume d_0 is +1 and evaluate the other possibilities for the adjacent data values. The resulting voltage when squared is the Eb for a 1Ω normalised system.

In assessing the total effect, there are 2^{N*2} possible states for the ±N adjacent data values either side of the intended sample point. To avoid the computation of all possible states we could use the central limit theorem to compute an effective ISI voltage from all of the adjacent bits by arguing there are independent random variable of a given standard deviation. This leads to the root sum square ISI voltage as:

$$isi(t) = \sqrt{\sum_{i=-\infty,\, i\neq o}^{\infty} s^2(t + i \cdot T_B)} \qquad 10$$

In practice, the index i would run for, say, ±6 bit periods since the time function is very small beyond this point. From this effective noise voltage, we could then compute the effect of this "interference" as adding to the noise power of the system. For the case of perfect timing, etc, we have the following:

$$\frac{E_B}{N_O} = \frac{s^2(0)}{\bar{n}^2} \qquad 11$$

Also we know that $s(0)$ = 1 so the noise power term n^2 = No/Eb, this leads to the total effective Eb/No for the case of additive noise as:

$$\left.\frac{E_B}{N_O}\right|_{effective} = \frac{s^2(t)}{\frac{N_O}{E_B} + isi^2(t)} \qquad 12$$

The numerator shows the effect of the reducing signal voltage due to timing errors while the denominator shows the added ISI voltage.

As for the PLL calculations for the QPSK system, we can argue that the FEC system will look at a number of coded symbols to decide on any one data bit. If the adjacent symbols were independent of the one considered, and the ISI term was moderate compared to the Gaussian noise, the noise power would simply add (eqn 12). However, the adjacent symbols the adjacent symbols represent a fraction of a bit (the code rate) are correlated with each other since they represent convolution sets of symbols for the data bits in the encoder shift register. As for the QPSK system we can modify eqn 12 to find the error probability by introducing K_{FEC} to account for possible degradation due to this effect:

$$Pef\left(\frac{E_B}{N_O}, t\right) \approx Pe\left(\frac{s^2(t)}{\frac{N_O}{E_B} + K_{FEC} \cdot R \cdot isi^2(t)}\right) \qquad 13$$

Here $Pe()$ represents the underlying function relating error probability with Eb/No. R is the code rate (223/512 for CCSDS) and K_{FEC} is likely to be in the region of 1-2, our calculations assume 2 as an upper bound.

To find the overall error rate for some distribution parameter T_E, we integrate the (error probability at timing offset *t*) * (probability of *t* occurring) to yield:

$$Perror\left(\frac{E_B}{N_O}, T_E\right) = \int_{\frac{-T_B}{2}}^{\frac{T_B}{2}} Pef\left(\frac{E_B}{N_O}, t\right) Pb(t, T_E)\, dt \qquad 14$$

Where *Pb()* is the Probability Distribution Function (PDF) of the timing errors, assumed to be reduced modulo T_B hence the integration range is $\pm 0.5 T_B$.

In the case of the baseband design, there are 3 possible sources of timing error:

- Static timing error. This should be zero in a DSP system with a “perfect” integrator and no significant Doppler shift.
- Uniform timing error. This arises from the time quantisation of the asynchronous system. Since the sample clock is asynchronous to the optimum timing point, it has a period of $T_S = 1/f_S$. The system will cycle by $\pm 0.5 T_S$ as the data is processed, assuming no static phase errors.
- Gaussian timing error. This is the result of estimating the correct timing point from the noisy signal.

All can be analysed by the above approach and is performed by the MathCAD file biterr3.mcd

~~1.1.1~~1.2.1 Static Bit Timing Errors

For the static bit timing errors, the function of eqn 12 is computed for different timing error values *t*. The required input Eb/No is then found by adjusting Eb/No until it yields the same error rate as the ideal system (using a standard root solving routine). The difference in dB(Eb/No) is the technology loss.

This was performed with a=0.7 for the HRIT system, and with a=1.0 for the LRIT system. The computed behaviour is shown in the following graph:

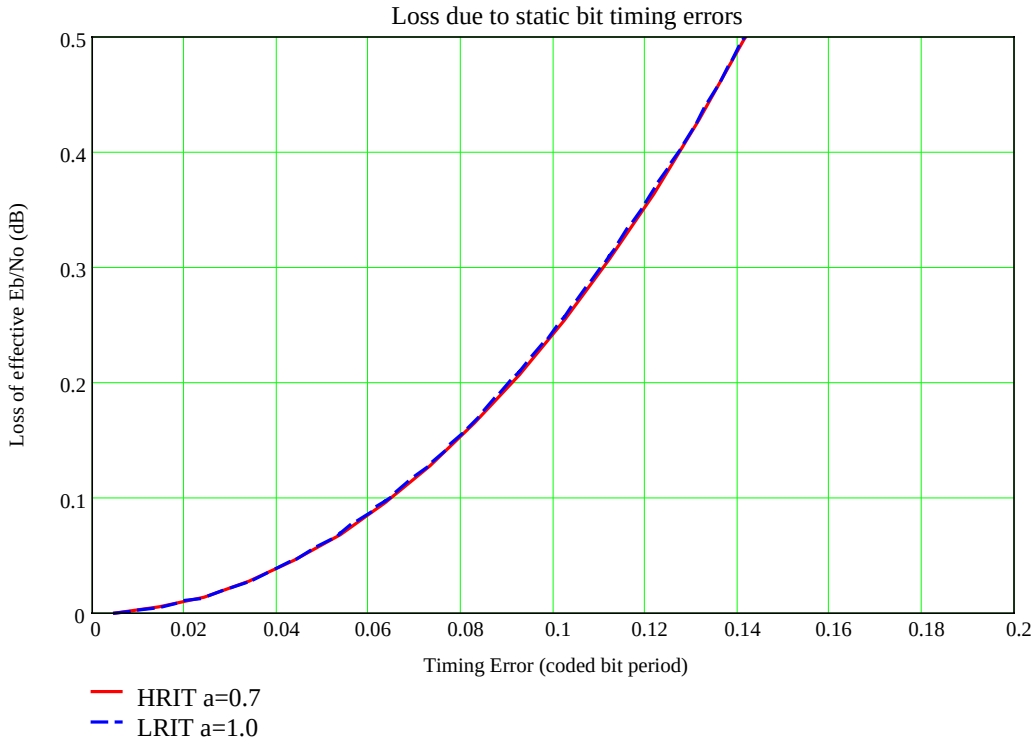

Figure 3: Loss Due to Static Bit Timing Errors

In Figure 3 (as for all of these calculations), the timing error is measured in fractions of a coded bit period. For the QPSK modulator in the HRIT system, the convolution symbols appear in parallel, so the actual time for one symbol is 1/(1E6*256/223) = 871ns. For the BPSK modulator in the LRIT system the convolution symbols appear in series, so the time for one symbol is 1/(128E3*2*256/223) = 3.4µs

In the case of the MSG, there is no significant Doppler shift and both the PSG and RX clock oscillators are expected to be very stable relative to the typical PLL bandwidths. The spacecraft translation has no effect on the bit rate, and the clock oscillators are at the bit rate, not multiplied up to L band, so LF phase noise should be negligible. Hence there is no dynamic reason for significant “static” (or low disturbance frequency) timing errors. Since we are implementing the bit timing recovery with DSP techniques, there should be no static errors. This is unlike analogue electronics where a DC offset might cause a small static error.

1.2.2 Uniform Timing Errors

For the uniform timing error problem, we model the system as having no static (average) error so the bit timing changes uniformly by $\pm T_E/2$ for the calculation. Since this is symmetric about zero, we have the PDF:

$$Pb(t, T_E) = \begin{cases} \frac{1}{T_E} & |t| < \frac{T_E}{2} \\ 0 & |t| \geq \frac{T_E}{2} \end{cases} \qquad 15$$

Taking the symmetry about zero in to account, and simplifying, the overall error rate is given by:

$$Perror\left(\frac{E_B}{N_O}, T_E\right) = \frac{2}{T_E}\int_0^{\frac{T_E}{2}} Pef\left(\frac{E_B}{N_O}, t\right) dt \qquad 16$$

The results of this calculation is shown in the following graph:

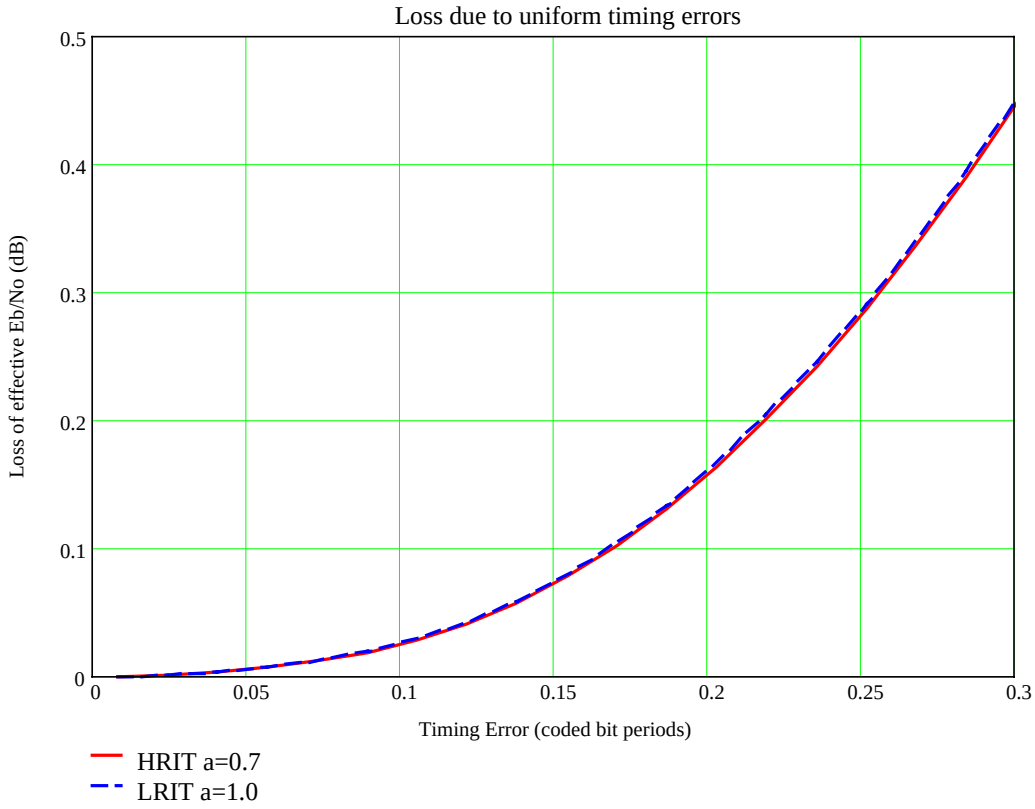

Figure 4: Loss Due to Uniform Bit Timing Errors

This is very important since is determines the resolution of the timing recovery system used. If we aim for a low loss due to the resolution of the interpolation system, for example 0.1dB, we require a uniform timing error of less than 0.16 symbol periods (T_E = 1/6.25). For a non-interpolated system, where the bit sync must select the nearest output of the matched filter, the sample rate would be 6.25 times the symbol rate, a considerable burden on the DSP software and on the number of taps for the FIR filter.

To avoid this, we can use an interpolation system. For a input sample rate to the interpolator of 2 samples/symbol, then the required resolution is about ¼ of the sample rate. To allow for the Gaussian errors present for the additive noise of the channel, a slightly higher specification would be prudent.

1.2.3 Gaussian Timing Errors

In this case we have a similar situation to the carrier recovery PLL. The bit timing errors have the same distribution, but for a range of ±½ symbol, rather than ±π radian. The PDF is given by:

$$Pb(t,\sigma) = \frac{\exp\left(\frac{\cos(2\pi t)}{(2\pi\sigma)^2}\right)}{I_0\left(\frac{1}{(2\pi\sigma)^2}\right)} \qquad 17$$

where σ is the standard deviation (RMS timing error, T_E) and I_0 is the modified Bessel function of the first kind and zero order. For small σ this approach is prone to math overflow and the approximation:

$$Pb(t,\sigma) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(\frac{-t^2}{2\sigma^2}\right) \qquad 18$$

is normally used. An appropriate threshold value for σ is 0.02 (equivalent to a PLL SNR_{Loop} of approximately 15dB). Above this value of σ the I_0 term is below 1.596E+26, within the range of single precision variables, and well within the range of double precision, therefore eqn 17 is appropriate. Below this threshold eqn 18 is a safer approach.

Double precision is used by MathCAD, however, there is the odd bug in the way MathCAD 6.0 detects overflow in some expressions.

Again, this function is symmetric about zero, and the Gaussian distribution is declining very rapidly by 8σ, so the computation is performed by:

$$Perror\left(\frac{E_B}{N_O}, T_E\right) = 2\int_0^{\min\left(8\sigma,\frac{1}{2}\right)} Pef\left(\frac{E_B}{N_O}, t\right) Pb(t, T_E)\, dt \qquad 19$$

This integral is inverted, for a given T_E, to find the Eb/No which yields the desired error rate. For σ above approximately 0.05, it is not possible to achieve the 5E-9 error rate of the MSG system and the iterative root solving routing fails to converge.

