LA-UR-09- 05170

Approved for public release; distribution is unlimited.

Title:	High-Efficiency Space-Based Software Radio Architectures & Algorithms (A Minimum Size, Weight, and Power TeraOps Processor)
Author(s):	Mark E. Dunham, Zachary K. Baker, Matthew W. Stettler, Michael J. Pigue, Eric N. Schmierer, John F. Power, Paul S. Graham
Intended for:	2009 International Conference on ReConFigurable Computing and FPGAs

Form 836 (7/06)

High Efficiency Space-Based Software Radio Architectures & Algorithms
A Minimum Size, Weight, and Power TeraOps Processor

Mark E. Dunham, Zachary Baker, Matthew Stettler, Michael Pigue, Paul Graham, and John Power
Los Alamos National Laboratory
Los Alamos, NM 87545
Email: mdunhamlanl.gov

Abstract

Los Alamos has recently completed the latest in a series of Reconfigurable Software Radios, which incorporates several key innovations in both hardware design and algorithms. Due to our focus on satellite applications, each design must extract the best size, weight, and power performance possible from the ensemble of Commodity Off-the-Shelf (COTS) parts available at the time of design. In this case we have achieved 1 TeraOps/second signal processing on a 1920 Megabit/second datastream, while using only 53 Watts mains power, 5.5 kg, and 3 liters. This processing capability enables very advanced algorithms such as our wideband RF compression scheme to operate remotely, allowing network bandwidth constrained applications to deliver previously unattainable performance.

1 INTRODUCTION & BACKGROUND

Since the original Adaptive Computing Systems Program at DARPA in 1994, Los Alamos (LANL) has been developing RF processing systems with FPGA devices, now known in the literature as Software Defined Radios (SDR). Originally designed with Altera and Xilinx first generation logic arrays, these receivers have grown with the chip families underlying them. While FPGA technology has now matured both architecturally and in process node to the point we prefer the term System-on-Programmable Chip (SoPC), the other components of an SDR have also gained in performance. In particular, modern Analog-to-Digital Conversion (ADC) technology now enables Direct Down-Conversion analog architectures due to the GHz input bandwidths available in 12 to 16 bit converters. Meanwhile point-of-load power conversion, high speed SRAM, and very large NVRAM technology all support extreme performance from the processing side. The net result is SDRs which are truly special purpose super-computers, and which can execute even the most complex algorithms in real time, over bandwidths meeting or exceeding those required for modern communications systems. As such the era of SDR has now arrived.

We focus on satellite-capable SDR designs, since the need for reconfiguration is greatest in systems which are no longer accessible after launch. This also supports a robust research effort in COTS Radiation Tolerance (COTS-RT) , which has also grown serendipitously with chip performance. Identifying promising COTS-RT parts is key to our success, since the Xilinx Virtex-4 is 1000 times more capable than the best RadHard micrprocessors available now. (cite?) In the last decade Rad Tolerance has advanced to the point that 400kRad Total Ionizing Dose (TID) is routine seen in commercial parts, and destructive latch-up is disappearing in many families. These phenomena are the result of process improvements intended to improve electrical performance, but the Rad Tolerance is a direct result as well. Remaining to be dealt with beyond TID and SEL, is SEE performance which in general is degrading as device feature size shrinks. The LANL Rad Tolerance program therefore devotes much of its R&D towards ways of mitigating SEE in modern electronics.

It is both a challenge and a blessing to the space electronics community that COTS electronics evolves at a rate of at least 2.5 generations/launch. In other words, by the time a new satellite design is delivered to the launchpad, the COTS technology will have advanced by at least 2.5 more generations. Here we describe a delivered receiver that is now 2 generations behind art, but also discuss laboratory work with current generation FPGA devices for the next build.

No SDR is complete without algorithms, which are often point designs to satisfy a specific communications need. LANL is more interested in general scientific applications that are flexible for a wide range of signal classes. Therefore we describe two such applications in some detail here, an incoherent wideband signal detector, and a fully coherent wideband transform/compression engine. Such algorithms form the building blocks of subsequent user-specific appli-

cations, or perform as general purpose science data gatherers. For instance, LANL has a well-known VHF lightning research program, based around global monitoring of severe storms via the GPS spacecraft. Radio-astronomy also forms another area of interest, especially radio bursts such as pulsars.

