DEPLOYMENT OF RE-ENTRANT WAVEFORMS

John D. Bard, Ph.D. (Space Coast Communication Systems, Inc., Melbourne FL, USA; jbard@spacecoastcomm.com); Sean Doyle (Space Coast Communication Systems, Inc; sdoyle@spacecoastcomm.com)

ABSTRACT

Software defined radio capitalizes on the notion that a single radio front-end supports many channels of receive/transmit. A software radio thus has a one-to-many mapping between a radio front-end and waveforms running in software. This architecture is a significant departure from the current software radio state-of-the-art and one that comes with new challenges and design protocols. Current software radios able to support “n” simultaneous channels include “n” modems, “n” digitizers and “n” RF front ends. Some implementations even have “n” black-side general purpose processors. This paper explores software radio in its consummate form; an opportunity to execute multiple independent channels on a single radio front end.

1. INTRODUCTION

As long as receive/transmit waveform functionality is tied in to a one-to-one mapping of software “channel” to radio hardware channel the economics of the software radio approach make it hard to compete with software developed for a single platform. Simply stated it is more expensive to design and build waveform software that can be run on many different platforms. Therefore, the economic benefit of the software radio approach is realized only after the initial deployment. Further this benefit is realized only if the entire system design accommodates the separation of radio platform versus waveform application. Additional complexity is incurred if the system supports the notion of multiple waveform instances mapping in to multiple virtual radio channels.

There are two important elements required to realize the one-radio-to-many-channels architecture. One is to create numerous virtual sources and sinks on a single physical radio platform. The other is to create waveforms that can be instantiated multiple times. A minimum condition for multiple instantiation is that the waveform software be written in a manner so as to be re-entrant.

The simplest approach to creating a re-entrant waveform is to altogether avoid the use of global variables. Global variables are memory locations that are accessed by more than one waveform. Thus the desired software radio architecture is immediately threatened by having a single global resource, i.e., the radio, being accessed by numerous channel objects which are not allowed to access global resources. A technique that circumvents this contradiction is to create a virtual channel on the radio where the waveform is allocated what looks to be a complete radio but in reality is only a spectrally-contained stream of in-phase and quadrature samples.

The desired configuration shall support the ability to dynamically allocate and de-allocate channels without the disruption of other channels already in operation. The software radio design shall accommodate different commercial chipsets. The waveform software shall run with no modification on one platform versus another. One platform might employ chipsets from Analog Devices, the other chipsets from Intersil. We shall impose the architectural constraint that all waveforms are fed from a single A-to-D. The single digitizer can feed multiple independent digital down conversion operations. Ultimately from the perspective of the waveform the number of digitizers in the radio system is a “don’t care”. We chose to intentionally impose a single hardware digitizer in order to insure the purity of the virtual radio channels.

We impose an additional constraint that the radio hardware has certain parameters that must be statically configured. That is once the radio hardware is initialized and execution of the run-time begins certain configuration parameters cannot be altered. For undisrupted operation these immutable parameters shall be set once in a manner so as to satisfy all anticipated modes of operation and combinations of waveforms. This prerequisite is not too restrictive in light of typical radio use cases. For example when one performs a particular mission the radio is typically preset with all the frequencies, cryptographic keys, user ID’s and other such network parameters. These presets, sometimes known as the Communications Plan of the Day (CPOD), are a way of life in the operation of any

Proceeding of the SDR 05 Technical Conference and Product Exposition. Copyright © 2005 SDR Forum. All Rights Reserved

network having multiple user nodes and multiple waveforms. Future JTRS Clusters will have tools available to perform this network and frequency planning but operationally each and every radio will be pre-configured for the task at hand. Another practicality driving this pre-configuration operation is the configuration of antennas, power amplifiers and receive-side low noise amplifiers. For example a radio may be configured for VHF line-of-sight and UHF SATCOM. Invariably this configuration requires a certain combination of line replaceable units (LRU's). Similarly our single radio hardware channel – perhaps analog and digital components - will need to be pre-configured to support the required "n" virtual radio channels. Virtual radio channels are a neat concept but underneath the software the reality of the physics of electromagnetism still needs to be addressed. This paper shall not shy from those requirements as pertains to feeding multiple IQ streams from a single sampled IF.

