32nd URSI GASS, Montreal, 19-26 August 2017

Analogue Design for Dynamic, Broadband Receivers

Antheun R. Botha(1), Jason Manley(1), Paul S. v d Merwe(1) and Abraham J. Otto(1)
(1) MESA Product Solutions, Stellenbosch, RSA, 7600, http://www.mesaproducts.co.za/

Abstract

RF measurement systems increasingly utilise the power of high-speed ADC technologies that offer broad instantaneous bandwidths while maintaining high dynamic-range performance. This paper evaluates the typical building blocks comprising a direct-sampling RF-receiver and provides an overview of the design considerations necessary to maintain analogue system performance.

1. Introduction

The electromagnetic-interference (EMI) environment is a constantly evolving sphere that continues to present a proving ground for radio-frequency (RF) instrumentation. The Square Kilometer Array (SKA), the largest and most sensitive radio-telescope project ever planned, has implemented its Field-Programmable Gate Array (FPGA) platforms in spectrum-monitoring stations. The Reconfigurable Open-Architecture Computing Hardware (ROACH)-2 board has been the cornerstone of recent MESA RF-receiver research and development, with the latest version being dubbed the Real-Time Transient Analyser (RTA-3) (Figure 1) [1]. A variety of studies have been conducted with the RTAs, including shielding-effectiveness [2], environmental-impact and propagation measurements [3].

Further RTA developments are underway to put the next-generation SKA-SA FPGA processing platform, the SKA Reconfigurable Application Board. (SKARAB), into service [4]. A SKARAB-based receiver will improve the system specifications and requires a re-evaluation of previous designs. It is the combination of a large instantaneous bandwidth, high-speed processing and high dynamic range (DR) that results in more accurate and efficient characterisation of unknown EMI signals. We outline the analogue-design considerations necessary to achieve these goals.

Figure 1. The RTA3. Photo courtesy of MESA Product Solutions.

2. Benefits of Higher Dynamic Range

A significant advantage of performing real-time, broadband measurements is that it enables characterising dynamic signals in both the Time (TD) and Frequency Domains (FD). This provides expedited insight into transient events that are typically obscured by time consuming conventional scanning devices that systematically outline the spectral contents of event sequences. In contrast, real-time captures represent unique events in both domains, and enables the extraction of statistical information. Transient signals may have low duty cycles and low spectral power densities, with energy spread over a wide bandwidth. However, a large instantaneous bandwidth allows the accurate rise-time capture of a single TD event. In the FD, averaging and dithering are common tools used to improve the DR. Together, they reduce or decrease the noise floor variance. But dithering distorts the TD envelope and is detrimental to the analysis of unknown transients. They are only useful tools where time-averaging is possible, i.e. non-transient signals.

3. Analogue Architecture and Dynamic Range

The RTA3 analogue-front-end layout is shown in Figure 2, featuring a low noise amplifier (LNA) at the input, followed by a switched-filter bank, level-setting and the ADC. This direct-sampling configuration features a 6-way filter selection, and each path is referred to as a frequency band. With direct-sampling, the filtered frequency band is sampled at the received RF. This layout was selected for its simplicity, since the ADC's analogue bandwidth was sufficient and it serves as an introductory example. Although the general design methodology discussed here also applies to topologies featuring frequency mixing, the additional mixing effects, such as intermodular distortion, will not be discussed.

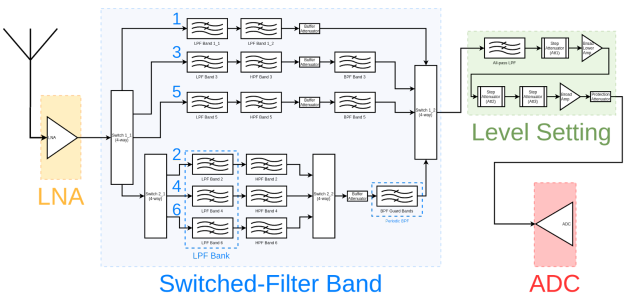

Figure 2. A direct-sampling front-end layout showing an LNA, 6-way filter selection, adjustable gain stage and an ADC.

Central to a real-time receiver is a high-speed ADC. In most cases, there is limited freedom in the selection of high-speed ADCs and it should define system performance in terms of the bandwidth and DR capabilities. So the signal-conditioning front-end (and also the digital-signal-processing back-end) must then be designed to avoid degrading this performance.

We will distinguish between three dynamic ranges (DR): 1) instantaneous, 2) input and 3) spurious-free DR (SFDR). Figure 3 illustrates the DR definitions described below.

We define the instantaneous DR to be the span between the maximum ADC input-power level and the noise floor. Practically, this is a measure of the difference between the smallest and largest signals that can be measured simultaneously. It is primarily determined by the resolution of the ADC, which defines the theoretical signal to quantisation noise ratio (SQNR) in decibels (dB).