The result of this calculation is shown in the following graph:

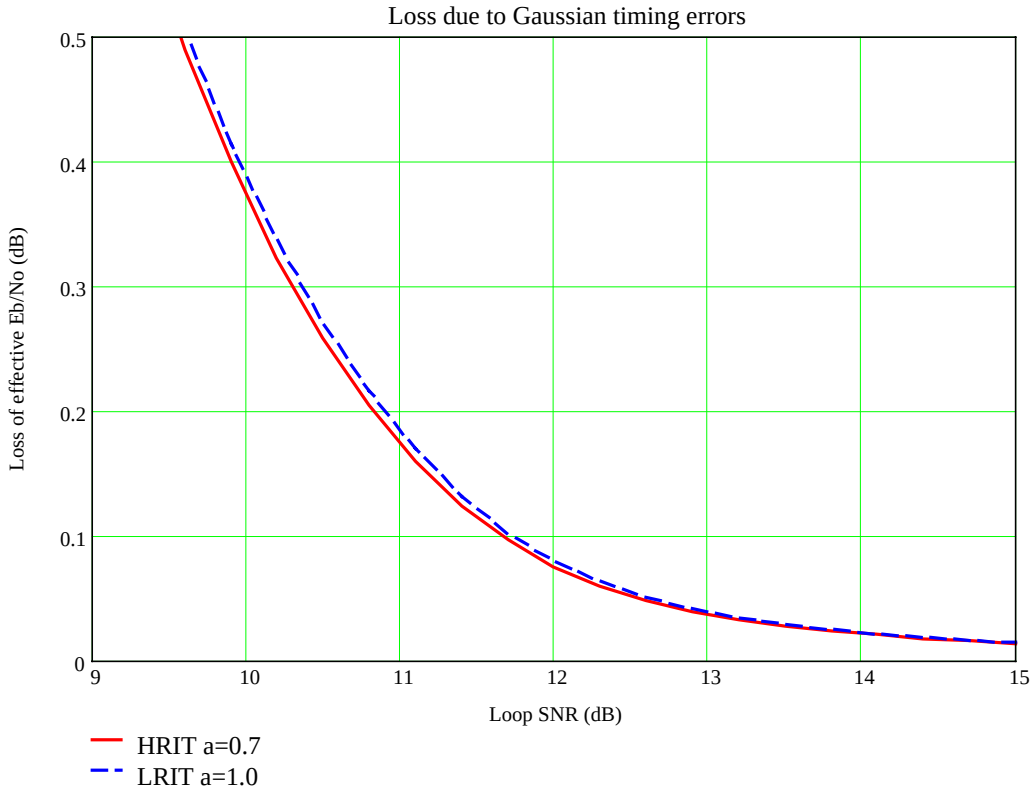

Figure 5: Loss Due to Gaussian Bit Timing Errors

Note the X-axis scale of Figure 5 is expressed in terms of SNR_{Loop}, compared to the RMS timing error in Figures 3 and 4. The much greater sensitivity to Gaussian timing errors is the direct result of the steep error rate against timing error behaviour. The small proportion of time spent at a few standard deviations produces the majority of the degradation in error rate.

To achieve the low overall technology loss value of 0.8dB, a sensible target is less than 0.2dB from the timing errors. From Figure 5 this requires a SNR of approximately 11dB (RMS timing error of 0.032 of the coded symbol time). The HRIT system with narrower baseband filter (a=0.7) is very slightly less sensitive to timing errors than the LRIT system (with a=1.0) in this case since the FEC system deals with the ISI quite well, the lower sensitivity is due to the slightly "rounder" HRIT system *s(t)* (see Figure 2).

Since operation at just above this figure results in a significant increase in loss, and there are uniform timing errors to be considered as well, a more prudent target might be 12dB. This is above the minimum for reliable operation (with a negligible probability of cycle slip), however, we must also consider the possibility of EDA element failure. This results in a periodic 3dB drop in the Es/No, therefore we might aim for a higher value, say 15dB, which translates to an RMS error of 0.02 bit periods. It is this specification that will determine the choice of bit sync PLL noise bandwidth once the data timing error system has been analysed.

If we now consider this in connection with the uniform errors, we have a mean square Gaussian term of 0.02^2 = 4E-4 and a maximum mean square value of 0.03^2 = 9E-4, so the mean square error from the uniform term should be less than 5E-4. This ignores the different probability distribution functions, but is a simple approach. For the uniform deviate of $\pm T_E$, the mean square error is $\frac{T_E^2}{12}$ so we have a maximum period of $T_E = \sqrt{12 \bullet 5.0E-4} = 0.078$ which requires a minimum of 13 interpolation points per bit. If we

use two samples per bit then we need at least 7 polyphase filters to achieve this performance. This is slightly pessimistic, however, it leads to a safe and reasonable design target.

1.3 Timing Estimation

1.3.1 Choice of Algorithm

The problem of extracting the timing information from the received baseband signal is an area which is still under active research. Unlike the carrier phase estimation system, this is strongly dependant on the baseband signal shape. There are various “ad hoc” structures that work well, even though they ware not derived systematically from an optimum solution. In some cases, the bit sync structure functions only for high bandwidth systems (the nearly “square” shape common on older systems, deep space systems, etc) and others that only work effectively only for narrow band systems, such as the raised cosine system of the MSG.

There is a Maximum Likelihood solution to the basic problem, unfortunately this approach invariably has a structure that is too complex for all but the lowest of data rates when implemented in real time software.

For the MSG link, the target for RMS timing stability is quite high (low σ, high loop SNR). However, the high stability of the received data stream, and the continuous nature of the transmission (so lock up time is not a significant issue), allow a narrow band PLL for tracking the data. The moderately high data rate for HRIT (for DSP software), together with the requirement for moderate cost, and the practical option to use a narrow band PLL system, all suggest that a simple sub-optimum timing estimator would represent a good compromise.

If, as we intend, the system is based on a VLSI DSP chip with the NCO, mixer, baseband filters and interpolator on chip, this relieves the DSP software of the most numerically intensive tasks. However, to make maximum use of this advantage the bit sync timing system should be based on matched filter samples only, not on any conjugate timing filter systems.

The implementation which we consider to be most appropriate is the “Zero-Crossing Detector” (ZCD) of (Gardner, 1988, p326, eqn 10.28) which requires two samples per bit, one at the optimum time (for data decisions) and one at the edge of the bit cell, when the 10 or 01 data pattern results in a zero crossing event.

The basic ZCD algorithm involves detecting a data transition from the optimum samples (at the current symbol, t = n and the previous symbol at t = n-1) and using this to direct the error voltage observed at the mid point (t = n-½):

$$v_E = x\left(n-\tfrac{1}{2}\right)\left[\operatorname{sign}\left(x(n)\right)-\operatorname{sign}\left(x(n-1)\right)\right] \qquad 20$$

where *x(t)* is the observed signal and noise from the matched filter and interpolator. The ZCD is simple to implement, although it requires two samples per bit, and the timing properties depend upon the shape of channel response function *s(t)*.

As for all NRZ bit timing circuits, if there is no data transition [*x(n)* same sign as *x(n-1)*] there can be little or no timing correction generated. This is not always true, for example, for non-minimum bandwidth signal formats such as Bi-ϕ-L (also known as SPL or Manchester coding) which have a transition in each bit cell. In these cases it is normal to achieve clock lock, although not clock phase determination, in the absence of data transitions. However, even for Bi-ϕ-L format the ZCD would not operate very well for clock recovery if there were few data transitions. Due to the randomisation of the MSG link (CCSDS, 1989) it is perfectly reasonable to assume a 50% data transition density, hence there should always be sufficient transitions for clock recovery.

The ZCD is a zero search algorithm, not a maximum finding system. This implies we must have a signal waveform, s(t), which is symmetric about zero time error otherwise the zero crossing data will be spread out and any resulting bias in the estimation would contribute a “static” timing error. For any true matched filter system this is correct, however if the TX signal is distorted, or the RX matched filter not ideally matched, the resulting s(t) may be asymmetric.

Another implication is there are two operating points as in the conventional carrier PLL, the wanted one which is stable, and the unwanted point which is unstable. Eventually the system will shift from the unstable point, but this could take a long time (known as “hang up”). Operation at the lowest acceptable SNR reduces this time, as can operation with a deliberate sweep / frequency offset until lock is properly detected.

A further attraction of the ZCD is the extension to complex baseband processing proposed by (Verdin, 1995, p2/1 to p2/7). They replace the sign(x) function used for real I and Q voltages when considered independently, with the “complex signum” function. This was not actually defined in the paper, but is the same idea extended to complex data:

$$\mathrm{sign}(z) = \frac{z}{|z|} \qquad 21$$

For a real valued signal, this is -1 for z < 0 and 1 for z > 0, and usually defined as 1 for z = 0 as well. For the complex value, this is the unit vector in the direction (same argument) of the original signal. The z = 0 case is a problem, typically this would be set to zero, or 1+j0 for compatibility with the real valued function, depending on how sign(z) is computed.

The importance of this “Complex ZCD” modification is the algorithm of eqn 20 can operate without carrier phase lock, provided the phase rotation is not significant over one bit period. This can greatly improve acquisition since the process can begin during carrier acquisition phase. It is possible that the simpler (and marginally better noise performance) of the original data-directed form of eqn 20 would be used when carrier lock is detected. The resulting equation is:

$$v_E = \mathrm{Re}\left\{ z^*\left(n-\tfrac{1}{2}\right)\left[\mathrm{sign}\left(z(n)\right) - \mathrm{sign}\left(z(n-1)\right)\right]\right\} \qquad 22$$

In this case, *z* is the complex baseband signal, given by *z* = *x* + *jy*, that is I + jQ, and * indicates complex conjugate.

There is the potential for odd synchronisation behaviour here, we could have both carrier and clock acquisition systems that can operate independently, but are better when operating together. It is claimed by (Verdin, 1995, p2/4) that the C-ZCD is free of hang-up due to the discontinuous S-curve at the unstable operating point, however, this was made based on zero noise analysis of the system. This is not true as they stated, but from our experience we are confident the proposed systems will lock reliably.

There are one sample per bit algorithms, such as the “Mueller and Müller” data-aided algorithm included in (Gardner, 1988, p332, eqn 10.34) and analysed in (Meyr, 1998, p196-205), however, our simulations for the excess bandwidth values used for the MSG link show the ZCD is better by around 10dB. In our opinion, the reduction in DSP software input sample rate is not enough to make up for the poorer performance of this method.

1.3.2 Analysis of Algorithm

If we had *a priori* knowledge of the data pattern, we can compute the basic detector shape for a 0-1 transition to get the synchronising curve (S-curve) of the detector from s(t) as:

$$u(t) = s\left(t-\tfrac{1}{2}\right) - s\left(t+\tfrac{1}{2}\right) \qquad 23$$

The S-Curve shows the error parameter (voltage for an analogue system, number for a DSP system) as a function of the error in the wanted parameter, in this case bit timing.

There is a factor of two from the [sign - sign] term of eqn 21/22, however, this is only for 50% of the data transitions on average, hence eqn 23. In the other cases (BPSK verses QPSK, etc) there are various other constants, but these can be found from simulation, etc, for the final system.

The resulting curve is only valid between the two bits used for the data decisions, and looks like:

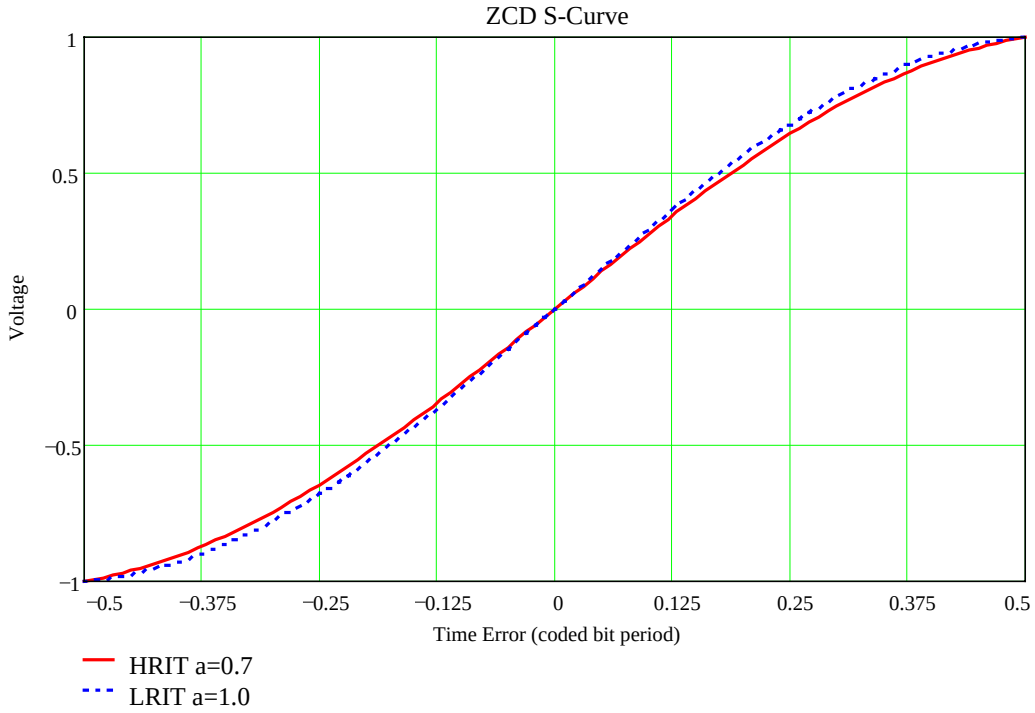

Figure 6: ZCD *a priori* Decision Timing Sensitivity Curve

It appears the S-curve is discontinuous at the edge (±0.5 bit error), however, this was based on the *a priori* knowledge of the data values, not on the data decisions that are actually possible. As the timing error approaches the edges of the region, both the ISI effects and the reduction in useful signal result in incorrect data decisions and a loss of detector response.

This S-curve shown is for the normalised system with s(0)=1, for these two cases the slope of the S-curve at t=0 is 2.8 (units of voltage / coded bit periods) for a=0.7 HRIT system and 3.0 for a=1.0 LRIT system. For the equivalent phase lock loop, one bit period represents 2π radian and the equivalent K_{PSD} are 0.45 $Vrad^{-1}$ and 0.48 $Vrad^{-1}$ respectively for the zero noise case.

In addition to the S-curve, we should compute the “self noise” (or data pattern induced jitter) term from the adjacent bits. The bits at ±0.5 are the S-curve generating terms and do not change for u(t) (other than reversing for a 1-0 transition rather than a 0-1). The other bits at integer multiples away generate the self noise voltage and the effect can be estimated as the root sum square value from:

$$n_S(t) = \sqrt{2\left[\sum_{i=1}^{\infty} s^2\left(t + \left(i + \tfrac{1}{2}\right)\cdot T_B\right) + \sum_{i=-1}^{-\infty} s^2\left(t + \left(i - \tfrac{1}{2}\right)\cdot T_B\right)\right]} \qquad 24$$

The factor of 2 accounts for the [sign - sign] term in eqn 20/22, this results in 4 times the variance (noise power) from $s^2(t)$, but this is only for the 50% of the time when data transitions occur. In practical calculations, the infinite summation limits would be reduced to around 6 since the contribution from s(t) beyond this region is negligible.