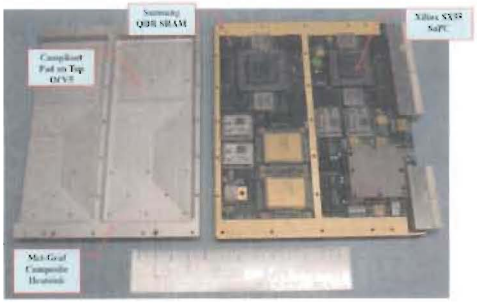

Figure 1. Real-time Signal Processor (RTSP) board with Xilinx Virtex 5 LX200 and SX55, seven banks of Quad Data-Rate (QDR) memory, and radiation hardened scrubbing and interface circuitry.

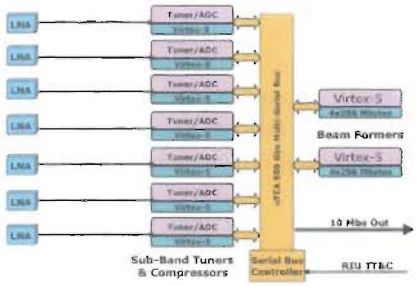

Figure 2. Multi-channel SDR architecture for potential applications such as beamforming.

2 DIRECT CONVERSION THEORY OF OPERATION

A key innovation of the LANL TeraOps SDR (T-SDR) is the use of direct conversion radio architecture. This feature is diagrammed in Figure 5, where it is seen that higher dynamic range is achieved concurrently with lower power, as long as analog tunability is given up. With input bandwiths

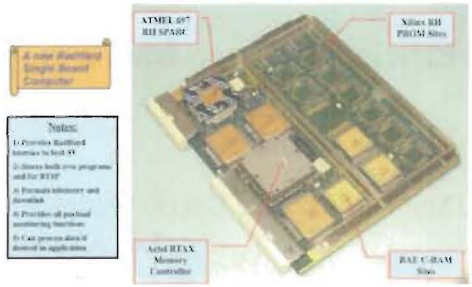

Figure 3. Flight computer board with radiation-hardened processor, SRAM, and Xilinx PROM programming circuits.

Figure 4. The use of commodity compact-PCI chassis allow for cheap and efficient development cycles.

exceeding 150 MHz in new designs on the bench, tunability is moved into the digital domain, while the analog section becomes wideband and fixed in frequency. This in turn simplifies design and manufacture, compared to traditional mixer-IF designs.

As 5 shows, Nyquist's original theory specifies the bandwidth supported by a particular sample rate, but not the center frequency, which can be any integer multiple of the sample rate which is within the input bandwidth of the analog system. Other terms for this process are used in the literature, including Nyquist conversion, and aliasing.

In order to better understand our efficiency gains, we will examine how Digital Intermediate Frequency (DIF) is created in the T-SDR receiver. Unlike most radios, T-SDR does not tune at all, it only switches between fixed bands. This is done in order to maximize dynamic range while simultaneously reducing analog power consumption. Since T-SDR covers 50 MHz bands, the need for tuning is greatly reduced, and we instead derive the wideband digital receiver

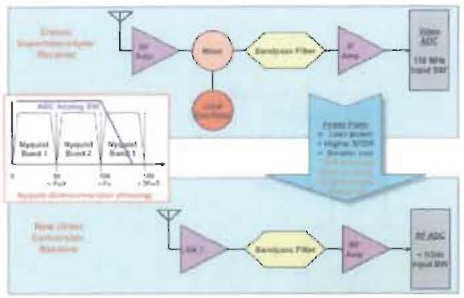

Figure 5. Comparing Direct Conversion to Super-Heterodyne Architectures

where tuning is accomplished as a signal processing process once the data has been collected.

However, the radio does digitally downconvert the input RF signal to baseband, as shown in the inset of Figure 5, entitled Nyquist Downconversion (Aliasing). It has long been recognized in the signal processing domain, that if an ADC has larger input bandwidth than half the sample rate, this bandwidth can be put to use mixing or downconverting the input to a lower frequency. In essence, the digital output stream always occupies the lowest Band 1, however the input may be positioned in any of the higher bands identically, other than an inversion of high and low frequencies in the even numbered bands. In the case of the Linear Technology LT2208 ADC, there are at least 16 possible Nyquist bands one can use, albeit with increasing degradation of effective bits in the ADC.