2. THE MODEL

A single digitally sampled data stream is input into numerous, parallel down-conversion and decimation engines – see **FIGURE 1**. This hardware feature is the genesis of the multiple virtual channels. The hardware might support the ability to start and stop individual sub-channels within the receive chain or it might not. Even so it might be possible under some circumstances for the software to make it appear that individual channels are being stopped and started.

A couple of features are worth noting in the figure. The pass-band may be down-converted from different frequencies. Say within the pass-band there are three adjacent modulated carriers separated by 1 MHz. Assume a 93 Msps sample rate of an IF centered at 70 MHz. The middle carrier is centered at 70 MHz. Thus each down-conversion frequency of the independent numerically controlled oscillators (NCO) shall be set to 24, 23 and 22 MHz.

Once down-converted, and perhaps in conjunction with the filtering feature, each signal is decimated to its baseband data rate. Commercial chipsets typically support numerous stages of decimation and filtering. Thus the figure which implies one stage of filtering and decimation might actually be implemented in multiple stages of filtering and decimation.

A channelization template is proposed in Reference [1]. The template is part of an overall transceiver definition that includes "all platform-specific hardware and software converting signal between the Antenna sub-system and *Modem* software" The template includes definition of the upper and lower bound rejection gains and slopes, width of the transition bands, pass band bandwidth ripple and tuning accuracy. This frequency information is supplemented with time domain definitions of maximum latency and delta group delay.

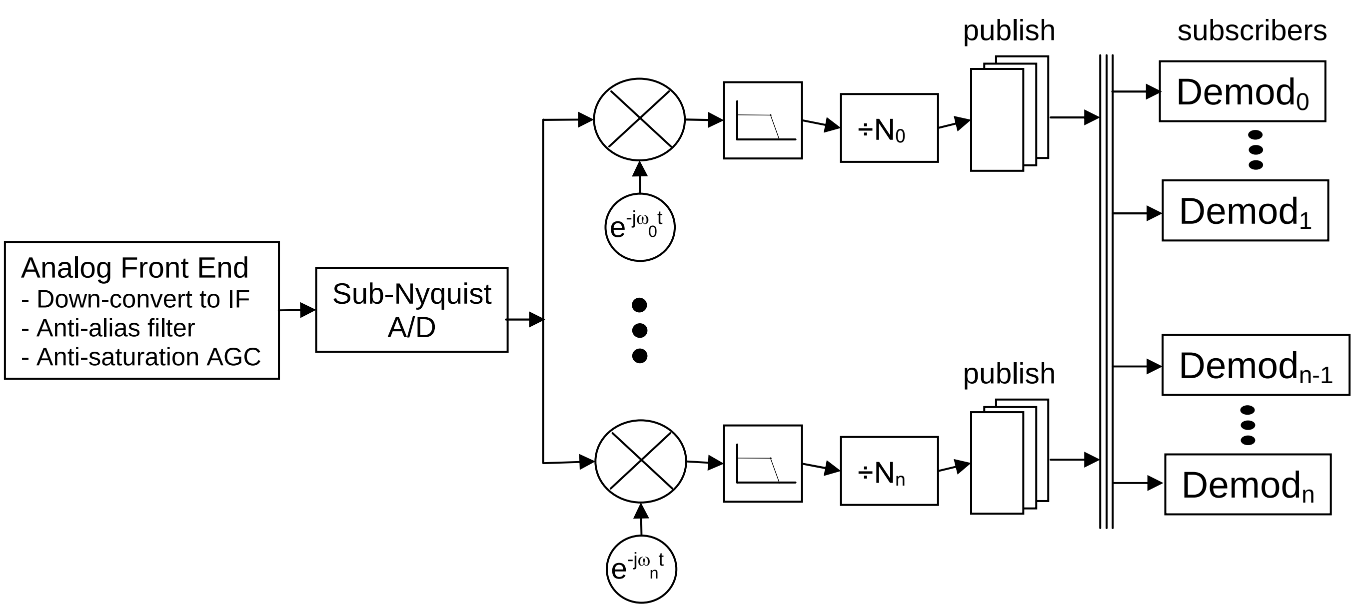

FIGURE 1 – MULTIPLE INDEPENDENT DOWN-CONVERSION STREAMS

 Copyright © 2005 SDR Forum. All Rights Reserved

Proceeding of the SDR 05 Technical Conference and Product Exposition. Copyright © 2005 SDR Forum. All Rights Reserved

In order to support hardware that has multiple filtering stages it should be possible to cascade these channelization templates. Also, from a software radio perspective, the channelization template should support all combinations of analog, digital hardware and software channelization implementations and operations. At a minimum the channelization template should be upgraded to include: 1) support for a decimation factor as used in multi-rate digital filters, 2) support for frequency translation to IF frequencies not just baseband and 3) support for the user to directly specify FIR or IIR coefficients. One might be tempted to add support for common filter types, for example elliptical filters (analog) or half-band filters (digital). This would be somewhat of an undertaking given the number of different types of filters that exist. Perhaps future extensions could include these specialized channel filters.