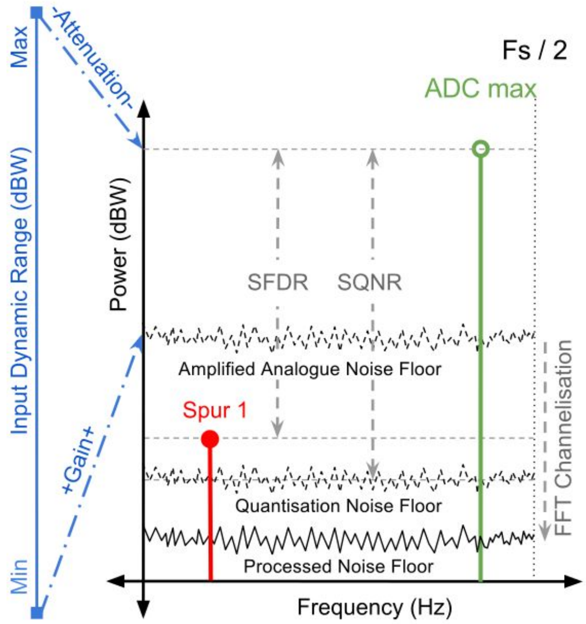

Figure 3. An illustration of the DR definitions present at the ADC. *ADC max* is the upper power limit for an uncompressed CW-tone input and *Spur 1* is the highest spurious product in the system. Note that the output noise from the analogue front end may be higher than the quantisation level with Spur 1 remaining observable above the processed-noise level. Spur 1 could be an input-related, aliased harmonic or self-generated RFI.

The input signal can be moved up or down, into the instantaneous-DR window, by adding attenuation or amplification to the signal (Fig. 3). In this way, we obtain the input DR: the allowable signal power levels that the system may observe without output compression, i.e. the range of power levels that the system can accurately measure. This is larger than the instantaneous DR.

The lowest detectable signal level is limited by the analogue noise floor, which is dependent on the system (T_{sys}) and ambient (T_{amb}) noise temperatures combined. The latter is inherent, whilst T_{sys} is dominated by the noise temperature of the LNA at the high-gain settings used during sensitive measurements. Equation 1 defines the analogue noise-floor level of the system at an input-power level [5]. The upper limit of the input DR is limited by and several dB below the LNA input 1 dB (IP1dB) compression limit, the maximum input-power level with acceptable accuracy.

$$P_{noise} = kT_{meas}B \quad (1)$$

$$T_{meas} = T_{sys} + T_{amb}$$

Note that the average noise-power level predicted by Eq. 1 is lowered due to the channelisation of the Fast Fourier Transform (FFT). The FFT effectively reduces the noise bandwidth (BW) from the instantaneous, analogue bandwidth of a frequency band to the frequency resolution of processed spectra. Equation 2 defines the decrease in the noise floor level shown in Fig. 3.

$$-10 \times log_{10}\left(\frac{BW_{Analogue}}{BW_{Processed}}\right) \quad (2)$$

The SFDR is lower than SQNR of the ADC and is provided by manufacturers as a metric for devices. SFDR is typically dominated by non-linearities in the response of the physical device. These unwanted effects may produce harmonic responses in the FD that are well above the processed noise floor. The harmonics may also extend well beyond the Nyquist band of interest, and alias repeatedly around the sampling frequency.

Practically, the SFDR is the largest range in CW-tone powers that can be unambiguously identified without being obscured by false components in output spectra. In general, the ADC's spurious response is complex and varies with frequency, amplitude and between Nyquist zones, making it difficult to characterise and calibrate. This motivates why it is essential not to introduce additional spurious components within the ADC SFDR.

4. Design Considerations and Limitations

The LNA is critical to maintain the sensitivity of the receiver. It is thus difficult to protect in sensitive systems, or those where a large input DR is required, and is often the most exposed component in the instrument.

Further, in band-selectable systems such as these, the LNA, out-of-band rejection filters and input RF switches must cope with the full Input DR over the full bandwidth at all times, while components after the filters are able to leverage the rejection offered by preceding parts and need only support the instantaneous DR in the band of interest.

Fortunately, antennas used on broadband systems are not likely to be efficient in the lower frequency ranges, where EMI is often ubiquitous. Nevertheless, many bands host an array of noisy transmissions (eg the UHF, VHF) while others may be relatively quiet (eg Ku-band). A monitoring system must support the coverage of both occupied and quieter spectral ranges.

The filtering section rejects out-of-band signals to avoid having unwanted spectral components aliasing into the sampled frequency range (i.e. passband). Both the out-of-band rejection and isolation specifications would ideally ensure that if a signal level is at the maximum input level in one frequency band, this level is reduced to below the minimum level of a second band observing smaller signals.

Figure 4 illustrates an example where Band 2 out-of-band rejection is not sufficient. This allows the maximum-level tone in Band 1 to leak into Band 2 as interference, aliasing into the SFDR region of the latter passband. The additional rejection required to avoid this is also shown, indicating that the rejection should ideally match the input DR. This also applies to non-reflective switch isolation.

