DESIGN OF RUGGED AND RECONFIGURABLE DIGITAL RECEIVERS

John Walker MSc BSC CEng MIEE (Pentland Systems Ltd., Livingston, Midlothian, UK; john.walker@pentlandsys.com)

ABSTRACT

The digital receiver is a crucial component in the Software Radio architecture. It must meet the various requirements of reconfigurability, modularity, flexibility and survivability to be suitable for use in rugged Software Radio applications. Pentland Systems have designed the RAD-2 to meet these requirements and this paper describes the main design choices made and experiences gained in the development process. The paper also covers the challenges of conduction-cooled design and the techniques that can be used to achieve maximum system performance with a limited power budget.

1. INTRODUCTION

A topic seldom covered in software radio forums is the design of rugged products specifically for use in harsh environments such as those encountered in jet aircraft, Unmanned Air Vehicles (UAVs), and all-terrain vehicles. The task of ruggedizing products is multi-faceted and various specialist techniques must be used to achieve maximum performance out of the system in the presence of limited power budgets and high shock and vibration.

Pentland Systems is a world-leader in the design and manufacture of signal acquisition hardware for use in sensor processing systems. We are focused on developing key enabling products for the emerging Software Radio market and exploiting our substantial experience in rugged product design.

Using the design of the RAD-2 Digital Receiver as a basis, this paper describes the challenges involved in design of a rugged, reconfigurable digital receiver for software radio applications.

2. DIGITAL RECEIVERS IN SOFTWARE RADIO

In Software Defined Radio (SDR), the digital receiver performs digitization at some point after the antenna; typically after wideband filtering, low noise amplification and down conversion to an Intermediate Frequency (IF). A proportion of the radio processing is carried out in the digital domain by software or firmware programmable elements. The transmit path can be considered the reciprocal of the received path, with the addition of a duplexer such that high-power transmissions are unable to damage the receiver subsystem.

In an ideal Software Radio (SR), the digitization takes place at, or very near, the antenna with all the radio processing performed by firmware and software in the digital domain.

Regardless of the extent of software penetration in the radio, the Digital Receiver must fulfil the following fundamental requirements:

- Reconfigurability – The functionality of the digital receiver can be reconfigured to implement new operational modes and to adapt to variable operational characteristics. Examples include selection of a different band or signal spectrum and frequency agility for security or to enhance performance.
- Flexibility – The hardware developed must be flexible enough to satisfy a wide range of operational requirements. For instance, a wide-band, linear front-end on the digital receiver is essential to maximize its usability in multiple applications.
- Survivability – The product must be able to survive demanding operational environments, characterized by wide temperature variances, high shock and vibration, high humidity and limited power budgets. These environments are common in platforms such as fast jets, UAVs and all-terrain vehicles.
- Modularity – The digital receiver component can easily be replaced or upgraded without affecting the remainder of the system. Robust logical and physical interface definitions are essential for this and the PMC specification provides an extremely good starting point to achieve such modularity.

3. DESIGN CONSIDERATIONS

Various key design choices affect a digital receiver's ability to fulfil these high-level requirements. This section covers the options in analog-to-digital conversion, sampling schemes, reconfigurable hardware and rugged product design available to the digital receiver designer. Details of the decisions taken in the design of the RAD-2 receiver are described.

