Smart Radio Using Multirate Filters to Process, Identify, and Reconfigure Itself to Receive and Demodulated Multi-Standard IRIG Signals

Robert Wade Lowdermilk
Signum Concepts
San Diego, California

Dragan Vuletic
Signum Concepts
San Diego, California

fred harris
CUBIC Signal Processing Chair
San Diego State University
fred.harris@sdsu.edu

ABSTRACT:

We present a detailed description of a DSP based radio that determines which of the multiple IRIG standard formats it is receiving and then reconfigures itself to demodulate the identified modulation format. The IRIG standard describes a set of serial time code formats for distributing time-of-day used in test ranges to time tag collected data. One of the modulation formats amplitude modulates a sine wave centered at one of 5-distinct carrier frequencies. Typical data rates and center frequencies are 100 pps on a 1.0 kHz carrier (Format B), 1 kpps on a 10.0 kHz carrier (Format A), and 10 kpps on a 100.0 kHz carrier (Format G). The radio we describe scans these frequencies, probing the observed signals to determine which band is occupied and then configures itself to the appropriate center frequency, bandwidth, and data rate.

1. INTRODUCTION

IRIG (Inter Range Instrumentation Group) Standard 200-95 defines a set of serial time code formats and modulation details for distributing time of day information. Time of day, including days, hours, minutes, seconds, tenth of seconds and hundredth of seconds, is standard time tag appended to collected data at various test ranges, ground tracking networks, and spacecraft and missile projects for data processing and reduction.

The time code contains 78 or 74 binary symbols encoded in Binary Coded Decimal (BCD) format. The Symbols are formatted by a set of rectangular pulses amplitude modulated on a sinusoidal carrier. The carrier frequency is an integer multiple of the modulating signal frequency. In Formats A, B, and G, this ratio is 10:1. The Pulse period and the carrier frequency are synchronized so that the leading edge of the rectangle pulse coincides with positive zero crossings of the carrier.

Three pulse types are defined for formats A, B, and G. These pulses are a Binary Zero "0", a Binary One "1", and a Position Marker "P". The Position Marker is used to separate the BCD Symbols. Due to the synchronization of the pulse edges and the carrier, the pulse types can be identified by the number of cycles per pulse segment. The binary zero is high for 2-cycles and is low for 8-cycles, the binary one is high for 5-cycles and is low for 5-cycles, and the position marker is high for 8-cycles and is low for 2-cycles. A detailed representation of these pulses is shown in figure-1. As seen, the pulses are non-return to zero with the ratio of high level to low level nominally set at 10-to-3 with acceptable ratios in the range 3-to-1 through 6-to-1.

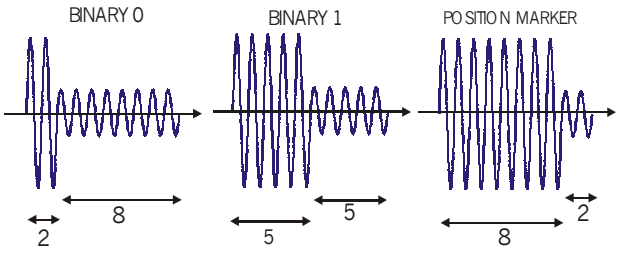

Figure 1. Modulated Pulses: IRIG Formats A, B, and G

The center frequencies and modulation rates for IRIG formats A, B, and G are shown below in table 1.

TABLE-1 Parameters for IRIG Formats

IRIG Format	Frame Rate	Carrier Freq	Bit Rate
A	10 per sec	10 kHz	1.0 kpps
B	1 per sec	1kHz	0.1 kpps
G	100 per sec	100 kHz	10 kpps

Earlier IRIG receiver equipment required the operator to manually select the format option the receiver was expected to demodulate. The system we describe in this paper uses a DSP engine to learn the signal parameters of the specific input signal by conducting a probe of the possible center frequencies. The probe conducts a hypothesis test using a multirate band-pass filter centered at the set of possible center frequencies, with bandwidths matching the corresponding data modulation bandwidth. In this mode the receiver searches a selected band for energy and if detected directs the receiver to a demodulation routine and if absent proceeds to test the next option.