3. DIGITAL CONSIDERATIONS

Suppose the hardware supports up to eight receive and transmit channels within a single IF band that is so many MHz wide. The device driver interface will allow an application to open, control and manage individual receive and transmit channels. **FIGURE 2** below illustrates an example interaction between the applications, device drivers and the hardware.

Each application requests or opens receive and transmits channels via calls to the device driver through the user API. The device driver attempts to satisfy the request by allocating the channel resources from the hardware card. If the resources are unavailable then the application is informed that the channels are unavailable.

The hardware will have statically and dynamically configured parameters based on the hardware capability. Configuration of the hardware that interrupts the flow of data between other applications will not be allowed. Thus, only those items that do not hinder the operation of the other open channels will be available for dynamic configuration. This restriction may lead to a statically defined channel configuration or channel plan. The hardware has a couple of items to be configured in order for the channel to perform the operations expected of the waveform application.

3.1 Board Controller

The board controller occupies so many MBytes of PCI memory. The driver shall make this address space visible and accessible in user space. This allows the user the ability to be able to configure buffers, set interrupt watermarks and even start and stop the DMA engine. Users should never attempt to access these functions directly.

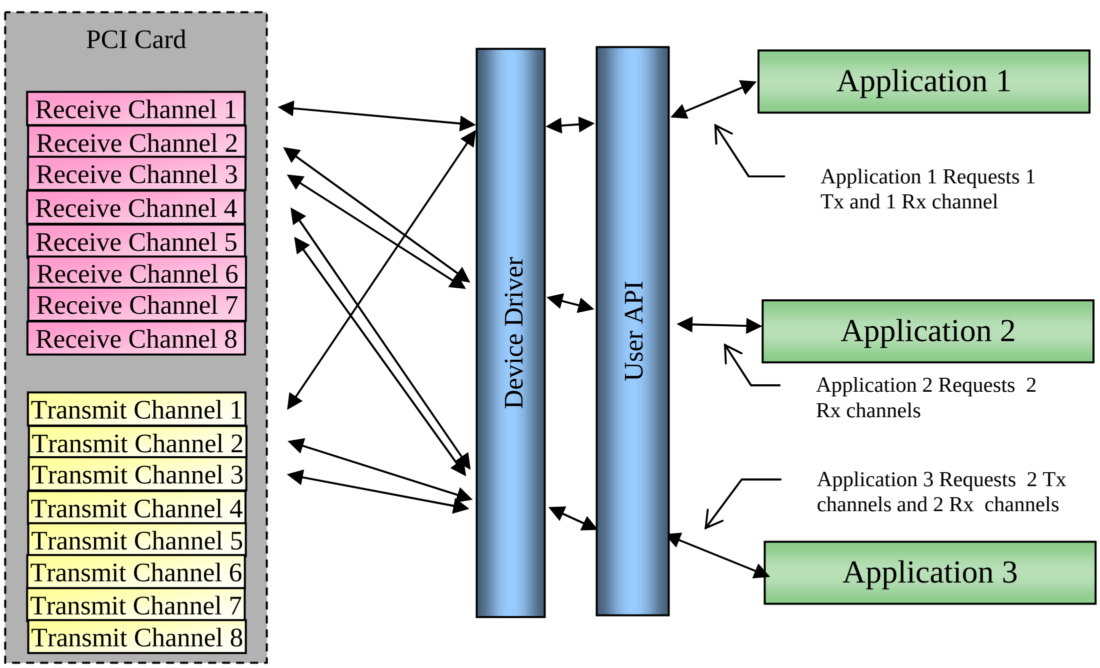

FIGURE 2 – MAPPING OF RESOURCES TO APPLICATIONS

 Copyright © 2005 SDR Forum. All Rights Reserved

3.2 Digital Upconvertor/Downconvertor

This is a portion of the PCI memory area that offers the ability to configure and control the up-conversion and down-conversion hardware. Settable parameters might include configuration of CIC filters, AGC settings, Numerical Controlled Oscillator (NCO) center frequency, and perhaps multiple stages of decimation and filtering.

