SAND2017-5567C

SOFTWARE DEFINED RADIO FOR STEPPED-FREQUENCY, GROUND-PENETRATING RADAR

Samuel C. Carey and Waymond R. Scott, Jr.

Georgia Institute of Technology
School of Electrical and Computer Engineering
777 Atlantic Drive NW, Atlanta, GA 30332-0250

ABSTRACT

Software defined radio (SDR) is a rapidly developing technology that implements signal processing components partially or completely in software. In this paper, SDR's potential as a platform for ground-penetrating radar (GPR) is explored. The stepped-frequency radar method is implemented using off-the-shelf SDR hardware and open-source software. SDR is typically designed for communications applications, so special consideration is necessary for remote sensing. Precisely timed commands achieve RF phase coherence as well as digital sample synchronization. A multi-tone digital signal takes advantage of instantaneous bandwidth, and a sequential tuning routine expands effective bandwidth to cover the 500-5000 MHz band. Initial design challenges are weighed against potential advantages of design flexibility and hardware versatility.

Index Terms— software radio, ground-penetrating radar, transfer function, UHD, coherency

1. INTRODUCTION

Ground-penetrating radar (GPR) has been shown to be an effective tool for detecting and locating buried objects such as utilities, archaeological artifacts, and landmines [1]. A GPR transmits an electromagnetic wave that is reflected by underground targets and then received again, as in fig. 1. Then, the round-trip time-delay of each reflection is used to determine the depth of each target.

Reflections can be caused by variations in subsurface permittivity or conductivity. While this sensitivity reveals various materials of interest, it can also produce false alarms when soil properties naturally fluctuate. A large effective bandwidth can help distinguish targets from false alarms by increasing the resolution of the transfer function. In high performance GPR systems, the ratio of the highest frequency to the lowest frequency of the operation band often exceeds 10.

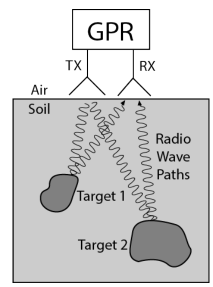

Fig. 1. A typical GPR scenario. The GPR hardware is held over the ground with its transmit (TX) and receive (RX) antennas directed into soil containing discontinuities in permittivity (targets). Radio waves from one antenna reflect from the discontinuities at different depths and are received by the other antenna.

Many different hardware configurations have been used to measure the transfer function between the transmitting and receiving antennas in a GPR. In pulsed systems, a very short time-domain pulse is radiated and received. Usually, the pulse is repeated many times, with the receiver sample time shifting slightly each pulse. The receiver can then build a representation of a single received pulse, with a resolution equal to the time shift.

Some GPRs have measured transfer functions by transmitting random or psuedo-random noise signals and correlating them with reflections [2][3]. Frequency-modulated continuous-wave (FMCW) GPR systems have used a linearly increasing transmitter frequency mixed with reflections [4].

A frequency domain approach is the stepped-frequency method. The GPR transmits a steady tone of a certain frequency and records the magnitude and phase of the reflection. Next, this measurement is repeated at equally spaced frequencies, scanning a certain band. A wider band increases range resolution, and measuring frequencies closer together decreases range ambiguity. The many tuning operations required for each scan can take considerable time, so a reasonable scan rate during a GPR survey can be difficult to achieve.

This work is supported by the U.S. Army REDCOM CERDEC Night Vision and Electronic Sensors Directorate, Science and Technology Division, Countermine Branch; the U.S. Army Research Office under Grant Number W911NF-11-1-0153; and by Sandia National Laboratories, a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

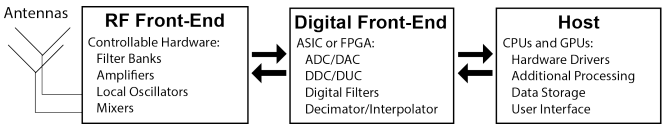

Fig. 2. A typical SDR. Components are divided into three categories, each representing stages along the processing chain.

Often, these frequency-domain samples are transformed into the time-domain to obtain the response versus time.

The methods above are generally implemented using custom hardware that is often difficult to design and build due to the extremely wide bandwidths required for performance. In this paper, an alternative approach of using software defined radio (SDR) for GPR is explored. The stepped-frequency radar method is implemented using off-the-shelf SDR hardware and open-source software.