Our receiver design exercises the ADC over its full datasheet range. As such, we have chosen 4 different operating bands, each exploiting a different portion of the ADC performance. One common theme found while working with our filter suppliers was that we needed 10 MHz transition zones in general, so we relaxed performance at the Nyquist frequency boundaries to 40 dB, as long as we were 80 dB down within another 5 MHz.

3 DETAILED ARCHITECTURE

In order to clearly outline the design aspects enabled by modern chip capabilities, we now describe each system component shown on Figure 6 in some detail.

- Input Antenna and Direct Conversion Receiver. Here a 16 bit, 120 Megasample/second ADC converts selected 50 MHz RF bands into Digital Intermediate Frequency (DIF) data streams, down-converting them at the same time it digitizes, in a process known as Direct Conversion. With Direct Conversion, no previous mixers, local oscillators, or IF amplifiers are needed. We simply amplify and filter the input RF, defining a 50 MHz band of interest. This Direct Conversion process results in much higher dynamic range due to the small number of active devices in the signal chain, while also saving significant electrical power for the missing components. As such, Direct Conversion is a major advance suitable for HF, VHF, and UHF receivers, and offering a one-time gain in both performance and power consumption. The output stream from our Direct Conversion Tuner-Digitizer (DCR) is a 16-wide parallel stream of data bits at the 120 MHz clock rate, leading to an aggregate data rate of 1.92 Gigabits/second.

  MRM is fed by low gain antennas which have nominal coverage from horizon to horizon at LEO. As such, 0 dBi gain is a good assumption, once losses are included to counter the basic directivity of 2 dB. We will use an active monopole antenna in the HF and low VHF where the T-SDR will be totally dominated by excess ambient noise, as shown below[1]. Noise temperature is easily converted into excess noise figure by normalizing to 300K, where kTB has the value of 114 dBm/MHz. Hence, at 100 MHz the total noise is 900,000K, or 3000 times more than at frequencies higher than 200 MHz. This means that ambient has an effective noise figure of 34 dB, before antenna performance is considered. If the antenna and LNA temperature is much lower than ambient, then only the ambient energy will be seen.

  Another factor confounding good signal reception at HF is that well matched antennas must subtend wavelength in some dimension, however at 50 MHz is already 3 meters, and keeps growing inversely with frequency. Hence the use of active antennas which obtain the same pattern as full sized elements, but which require matching amplifiers to deal with the very high impedances present, is very attractive at HF and to a lesser extent at VHF. MRM will use an advanced active antenna resonant at 175 MHz, which will not degrade the noise performance coming from natural sources as shown above.

- The DIF output of the DCR is fed to a Real Time Signal Processor, containing 2 each Xilinx Platform FPGA, and their associated SRAM cache memories. The 4th generation Xilinx devices come in 3 types, optimized for signal processing (SX series), for communications (FX series), and for logic (LX series). Based on past experience in building software radios at LANL and other DOD entities, a concate-

[1] CCIR Reports 341-6, 258-5, & 670-1, International Telecommunication Union, Geneva, SW.