Some hardware allows the output of independent channels driven from the same source to be multiplexed in a polyphase configuration. This allows the digital horsepower of the independent channels to be combined thus increasing throughput.

3.3 DMA Interrupt Processing

Some hardware has DMA master capability. In this case the DMA engine runs autonomously. On the downlink the card will send an interrupt to the user's processor to indicate that a transfers has just completed and that the processor had better get the baseband I's and Q's out of harms way before it gets over-written.

This underlying hardware construct implements a push convention. If the card cannot DMA master then the processor needs to pull the data from the card. The specifics of this transaction should be hidden from the user. This implies another abstraction layer above the device driver – the User API

3.4 User API

The user API provides access to the low-level driver functionality without the user knowing all of the driver intricacies and details. It provides a level of hardware abstraction in hopes of insulating the user to future hardware changes.

Underneath the User API is an agent that is hardware specific – **FIGURE 3**. This agent shall be coded to the meet the specific interface requirements of the hardware by acting as a translator from the User API to the target hardware. An example of this might be coefficients for a FIR. A generic API would allow the user to specify coefficients in floating point whereas the underlying hardware might require 12-bit coefficients with an LSB of 2^{-15}. This platform specific operation would be carried out by the agent.

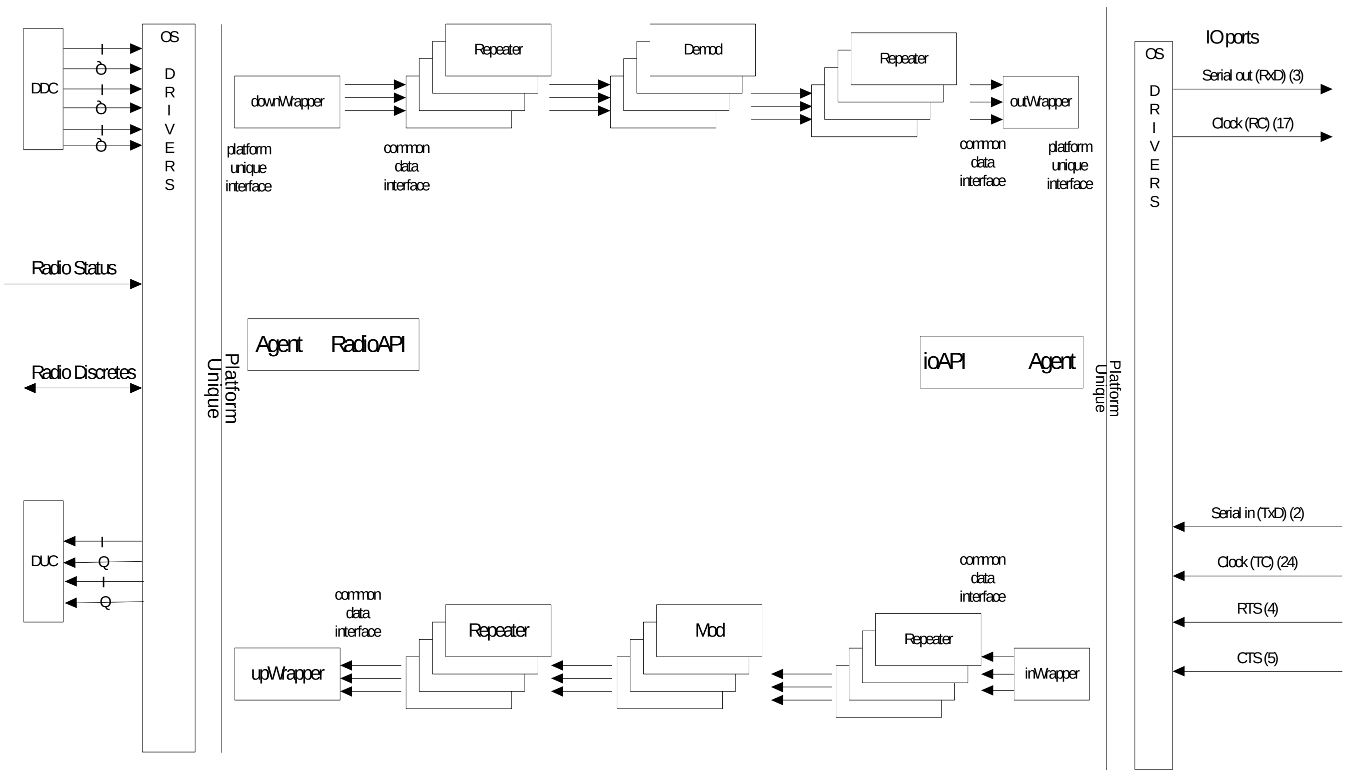

FIGURE 3 – USE OF A HARDWARE DEPENDENT AGENT

One of the means of increasing the portability of waveform software and also a means of simplifying the task of creating re-entrant waveforms is to allow as much of the waveform processing as is possible to occur without regards to the real-time. We assume that processing power is more than sufficient that everyone gets a chance to run in a timely manner with room to spare. The bulk of the waveform processing software is thus written as an application that operates on a stream of data as if it were coming from a asynchronous file or a socket instead of a real-time radio front end. This very powerful separation of the waveform from its time-critical environment is yet another job for the agent.

4. ANALOG CONSIDERATIONS