An SDR is a signal processing chain that converts between physical radio waves and digital data. An ideal SDR would implement as much of this chain in software as possible, receiving and transmitting by sampling and synthesizing RF signals, with the ADC and DAC connected directly to the antenna. Currently, this is not technically feasible. Instead, components are implemented either as RF circuits, specialized digital circuits, or software, as shown in fig. 2. The term "software defined radio" refers to all three of these stages working together, although it may also refer to just the front-ends which are typically sold with the expectation that the user will provide their own host computer. Today, users with little analog design experience can purchase a generic, off-the-shelf SDR hardware unit, download compatible software, rapidly prototype almost any radio system in software, and immediately test it in the real world.

This new process increases efficiency at every stage of radio development. For example, theoretical models of hardware and signal channels are easily replaced with their physical versions. Hardware fabrication is replaced with software compilation, dramatically reducing the costs and delays of design iteration cycles. The development process can also benefit from unprecedented ease of collaboration, where developers can share new improvements online and ensure that they are working on the most recent versions with their colleagues. Development can even continue after device deployment, via remote software updates. Because the capabilities of each device are increasingly broad, the market can converge to fewer devices and allow economies of scale to reduce prices.

The capabilities of SDR appear to meet the requirements for several GPR methods, suggesting that a GPR could be implemented using SDR, and reap the described advantages of SDR. Various radar methods have already been discussed and implemented using SDR [5][6]. Ground-penetrating radar in particular has also been discussed [7], but a search for actual implementations of SDR GPR revealed no examples.

2. APPROACH

2.1. Hardware and Software

For this project, Ettus UBX-160 RF front-end daughterboards have been chosen for their relatively high bandwidth of 160 MHz, as well as their ability to synchronize phase between multiple local oscillators (LOs) [8]. Two UBX-160 units are housed in an Ettus X310, which serves as the digital front-end of the system. The X310 was chosen for its compatibility with the UBX-160, as well as its relatively large, built-in Xilinx Kintex-7 field programmable gate array (FPGA). Using separate daughterboards for transmitting and receiving improves isolation between the two RF channels.

Ettus devices are supported by the open-source UHD driver library. A GNU Radio application uses UHD to interface with the X310 and processes the data. Results are plotted in MATLAB.

2.2. Synchronization

For a fixed range profile, repeated measurements should yield repeated results, within a margin of error. This system must carefully control two major factors in order to achieve this consistency.

First, the RF front-ends of the transmit (TX) and receive (RX) chains must be synchronized. Each UBX daughterboard has separate LOs for TX and RX. For communication applications, this is useful for transmitting and receiving in multiple bands simultaneously. However, this application requires that both LOs on both UBXs be tuned to the exact same frequency, with a consistent phase offset between them. Fortunately, the X310 provides the option of feeding all LOs with the same clock, achieving frequency synchronization. The UBX daughterboards also allow phase synchronization using timed tune commands [8].

Second, the transmitted and received digital signals must be aligned in time. By default, if a user commands that a burst be transmitted and received, both actions will be completed at unpredictable times in the future, depending on numerous unpredictable factors such as scheduling and buffering. This alignment uncertainty renders any measurements of range meaningless. Fortunately, UHD allows users to timestamp transmission and reception commands. The application sends a burst to the SDR where it waits in a buffer for a specific time in the future. The user also sends a request to receive a burst to the SDR where it waits for the same time. When the time comes, the first samples of the TX and RX bursts are read and written on the same clock cycle, or a fixed number of clock cycles apart.

Breaking operation into bursts has another advantage. While the X310 is designed to transmit and receive at 200 MS/s (200 million complex samples per second), continuous operation requires the designer to choose between lowering the bandwidth to a few megahertz and upgrading host hardware to keep up. With bursts however, the host can simply use the breaks in between bursts to catch up, if buffers are sufficiently large.

2.3. Band Coverage

The X310's maximum sampling rate limits the digital bandwidth of the SDR to 200 MHz. The overall bandwidth of the system is further constrained by the 160 MHz RF bandwidth of the UBX-160 front-ends. Since this is much narrower than the desired 4500 MHz-wide band, the large band is broken into smaller bands which are scanned one-at-a-time.

The front-ends are sequentially tuned to the center of each small band. These centers are chosen in increments of 150 MHz to reduce artifacts from fractional-N tuning and avoid the edges of the RF filters. At each center frequency, the X310 transmits and receives a 200 MS/s signal containing 8 tones, equally spaced in frequency on either side of the center to cover that band. This allows 8 GPR frequencies to be measured at the same time, resulting in an 8-fold speed increase over a single-tone, hardware GPR.