Since the center frequencies and data rates of the options being examined are related by factors of 10, a single 10-to-1 down sampling filter is applied sequentially to reduce the bandwidth and sample rate and thus permit the entire search to be conducted at the same fixed ratio of sample rate to bandwidth. Following design guidelines formed in other multirate applications we use

the resampling operations to translate, by aliasing, successive center frequencies to a common processing frequency. Further, we use the resampling process to aid in the timing recovery process by assigning phase control of the resampling operation to the timing acquisition loop. The loop advances and retards the phase of the down sample counting clock to align the rising edge pulse boundary with the output sample clock.

2. MULTIRATE FILTERING AND TRANSLATION

The center frequencies of the formats to be processed are 100 kHz, 10 kHz and 1 kHz. The sampling rate for the receiver was selected to be 800 kHz so that the timing recovery loop would have 8-samples per cycle of the highest frequency for ease of timing alignment. The timing resolution requirements actually required 16-samples per cycle and rather than operate the ADC at the higher rate we raise the sample rate with a 2-to-1 interpolator. Other combinations of over sampling and interpolating filters were considered and rejected in favor of the 800 kHz sample rate.

The first processing task is the hypothesis test of IRIG format G, the 100 kHz center frequency signal. A complex band pass filter centered at 100 kHz, 1/8-th of the input sample rate, processes the signal. The output of this filter is down sampled 10-to-1 to obtain a sample rate of 80 kHz, which is a lower sample rate than the 100 kHz center frequency. The center frequency is thus aliased to 100 kHz –80 kHz or 20 kHz, which is one quarter of the 80 kHz output sample rate. Spectra at the quarter sample rate can be trivially down converted to baseband since there are no multiplications required to apply the heterodyne exp(j n π/2). The spectral response of the filtering and resampling operation is shown in figure 1.

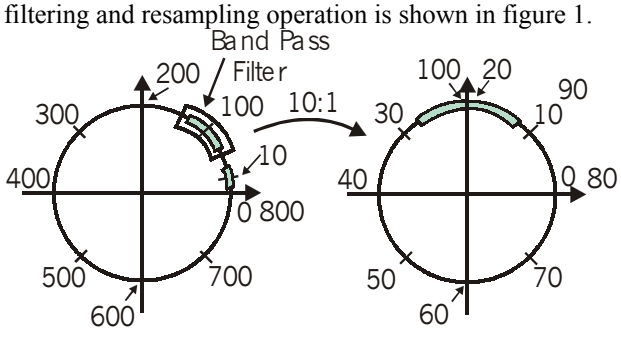

Figure-1. Spectral Position of Filter (100 kHz)at Input Sample Rate (800 kHz) and Spectral position of Filter at Re-Sampled Output Rate (80 kHz).

The band pass filter accomplishes the filtering and the 10-to-1 resampling simultaneously by being configured as a 10-stage polyphase filter. In this configuration, the order of filtering and down sampling is reversed so that the filtering operation occurs at the output rate rather than at the input rate. This interchanged, valid under conditions known as the "Noble Identity", is illustrated in figure 2. The Noble Identity recognizes that M-units of delay at the input clock are the same as 1-unit of delay at the output clock after an M-to-1 down sample operation. Applying the Noble Identity to cast a filter followed by a down sampler to a sum of sub filters operating at the output rate rather than at the input rate is shown in equations 1-4.

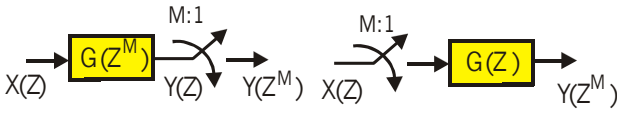

Figure 2. Noble Identity Filter and Down Sample Equivalent to Down Sample and Filter

$$H(Z) = \sum_{n=0}^{N-1} h(n)\, Z^{-n} \tag{1}$$

$$H(Z) = \sum_{r=0}^{M} \sum_{n=0}^{\frac{N}{M}-1} h(r+nM) Z^{-(r+nM)} \tag{2}$$

$$H(Z) = \sum_{r=0}^{M} Z^{-r} \sum_{n=0}^{\frac{N}{M}-1} h(r+nM) Z^{-nM} \tag{3}$$

$$H(Z) = \sum_{r=0}^{M} Z^{-r} H_r(Z^M) \qquad \text{where} \quad H_r(Z^M) = \sum_{n=0}^{\frac{N}{M}-1} h(r+10n) Z^{-nM} \tag{4}$$

Figure 3 presents the structure formed by applying the noble identity to the down sampling band pass filter. Here G(Z) is the Z-transform of band pass version of the filter H(Z) with g(n)=h(n)exp(j (0.1) 2π n). The Noble identity can be applied to pass band filters as well as base band filters.