This can be computed for time errors of ±0.5 bits which correspond to the operating range of the S-curve. This results in:

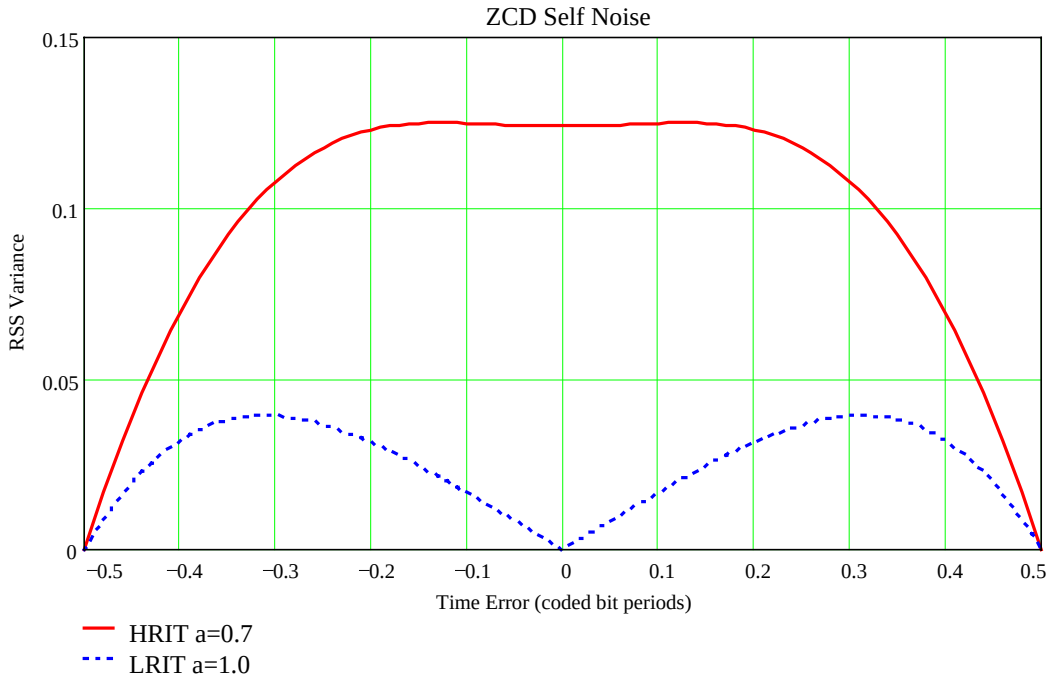

Figure 7: ZCD Self Noise Characteristics

It is interesting to note the LRIT system, with an excess bandwidth of 100%, has no self noise in this configuration at zero timing error. This is due to the s(t) waveform having two sets of zero crossings, both the Nyquist channel design at multiples of 1 bit period, and half way between them. This effect can be seen from a close inspection of Figure 2.

The HRIT system has relatively constant noise before it declines toward the ±0.5 bits edge of the detector. The decline to zero at ±0.5 bits is not useful, the effective S-curve becomes zero at this point. This decline is due to approaching the zeros of the s(t) shape, exactly the characteristic for zero ISI that was originally required for the data transmission. It must be noted that the real “self noise” is not zero here, since the decision errors result in a different source of disturbances from the wanted signal.

To compute the effective timing jitter, we need to know the S-curve and the noise contribution. We have the “self noise” which introduces data pattern jitter at all input SNR values, and we have the Gaussian noise due to the receiver and link noise contribution.

The RMS noise voltage from the wanted signal (this depend upon the energy per symbol to noise density ratio) in the normalised system, s(0)=1, is:

$$v_N = \sqrt{\frac{N_O}{2E_S}} \qquad 25$$

This is the standard deviation of the additive noise term.

Since there appears to be no comprehensive analytic analysis of the different bit timing detectors, the performance was computed by Monte Carlo simulation. A cyclic 1024 bit PRBS with either real-only data (for BPSK) or complex data (for QPSK) was used, this was “raised cosine” filtered and stored in an array with 32 samples per bit cell. A test signal with additive Gaussian noise, standard deviation from eqn 25, was generated and filtered by a close approximation to the RRC filter of the RX. The slight correlation between successive noise samples is important in accurate modelling of the timing error behaviour,

particularly the “mid bit cell” transition detectors. This was applied to the timing detector and the results stored for statistical analysis.

The resulting timing error signal was accumulated for 10^6 points (only $5*10^5$ for the S-curve graphs) and the mean and variance calculated. The 95% confidence range is ±0.2% of the variance observed, however, the detector sensitivity is calculated by numerical differencing, so the effect of the variance is much larger.

For a numerical example, the 0dB BPSK curve has mean values of approximately ±0.08 at the points for differencing (±1/32 bit periods about the 0.5 bit point). The variance is about 0.9, so the standard deviation of the mean is around $2*0.9/10^{6/2}$ = 1.8E-3. When compared to the 0.16 value expected from differencing, this is 1.1%, or a 95% confidence range of 2.3%. At lower Es/No, and for higher loss systems, the simulation uncertainty is increased.

We make the assumption in the following analysis that the timing error will be distributed in a Gaussian manner. This is reasonable if the data is essentially random and the loop bandwidth is narrow compared to the data rate. Even the modest non-Gaussian nature of the self noise terms can be ignored in this case. For example, consider that most low frequency laboratory noise generators use a LPF version of a binary PRBS.

Once we have the mean (the timing error voltage for the PLL) and the variance (the additive noise contribution) we can compute an equivalent noise term for the bit sync PLL. This is calculated from:

$$\frac{N_O}{2E_S} Loss = \frac{v_n^2}{K_{PSD}^2} = \frac{\mathbf{Var}\langle v_E \rangle}{\left[\frac{\frac{\partial}{\partial t}\mathbf{E}\langle v_E \rangle}{2\pi}\right]^2} \qquad 26$$

The left hand side has the noise variance from the pure carrier tracking PLL together with a “loss” term to account for the difference between the timing estimation and the ideal carrier phase estimator. The factor of two is for direct comparison with the phase estimation case, half of the additive noise in in-phase with the carrier and produces a disturbance while the other half is in quadrature and produces virtually no response.

The centre equation contains the measured parameters of interest. The noise power includes both additive and self noise. The K_{PSD} indicates the relative magnitude of this noise, expressed as an angle.

The right hand side of eqn 26 shows how this is determined. The numerator is the noise power (variance) from the detector, calculated from the squared error about the mean value. The denominator finds the expected bit timing detector sensitivity in volts per radian (the factor of 2π converting from “per bit” to radian), when squared this represents the available corrective signal power from the detector.

Although it is not shown in eqn 26, K_{PSD} (or v_n, depending on the system details) is dependant on the signal voltage and here we are considering the normalised system. Also the “noise density” is considered as spread evenly over the detector sampling frequency range from the one centre sample per bit used, which is the case in the ZCD studied.

If we average *N* samples from the timing error detector, we reduce the variance by a factor of *N*. If we consider the equivalent continuous time filter with unity DC gain, the frequency response is given by:

$$H(f) = \int_{-\infty}^{+\infty} h(t) e^{-j2\pi ft} dt = \int_{-\infty}^{+\infty} \frac{1}{NT_B} \mathrm{rect}\left(\frac{t}{NT_B}\right) e^{-j2\pi ft} dt = \mathrm{sinc}(\pi f NT_B) \qquad 27$$

Where T_B is the bit period as before and the *rect()* function is 1 for -0.5 to +0.5, and zero elsewhere. The double sideband noise bandwidth of the (loop) filter, B_L, can be found from:

$$B_L = \frac{\int_{-\infty}^{+\infty} |H(f)|^2 df}{|H(0)|^2} \qquad 28$$

And using the relationship:

$$\int_{-\infty}^{+\infty} \mathrm{sinc}^2(\pi f X) df = \frac{1}{X} \qquad 29$$

this leads to:

$$B_L = \frac{1}{NT_B} \qquad 30$$

Using this noise bandwidth to bit averaging equivalence, and eqn 26, we have:

$$\bar{\theta}^2 = \frac{1}{N}\frac{v_n^2}{K_{PSD}^2} = B_L T_B \frac{v_n^2}{K_{PSD}^2} = B_L T_B \frac{N_O}{2E_S} Loss \qquad 31$$

The term $B_L T_B$ is effectively normalising the loop noise bandwidth to the data rate, since the variance term used is “per bit” of the data stream. Using our definition of loop SNR this leads to:

$$SNR_{Loop} = \frac{1}{2\bar{\theta}^2} = \frac{E_S}{N_O B_L T_B}\frac{1}{Loss} \qquad 32$$

Note that E_S/T_B is the carrier power C, we have:

$$SNR_{Loop} = \frac{C}{N_O B_L}\frac{1}{Loss} \qquad 33$$

which allows a direct comparison with conventional carrier tracking PLL systems, providing the cyclo-stationary noise (noise density as a function of timing error) and the possibly non-sinusoidal PSD characteristics are noted.

Looking at eqn 26, we can determine the loss term from simulation as:

$$Loss = \frac{v_n^2}{K_{PSD}^2}\frac{2E_S}{N_O} = \frac{\mathbf{Var}\langle v_E \rangle}{\left[\frac{\frac{\partial}{\partial t}\mathbf{E}\langle v_E \rangle}{2\pi}\right]^2}\frac{2E_S}{N_O} \qquad 34$$

This is analogous to the loss from the regeneration non-linearity in the carrier tracking PLL. In the bit synchroniser the loss is generally much greater and dependant on the signal shape, etc, which is not the case in the carrier phase estimator system. From these, we can determine a suitable loop noise bandwidth from:

$$dB(SNR_{Loop}) = dB(C/No) - dB(PLL\ Noise\ Bandwidth) - dB(Loss) \quad 35$$

and hence the PLL noise bandwidth from:

$$dB(PLL\ bandwidth) = dB(C/No) - dB(Loss) - dB(SNR_{Loop}) \quad 36$$

1.3.3 Performance Bound

It is of interest to consider the fundamental limit of performance for any system as a measure of the quality of the practical approach. For the bit timing recovery we present the following results in a similar manner to the above for a comparison with the well developed carrier phase tracking application of the PLL.

We have the total noise voltage as the sum of the additive noise and the self noise. This results in the equivalence of $v_n^2 = \frac{N_O + N_S}{2E_S}$ for the normalised s(0)=1 system so we can rearrange eqn 31 as:

$$\overline{\theta}^2 = B_L T_B \frac{N_O + N_S}{2E_S K_{PSD}^2} = \frac{B_L N_O T_B}{2E_S}\left(1 + \frac{N_S}{N_O}\right)\frac{1}{K_{PSD}^2} \quad 37$$

Noting that $E_S/T_B = C$, the carrier power, it follows that:

$$\overline{\theta}^2 = \frac{B_L N_O}{2C}\left\{\left(1 + \frac{N_S}{N_O}\right)\frac{1}{K_{PSD}^2}\right\} \quad 38$$

A comparison of this with the usual pure carrier tracking PLL formula shows the term in the curly brackets representing a “loss” term against some ideal. Looking at this loss term, the effect of high SNR operation (where N_O is small) is an increase in phase variance due to the self noise term N_S.

Clearly a system with a steep delay discriminator response (K_{PSD} high for a given signal level) will be better, providing this is not achieved at the expense of the additive noise term. In eqn 38 it must be remembered that K_{PSD} has already been normalised to the E_S signal level.

There is an upper limit to the performance of any synchroniser in terms of the signal to noise ratio and the observation time interval (the loop bandwidth, etc). This is known as the Cramér-Rao bound and is based on the assumption that we know everything about our signal, except for the parameter of interest, in this case the timing error. For this case the limit is given by (Meyr, 1998, p331, eqn 6-28) and after rearrangement becomes:

$$\overline{\theta}^2 = \frac{B_L N_O}{2C}\left\{\frac{\int_{-\infty}^{+\infty} |H(f)|^2 df}{T_B^2 \int_{-\infty}^{+\infty} f^2 |H(f)|^2 df}\right\} \quad 39$$

The curly brackets of eqn 39 can be compared to those for eqn 38. Here the numerator normalises to unit signal power, and T_B normalises the denominator to unit bit rate. Essentially this is the distribution of signal power with frequency, the greater the 2nd moment, the higher the potential timing accuracy. It

should be noted that this “loss” may be less than 1 for very wideband (relative to the bit rate) signals, unlike the loss in a carrier tracking PLL.

For the raised cosine system this loss term can be evaluated by numerical integration and is shown in Figure 8:

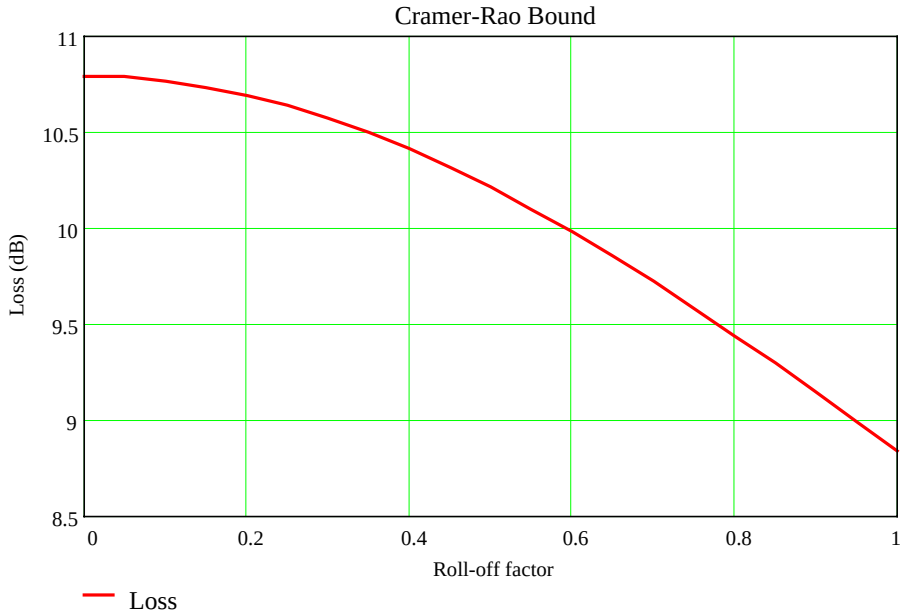

Figure 8: Cramér-Rao Bound for Raised Cosine Signals

For the MSG system, the limit for HRIT with a=0.7 is 9.73dB and for LRIT with a=1.0 the limit is 8.84dB. This limit is unlikely to be reached since it assumes prior knowledge of the data sequence, however, it represents a fundamental point of comparison for the practical systems.