A zero frequency tone is not used, which avoids the DC-offset inherent to direct-conversion receivers. Hardware GPRs must use either careful calibration or a superheterodyne architecture in order to achieve this same isolation.

IQ-imbalances in the quadrature DACs and ADCs can cause tones on either side of the center frequency to be partially mirrored and coupled to the other side. While UHD includes routines to partially calibrate for this, the effect is further diminished by shifting all tones in frequency by a quarter of the tone spacing. This leaves tones equally spaced from both intentional tones and their mirrored versions.

Each tone is also shifted in phase to minimize the maximum amplitude of the combined signal, thus allowing a higher average power to be sent through the amplifiers [9]. Once a DFT samples the magnitude and phase of the reflected tones, these same phase offsets can be subtracted out.

Once samples from the entire band are collected, the host performs an IDFT and records the result. See Section 3 for plots of this process in action. The entire scan is repeated for each new location in the survey.

2.4. Speed

By default, UHD chooses between two LNAs in the UBX RX chain depending on whether the center frequency is above or below 1500 MHz [10]. Each time an LNA is selected, there is a significant warm up period. Therefore, the UHD code has been modified and recompiled such that only the upper band LNA is used for both bands. This allows the scan speed to be increased without sacrificing gain stability.

Although the 4500-MHz-wide band could be covered in only 30 sub-bands, 32 are used in order to round up the number of samples to a power of two (256). Considering only the time needed to transmit 16,384 sample bursts and retune the LOs, the entire band could be theoretically be scanned in as little as 5.8 ms [11]. This speed might be approachable if command and processing operations were implemented on the X310's FPGA instead of the host. For this implementation, however, communication over ethernet increases the scan time to about 100 ms.

3. TESTING

Using the methods described above, a basic, stepped-frequency radar system has been developed. 256 samples are taken 18.75 MHz apart for each frequency sweep, approximately covering the 500-5000 MHz band. The system was tested using a loopback from one UBX daughterboard to the other, through a coaxial cable, attenuation, and a line stretcher device, as shown in fig. 3. The line stretcher is a telescoping, air-filled waveguide used to precisely adjust the path length of the signal.

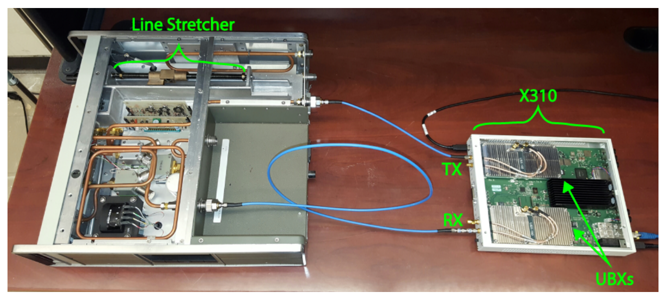

Fig. 3. Two UBX daughterboards inside of an X310 (top panel removed) transmit and receive signals through coaxial cables, attenuation, and a line stretcher. Attenuators are connected in series at the RX port. The line stretcher is housed with other electronics which are not used.

Figure 4 shows the magnitude and phase of the discrete frequency response of the loopback. The vertical lines separate the samples into the bands of 8 tones that were collected at each center frequency. The bump-like shape of each band is due to the 160 MHz low-pass filters of the transmitting and receiving RF front-ends.

After the frequency response is saved and divided from subsequent measurements, the magnitude response and phase response become approximately flat lines at zero, and the impulse response becomes approximately an impulse centered at a time of zero. This calibrated response is what one would expect from an ideal loopback of zero-length.

In fig. 5, after calibration, the line stretcher has been adjusted to extend the path length by 30 cm, resulting in a corresponding shift of the impulse in time. One can see that the phase across bands is linear and continuous, indicating that LO synchronization was successful. The change in path length can be divided by the the change in time of the impulse response's peak to calculate the velocity of propagation

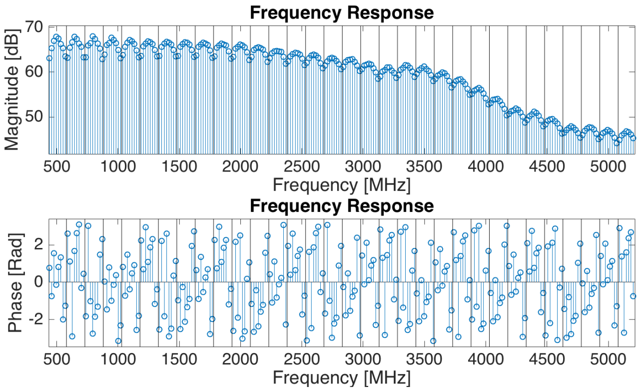

Fig. 4. Uncalibrated scan of frequency response.

in the line stretcher to be 1.0001 times the speed of light in a vacuum. Since the line stretcher is air filled, this computed velocity is essentially correct.