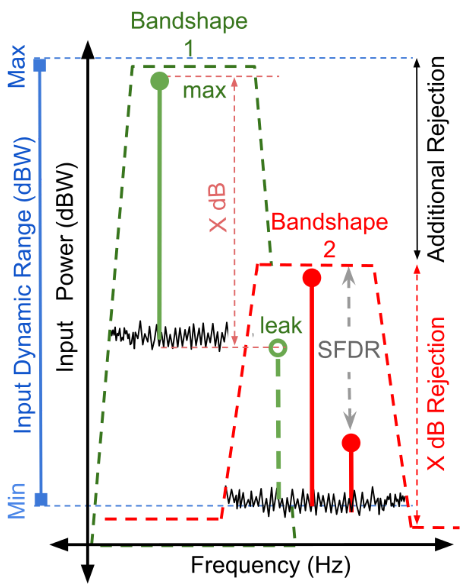

Figure 4. An illustration of adjacent frequency-band measurements showing bandshape responses at calibrated (input) levels and unwanted leakage limiting the SFDR.

A very stringent out-of-band rejection specification can reduce the realisable bandwidth, will increase filter complexity and is likely to widen the roll-off margins required to avoid SFDR limitations due to aliasing. In order to maintain realistic filter and switch specifications, the input-DR may need to be reduced. If the input DR exceeds the rejection and isolation, a protocol may be put in place to track the largest out-of-band spectral components present during measurements and warn the user against DR limitations.

The total gain and attenuation must be specified to allow optimal sampling. The ideal is to use the entire instantaneous DR to maximise system performance. However, it is important to consider that the noise levels described in section 2 represent average-power estimates. Noise voltage is a stochastic, stationary signal with amplitudes defined by a bell-curved normal distribution. Given enough time, this random signal will exceed the linear-amplitude ranges of the analogue components or ADC.

The gain distribution throughout the front-end must ensure that the occurrence of non-linear voltage levels is very low and has a negligible effect on the accuracy of measurements. As an example, the maximum gain for the front-end in Fig. 2 amplifies the analogue noise floor to the centre of the instantaneous DR. For a 12-bit ADC, there remains a ~30 dB margin between the noise floor and maximum input-power levels. The occurrence of overrange in this case is negligible, whilst small signals can still be accurately characterised.

For a system implementing adjustable attenuation (amplification) for level-setting, an attenuation (gain) mapping is required. This mapping describes the division of attenuation (gain) among the adjustable components. Firstly, the headroom must be maximised, which is the margin between input noise floor and the lowest IP1dB. This can be done with an algorithm incorporating cascaded headroom calculations (Fig. 4). The available headroom is also paramount to the SFDR performance and amplifiers must be sufficiently buffered. This will also maintain a safe margin from the noise floor.

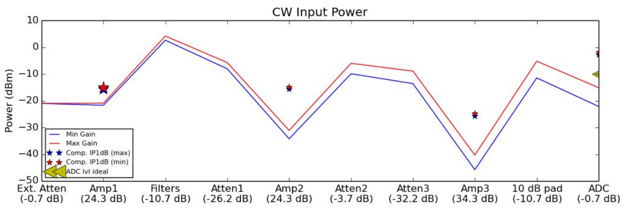

Figure 4. Plot of cascaded RF power generated by an attenuation-mapping algorithm. Note that the system gain is neutralised with balanced attenuation.

The assignment of attenuation will start from the attenuators closest to the ADC, buffering the final amplifier. (Gain reduction starts at the first amplifier after the LNA.) It is good practice to add enough attenuation to allow the gain of each amplifier to be neutralised. This retains full control of the headroom throughout the signal path. Excess attenuation is also useful to protect the system in unknown environments.

5. Conclusions

A broadband, high-DR receiver is capable of performing measurement in a unique fashion, especially when compared to scanning systems. The analogue design of a broadband receiver originates with the selection of a high-performance ADC. The DR performance of the sampler determines the required total gain and attenuation, as well as filter and switch characteristics. The properties of both analogue-signal levels and the digitised representation of processed data must be considered when interpreting the system response. There are several guidelines for assigning gain, attenuation and filter specifications that simplify the base design. Although most complex parameter variations are ignored, this allows the designer to get an overview of the potential performance characteristics and requirements for a broadband receiver front end.

6. References

1. MESA Product Solutions, [Online]: http://www.mesaproducts.co.za/, Last accessed January 2017

2. J. A. Andriambeloson, H. C. Reader, A. R. Botha, "Time-domain nested-enclosure technique for shielding effectiveness characterisation", proceedings for Electromagnetics in Advanced Applications (ICEAA) 2013 International Conference on, pp. 901-904, 2013.

3. A. J. Otto, R. P. Millenaar and P. S. van der Merwe, "Characterising RFI for the Square Kilometre Array", IEEE proceedings of Radio Frequency Interference (RFI), Socorro, New Mexico, U.S.A, 17-20 October 2016, pp.81-85, DOI: 10.1109/RFINT.2016.7833536

4. R. P. Millenaar and A. J. Otto, "Innovations in instrumentation for RFI monitoring", IEEE proceedings of Radio Frequency Interference (RFI), Socorro, New Mexico, U.S.A, 17-20 October 2016, pp. 65-68, DOI: 10.1109/RFINT.2016.7833533

5. David M. Pozar, "Transmission lines and microwave networks," *Microwave and RF design of wireless systems.*, 1st ed., New York: Wiley, 2001, pp. 89–91.