

FPGA's Role in the Development of Synthetic SARs

The old adage "big things come in small packages" can accurately describe a Sandia National Laboratories effort to make synthetic aperture radars (SARs) smaller, lighter, and cheaper.

|By Dale Dubbert,
George Sloan,
and Armin Doerry

gathering, surveillance, reconnaissance, and precision weapons guidance. Miniature radar platforms include applications as varied as small UAV's, surveillance aircraft, and precision guided munitions, which must find their target in environments where a GPS is ineffective.

However, component miniaturization is only half of the solution. Optimal miniaturization requires a complete system redesign with consideration for modern devices, materials, processes, and algorithms.

An important ingredient in the miniSAR recipe is the extensive use of high-throughput digital signal processing functions as implemented using the latest high-speed, high-density FPGA technology. The modern FPGA affords miniaturization through the migration of former analog functions into the higher performance, programmable, and more robust digital domain. Additionally, advanced signal generation and processing techniques such as QECDWS, digital quadrature demodulation on receive, and variable bandwidth IF filtering all contribute to the simplification of the associated analog electronics that comprise the radar receiver/exciter (RX/EX) subsystem. Not only do FPGAs enable analog-to-digital migration, but they are also the perfect venue for the implementation of systolic architectures for real-time SAR image formation processing.

Because of their highly parallel nature, a 10 to 20 times increase in SAR processing density using modern FPGAs, as compared to current multi-node RISC processors, such as the PowerPC is expected.

Historically, FPGAs were much harder to program than a processor, but in recent years EDA tools (for this project Simulink from The MathWorks and System Generator for DSP from Xilinx) have made FPGA technology accessible to signal processing engineers.

The heart of miniSAR is the REA, which is a standard 3U compact PCI format with 8 module slots (See Figure 1). The REA performs all signal generation, receiving, processing, and control functions for the radar. Two custom high-speed digital modules form the foundation of the radar RX/EX subsystem, housed in the REA. These are the QECD-

GLOSSARY OF ACRONYMS

ADC - Analog-to-Digital Converter

APF - Azimuth Pre Filter

DAC - Digital-to-Analog Converter

DSP - Digital Signal Processor

EDA - Electric Design Automation

ENOB - Effective Number of Bits

FGPA - Field Programmable Gate Array

GPS - Global Positioning System

IF - Intermediate Frequency

LO - Local Oscillator

LTCC - Low-Temperature Co-fired Ceramics

LUT - Look-up Table

MAC - Multiply and Accumulate

PEC - Phase Error Correction

QECDWS - Quadrature Error-Corrected Digital Waveform Synthesis

REA - Radar Electronic Assembly

RISC - Reduced Instruction Set Computing

RAM - Random Access Memory

ROM - Read-only Memory

SAR - Synthetic Aperture Radar

SNR - Signal-to-Noise Ratio

UAV - Unmanned Aerial Vehicle

The coherent nature of the SAR signal combined with sophisticated processing algorithms presents an opportunity to exploit techniques such as high-sensitivity change detection, ultra-precise 3D terrain mapping, and moving target indication.

A new SAR program, under development at Sandia National Laboratories,

has the potential to create a new class of SAR — SAR systems that aren't large, heavy, power-hungry, and expensive. This "miniSAR" program is aimed at reducing the weight of its airborne sensor from 120 pounds to 30 pounds. The prototype is scheduled to fly in early 2005.