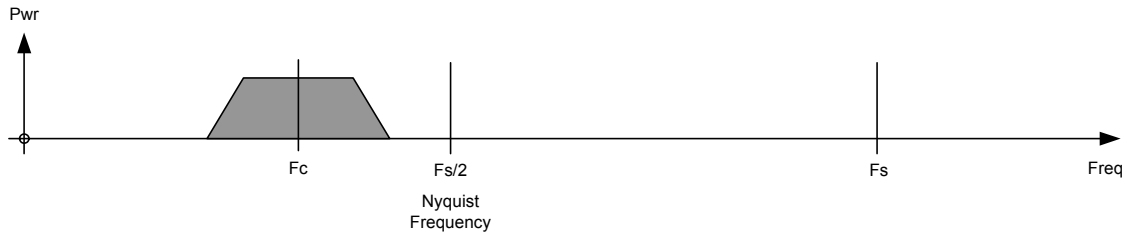

Figure 1 - Oversampling Methodology

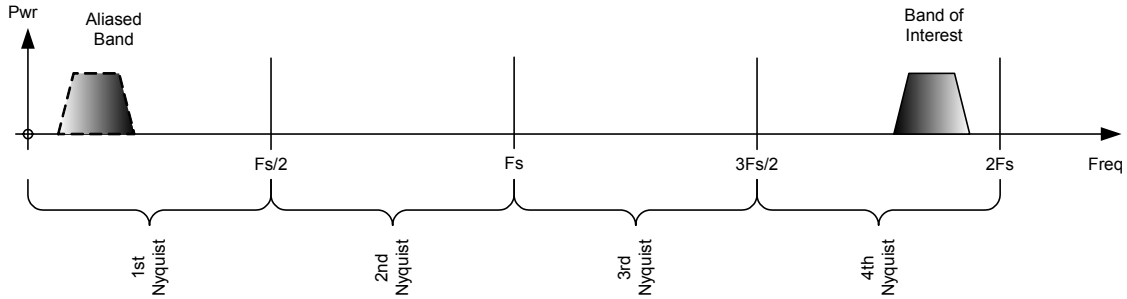

Figure 2 - Undersampling Methodology

3.1. Analog to Digital Conversion

What is the best choice of converter implementation for Software Radio? It depends on the application requirements, on the level of flexibility desired by the system architect and on the restrictions imposed by the operational environment. See [1] for details on data conversion approaches in Software Radios.

In the RAD-2, the decision was taken to use the AD6645 converter [2], which offers 14-bits of resolution with a maximum sampling rate of 105 MSPS. As well as high SNR for oversampling, this converter is also ideal for undersampling applications, as it achieves in excess of 100dB multi-tone Spurious Free Dynamic Range (SFDR) in the second Nyquist band. See section 3.2 of this paper for a detailed description of oversampling and undersampling.

Whilst this paper concentrates chiefly on the receive path, the complexities involved in the transmit path should not be dismissed. The main challenge in high-performance radio transmit applications is the reconstruction of the analog signal from the digital equivalent. Digital interpolation filters, high-speed DACs and analog reconstruction filters are the fundamental components required to reconstruct the analog signal.

Further information on A/D conversion architectures can be found in [3].

3.2. Sampling Schemes

In essence, there are two possible options for signal acquisition in a digital receiver: oversampling and undersampling. In oversampling, the sampling frequency is at least twice that of the highest frequency in the band-of-interest. In undersampling, use is made of the 'aliasing' effect of discrete sampling to acquire signals much higher than the sampling frequency. In undersampling, the sampling frequency must be at least twice the bandwidth of the band-of-interest.

3.2.1.Oversampling

Given an IF or carrier frequency that is less than half the maximum sampling rate of the ADC, an oversampling method can be adopted to digitize the signal bandwidth. In this method, the sampling rate adopted must be at least twice the highest frequency to be acquired:

$$Sampling\_Frequency \geq 2 \times Max\_Frequency \qquad (1)$$

Oversampling systems generally sample at 2½ to 3 times the maximum signal frequency, which relaxes the requirement on the anti-aliasing filter. Tight anti-alias specifications can lead to distortion of the upper frequencies in the band-of-interest and can be extremely costly to design and manufacture. They are also susceptible to temperature variance. Figure 1 depicts the oversampling method.