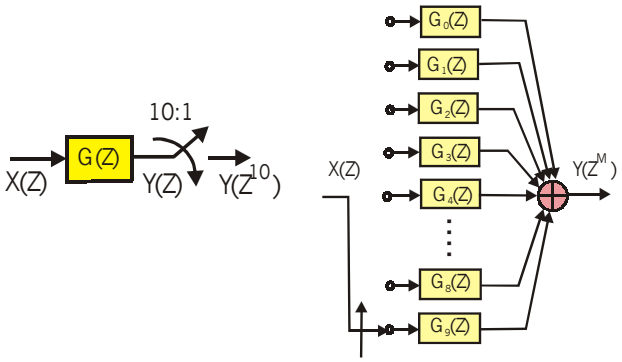

Figure-3 Partition of Band Pass Filter and 10-to-1 Down Sampler into 10-Stage Polyphase Filter

The energy content at the output of the band pass filter is tested for presence of a signal. The sample by sample energy is obtained by squaring the real signal, or by the sum of squares of the down converted and filtered base band signal, or by the sum of squares of the analytic signal formed by a Hilbert Transform filter. The sample energy is averaged in a leaky integrator to obtain a smoothed estimate of the signal energy. The sample energy, the averager output, is compared with a threshold. If the energy level exceeds the threshold, the signal is passed on to the demodulation and timing acquisition routines. If the energy level does not cross the threshold, processing of the next frequency band is initiated.

2. BASEBAND MULTIRATE FILTERING

If the energy detector test for IRIG-G declares that there is no signal present the search moves forward to the next option, which specifically is a test for the IRIG-A signal. The IRIG-A signal center frequency is 10 kHz, one-tenth of the IRIG-G center frequency. If we reduce the bandwidth by a factor of ten and simultaneously reduce the sample rate 10-to-1 the IRIG-A center frequency relocates to one-eighth of the new sample rate, a condition matching the ratio of input sample rate to the IRIG-G center frequency. Processing for the IRIG-A signal is then identical to that for the IRIG-G signal, but is conducted at the reduced sample rate. The effect on the input spectrum of the low pass filtering and the 10-to-1 down sampling is shown if figure 4. We see this effect is similar to the effect demonstrated in figure 1 for the band pass filtering and 10-to-1 down sampling. The primary difference here is that the frequency translation from fs/80 to fs/8 does not experience a wrap-around from the second to the first Nyquist zone as did the band pass spectrum of figure 1.

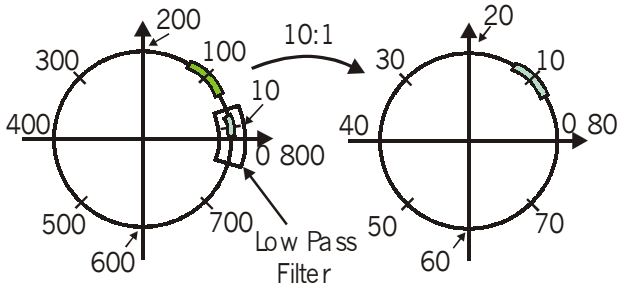

Figure4. Spectral Position of Low Pass Filter (10 kHz) at Input Sample Rate (800 kHz) and Spectral position of Filter at Re-Sampled Output Rate (80 kHz).

The low pass filtering and 10-to-1 down sampling is seen to be a pre-processor that moves the next IRIG frequency band to the proper spectral location to be processed by the IRIG detector functions. The spectral shift is accomplished by the simple expediency of filtering and down sampling. The signal processing flow required to examine the three IRIG frequency bands is shown in figure 5.