1.4 LRIT BPSK System

1.4.1 Simulation Results

If we start with the BPSK system with carrier lock, we can use the data-directed approach on the I channel data only. This results in the following curves (where the centre of the graph, Timing Position = 0.5, represents the “zero error” desired position):

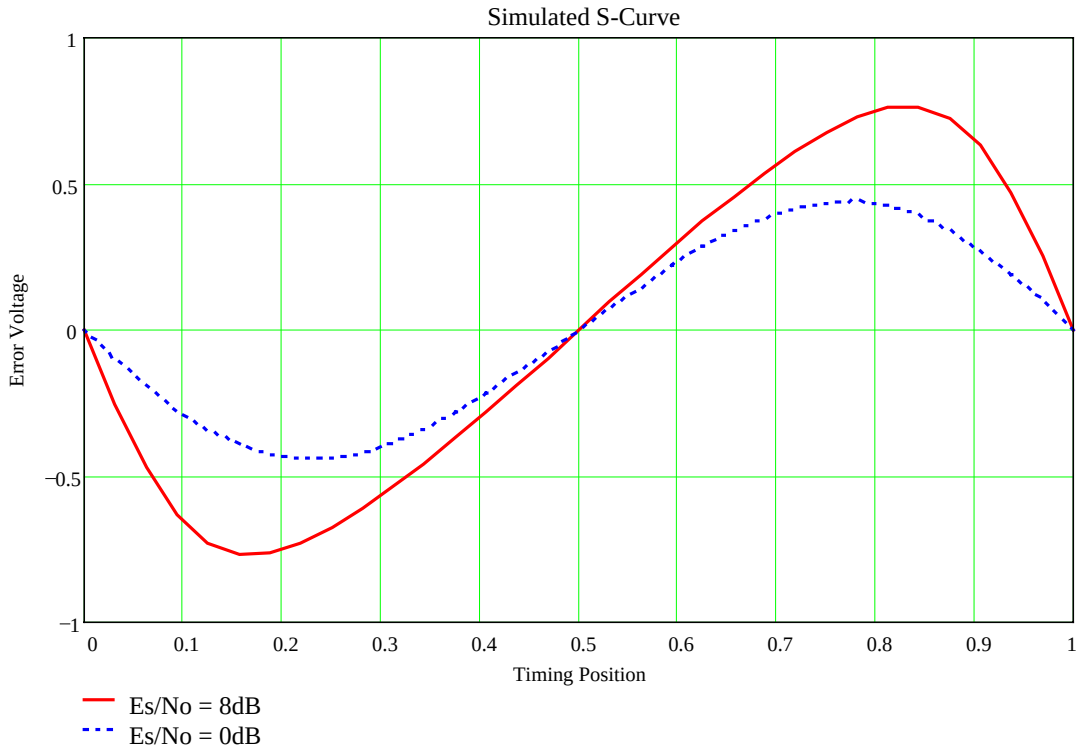

Figure 9: BPSK Data-Directed S-Curve

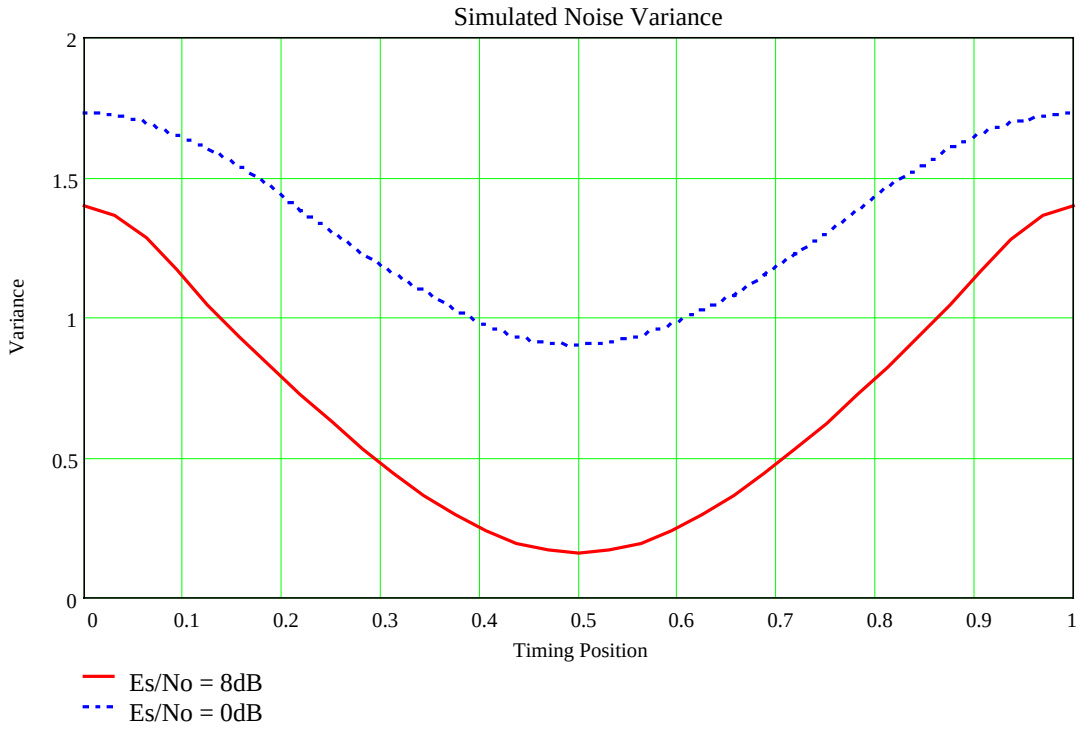

Figure 10: BPSK Data-Directed Variance

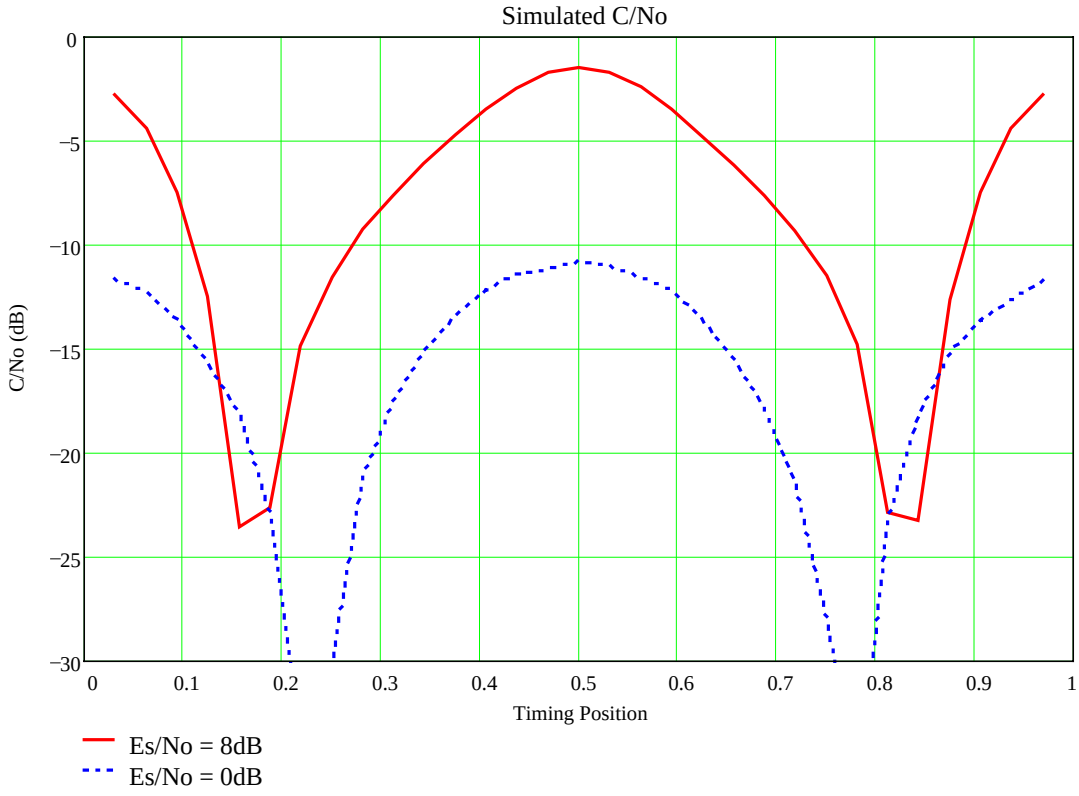

Figure 11: BPSK Data-Directed C/No

In Figures 9-11 the simulation is performed for two input Es/No values, the high SNR situation with 8dB, and the normal situation with Es/No = 0dB. The theoretical value for MSG is -0.8dB, but there is an additional margin of 0.8dB for technology loss in the user station requirements. In performing the loop calculations, only 0.6dB is allowed since the demodulator losses are set to be less than 0.2dB (much less in practice, however).

In Figure 9 the effect of noise on the data decisions can be observed as a reduction in the system sensitivity, this is most pronounced away from the centre of the operating range (timing position of 0.5 indicates the sampling is at zero error). The *a priori* decision sensitivity of this detector is 3.0 and in the high SNR case the simulated slope is close at 3.04 Volt/bit. At the operating point of 0dB this has reduced to 2.4, a decrease of 20%.

Also apparent in Figures 10 and 11 is the increase in detector noise as the system moves away from the optimum point. There is a contribution from the self noise term, but this is very small, the majority of the variance has two sources:

- An increase in data decision errors, reducing the detector response and adding signals from the non-zero crossing bit cells.
- The error signal has two values, the timing error estimate from the data transitions, and a zero value for no transition. Even with no noise, this leads to random variations in the mean value from the random data pattern.

Clearly we must operate close to the zero error (the ½ bit point) to minimise both sources of noise, however, this is essential to minimise technology loss in any case.

The effective C/No (for unity symbol rate) is plotted in Figure 11 from eqn 26 by numerical differencing and smoothing of the simulated data points. This explains the restricted range of the x (time) axis. At low Es/No this shows the noise of the simulation and the overall accuracy decreases.

For other input Es/No, the simulation results for this data-directed algorithm, and the complex sign version, are summarised in Table 1 and illustrated in Figure 12. This table was calculated separately from the S-curves and used a greater number of simulation points to reduce the uncertainty.

Input Es/No	Data-Directed ZCD Loss	Complex ZCD Loss
8dB	9.4dB	9.6dB
6dB	9.5dB	9.8dB
4dB	9.7dB	10.2dB
2dB	10.1dB	10.9dB
0dB	10.9dB	11.8dB
-2dB	11.9dB	13.1dB
-4dB	13.1dB	14.6dB

Table 1: BPSK Timing Detector Loss

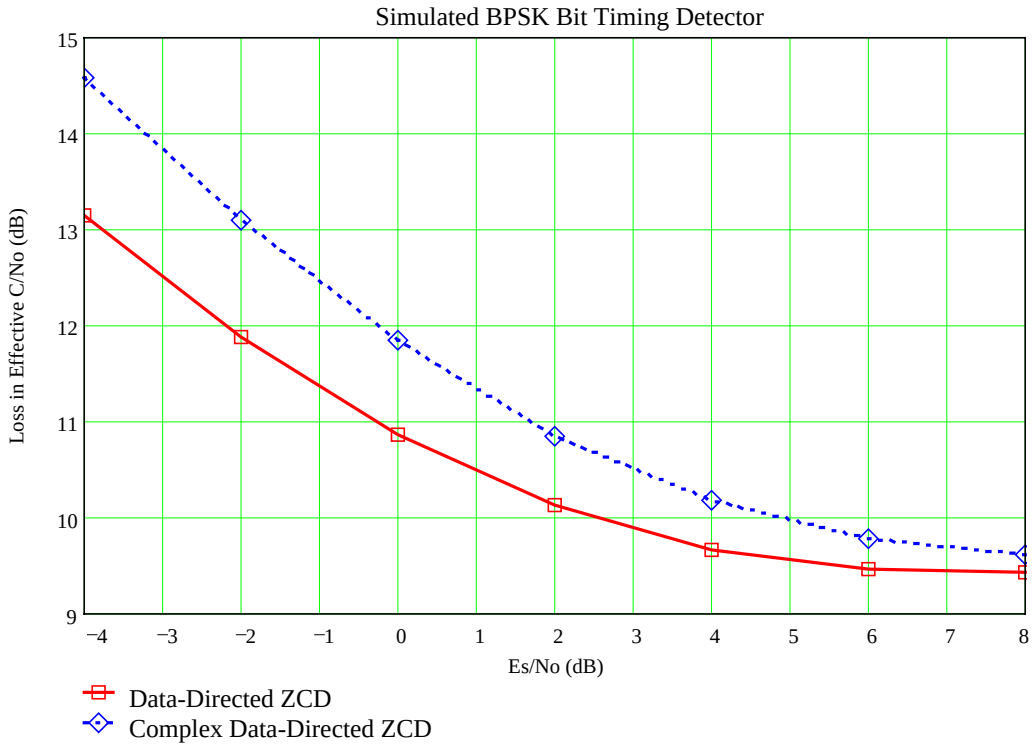

Figure 12: BPSK Data-Directed Loss

For the data-directed version, the loss is around 9.4dB for high input SNR. As a general rule, the loss term increases again for higher Es/No since the self noise term (eqn 24) dominates over the additive Gaussian noise, however, for the particular case of LRIT (a=1.0) and zero timing error, the self noise is term is zero. An additional point about the ZCD is the decreases in performance as the excess bandwidth factor *a* is reduced. In the case of LRIT, we have the best waveform for this type of timing error detector.

When we compare the lower bound on loss of 8.84dB, we conclude that for LRIT systems the performance at high SNR of 9.4dB is sufficiently close to the theoretical limit to be acceptable. At the operating point of around -0.2dB (Es/No of -0.8dB and 0.6dB technology loss margin, after allowing no more than 0.2dB for demodulation degradation) the loss is approximately 11dB, only 2.2dB away from the theoretical limit. When considered together with the simplicity of the ZCD structure, and the ability to use a only the matched filter outputs and not any conjugate timing error filter, it clearly makes this the optimum performance to cost solution.

1.4.2 Calculation of PLL Parameters

From Figure 12, the loss at Es/No = -0.2dB appears to be 11dB. We are operating with a symbol rate of 293.9E3 and so the loop noise bandwidth is given by:

$$dB(B_L) = \{dB(\text{Symbol Rate}) + dB(Es/No)\} - dB(SNR_{Loop}) - dB(\text{Loss at Es/No})$$

$$= (54.7 - 0.2) - 15 - 11 = 28.5dBHz = 708Hz \qquad (40)$$

If we use a damping factor of 1.14, this requires a loop natural frequency of 83Hz.