The Benefits of Miniaturization

SAR miniaturization has far-reaching benefits when applied to intelligence

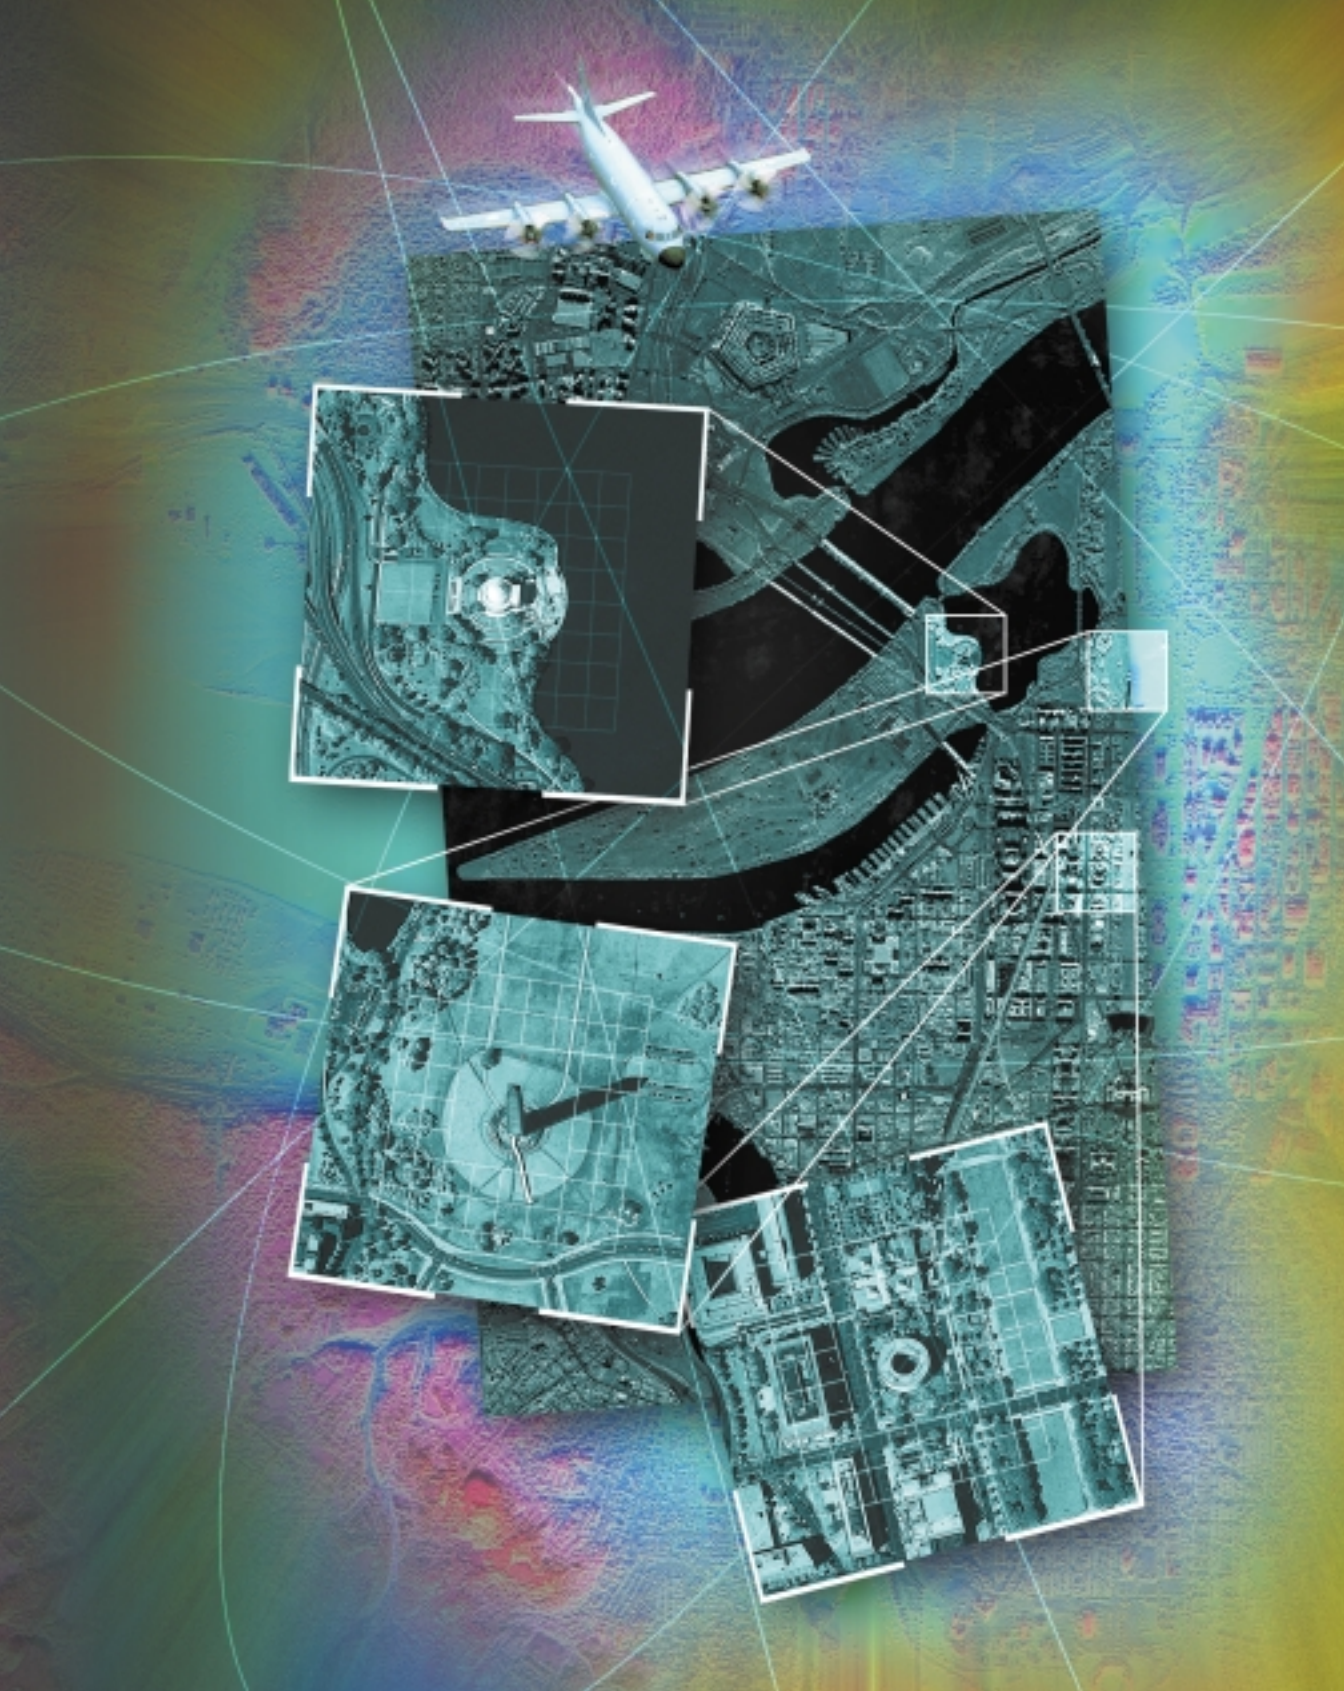


Illustration by Bill Yermal. Images courtesy of Sandia National Laboratories and Space Imaging.

WS and the digital receiver (DRX). Both the QECDWS and DRX perform critical radar signal generation and signal processing functions in modern FPGAs.

The QECDWS Model

The QECDWS is a 1.2 GHz synthesizer with dual DACs to produce quadrature outputs I and Q. The I and Q pair enables significant simplification in the analog

hardware required to convert the synthesized linear-FM “chirp” waveform to the desired microwave frequency. Also, the realizable bandwidth of the synthesized waveform is double that of a single-ended DWS, which is a huge benefit for wide bandwidth, ultra-fine resolution radars such as the miniSAR.

The major pitfall of quadrature synthesis is associated with the RF upcon-

verter, consisting primarily of a quadrature (single side-band) mixer, which exhibits poor quadrature image and mixer LO rejection. The image and LO leakage result in unwanted spurious energy in the radar waveform, which directly translates to a dynamic range reduction in the radar product (i.e., bad images).

This is where the FPGAs power and flexibility shine. First, the quadrature

image is removed by pre-distorting the amplitude and phase of the synthesized waveform. This is done using QEC loops in the synthesizer (See Figure 2). The entire synthesizer, with the exception of the DACs, is implemented in a 3-million gate FPGA operating at a clock frequency of 150 MHz. The basic chirp synthesizer consists of two accumulators, one for frequency and one for phase.

In a garden variety direct digital synthesizer, the output of the phase accumulator is fed directly into a sine ROM LUT, whose output drives a DAC. For quadrature synthesis, two ROMs (sine and cosine) feed dual DACs to produce a quadrature pair (I and Q). The quadrature amplitude and phase pre-distortion, to null the unwanted quadrature mixer image, is implemented using QEC RAM as a LUT.

The RAM address is generated using the instantaneous frequency output f_i from the frequency accumulator. For the quadrature phase correction, phase coefficients κ_{QEC} in the QEC phase RAM are selected and added to the output of the phase accumulator prior to the cosine ROM. For quadrature amplitude correction, a QEC magnitude RAM is addressed by f_i , generating amplitude coefficients κ_{QEC} , which are used to scale the amplitude waveform via a digital multiplier, either at the sine ROM or cosine ROM output. Which multiplier is active depends on the sign of the correction coefficient κ_{QEC} for a given f_i .