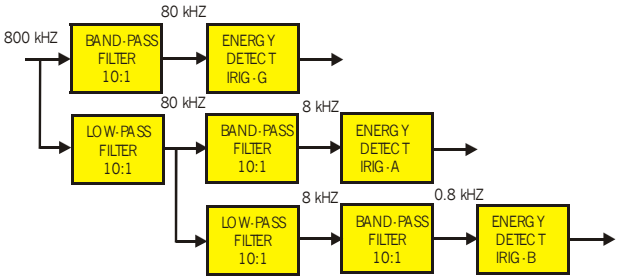

Figure 5. Processing Blocks Comprising 10-to-1 Band Pass Filters and 10-to-1 Low Pass Filters for Spectral Isolation and Spectral Translation

We note that the two filtering operations suggested by figures 1 and 4 actually use the same prototype baseband filter weights that are heterodyned to the one-eighth center frequency for use as a band pass filter and are used directly as a low pass filter. In a third generation design, both filtering task are performed simultaneously in a single eight-stage polyphase filter that down samples 10-to-1. The multiple aliases are separated with a set of post filter phase-rotators as opposed to using separate filters with the phase rotators embedded in the filter weights. The signal processing flow required to examine the three IRIG frequency bands in the third generation design is shown in figure 6.

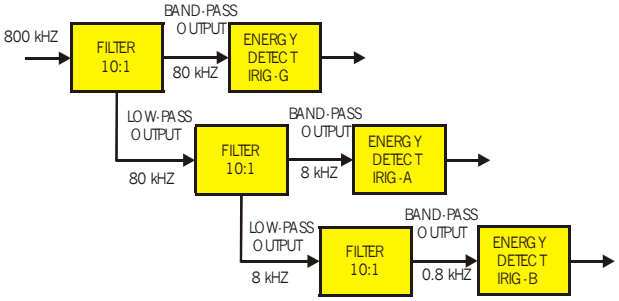

Figure 6. Third Generation Processing Blocks Comprising 8-stage Polyphase filter with 10-to-1 Re-Sampler to Simultaneously form Outputs from Low-Pass and Band Pass Filters

3. TIMING SYNCHRONIZATION

Part of the signal-processing task in the receiver is to align the processing clock with the bit boundaries of the modulation signal. We noted in the introduction that the modulation signal is locked to the frequency and phase of the sinusoidal carrier. The key consideration is that the leading edge of the modulation pulse coincides with a

positive going zeros crossing of the carrier signal and that there are ten cycles of the carrier wave per modulation pulse. The selected sample rate was 800 kHz so that the IRIR-G format, at center frequency of 100 kHz, has 8-samples per cycle. Timing specifications required that the discrete signal samples be aligned with the carrier zero crossings within 1.0 µsec. The time interval between 800 kHz input samples is 1.25 µsec so that the maximum time offset between a zero crossing and the nearest time sample is 0.625 µsec. As a safety margin we used a 1-to-2 interpolator increase our timing precision to (1/4)*1.25 µsec or 0.3125 µsec. The interpolator is used to output data samples at twice the input sample rate, but is used in a 2-to-1 down sample mode in which it either outputs the original input samples or new interpolated samples residing midway between the input sample positions. The timing relationship between the carrier, the ADC samples of the carrier, and the 2-to-1 interpolated samples is shown in figure 7. In this figure we see that the sample closest to the positive going zero crossing is the interpolated sample immediately to the right of that zero crossing.

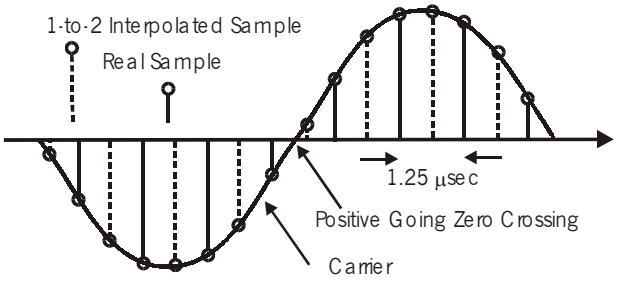

Figure 7. Carrier, Positive Going Zero Crossing, Real Samples, and 1-to-2 Interpolated Samples

The receiver identifies which of the sixteen possible samples per sine wave cycle is closest to the positive going zero crossing . It does this in two passes, first examining the eight ADC samples and then the eight interpolated samples. Eight successive sine wave samples are commutated to eight averaging filters to determine the average value of each sample location. The filter bank is scanned to find two successive filters with a sign change, indicating the positive going zero crossing and the filter with the smallest average is the best approximation to the location of the positive going zero crossing. The process is repeated for both output phases of the 1-to-2 interpolator to determine in which samples set the sample closest to the positive going zero crossing resides. A block diagram representing the timing recovery process for IRIG-G is shown in figure 8.