We must also consider the effect of an EDA element failure. This would periodically reduce the input Es/No by 3dB. Using the graph of Figure 12, the loss at Es/No of -3.2dB is about 12.6dB. Considering the previous PLL bandwidth, the equivalent SNR is given by:

dB(SNR) = {dB(Symbol Rate) + dB(Es/No)} - dB(B_L) - dB(Loss at Es/No)

= (54.7 - 3.2) - 28.5 - 12.6 = 10.4dB 41

At 10.4dB the mean time to cycle slip is approximately:

$$T_{AV} \approx \frac{2}{B_L}\exp\left(\pi \cdot SNR_{Loop}\right) = \frac{2}{708}\exp\left(\pi \cdot 10.96\right) = 2.5E12 = 80000\,years \qquad 42$$

This effectively zero probability of cycle slip indicates the PLL will track through EDA failure without any problems. Of course, we are assuming the bit timing detector is sufficiently close to the sine detector used for the conventional PLL analysis that created eqn 42, and the extrapolation to high SNR is reasonable. However, even an error of a factor of one thousand (just under 1dB in SNR) would not change the conclusion.

The symbol rate from the PGS is known to high accuracy, almost certainly limited by the stability of the receiver's data sample clock. We would aim for better than ±25ppm for the sample clock as a general specification which is higher than most general purpose crystal clock modules. This translates in to an error of ±6.4Hz for the LRIT symbol clock. This is much less then the PLL bandwidth so the lock up time should be a few periods of the natural frequency (perhaps 100ms or so). The tuning range of the bit clock would be comparable to the natural frequency to allow reasonable phase slew rate, so a bit sync NCO range of ±100Hz is not unreasonable.

1.4.3 Conclusions

For data tracking the application of Gardner's ZCD (eqn 20) to the LRIT BPSK bit sync and a PLL with a natural frequency of 83Hz and damping factor of 1.14 should function almost perfectly even in the event of EDA failure.

Considering the short lock up time for the LRIT BPSK demodulator the additional time for bit lock is insignificant and the marginal advantage of the Complex ZCD in this area is not really justified.

As for the carrier PLL design, the impact of time delay in the PLL system is important and a similar specification of no more than 100µs is reasonable. This applies for the proposed system with an external NCO and interpolator, if this was in the DSP software there would be zero time delay.

1.5 HRIT QPSK System

1.5.1 Simulation Results

If we start with the QPSK system with carrier lock, we can use the data-directed approach on both I & Q channels. This results in the following curves:

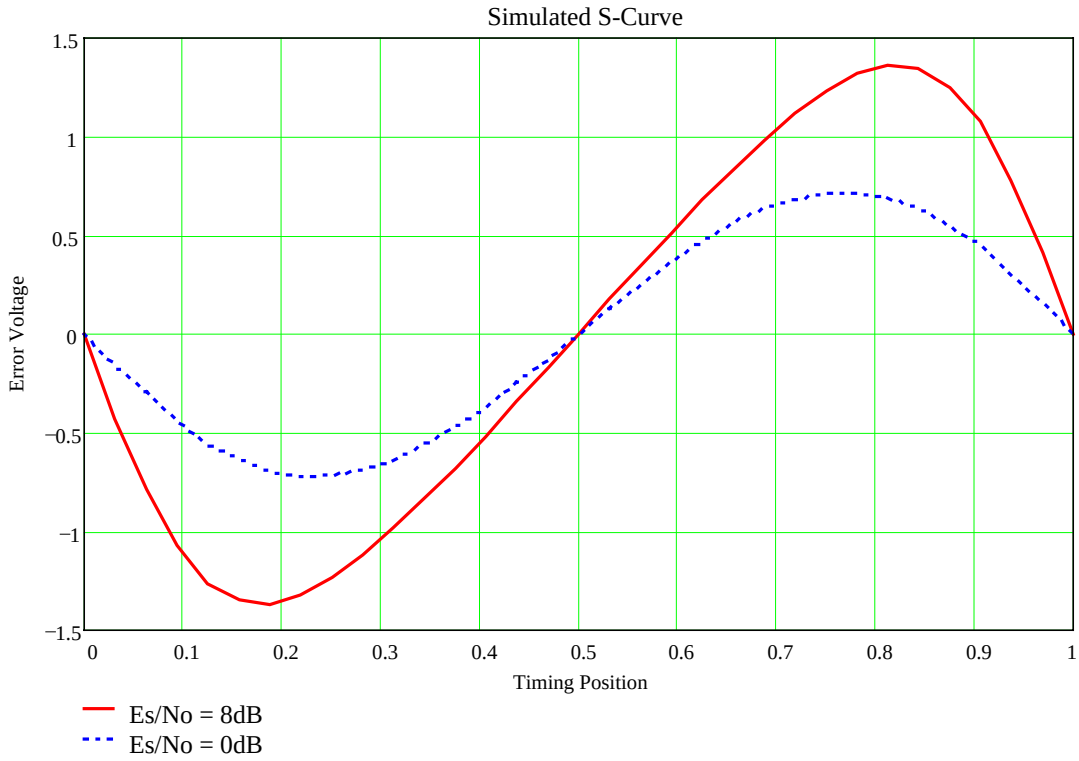

Figure 13: QPSK Data-Directed S-Curve

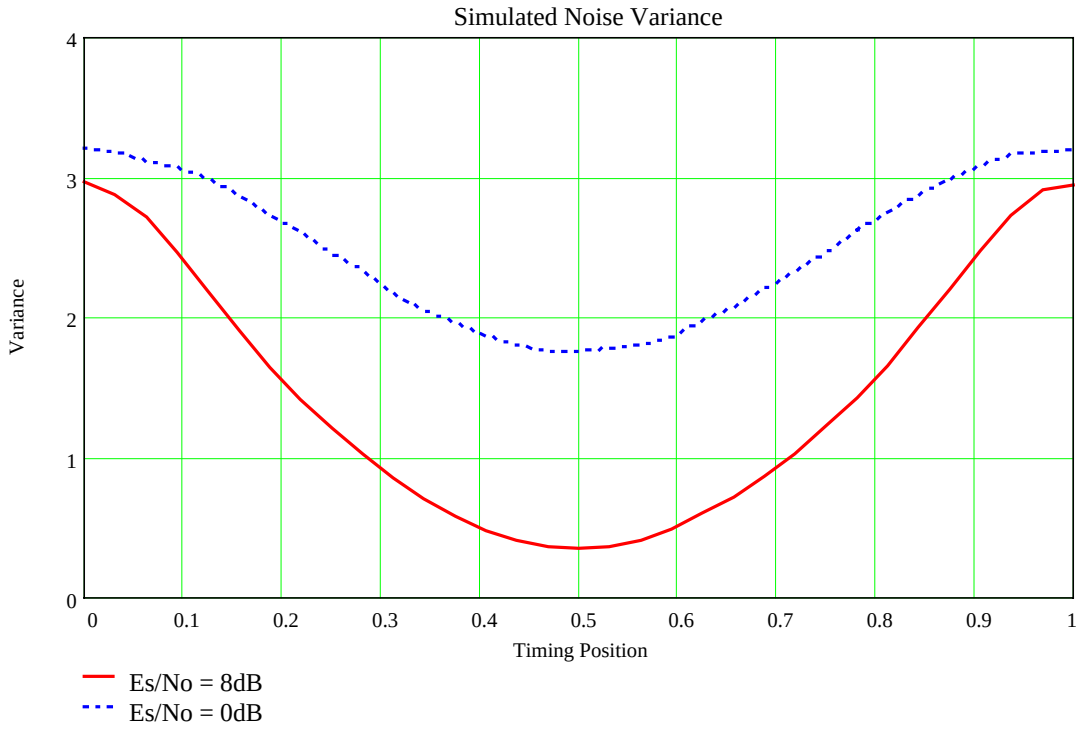

Figure 14: QPSK Data-Directed Variance

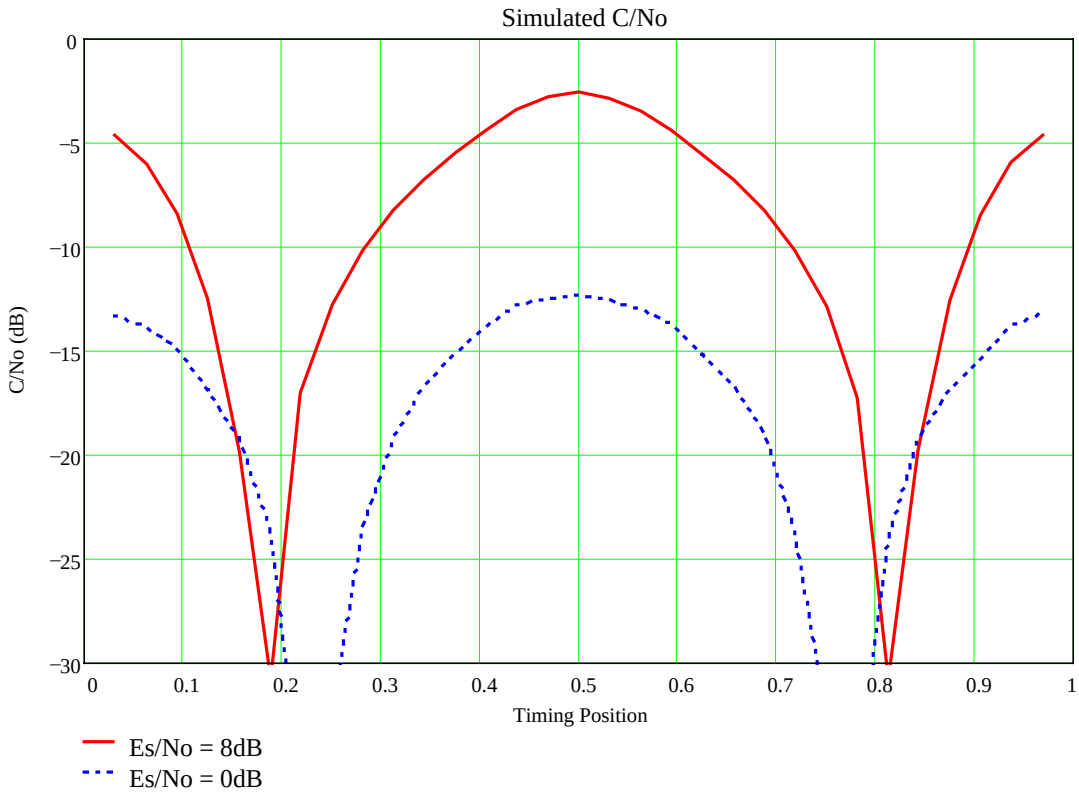

Figure 15: QPSK Data-Directed C/No

In Figures 13-15 the simulation is performed with a=0.7 for two input Es/No values, the high SNR situation with 8dB, and the normal situation with Es/No = 0dB, as for the LRIT case.

The *a priori* decision sensitivity of this detector is 2.8*2 = 5.6. The factor of two accounts for the two channels (I & Q) contributing decisions. In the high SNR case, the simulated slope is close at 5.5 Volt/bit. At the operating point of around 0dB this has reduced to 4.0, a decrease of 28%.

The results for other values of Es/No, and for the C-ZCD, are given in Table 2 and illustrated in Figure 16.

Input Es/No	Data-Directed ZCD Loss	Complex ZCD Loss
8dB	10.6dB	11.4dB
6dB	10.5dB	11.3dB
4dB	10.7dB	11.4dB
2dB	11.4dB	11.7dB
0dB	12.3dB	12.3dB
-2dB	13.5dB	13.2dB
-4dB	14.9dB	14.4dB

Table 2: QPSK Timing Detector Loss

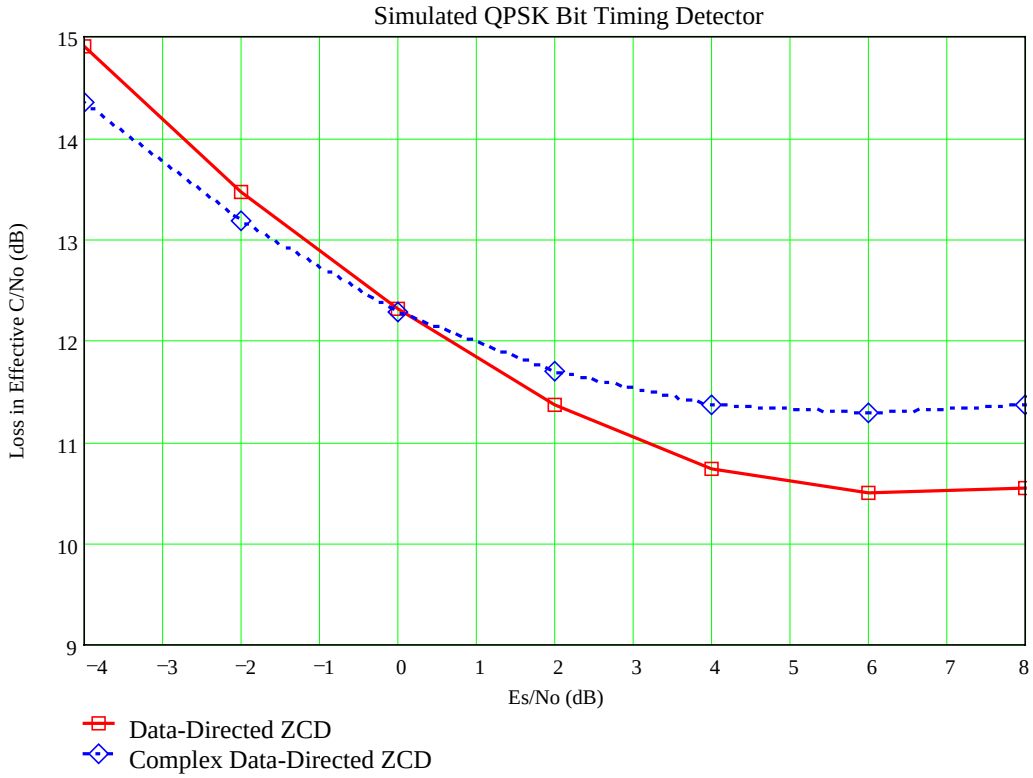

Figure 16: QPSK Data-Directed Loss

In calculating the QPSK results, these is a 3dB factor to be noted due to the computation on each pair of symbols for each timing detector result. Once this factor of 3dB is allowed for, the loss is slightly higher than for BPSK, but of the same order. The other factor which is different is the a=0.7 used for the HRIT QPSK calculations.