Typical RF environments are crowded with users. The wideband IF frequency that is presented to the digitizer might have hundreds of other carriers in the pass-band. In legacy radios a band-pass filter of the appropriate center frequency with a bandwidth identically matched to the signal-of-interest (SOI) is what is fed to the down-converter. Software radio offers many, many flexible and improved ways of doing RF receive and transmit, but a matched filter at the input to the down-converter is simply optimal and software radio cannot beat that.

With some understanding of the RF environment in which we will find our signal we can plan our filtering and decimation to isolate as best as possible the SOI and under most conditions our performance will be at least as good as the legacy radio. Given a peak power to average power ratio of 3, the relationship between dynamic range and the number of digitizer bits is well known [2].

$$\text{dynamic range (dBW)} = 20\log_{10}(2^n) = 6.02B \quad (1)$$

The relationship shows that the digitizer measures voltage not power. Let's consider three scenarios separately as they relate to the number of bits in the digitizer and the ability to recover the SOI. First, we consider "m" carriers of equal power within the pass-band. How many digitizer bits are required to allow the reasonable extraction of our SOI. Next we consider the effect of range. Nearby transmitters will tend to overwhelm the digitizer whereas our SOI might be many kilometers away. Finally we consider the effect of wideband AWGN as it steals digitizer bits from the SOI.

4.1 Implementation Loss due to Quantization

We first derive the number of bits required to reasonably recover the signal-of-interest. A PSK simulation was written in which a series of random bursts were generated in the presence of AWGN and then perfectly up-converted to some random IF frequency. The signal was then quantized to a certain number of bits and then down-converted back to baseband. Data decisions were computed on the quantized baseband signal and bit error rate computed. The FIGURE 4 shows PSK implementation loss computed for a signal that has been quantized to 3 bits before down-conversion.

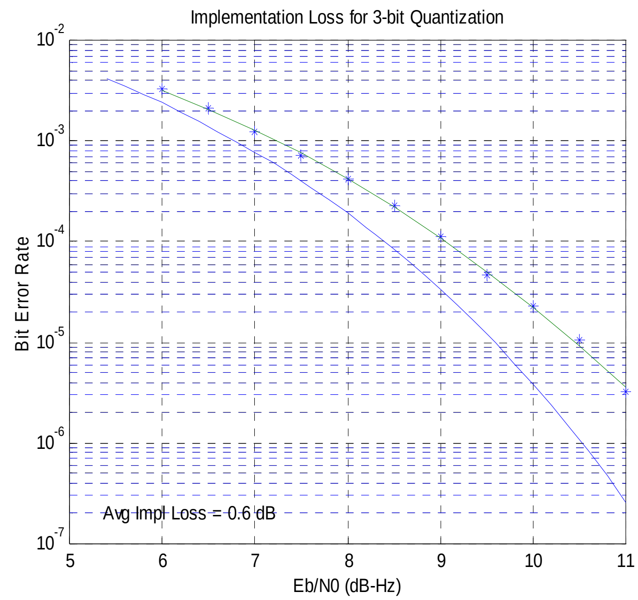

FIGURE 4 – QUANTIZATION LOSS

The down-conversion was performed using perfect knowledge of the IF frequency and phase. Of course in a real digital down-converter the down-conversion is itself quantized. We assume that arithmetic precision within the digital down-converter is significantly better than the 3 bits used by the SOI and thus offer a negligible contribution to error. The figure shows that with 3 bits of pre-down-conversion quantization, implementation loss goes from 0.3 dB at 10^{-2} all the way to 1.0 dB at 10^{-5}. For convenience sake we adjudicate that FOUR bits provides for reasonable signal recovery.