The timing recovery process for IRIG-A and IRIG-B differs from the process for IRIG-G. This is because the IRIG-A and IRIG-B options have access to a finer resolution timing increment through phase control of the

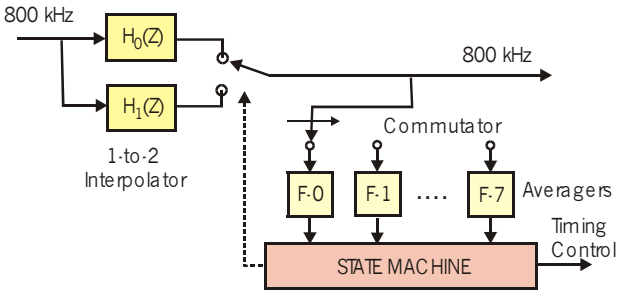

Figure 8. Block Diagram of Timing Recovery Process for IRIG-G to Determine Input Sample Closest to Positive Going Zero Crossing

10-to-1 low pass down sampling filters. The effective granularity of the IRIG-A sample clock is 160 samples per cycle rather than the 16 samples per cycle of the IRIG-G option. The 16-to-1 resolution is accessible in the IRIG-A in the same manner that we operated the IRIR-G timing recovery process. Initially we use the commutated eight averaging filters and the 1-to-2 interpolator to access the mid sample points. Finer resolution is obtained by slipping phase one count at the input to the polyphase 10-to-1 low-pass filter. The phase-slip occurs by inserting a no-op in the 10-port commutator distributing samples to the polyphase filter bank. The IRIG-B option uses two 10-to-1 down sample filters to realize an effective granularity of its clock of 1600 samples per cycle. The timing recovery for the IRIG-B accesses this granularity by performing the same phase slippage at the input to the first of the pair of low-pass 10-to-1 polyphase filters.

A block diagram representing the timing recovery process for IRIG-B is shown in figure 9.

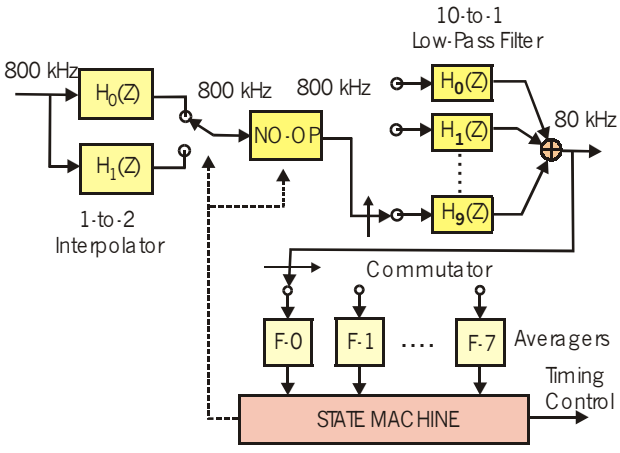

Figure 9. Block Diagram of Timing Recovery Process for IRIG-A to Determine Output Sample Closest to Positive Going Zero Crossing

4. SAMPLED DATA CARRIER ALIGNMENT

The ADC sampling clock is set to a value of 800 kHz, a value nominally 8 times the IRIG-G carrier, 80 times the IRIG-A carrier, and 800 times the IRIG-B carrier. Due to oscillator tolerances at the modulator and at the demodulator, the frequency of the sampling clock will exhibit small offsets that have to be removed by VCO control of the local oscillator. An alternate scheme is to use a free running sampling clock and interpolate from the sample times offered by the clock to the sample times synchronous with the received carrier. We elected, for this design, to modify the clock by controlling the VCO control voltage. The block diagram of the systems required to control the frequency of the sample clock to place the incoming carrier at the quarter sample rate is shown in figure 10.