In the QPSK case the complex variation of the ZCD is slightly poorer at high Es/No, where the phase estimation errors (the argument of the complex sign function) lead to greater detector noise. This is similar to the demodulator phase noise problems. In the BPSK case any phase errors in the complex sign only reduced the wanted signal, however, in the QPSK case this introduces crosstalk noise in the timing signal as well.

At low Es/No values the decision errors in the data-directed case become significant and performance of the two are similar. At very low Es/No (less than 0dB) there is a minor advantage with the CZCD, this is attributed to the smaller magnitude of incorrect data direction in certain cases of decision errors. (In fact there is a variation of the ZCD not considered here using a Maximum Likelihood non-linearity with the best characteristics of both systems).

1.5.2 Calculation of PLL Parameters

At the expected operation point of Es/No = -0.2dB, the loss is approximately 12.4dB, when compared to the bound of 9.73dB this is 2.7dB poorer, not a serious loss.

Using these figures, and the equivalent serial coded symbol rate of 2.296M/sec (the appropriate figure for the above presentation of the simulated data) we have:

$$dB(B_L) = \{dB(\text{Symbol Rate}) + dB(Es/No)\} - dB(SNR_{Loop}) - dB(\text{Loss at Es/No})$$

$$= (63.6 - 0.2) - 15 - 12.4 = 36dBHz = 3981Hz \qquad 43$$

If we use a damping factor of 1.14, this requires a loop natural frequency of 467Hz.

As a possible saving in CPU effort, only the I or Q channel could be used, requiring a PLL with half of the noise bandwidth (f_N = 233Hz). This is not the best option, but it would be acceptable if both uplink modulator, and receiver demodulator, have DSP based systems with zero timing error (phase shift) between the I and Q channels. An additional attraction is the simplification in DSP code, the same code (with slightly different constants) being used for both LRIT and HRIT.

Again we must consider the effect of an EDA element failure. This would periodically reduce the input Es/No by 3dB. Using the graph of Figure 16, the loss at Es/No of -3.2dB is about 14.3dB. From this we have:

$$dB(SNR) = \{dB(\text{Symbol Rate}) + dB(Es/No)\} - dB(B_L) - dB(\text{Loss at Es/No})$$

$$= (63.6 - 3.2) - 36 - 14.3 = 10.1dB \qquad 44$$

At 10.1dB the mean time to cycle slip is approximately:

$$T_{AV} \approx \frac{2}{B_L}\exp\left(\pi \cdot SNR_{Loop}\right) = \frac{2}{3981}\exp\left(\pi \cdot 10.23\right) = 4.6E10 = 1400\,years \qquad 45$$

Considering the EDA element is 1/32 of each spin, cycle slip is effectively a zero probability even under EDA failure.

We could consider the CZCD option. Looking at Figure 15, the loss at Es/No = 0dB or below is very similar to the data-directed system, hence the PLL parameters are essentially the same. Since the maximum carrier frequency error is 150kHz (±75kHz specified and the opposite for the NCO during the worst case search start value), the phase rotation over one symbol period is given by:

150E3 * 2π * 0.871E-6 = 0.82 radian = 47°

The CZCD would probably allow bit sync lock to occur, and certainly when the frequency error decreased to 30kHz or similar. This would increase the useful C/No for the carrier PLL by 3dB and hence increase the possible sweep rate. Although this is initially attractive, the saving of perhaps 3 seconds in acquisition time must be weighed against the extra complexity for the complex sign look-up table and the

corresponding subtract and multiply operations. The data-directed algorithm can be implemented by with only two sign tests and an add/subtract operation per I/Q channel used.

It is worth observing that our previous generations of analogue and digital (synchronous sampling) bit synchronisers would lock at Es/No below 0dB, even though they were designed for uncoded (Es/No around 11dB) applications. The use of VCXO based PLL systems forced comparable loop natural frequencies to the calculated MSG values. In addition, it was observed that these older bit synchronisers were able to achieve synchronisation in the absence of carrier phase lock. We can expect the asynchronous digital systems to do the same, although the CZCD would perform much better in this respect.

As for the LRIT system, we know the bit clock rate to high accuracy, applying a specification of ±25ppm for the symbol clock we have ±29Hz. Again, a higher tuning range for the NCO is normal to allow a fast phase slew rate when adjusting to the correct time so a tuning range of ±250Hz is not unreasonable. Since the frequency error is less than the loop natural frequency, the time to bit sync lock is a few PLL “periods”, in this case only tens of milliseconds.

1.5.3 Conclusions

We conclude the use of the ZCD for the HRIT QPSK system is also the best compromise. Although the CZCD would allow a higher carrier PLL sweep rate, the complexity is not really justified. If fast acquisition were really important we would attempt a frequency discriminator or FFT (maximum likelihood) analysis of the frequency error.

The bit sync PLL for HRIT could use both I and Q data with a PLL with a damping factor of 1.14 and a natural frequency of 467Hz. This would function correctly in the presence of EDA element failure. In this case the system is processing two coded symbols (E_S) of signal energy per QPSK “symbol” received.

A minor reduction in CPU load could be achieved by using the data-directed ZCD on the I channel only, the same as for BPSK. In this case there is half of the signal energy available and hence the PLL noise bandwidth must be reduced by a factor of two, leading to a PLL natural frequency of 233Hz.

As for the carrier phase tracking system, the effect of time delay in the PLL needs to be considered (since our proposed solution has the interpolator in separate DSP hardware from the ZCD software). From the carrier PLL analysis, the higher PLL bandwidth of 467Hz needs care to ensure the overall time delay is kept to a minimum, ideally 50µs or less. This is an important factor in the DSP hardware, processor choice, and software architecture. The use of the lower natural frequency (single I channel at 233Hz) would ease this real-time requirement to a similar value as the carrier PLL case requires.

1.6 Bit Synchroniser Summary

The bit synchroniser characteristics are summarised in the table below. Due to the choice of interpolator (the 32 step per sample of the HSP50214) the loss term is essentially due to the Gaussian timing errors for the 15dB SNR_{Loop} used in the design. This assumes there are no ‘static’ errors, for example from an asymmetric (distorted) transmitted signal, etc.

Parameter	HRIT	LRIT
Timing Error Algorithm	ZCD (on I data only)	ZCD
PLL Natural Frequency	233Hz	83Hz
PLL Damping Factor	1.14	1.14
Bit Rate Frequency Range	±250Hz	±100Hz
Technology Loss	< 0.025dB	< 0.025dB
Operation with EDA failure	Yes	Yes

Table 3: Bit Synchroniser Characteristics

1.7 Implementation of Baseband Filters

1.7.1 Proposed Hardware Solution

The method proposed for the receiver is to use an interpolation system to achieve finer control of the optimum signal sampling point. Due to the Harris HSP50214 hardware used, this involves decomposing the matched filter in to a fixed rate FIR filter and an interpolation filter (both present on the one VLSI chip).

We have obtained the coefficients of the interpolating polyphase filter from the Harris technical support service. This has 192 coefficients arranged so the interpolator uses 6 from 32 possible subsets. This results in a uniform symbol timing error of 1/32 per sample (1/64 per bit) which is better than the minimum found when discussing the technology loss from Gaussian bit timing errors.

The overall time domain response is shown in Figure 17:

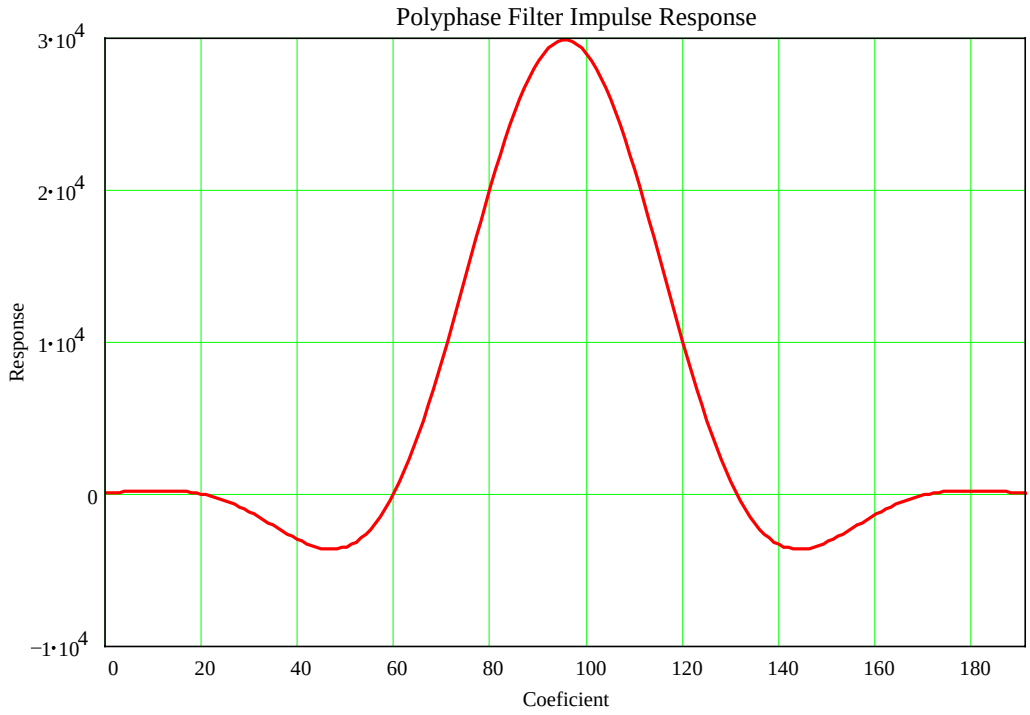

Figure 17: HSP50214 Polyphase Filter Impulse Response

1.7.2 Computation of Optimum Filter

Considering the bandwidth determining signal processing stages involved in the receiver, there is the 1st IF bandpass filter, then through HSP50214 there is the decimation filter, programmable FIR filter and finally the polyphase interpolation filter. The HSP50214 has other sections, but these are bypassed in our configuration.

If we apply the matched filter theorem of eqn 2, the optimum overall filter frequency response is:

$$H_{IF}(f)H_{CIC}(f)H_{FIR}(f)H_{PP}(f) = H_{RX}(f) = H_{TX}^{*}(f) = \sqrt{S(f)} \qquad 46$$

where H_{IF} is the response of the IF bandpass filter(s) (from the low pass equivalent model), H_{CIC} is the "Cascaded Integrator Comb" filter used for the main sample rate decimation in the HSP50214, H_{FIR} is the programmable FIR filter, and H_{PP} is the polyphase interpolating filter.

The overall objective of the filter design is to find the coefficients for H_{FIR} which result in the optimum overall receiver frequency response $H_{RX}(f)$. We can perform the required deconvolution of the filter responses in the frequency domain by:

$$H_{FIR}(f) = \frac{H_{TX}^{*}(f)}{H_{IF}(f)H_{CIC}(f)H_{PP}(f)} \qquad 47$$

The required time domain response, $h_{FIR}(t)$, (and hence the FIR coefficients) can be found from the inverse Fourier transform of $H_{FIR}(f)$, however, is not quite so simple because the FIR filter has a limited number of taps and so we must truncate $h_{FIR}(t)$ to only a small number of bit periods.

The key issue with the design of the matched filter is not just the overall frequency response, but also the overall impulse response. The important parameter is the ISI term of the overall system since the additive noise from adjacent bits may result in greater loss than the slightly sub optimum filtering of channel noise. Unfortunately it is the just frequency response which most existing filter literature, and filter CAD packages, aim to optimise.

The calculation of the loss introduced by imperfections in the matched filter can be performed by assessing the total ISI contribution at the sampling point from the past and future data bits. If the TX and RX filters are exact, and the channel linear, then this should be zero. Unfortunately, both channel errors and RX matched filter errors may be present and these give rise to ISI. The main problem in connection with the RX design is one of optimising the available filter architecture to minimise the ISI and hence the technology losses. The two practical problems facing the design are:

- The RX hardware has a limited number of FIR filter taps, this results in a truncated time domain filter impulse response and hence a non-ideal characteristic.
- The spacecraft transponder, and possibly the TX modulator or uplink amplifier, are not perfect and so the signal received will not be the ideal case of undistorted signal and additive noise.

The first problem, that of the practical filter constraints, can be addressed by a number of design and/or optimisation methods. The objective being to minimise the technology loss, a combination of frequency response errors and ISI.

That leads to the second problem, if the system is to be optimised for the real operational situation the actual received signal must be known. This is not available, however, they have been accounted for since the losses associated with the TX area have been included in the link budget by assuming a perfect RX. The practical RX losses have then been accounted for to make up the total link budget, this leads to the 0.8dB technology loss target for our design.

Although it is not possible to recover all of the degradation from the non-linear amplifier(s), it might be possible to recover some of the losses and possibly allow operation with an output level closer to saturation. Unfortunately, it is not just amplitude distortion losses that need to be considered, but the AM/PM conversion of the HPA would probably need compensation as well and this could require a complex matched filter (with the I & Q channels cross coupled). This is possible, but for the chosen

hardware (Harris HSP50214), and the MSG data rates, the on-chip FIR filter is too small in this configuration.

1.7.3 Effect of Filter ISI

If we consider the ISI target for the design, we need to analyse the impact on the BER from this effect. Again we use the same approach that was applied to the bit synchroniser timing error losses by considering the ISI power as an additive "noise" to be considered with the thermal noise from the channel. This is reasonable if the ISI waveform consisted of a number of similar independent voltages from the past and future symbols, then the central limit theorem leads to a Gaussian distribution and a root sum square addition is appropriate. This results in:

$$\frac{isi^2}{E_S} = \frac{\sum_{i\neq 0} s^2(i \cdot T_B)}{s^2(0)} \qquad 48$$

Where *isi* is the additive voltage due to the received signal from all bits other than the current bit. The ISI power is considered relative to the symbol energy E_S since this is the magnitude of the wanted signal $s^2(0)$

To account for the possibility that the ISI term is dominated by one or two large values (hence a non-Gaussian PDF), and for the Viterbi decoder that looks at a number of adjacent coded symbols to decide on the value of any particular bit, we could introduce a correction factor K_{FEC} to the additive power as before. This is likely to be in the range of 1-2 so taking an upper value of 2 should lead to a conservative design.