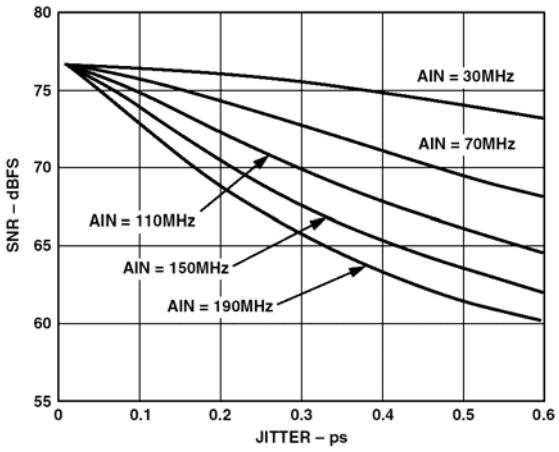

Figure 3 – SNR vs. Aperture Jitter

3.2.2. Undersampling

With undersampling, the principle is to bandlimit the desired frequency range prior to conversion and utilize the 'folding' phenomenon of the Sampling Theorem to digitize the band-of-interest. In undersampling, the restriction is that the sampling rate must be at least twice the bandwidth of the signal:

$$Sampling\_Frequency \geq 2 \times Bandwidth \qquad (2)$$

The same recommendation as for oversampling applies here to relax the specification on anti-aliasing filters. Frequencies in excess of the Nyquist Frequency are 'folded' down into the 0 to Fs/2 range according to the mapping in the table below. When choosing the sampling frequency for an oversampling system, the designer must be careful to avoid spurious aliasing signals from interfering channels.

Nyquist Region	Frequency Band	Aliased Band	Orientation
1st	0 to Fs/2	0 to Fs/2	Correct
2nd	Fs/2 to Fs	Fs/2 to 0	Flipped
3rd	Fs to 3Fs/2	0 to Fs/2	Correct
4th	3Fs/2 to 2Fs	Fs/2 to 0	Flipped
…	…	…	…

Note that in odd-numbered Nyquist regions the frequency orientation is 'flipped'. In Figure 2, the theoretical aliasing of the band-of-interest in the 4th Nyquist region down to the 1st region is shown.

Using undersampling, the RAD-2 can acquire VHF frequencies with a maximum sampling rate of only 105MHz. For instance, given a signal with center frequency 130MHz and bandwidth 10MHz, a sampling rate of 75MHz will down-convert the signal band to central frequency 20MHz with an inverted 10MHz bandwidth. Front-end anti-aliasing filters would only have to select the 4th Nyquist region prior to digitization to avoid aliasing. Post-digitization filtering and down-conversion would allow baseband processing of the signal.

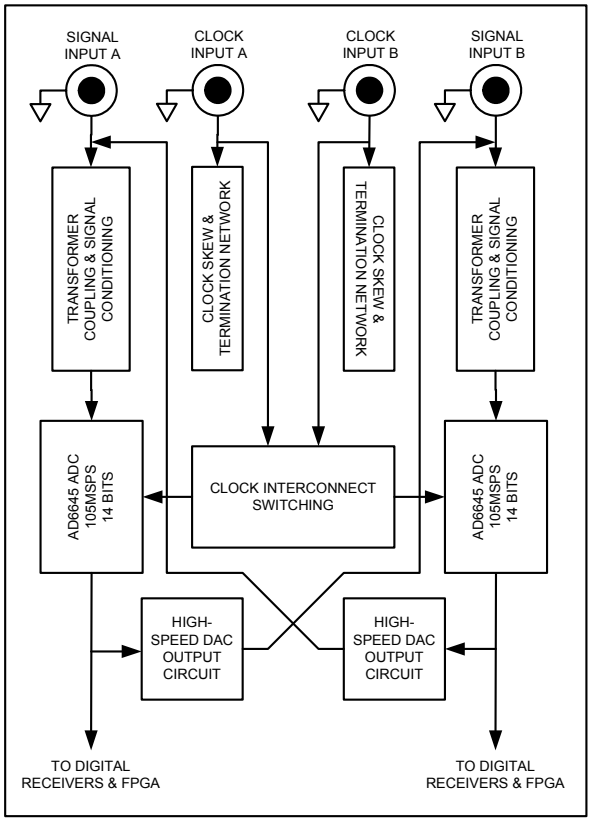

Figure 4 – RAD-2 Front-End Architecture

3.3. Performance Specifications

3.3.1 Aperture Jitter

Aperture jitter is caused by sample-to-sample variations in the clock source. The effect of this jitter is directly translated into aperture error, which is an error in the converted signal level. This effect is magnified significantly at higher analogue frequencies, such as those experienced at the IF of a radio receiver, which have a consequently higher slew rate.