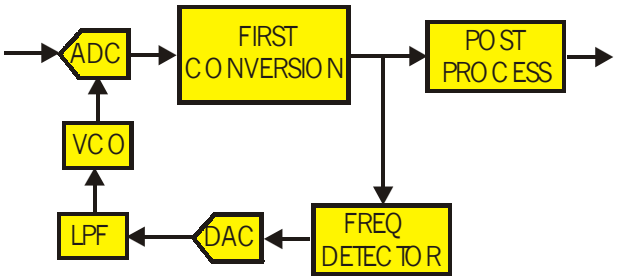

Figure 10 Control Loop to Adjust Sample Frequency so that First Conversion Carrier Resides at Quarter Sample Rate

The task of forming an error proportional to the frequency offset is performed as the standard imaginary part of the conjugate cross product of the current I-Q pair and the previous I-Q pair of the down converted carrier as shown in figure 11. simple.

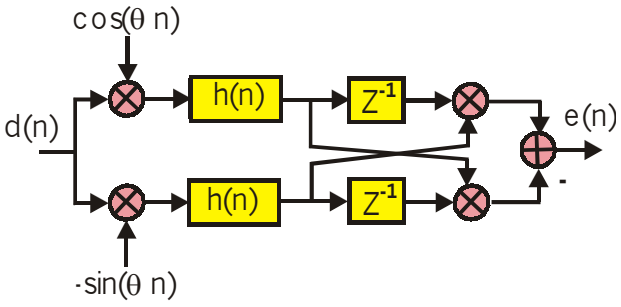

Figure 11. Block Diagram of Frequency Error Detector for Carrier Recovery Loop.

The processing required for the frequency error detector is particularly simple because the center frequency of the processed carrier is at the quarter sample rate and the cosine and sine heterodynes to down convert the signal can be performed without multiplies. Similarly, the low-pass filters following the down conversion can be CIC (cascade integrators and comb) filters or re-sampled versions known as the Hogenauer filter. The simplified block diagram of the frequency error detector is shown in figure 12.

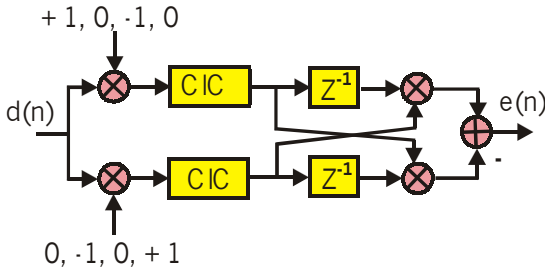

Figure 12 Frequency Error Detector: Trivial Down Converter, Simple Filter, and Imaginary Part of Conjugate Cross Product

The CIC filter is equivalent to a cascade of one or more, boxcar integrators but implemented by manipulation of the boxcar's Z-transform. The Z-transform is partitioned and realized in separate cascade segments. The p-stage versions of the boxcar and CIC have the same Z-transform shown in equation 5. The also have the same spectral response, the periodic extension of sin(x)/x, the Dirichlet kernel shown in equation 6.

$$H(Z) = \left[\frac{(1 - Z^{-M})}{(1 - Z^{-1})}\right]^{P} \qquad (5)$$

$$H(\theta) = \left[\frac{\sin(M\frac{\theta}{2})}{\sin(\frac{\theta}{2})}\right]^{P} \qquad (6)$$

Both the boxcar and the CIC filters require M-address locations for data storage. The attraction of these filters, over standard FIR filters is that they are multiply free. The advantage of the CIC over the boxcar is the dramatic reduction in additions required to implement the filter. An M-tap boxcar requires M-additions while an M-tap CIC requires 2 additions. The block diagram of the two-stage, M-tap CIC filter is shown in figure 13.