We define the technology loss as the increase in E_B (which is identical to the increase in E_S) required to achieve the same error rate as the perfect system could:

$$Loss = \frac{E'_S}{E_S} \qquad 49$$

Where E'_S is the input symbol energy required for the practical system. By considering the ISI as an addition to the thermal noise N_O present, we have:

$$\frac{E'_S}{N_O + K_{FEC}\frac{isi^2}{E'_S}} = \frac{E_S}{N_O} \qquad 50$$

This can be rearranged to yield the loss as:

$$Loss = \frac{1}{1 - K_{FEC} \cdot R \cdot isi^2 \frac{E_B}{N_O}} \qquad 51$$

This equation shows the dependence on Eb/No and the code rate *R*, systems that operate at low Es/No have a greater tolerance to ISI simply because they have a large amount of thermal noise present normally, and the ISI becomes a small proportion of this. This equation is plotted in the figure below for the uncoded system, the R=1/2 Viterbi decoder, and for the CCSDS system at an error probability of 5E-9, as used for the MSG.

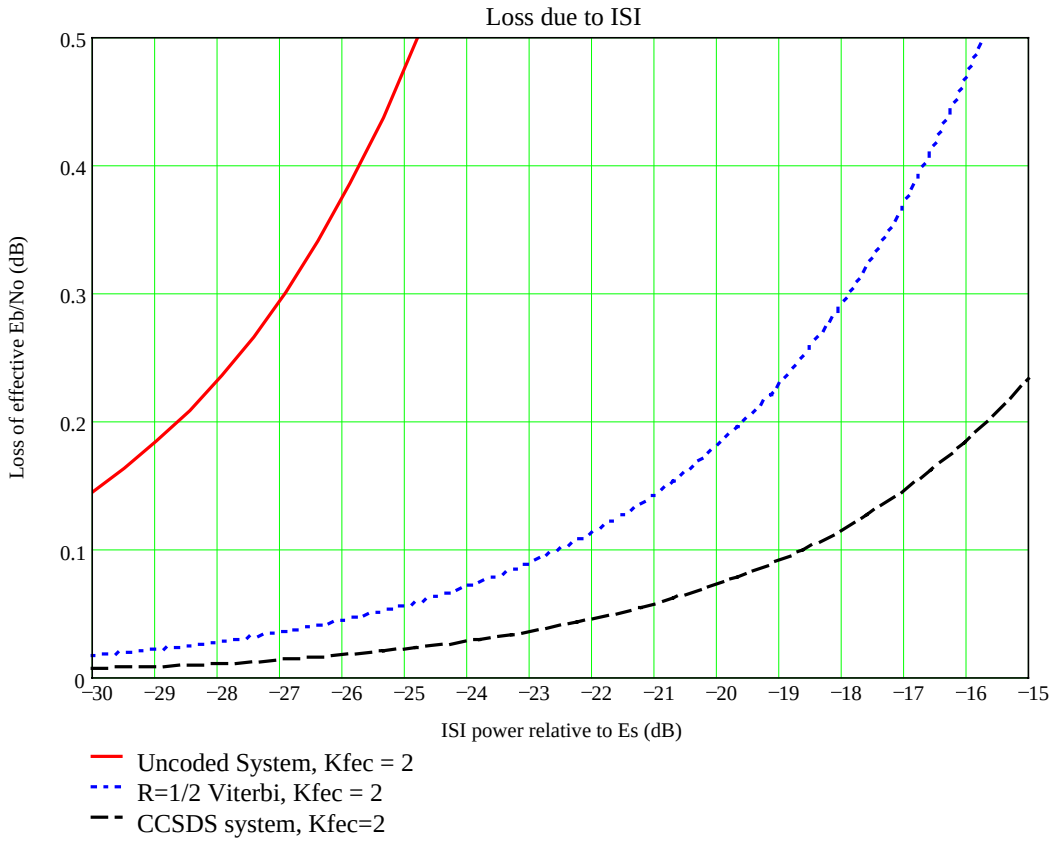

Figure 18: Effect of ISI on Technology Loss

From the overall technology loss budget, we aim for less than 0.2dB from the matched filter system. From Figure 18 this requires the ISI power to be less than –19.2dB with respect to the wanted signal. This is about 1.2% for power, or 10.9% for signal voltage, and is not an unreasonable target. To allow for any errors in accounting for the other filters in the system (principally the IF bandpass filter, which is an analogue LC design and manually adjusted for each RX to slightly different values) and for the frequency response errors it might be prudent to aim for –22dB or better (less than 0.1dB loss). As a general comment, it can be seen from Figure 18 that the concatenated CCSDS system (with an E_S/N_O of typically -0.2dB) is very useful in reducing the influence of ISI on the overall technology loss associated with the matched filter design.

1.7.4 Hardware Configuration

Achieving the ISI target requires a study of the impact of a short time domain window on the design of the RX matched filter and any optimisations possible. The length of the FIR filter is constrained by the choice of hardware. The chosen device, the Harris HSP50214 has up to 255 FIR taps but these are run by a DSP state machine. The limit is a maximum of 2 taps per processor clock cycle (35MHz in our design) if a real, even, and symmetric filter configuration is used, plus an additional two cycles (presumably to initialise the filter).

The number of times this FIR filter is evaluated depends on the choice of decimate rates and the configuration of the output polyphase filter. There are a considerable number of inter related constraints on the design and our design is based on the application note for the HSP50214 (Harris Corp., 1997b, AN9720). The chosen configurations are listed in Table 4 below. It is important to note that we have found this device to be only just suited to the HRIT data rate.

Parameter	HRIT	LRIT	Comments
A/D (Input) Sample Rate	35MHz	35MHz	CLKIN = 52MHz maximum. Need not be synchronous with the processing clock. Our choice based on external EMC and complexity considerations.
Processor Clock	35MHz	35MHz	PROCCLK = 35MHz maximum.
Signal Bandwidth	976kHz	293.9kHz	$F_{MAX} = (1+a)/(2*T_S)$
CIC Decimation Factor	5	14	R_{CIC} = CLKIN / F_{CIC} Must be integer, range 4-32.
CIC Output Rate	7MHz (10% below ideal)	2.5MHz	F_{CIC} = CLKIN / R_{CIC} Ideally be 8 * F_{MAX} for -84dB alias rejection.
Halfband decimation filters	Bypass	Bypass	Not used.
FIR Decimation Factor	2	2	$R_{FIR} = F_{CIC} / F_{FIR}$ Must be integer.
FIR Output Rate = F_{FIR}	3.5MHz (10% below ideal)	1.25MHz	Ideally 4 * F_{MAX}, maximum limited to PROCCLK / 6 = 5.833MHz
Number of FIR taps	16	52	N_{TAP} = floor(PROCCLK*R_{FIR}/F_{CIC} - R_{FIR})*2 for symmetric, even, real filters.
FIR Time response	2.286µs 2.6 symbols	20.8µs 6.1 symbols	N_{TAP} / F_{CIC}
Polyphase Decimation Factor	1.52	2.13	F_{FIR} / F_{SAMP}, non-integer adjusted for bit sync, range 1-4
Polyphase Output Rate	2.296MHz	587.8kHz	F_{SAMP} = 2 Sample / Symbol exactly.
Halfband interpolation filters	Bypass	Bypass	Not used.

Table 4: HSP50214 Configuration Parameters

From this table, the area of concern is the HRIT matched filter. Not only is the filter only 2.6 symbols in total period (±1.3 symbol), but this is also the system with the narrower bandwidth (a=0.7) hence this has longer time domain “tails” in the optimum filter response. The LRIT case, with 6.1 symbols and a roll off factor of a=1.0, should result in a near perfect filter design.

1.7.5 Computation Procedure

The computation of the FIR coefficients is performed by the MathCAD files `rcos8lrit.mcd` and `rcode8hrit.mcd` which implements the “FFT window” design method, with modifications. In principle, this works by computing the compensated FIR frequency response using eqn 47 and then using the inverse FFT to convert this to the time domain. The CIC filter response and the polyphase interpolator are taken from the Harris data.

The IF filter model is a non-causal version of the analogue IF filter frequency response, this removes the group delay to simplify the time alignment for the computation. The actual IF design has a constant group delay up to the -3dB point so this is a reasonable approximation. Compensation for any errors beyond this frequency is not required (since *H(f)* = 0 in this region), although it is possible by using a non-symmetric FIR design (unfortunately this would reduce the number of FIR taps in the HSP50214 by a factor of 2).

The resulting time domain response is then reduced to the appropriate finite number of filter taps by some “window” function. The DSP literature is full of FFT examples and window designs, however, we found that in most cases a rectangular window yielded the best overall design.

The truncated filter is converted back to the frequency domain and the resulting frequency domain error can be applied to pre-distort the frequency response for the next time domain waveform, this attempts to reverse the effect of the window on the filter, in a crude (but simple) manner. In some cases this improves the overall performance, but not always. The computation is evaluated with both options and the best chosen. A brute force optimisation based on multi-dimensional minimisation was considered, but the final design (for a=0.7 to 1.0) achieved high performance with this simple window method and the optimiser was not required.

The resulting design is compared to the ideal by the ISI computation of eqn 51 from the total system time domain response. This is computed from the inverse Fourier transform of the equivalent received signal spectrum:

$$S(f) = H_{TX}(f) \cdot \left\{ H_{IF}(f) H_{CIC}(f) H_{FIR}(f) H_{PP}(f) \right\} \qquad 52$$

where the term in the curly brackets represents $H_{RX}(f)$. The computation of the ISI term is performed by cubic spline interpolation of $s(t)$ to find the value at exactly integer points, since the FFT sampling is not on integer multiples of the symbol rate. This is slightly sub-optimum, however, it is easy to implement (in MathCAD when compared to the exact method of time shifting by phase rotation in the frequency domain and FFT converting back to the time domain for each time offset) and the magnitude of the error is insignificant in this case, where several points per symbol are computed by the FFT.

In addition to the ISI losses, there is also the loss of effective S/N from the non-ideal frequency response to be considered. The classic derivation of the matched filter theorem (for example, Cooper, 1988, p88-91) takes the frequency domain expression for the filter output SNR at the sample instant. This, with optimum sample time t=0 for simplicity, is given by:

$$SNR = \frac{\left| \int_{-\infty}^{+\infty} H_{TX}(f) H_{RX}(f) df \right|^2}{\int_{-\infty}^{+\infty} N(f) \left| H_{RX}(f) \right|^2 df} \qquad 53$$

The numerator computes $s^2(t=0)$ from the transmitted signal spectrum and the received filter frequency response. Essentially, the integral evaluates the inverse Fourier transform for time $t=0$ only. The denominator evaluates the noise power from $N(f)$, the noise power spectral density function, and the power transfer function of the receiver filter.

If the noise is white ($N(f)$ = constant), the maximisation of eqn 53 yields eqn 2 and, for real signals, in the time domain this is the well known time reversal of the transmitted signal to find the optimum impulse response for the receiver's matched filter.

By evaluating eqn 53 numerically both for the ideal case, and for the actual filter response found, the loss in effective SNR can be computed. This is added (in dB) to the loss from the ISI term to estimate the overall deviation from ideal case due to the filter design. This was computed based on the assumption that the noise is white. If interference is present, then it could be accounted for by evaluating eqn 53 with the appropriate power spectrum.

1.7.6 Results for FIR Design

1.7.6.1 HRIT Design

The result of this operation produces the filter characteristics shown. In the case of the filter coefficients, these have to be converted to the nearest integer, based on the 22-bit size used in the HSP50214. The resulting file HRIT_FIR.PRN has the 16 coefficients for the filter (these are symmetric, so only 8 are needed).

The computed loss due to ISI and frequency response errors is 0.004dB, however this takes no account of alias profile noise and quantisation errors elsewhere in the system, these are very small, however,

compared to this impressive value for loss they might be significant. A prudent analysis would simply specify the filter design as better than, say, 0.05dB for the computation of the overall technology loss performance.

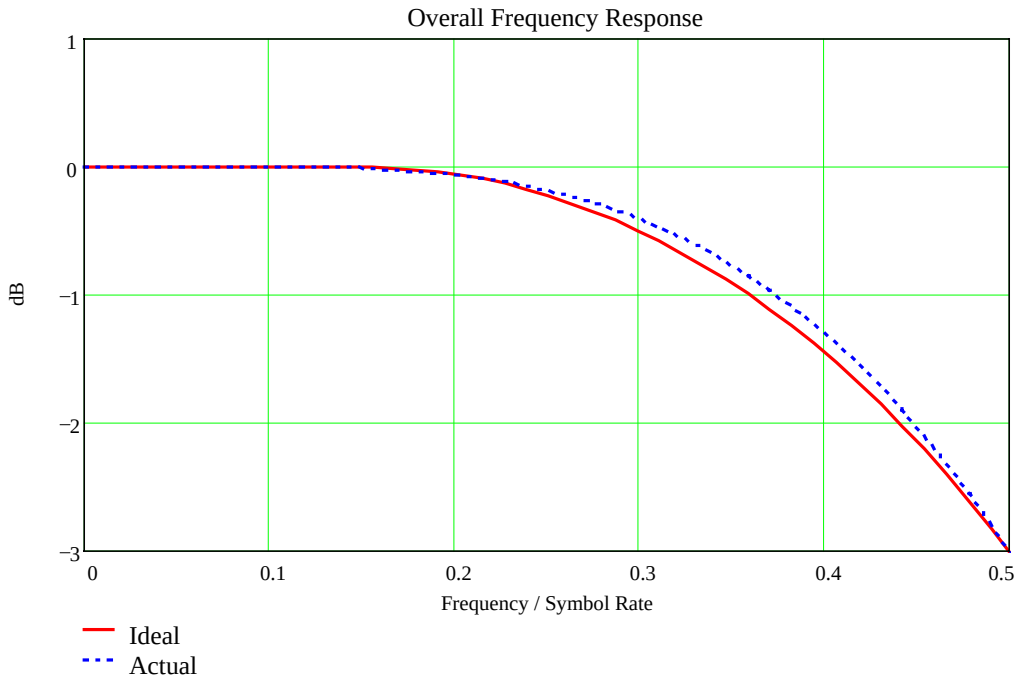

Figure 19: HRIT Matched Filter Frequency Response - Passband

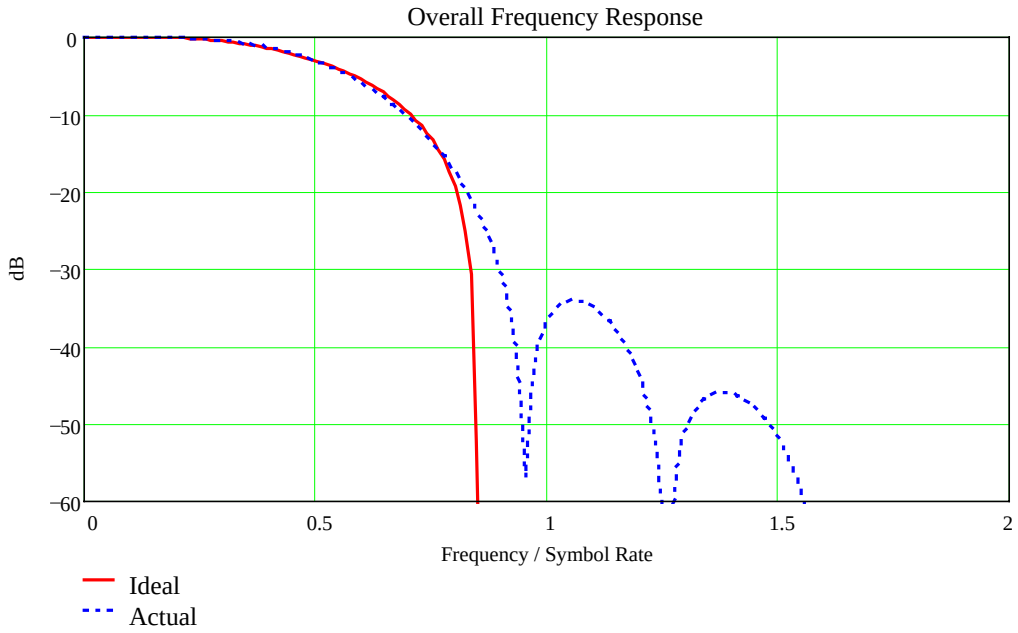

Figure 20: HRIT Matched Filter Frequency Response - Stopband

1.7.6.2 LRIT Design

The file LRIT_FIR.PRN has the 52 coefficients for the filter. The computed loss due to ISI and frequency response errors is 0.0003dB. Again, a prudent analysis would simply specify the filter design as better than 0.05dB.