The effect of aperture jitter on SNR can be seen from Eq. 3 where F_A is the analog input frequency and t_j is the aperture jitter in seconds (for brevity, terms for average Differential Non-Linearity (DNL) error and thermal noise have not been included). The graph in Figure 3 shows the

resulting effect on the SNR performance for various analog input frequencies.

$$SNR = 1.76 - 20\log(2\pi F_A t_j) \qquad (3)$$

Purity of the clock source is, therefore, absolutely essential in direct conversion systems where the receiver acquires data directly from the IF carrier. Further information on Aperture Jitter can be found in [7].

3.3.2 Dithering

Dithering is an extremely powerful aid for increasing the SFDR performance of the radio digital receiver. Dithering is the insertion of pseudo-random noise into the analog input prior to signal conversion. Repeated errors (such as spurious signals) reduce the SFDR performance and the pseudo-random noise randomizes these errors by turning them into wideband noise. The calculation of required dither signal power is based on the ADC resolution and number of ADC conversion stages. Careful choice of the dither signal power is important to maximize the increase in SFDR while minimizing the impact to SNR. Figures 5 and 6 show the effect of dithering on spurious signals.

The RAD-2 includes two AD9744 D/A converters, as shown in Figure 4, which can be used to inject dithering signals into the analog signal prior to A/D conversion.

These DACs can also be used to enhance the resolution of the ADC process. By injecting a known pseudo-random noise stream into the input signal prior to conversion the probability of signal levels occurring at quantization boundaries is reduced (such signal levels increase the noise floor). The injected noise is then removed digitally afterwards, resulting in higher SNR and effective resolution. Previous experience at Pentland of this technique has shown a 12-bit converter can provide the effective resolution of a 16-bit converter.

Note that the DACs can also be used to generate test tones for injection into the analog inputs to test the RAD-2 subsystem. The test tones can be selected by digital down-conversion and the acquisition performance verified in-system. This capability is extremely important in high-integrity systems where constant system availability is necessary.

4. ASPECTS OF RUGGED DESIGN

"Ruggedizing" the digital receiver to survive tough environmental conditions presents unique challenges: heat, vibration and humidity are the three major contributors to electronic failure, and the remedies for these are often conflicting.