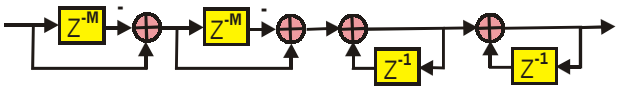

Figure 13. Two Stage, M-Tap, Cascade-Integrator-Comb, (CIC) Filter

5- CONCLUSIONS

We have presented important segments of an intelligent radio receiver that demodulates a subset of IRIG Format signals. The signal set we examined, IRIG-G, IRIG-A, and IRIG-B have a common structure, which allowed us to design the receiver for the highest frequency signal and then by simple signal processing manipulations scale and slide lower frequency bands into the processing stream designed for the high-band. In the processing of the high band we used spectral aliasing to move the signal band center from one digital frequency $f_1/8$ to a second digital frequency $f_2/4$. In this case, the resampling violated the Nyquist criterion for the center frequency but not for the bandwidth. Consequently, the spectrum in the second Nyquist wrapped into the first Nyquist zone.

The processing to move the next spectral band into the previously processed spectral band also involved resampling and spectral aliasing. Here the signal at center frequency $f_1/80$ is aliased to $f_2/8$ by a 10-to-1 down sampling. For this case there is not a transfer of spectrum from one Nyquist zone to another.

Resampling was also invoked in the timing recovery modes of the receiver. A 2-to-1 interpolating filter is used in the high band (IRIG-G) to decrease the timing granularity as we locate the data sample closest to the positive going zero crossings of the carrier frequency. In the lower bands, the 2-to-1 resampling filter is augmented by phase control of the commutator delivering samples to the 10-to-1 down sampling low-pass filter.

Resampling is an option we discussed in the implementation of the low-pass CIC filter used in the frequency error detector that guided the Costas loop adjustment of the sampling clock. Another resampling options we considered for this receiver was replacement of the Costas loop used in analog frequency control with a set of digital interpolators to effect arbitrary resampling of fixed sample rate data from input clock positions to output samples synchronized to carrier frequency and phase. Since the carrier and the modulation signal are phase and frequency locked, this option would replace both the carrier alignment and the timing recovery processes.

The common thread in the design of this receiver is the use of multirate processing to move spectrum by aliasing, to reduce workload by operating filters at reduced output rate rather than higher input rate, and to obtain higher resolution time increments by move time samples with M-path interpolators.

6-REFERENCES

[1] fred .j. harris and Chris Dick, "Performing Simultaneous Arbitrary Spectral Translation and Sample Rate Change, in Polyphase Interpolating or Decimating Filters in Transmitters and Receivers", 2002-Software Defined Radio Technical Conference, San Diego, 11-12 November 2002.

[2] Steve Scott and fred harris, "Flexibility, Performance and Implementation Advantages of Recursive, Linear and Non Linear Phase, Polyphase Filters in Transmitters and Receivers", 2002-Software Defined Radio Technical Conference, San Diego, 11-12 November 2002.

[3] fred harris, "A Fresh View of Digital Signal Processing for Software Defined Radios, Part I" and "Part II". International Telemetering Conference, San Diego, CA, 21-24, October 2002

[4] Michael Rice and fred harris, "Polyphase Filter Banks for Symbol Timing Synchronization in Sampled Data Receivers", Milcom-2002, Anaheim, CA, 7-10 October 2002

[5] fred harris, "On Multichannel Receivers with Arbitrary, Uncoupled Selection of Channel Bandwidth, Channel Spacing, and Channel Sample Rate", 2002 International Symposium on Advanced Radio Technologies (ISART), Mar. 4-6, 2002, Boulder Colorado

[6] Michael Rice and fred harris, "Multirate Digital Filters for Symbol Timing Synchronization in Software Defined Radios", IEEE Journal on Selected Areas in Communications, Signal Synchronization in Digital Transmission Systems, December 2001, Vol. 19, No.12, pp.2346-2357.

[7] E. B. Hogenauer, "An Economical Class of Digital Filter for Decimation and Interpolation", IEEE Trans., Acoust., Siignal Processing, vol. ASSP-29, pp. 155-162, April 1981.

[8] Heinrich Meter, Marc Moeneclaey, and Stefan Fechtel, "Digital Communication Receivers", John Wiley & Sons, 1997.

[9] Michael Rice, Chris Dick, and fred harris, "Maximum Likelihood Carrier Phase Synchronization In FPGA-Based Software Defined Radios", ICASSP-2001 (International Conference on Acoustics, Speech and Signal Processing) Salt Lake City, Utah, 7-11 May 2001