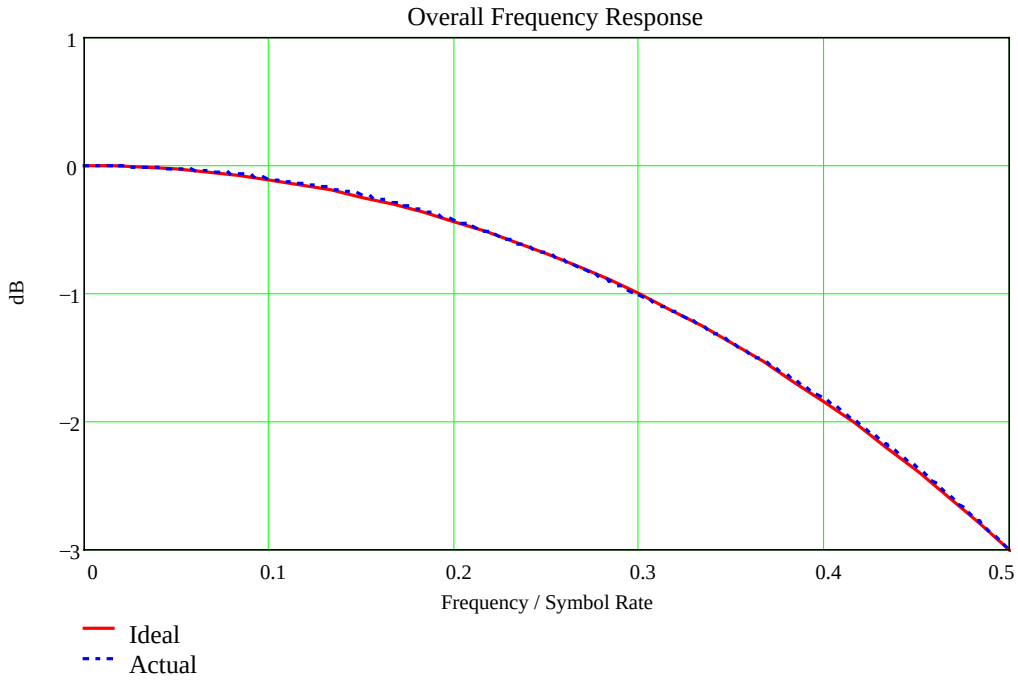

Figure 21: LRIT Matched Filter Frequency Response - Passband

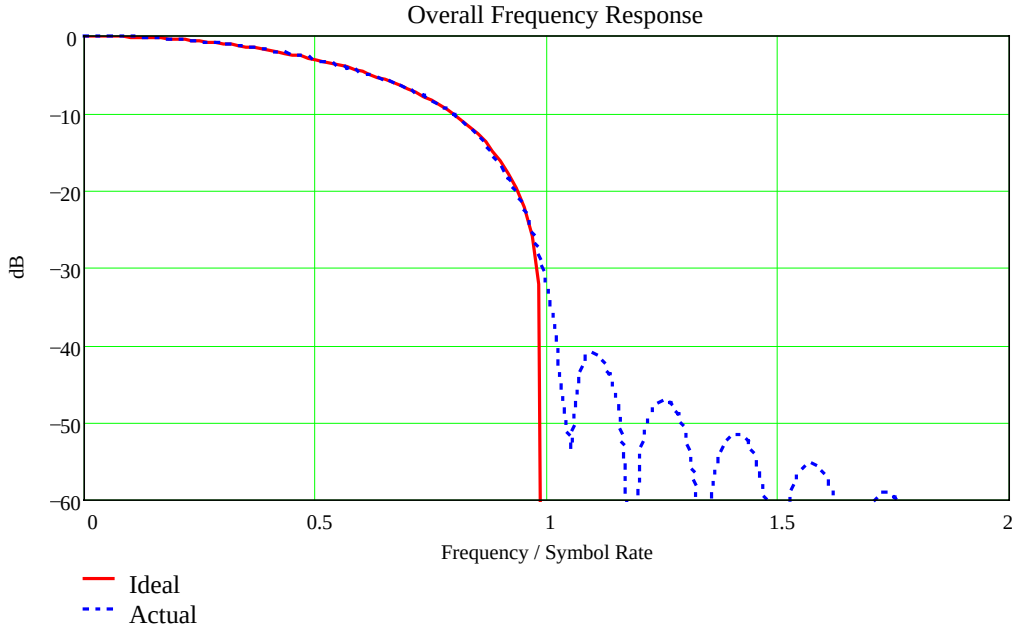

Figure 22: LRIT Matched Filter Frequency Response - Stopband

1.8 Design of the IF Bandpass Filter

The calculation of the filter performance requires the IF filter to be specified. Ideally, this filter would have a negligible effect on the data so variations in filter behaviour due to initial adjustment, temperature, ageing, etc., could be ignored.

A further consideration is the bandpass to lowpass transformation, if the signal is not placed "symmetrically" in the low pass equivalent design there will be an I-Q crosstalk degradation due to the complex time domain response. This is important since the classical filter synthesis method starts with the lowpass model and transforms to a bandpass design, then to coupled resonator values for an LC implementation. If the fractional bandwidth is too large then the resulting filter, which has geometric symmetry (frequency on a logarithmic scale), is not very symmetric in the arithmetic (lowpass, linear frequency axis) sense. Another effect is the frequency uncertainty in the downlink, this results in IF frequency centring errors and these must be minimised.

These requirements are coupled to the need for image channel and A/D aliasing rejection. A compromise is needed which can be realised in low cost electronics without excessive alignment time in production. Our design is based on a 4^{th} order 2MHz lowpass Bessel (maximally flat) model, the pole locations, frequency response and LC design method can be found in any of the textbooks on classical filter synthesis, for example (Zverev, 1967).

This filter has the pole locations (for the 1 rad s^{-1} prototype) located at:

-0.9952088 ±j1.2571058

-1.3700679 ±j0.4102497

The frequency response in our application is shown in the figure below

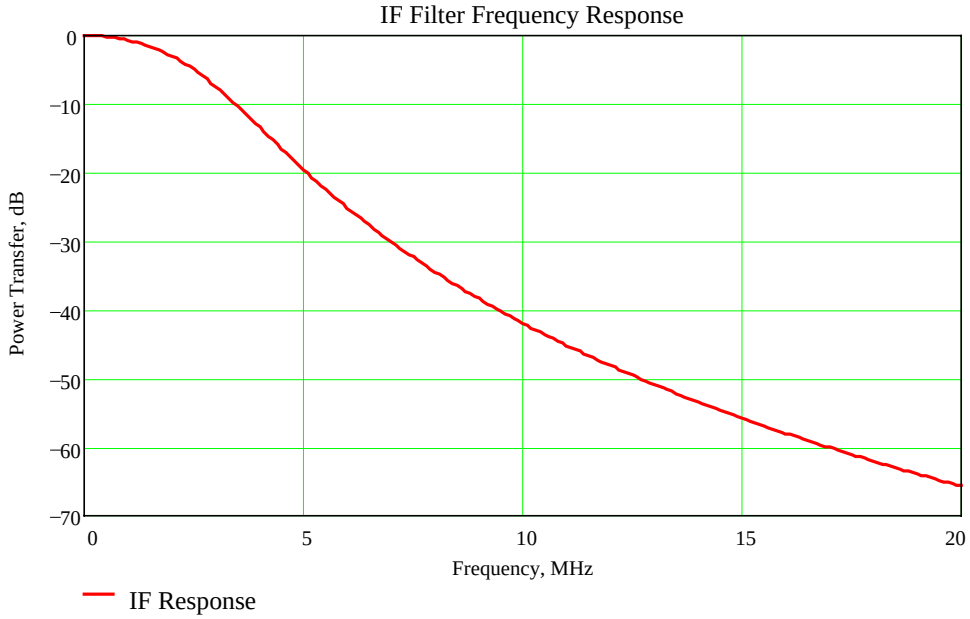

Figure 23: IF Filter Frequency Response

This lowpass model is then transformed to be centred on the 1^{st} IF frequencies. These IF frequencies are calculated based on the 1^{st} downconverter design.

The resulting filter has a 4MHz -3dB bandwidth centred on the appropriate IF. This is fixed, during the design study it was apparent that direct interchange of the IF and A/D stages between the HRUS and LRUS was of no real advantage. The use of fixed IF filters optimised for the individual channel allows for the highest performance for a particular input stage technology. The additional investment in a fully programmable wideband design was not a requirement.

The choice of the 8.75MHz 2nd IF is based on the optimum alias free range for the 35MHz sample clock. The HSP50214 can support higher sample rates, however, the hardware design is simplified by the use of a common clock for the A/D and the HSP50214 processor clock. It would be possible to use any of the odd multiples of $f_s/4$ as the 2nd IF, however, the aliasing and the image rejection properties are the same.

The use of sampling at the 1st IF to alias directly down to the required range for the NCO/Mixer of the HSP50214 was considered, however, the frequencies from the 1st IF are not well suited to the HSP50214 processor clock and the resulting A/D performance is significantly poorer (since the input bandwidth and aperture time need to be suited for 140MHz, not 9MHz operation in this case). Hence the use of a 2nd conversion stage, but with a simple lowpass filter to remove the mixer's "double frequency" output terms.

Since the HRIT DSP filter has good rejection beyond 1.5MHz, the rejection of alias or image channel is most important in this range, hence the performance that matters is the value at 2*IF-1.5 = 16MHz, and the chosen design has better than 55dB rejection at this point. A narrower filter could be used, however, the choice of 2MHz was based on a practical Q factor of 140 for the inductors used, which results in around 5dB insertion loss. In addition, this relatively wideband design avoids any significant losses if the bandwidth changes, or centre frequency shifts, by less than 10%.

It is intended to use the same basic design for HRIT and LRIT, but with minor changes to the different 1st IF centre frequencies. This allows almost identical hardware saving on development cost and effort.

The overall design result is summarised in the table below:

Parameter	HRIT	LRIT
RF Downlink	1695.15MHz	1691.00MHz
1st LO Frequency	1553.5MHz	1553.5MHz
1st IF Frequency	141.65MHz	137.5MHz
1st IF Filter Bandwidth	4MHz	4MHz
1st IF Filter Type	4th Order Bessel	4th Order Bessel
Image Rejection	> 55dB	> 55dB
2nd LO Frequency	132.9MHz	128.75MHz
2nd IF (A/D Input)	8.75MHz	8.75MHz
A/D Sample Clock Rate	35MHz	35MHz
Number of FIR taps	16	52
Total Filter ISI	-61dB	-64dB
ISI Loss	5E-6dB	2.4E-6dB
Frequency Response Loss	0.004dB	2.9E-4dB
Total Loss Assumed	< 0.05dB	< 0.05dB

Table 5: IF Frequency and Filter Design

1.9 References

CCSDS, 1989	*Advanced Orbiting Systems, Network and Data Links*	CCSDS 700.0-G-2, Green Book, Issue 2, Consultative Committee for Space Data Systems, NASA
Gardner, F. M., 1988	*Demodulator Reference Recovery Techniques suited for Digital Implementation*	ESTEC final report, contract 6847/86/NL/DG
Cooper, G. R., and McGillem, C. D., 1986	*Modern Communications and Spread Spectrum*	McGraw-Hill International, ISBN 0-07-100159-X
Meyr, H., Moeneclaey, M., and Fechtel, S. A., 1998	*Digital Communication Receivers: Synchronisation, Channel Estimation and Signal Processing*	John Wiley & Sons, New York, ISBN 0-471-50275-8
Verdin, D., and Tozer, T. C., 1995	*Symbol Timing Recovery for M-PSK Modulation Schemes Using the Signum Function*	New synchronisation techniques for radio systems, colloquium, London, 27 November, 1995, p2/1-2/7 Institution Of Electrical Engineers
Harris Semiconductors Corporation, 1997a	*HSP50214 Programmable Downconverter*	Harris Semiconductors Corporation, Data Sheet file number 4266.1, June 1997. http://www.semi.harris.com/data/fn/fn4/fn4266/fn4266.pdf
Harris Semiconductors Corporation, 1997b	*Calculating Maximum Processing Rates of the PDC (HSP50214)*	Harris Semiconductors Corporation, Application Note AN9720, June 1997. http://www.semi.harris.com/data/an/an9/an9720/an9720.pdf
Zverev, A. I., 1967	*Handbook of Filter Synthesis*	John Wiley & Sons, New York. ISBN 0-471-98680-1