4.2 Number of Carriers versus Number of Bits

We consider "m" carriers of equal amplitude contained in the pass-band. Exactly one of those carriers is our SOI. Consider the sum of "m" independent sinusoids.

Proceeding of the SDR 05 Technical Conference and Product Exposition. Copyright © 2005 SDR Forum. All Rights Reserved

$$\sum_{i=1}^{m} \cos(2\pi f_i t + \phi_i) \qquad (2)$$

Intuition tells us that at some instant all of the sinusoids will be exactly in phase and the instantaneous amplitude will be 1+1+1+ … "m" times or exactly "m". Intuition also tells us that if the carriers have no harmonic relationship to one another - more specifically their phases are randomly related – then the likelihood of even exceeding "m/2" is small. In fact, Monte Carlo simulation of this scenario shows that the sum of the random sinusoids rarely exceed "m/2". We can compute that in using m/2 as a set-point we will saturate no more than 1 second out of 4200 seconds with all transmitters uniformly transmitting.

Assume the front-end AGC keeps the incoming signal in the range [-1, +1], i.e., full-scale at the A-to-D. We know that "m" unity amplitude carriers have an amplitude of m/2, thus with the gain applied for normalization our SOI now has an amplitude of "2/m". Consider the following equation:

$$\text{Quantization} = 2^{-(n-1)} \qquad (3)$$

Equation (3) simply says that, for instance, n = 3 bits of digitizer give a quantization of 0.25. To compute the bits required to comfortably recover "m" carriers we first compute the quantization level just underneath our SOI. Call that the first bit, subtract three for a total of 4 bits to represent the quantization required for our signal and then set equal to (3)

$$Int_{-\infty}\left\{\log_2\left(\tfrac{2}{m}\right)\right\} - 3 = Q = 2^{-(n-1)} \qquad (4)$$

where, the integer function truncates towards minus infinity. After some manipulation, the following rule of thumb is derived:

$$\text{Int}(\log2(m)) + 4 = \text{number of bits} \qquad (5)$$

For a simple example, to extract one carrier from among 30 of equal amplitude in the pass band requires at least an 8-bit digitizer.

4.3 The Effect of Range

Since power distinguishes at the range squared, at the digitizer, voltage is inversely proportional to range. Consider a tactical battlefield scenario. We presuppose a multiplicity of transmitters scattered between 10 meters and 10 kilo-meters. For a range that varies by a factor of 1000 that is approximately 2^{10}, thus 10-bits are required to accommodate that dynamic range. If you wanted to receive the signal from the distant transmitter with negligible implementation loss due to quantization another 4 bits would be required for a total of 14-bits; state of the art for A-to-D's sampling at IF frequencies. This scenario is dramatically over-simplified and does not adequately cover the case of a distribution of "m" transmitters.

4.4 The Effect of AWGN

The effect of additive white Gaussian noise can be addressed in a manner similar to "n" independent carriers. For all practical purposes signals not of interest are noise. The distribution of instantaneous voltage of wideband noise at the digitizer is Gaussian. We can develop a threshold and say that the wideband noise can push the digitizer into saturation no more than x percent of the time. A detailed analysis is deferred to a later work.

5. CONCLUSION

When one considers the engineering facts that must be solved in order to run multiple waveforms off of a single digital IF, the problem is more than avoiding the use of global variables and the allocation of resources. Specific consideration must be given to the construction of the radio API so that it is possible to instantiate and destroy waveforms on one virtual channel without loss of data on other virtual channels. Furthermore the software radio is not exempt from the laws of physics and careful design of analog components prior to the digitizer is requisite so as not to overwhelm the A-to-D.

6. REFERENCES

[1] E. Nicollet, and C. Serra, "SCA 3.1 – Transceiver sub-system – Version 0.2, "*SDR Forum Members Only*", Jan 2005, p. 7

[2] Jeffrey H. Reed, *Software Radio: A Modern Approach to Radio Engineering*, Prentice Hall, Upper Saddle River, New Jersey, 2002, p 202.

Proceeding of the SDR 05 Technical Conference and Product Exposition. Copyright © 2005 SDR Forum. All Rights Reserved