An additional correction loop, referred to as the common-mode PEC, applies a phase predistortion to both the sine and cosine branches of the synthesizer, via coefficients ϕ_{PEC} loaded in the PEC RAM. The PEC is used to remove the effects of frequency-dependent, non-linear phase in the RF and microwave portion of the radar RX/EX. The result is a sharper radar impulse response, which directly translates to improved image acuity.

As for the undesired mixer LO leakage, precise DC offsets (I_{DC} and Q_{DC}) are applied prior to the DAC. Proper selection of I_{DC} and Q_{DC} have been shown to reduce the undesired LO leakage from approximately -30 dBc to less than -60 dBc at a center frequency of 3.3 GHz. The downside of this technique is the



Ku-Band SAR image of Washington, D.C., created in cloud cover.

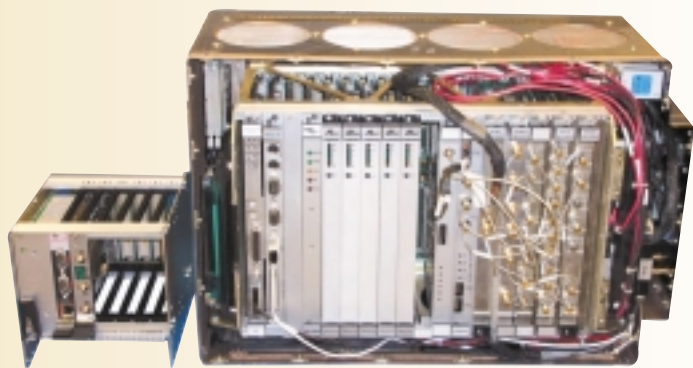


Figure 1. “miniSAR” radar chassis at 10 pounds (left) next to current state-of-the-art radar chassis at 60 pounds (right).

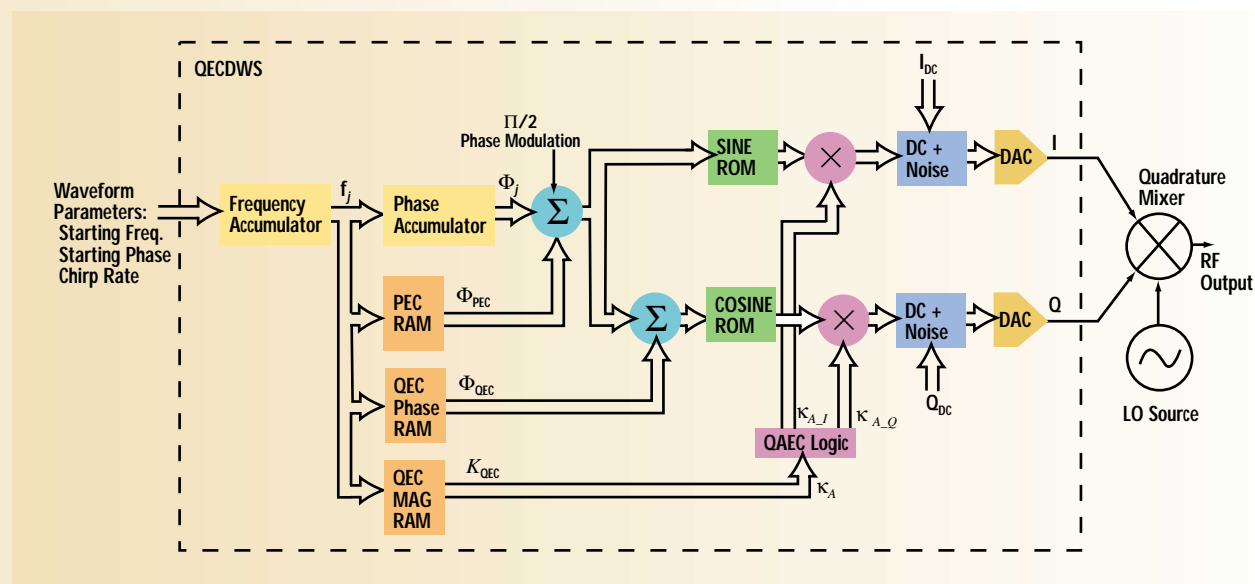


Figure 2. QECDWS simplified block diagram with quadrature mixer.

requirement of a high precision DAC to properly convert the precise digital DC value (in the micro-volt range) to analog. For high-speed DACs (>1 GHz), the precision or ENOB is not adequate. The solution is to apply a pseudo-random noise dither to the DC value prior to the DAC. The noise dither allows the effective precision of the applied DC terms, as integrated over several DAC samples, to be significantly less than the quantization precision of the DAC.