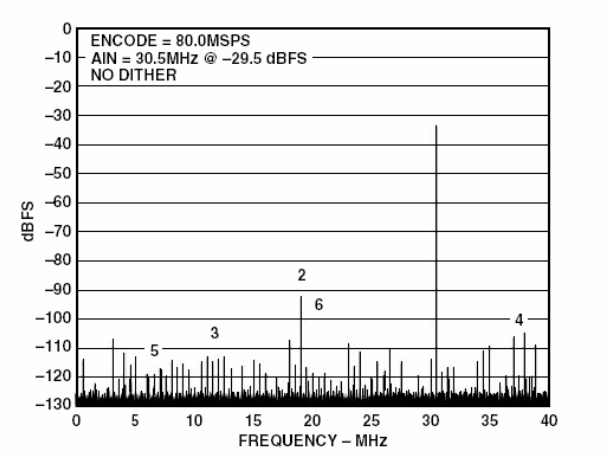

Figure 5 – 1M FFT without Dithering

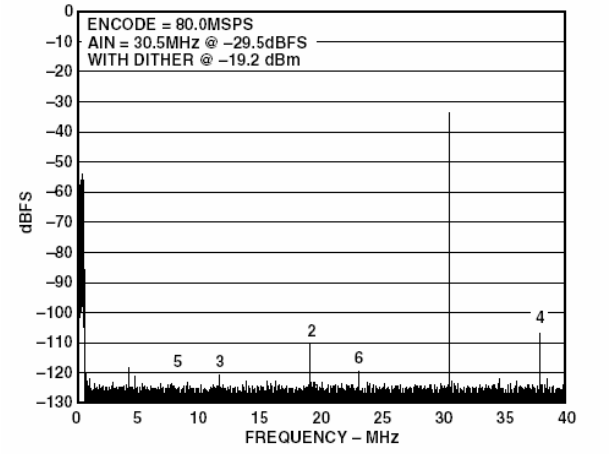

Figure 6 – 1M FFT with Dithering

4.1. Cooling

Electronic components are not 100% efficient in energy conversion; heat is generated by active components whilst in operation. Assuming that the heat source remains constant, the temperature within the component will continue to rise until the rate of heat being generated is equal to the rate of heat flowing away from the component. In restricted environments, the system architect is sometimes unable to use natural or forced-air convection to cool the electronics and must rely on conduction-cooling, where the heat is drawn out through the mechanical assembly.

The ability of a material to conduct heat depends on the physical properties of the material and on the geometry of the atomic structure. Eq. 4 gives the basic equation for heat flow where Q is the power dissipation, K the thermal conductivity, A the cross-sectional area, ΔT the temperature difference and L the length of flow.

Figure 7 – Air Transportable Racks

$$Q = KA\frac{\Delta T}{L} \qquad (4)$$

In light of Eq. 4, high power consumption components are situated close to the PCB edges (the heat transport interfaces) to minimize heat build-up in the PCB. Metallic PCB cores (thin layers of high thermal conductivity metals such as copper) are also used to transport heat away from the components. The PCB layers are mounted either side of the metallic core and thermal vias (through-holes) act as conduits for the heat from the component to the central core.

Difficulties occur when adjacent materials have different Thermal Coefficients of Expansion (TCEs). Consider components mounted on a PCB; when heated the adjacent materials will expand by different amounts, resulting in stresses on the component leads and possible fracture. Consequently, the designer must choose component and PCB materials with as similar TCEs as possible. However, as this is not always possible, the complementary route is to keep the temperature as low as possible which minimizes the stress on component legs. Underfill technology is particularly suitable for this purpose. Underfilling is designed to make flip chip assembly tough and reliable by filling the gap between the die and substrate with an epoxy or other adhesive, as shown in Figure 8. This gives three advantages: increased mechanical strength, reduced effects from moisture or other contaminants, and increased thermal conductivity. For fully rugged environments, BGAs on the RAD-2 digital receiver will be underfilled to protect the unit from the high stresses in shock/vibration and thermal expansion.

However, in some components, such as the Xilinx FPGA, the chip is attached to the top of the package, rendering underfill ineffective as a heat transfer channel. In the RAD-2, a top-mounted heat sink is used to extract heat from the top of the component to high-capacity thermal vias on the PCB surface.

A common method with larger form-factor PCBs, such as VME cards, is to machine a solid metal 'cover',

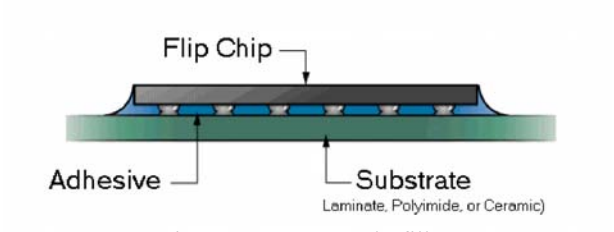

Figure 8 – BGA Underfill

which is designed to match the component profiles. The intervening space between the component and metal cover is then filled with a thermal transfer material to increase the thermal conductivity, thereby keeping the component at a lower temperature.

Conduction-cooled boards are mounted within Air Transportable Racks (ATRs) for use in the final platforms. For examples of ATRs see Figure 7 above and for further information see [6]. To achieve a good thermal contact with the rack assembly in an ATR, wedge-lock retainers on the card edges are used to fix the cards into place. These retainers exert a high pressure onto the rack slot, which results in increased thermal conductivity and also provides a high retention force to secure the card units under high levels of shock and vibration.

4.2. Conduction-Cooled PMCs

The RAD-2 has been designed compliant with the Conduction-Cooled PMC standard [4]. This standard shortens the normal PMC mezzanine PCB length to 143.75mm so the PCB can fit on a Conduction-Cooled VME baseboard [5]. This form-factor is a widely adopted standard in the Commercial Off-The-Shelf industry, where adherence to a standard physical specification results in a product usable in multiple applications.