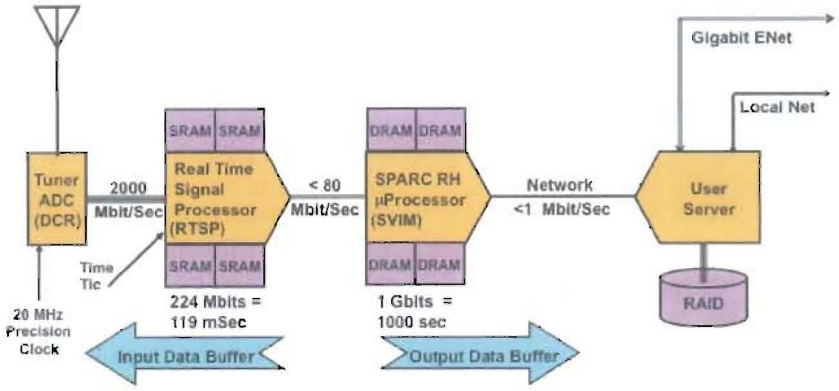

Figure 6. System Architecture for MRM system

nated LX device after an SX device offers the optimum combination of multiply-accumulate-heavy capability in the 4VSX55 device, followed by general purpose logic-heavy capability in the 4VLX160 device. Ground benchmarks of these devices have already shown the ability to process 10 Gigabits/second of data, therefore we expect no issue with accepting our 1 Gigabit/second stream. In the space allocated to the RTSP we have been able to accommodate 7 each synchronous cache RAM memories, each containing 32 Megabits of data organized 32 bits wide by 1 Meg deep. These cache Ram enable memory intensive operations such as FFT, within the Xilinx devices. However it will still be possible to define algorithms which cannot execute within the devices in real-time, or within power consumption limits. We predict the Virtex-4 devices offer 30 60 GigaOps/Watt processing capability, which should allow a 16K FFT to be executed with 2 watts of power in the Xilinx FPGA. Data exits the RTSP via a PCI bus interface managed by an ACTEL RTAX Rad-Hard PLA, and theT-SDR requirement is only that 10 Mbytes/second be transferred over the PCI bus with its theoretical 120 Megabytes/second capacity.

- Following the RTSP a Satellite Vehicle Interface & Memory (SVIM) card receives the PCI data, and temporarily stores it in main memory of a SPARC 32 bit integer processor via DMA. The SPARC is a European Space Agency standard design, and is Rad-Hard for LEO and Geosynchronous orbit conditions. This SPARC processor also maintains a processed data buffer 1 Gigabit deep. This data buffer size was chosen as the maximum size that can be downlinked at a 1 Megabit rate, within one orbit of a LEO satellite in continuous contact with the ground.

- The SVIM also provides the Rad Hard and non-volatile storage of all software for theT-SDR payload. We have provided 3 separate non-volatile banks of memory:

  - 256 Kbyte of Start-Up ROM (SUROM) are provided to the SPARC alone, so that critical OS level functions are stored in assured RadHard memory. Our intent is that all elements of software required to decompress and load FPGA bitstreams will also be contained in this SUROM.

  - A bank of Xilinx PROM 48 Megabits deep, with hard ECC, which can only be reprogrammed once on the ground, through theT-SDR test port. Write enable is controlled by the test port, and is disabled when test port is not engaged. Lossless compression will be used in order to reduce bitstream image size by 4X or more, depending on the application. The decompression will need to be performed in the SPARC each time a data block is read.

  - A bank of C-RAM TBD Megabits deep, with hard ECC, which can be reprogrammed at will,

including on-orbit. This is the target memory for any subsequent reprogramming after launch. It also must contain the SPARC application code associated with the FPGA application, and will will also be losslessly compressed in the same format as the PROM.

- The SPARC processor is best in class, with about 100 Mips/watt capability, and will have a 100 MHz core clock. It is supported by the following user memories:
  - SRAM of 6 Mbytes, organized as 3 SRAM banks of 512Kx32, data bus width is 32 bits.
  - SDRAM of 128Mbytes, organized as 2 SDRAM banks of 16Mx32, data bus width is 32 bits.
  - SUROM of 256Kbytes, organized as 2 ROM banks of 32Kx32, data bus width is 32 bits
  - RTAX FPGA registers, total size TBD, as Atmel uP memory mapped I/O, address buss width 28 bits, data buss width 32 bits
- The SVIM is tasked with 3 major functions: responding to and transmitting telemetry, retrieving configuration data dynamically from C-RAM memory to support SEU mitigation, and operating the downlink with any associated data compression or thinning operations. The SPARC is capable, but will have to assign interrupt priority to these tasks in normal operations. As such, the majority of reprogrammability is assigned to the Xilnx devices, unless one of the 3 primary SPARC functions can be de-emphasized in a particular mission. An example of this is during data reception from the ground, where we do not expect to need to decompress any on-board code. In that case we my operate an AES-256 decryption module in place of the decompressionPrepare Your Paper Before Styling

4 RESULTS & PERFORMANCE TESTING

Here we present a few of the validated performance benchmarks of the existing T-SDR. (Zack?)