4. CONCLUSIONS

This prototype demonstrates a viable approach to implementing a wide-band GPR using SDR. By carefully choosing the UBX daughterboards and using special synchronization commands, the stepped-frequency radar method has been used to generate a stable, realistic time-domain impulse response of a physical simulated target.

Following the basic validation in this paper, this GPR system was tested with actual antennas, soil, and targets by Carey [11]. The system was also extended to locate modulated scatterers with excellent clutter rejection. With these proven capabilities, this GPR system could be very helpful for a variety of sensing applications.

5. REFERENCES

[1] David J. Daniels, Ed., *Ground Penetrating Radar*, The Institution of Electrical Engineers, Michael Faraday House, Six Hills Way, Stevenage, Herts., SGl 2AY, United Kingdom, 2 edition, 2004.

[2] Guangyou Fang Qunying Zhang, Shengbo Ye, "Design and testing of a pseudo random coded gpr for deep investigation," in *IEEE International Geoscience and Remote Sensing Symposium*, 2016.

[3] Ram M. Narayanan, Yi Xu, Paul D. Hoffmeyer, and John O. Curtis, "Design, performance, and applications of a coherent ultra-wideband random noise radar," *Optical Engineering*, vol. 37, no. 6, pp. 1855–1869, 1998.

[4] Mariusz Zych, "Measuring experiment of FMCW ground penetrating radar," in *Radar Symposium (IRS), 2011 Proceedings International*, 2011.

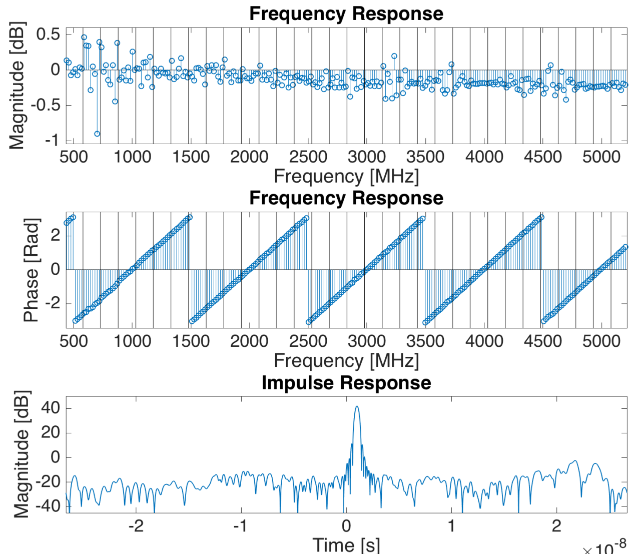

Fig. 5. Calibrated scan of frequency response and resulting impulse response, after a path length extension of 30cm. The impulse response is oversampled by a factor of 256 and has a peak located at 1.001×10^{-9} seconds.

[5] Sandra Costanzo, Francesco Spadafora, Giuseppe Di Massa, Antonio Borgia, Antonio Costanzo, Gianluca Aloi, Pasquale Pace, Valeria Loscrì, and Hugo O. Moreno, "Potentialities of usrp-based software defined radar systems," *Progress In Electromagnetics Research B*, vol. 53, pp. 417435, 2013.

[6] "GNU Radio Radar Toolbox," `https://grradar.wordpress.com/about/`, Accessed: 2017-4-6.

[7] Jonathan Ralston and Chad Hargrave, "Software defined radio: An open source platform for prototype GPR development," in *14th International Conference on Ground Penetrating Radar (GPR)*, 2012.

[8] "Device synchronization," `https://files.ettus.com/manual/page_sync.html`, Accessed: 2016-11-29.

[9] Waymond R. Scott, Jr., "Efficient drive signals for broadband cw electromagnetic induction sensors," in *IEEE International Geoscience and Remote Sensing Symposium*, 2013.

[10] Ettus Research, *UBX USRP Daughterboard, 10MHz-6GHz*, 9 2015, Rev. C.

[11] Samuel C. Carey, "Software defined radio for stepped-frequency, ground-penetrating radar," M.S. thesis, School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, April 2017.