However, the CC-PMC specification assumes that all connections are made through the four connectors located at the bottom of the card (Figure 9) and does not include scope for front-panel connectors. In the design of the RAD-2, this shortcoming of the specification was recognized as it severely affects signal routing on the PCB. While no standard approach to circumventing this problem has been developed, the RAD-2 can accommodate either a cut-out in the front panel of the base board, or inversion of the connectors to allow signal routing from the side or back of the rack. To aid the development of rugged digital receivers for Software Radio applications, the ANSI/VITA standards should be amended to allow for RF signal connections, either through the rear of the rack or through the front.

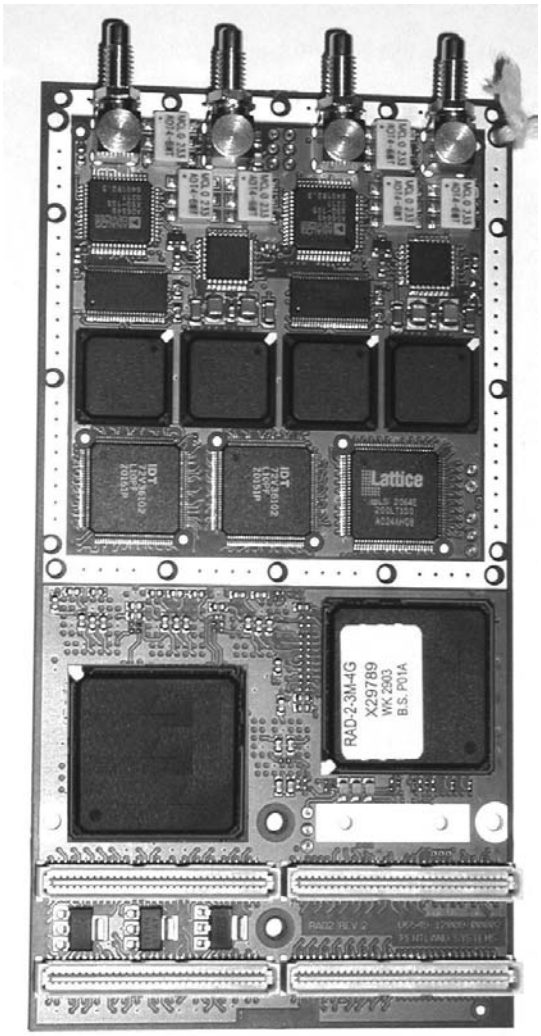

Figure 9 – RAD-2 PMC Module

5. CONCLUSIONS

The role of the digital receiver in a Software Radio system has been described and the four high-level requirements of reconfigurability, flexibility, modularity and survivability outlined. A summary of the challenges encountered in the design of the RAD-2 digital receiver has also been given. It should be noted that the complexity of these design choices is substantial and this paper is meant to serve only as a brief introduction.

The lack of front-panel connections in the CC-PMC standard was identified as a serious shortcoming and a recommendation has been made that the problem be addressed by an amendment to the standard and/or related standards.

6. REFERENCES

[1] B. Brannon et al, "Data Conversion in Software Defined Radios", *Software Defined Radio*, Part I, pp99-125, 2002

[2] *AD6645 Datasheet*, Analog Devices Inc., http://www.analog.com/

[3] N. Kularatna, "Understanding A/D Converter Architectures", *IEE Electronic Systems & Software*, Vol 1, Issue 4, pp20-23, Aug/Sep 2003

[4] *American National Standard for Conduction Cooled PMC, ANSI/VITA 20-2001*, Approved August 31 2001, VMEbus International Trade Association

[5] *IEEE 1101.2 Standard for Conduction-Cooled VME*, IEEE

[6] T. Perazzo, "Cooling Ruggedized Systems by Conduction", Application Note, *APW Electronic Solutions*

[7] *AN501 Aperture Uncertainty and ADC System Performance*, Analog Devices Inc., http://www.analog.com/

(Figures 3, 5 and 6 reproduced courtesy of Analog Devices Inc.)