5 WIDEBAND COMPRESSION APPLICATION

The general term for this problem is data compression and expansion, or "companding." Application 3 deals with a "lossy" companding algorithm, which is suited to the transmission of RF signals. Most RF signals are modulated as the mathematical product of the information and an underlying carrier., where the strength of the received signal can vary dramatically. Thus, the great majority of practical

Figure 7. Sealed satellite module

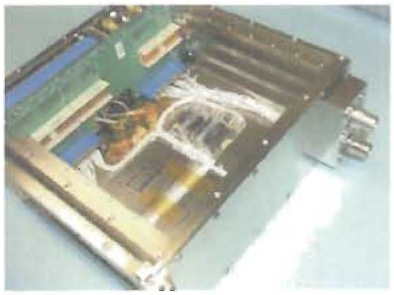

Figure 8. Chassis and power supplies for system

sensor applications are multiplicative modulation processes, and only require constant relative precision. This implies that logarithmic representations may be used without loss of resolution, as long as they are properly scaled.

The basic block diagram of Application 3 is illustrated in Figure 10

The basic process consists of applying a logarithmic transform to the output of a Fourier transform. The purpose of the Fourier transform are to convert convolution processes to multiplication in frequency space, and to separate various signals that overlap in the time series, but which are separate in the frequency space. The log transform then allows the multiplicative signals still contained in each frequency sample to be treated as additive superposition.

The 120 Mega-Sample input data is subjected to a 1/3 overlap delay (1/3 of the FFT size of 16k samples) to produce two data streams. The allows the output to completely cover the input signal after the windowing function is applied. Without the overlap delay, signal occurring at the edge of the transform might be masked and inadvertently destroyed. The two signals are then fed through a window multiplier to reduce spectral leakage and to allow near-perfect reconstruction. The outputs of the windowing oper-

Figure 9. System view of analog bandpass filters and analog-to-digital conversion

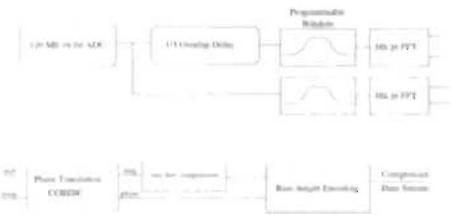

Figure 10. Application architecture for Gabor transform-based wideband compression system.

ator are suitably delayed such that both signals can be presented to the FFT simultaneously. This is necessary because the system uses a single FFT to compute the transform of both streams, reducing the area and power requirements for the large 16k point FFT. Because the inputs to the FFT are real, the output of the FFT is separable into two unique frequency transforms.

The output of the FFT is then fed into a phase translation CORDIC. This is a pipelined core that converts the complex FFT output into corresponding polar representation consisting of magnitude and phase.

Our preliminary results indicate that the V5SX95 with the 16k FFT and CORDIC cores can operate at roughly 317 MHz, occupying about 18% of the device. This benchmark will be tested in our NextRF module using a VirtexSX95T device in place of the current Virtex-4 devices in MRM. In this instantiation, the compressor runs within the receiver head, and delivers a serial stream of digital data to post-processors which combine and further reduce the signal size into a final output datastream.

6 SUMMARY AND WORK IN PROGRESS

Even as the T-SDR was developed, new COTS hardware appeared that dramatically changed architectural possibilities again. The two most significant developments for LANL were:

- advent of Multi-Gigabit Serial cores in FPGA, DSP, and micrprocessors. These hard IP blocks enable all-serial modular systems
- The development of silicon-on-insulator ADC such as the TI5424, which offer very strong Rad Tolerance while providing input bandwidths exceeding a GHz and sample rates ¿ 400 Mss. This enables continuous band coverage from Direct Conversion architectures.

Together, these key evolutions led us to consider multi-channel SDR architectures for MIMO and other spatial diversity applications. One such architecture now under active pursuit at LANL is shown below.

Acknowledgments

The LANL team gratefully acknowledges long term support from DOE NA-22, DARPA, and other DOD agencies in development of all our SDR work. Only through such long-term support can a new capability be fully developed. We also the collaboration of our commercial partners, Xilinx, Lockheed-Martin, BAE Systems, Linear Technology, and Texas Instruments.