For radar applications, this works well since many DAC samples (tens of thousands) are integrated in a single radar pulse.

The result of the corrections as applied to the QECDWS is a spectrally pure output at the quadrature mixer. For radar systems with relatively narrow to moderate bandwidth requirements, no additional bandwidth multiplication is needed. Thus, a QECDWS in concert with a single quadrature mixer and LO source can produce the necessary microwave radar signal directly. For wide bandwidth systems, a frequency multiplication stage is added to achieve the required RF bandwidth and center frequency (3 GHz centered at 16.7 GHz for the miniSAR).

All QECDWS coefficients are determined using built-in calibration, allowing an end-to-end radar loop to be utilized to determine the optimal coefficients required to maximize the dynamic

range of the received and processed radar signal.

There are multiple designs for SARs and designers can approach SAR devices from a number of directions. For Sandia-built SARs, the use of dynamic programmability of the synthesizer parameters is the preferred approach. The reason behind this particular approach is that by applying pulse-to-pulse changes to the fundamental waveform parameters (phase, frequency, and chirp rate), radar data can be collected in a way that simplifies the ensuing image formation process. This simplification leads to fewer operations per image. Combined with a high compute-density processing engine, the result is a two-fold reduction in the size and power of the image formation processor.

The DRX Module

Another significant element in the high-speed digital subsystem, which further improves SAR design, is the DRX module (See Figure 3). The input of the SAR DRX is a 300 MHz IF, which is sampled by an 8-bit ADC operating at a 1.2 GHz clock frequency. For certain hardware realizations, the analog front end of the receiver, which produces the 300 MHz IF, includes a band-limited noise source to insure that the ADC quantization errors are randomized. This is extremely important in a high-performance SAR where 80 to 100 dB or more of SNR improvement is applied during image formation. The noise, or dither, source is spectrally separate from the signal, allowing the FPGA-based DSP to remove the noise after the ADC. The benefit is as much as an 11 dB increase in the image output SNR.

The DRX employs many functions that either were analog, or simply never existed in earlier generation radar systems. Good examples of DRX functions that have migrated from analog to high-performance reconfigurable logic are the quadrature demodulator and the first stage range pane filter.

The quadrature demodulator in Figure 3 accepts the 1.2 GHz ADC data and performs both a demodulation and decimation filtering function, producing quadrature (I and Q) outputs at a 300 MHz rate.

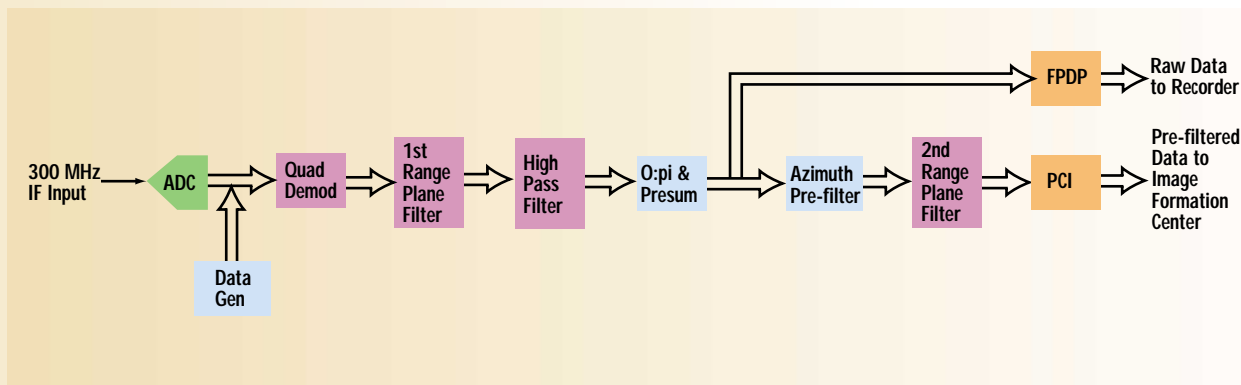


Figure 3. DRX simplified block diagram.

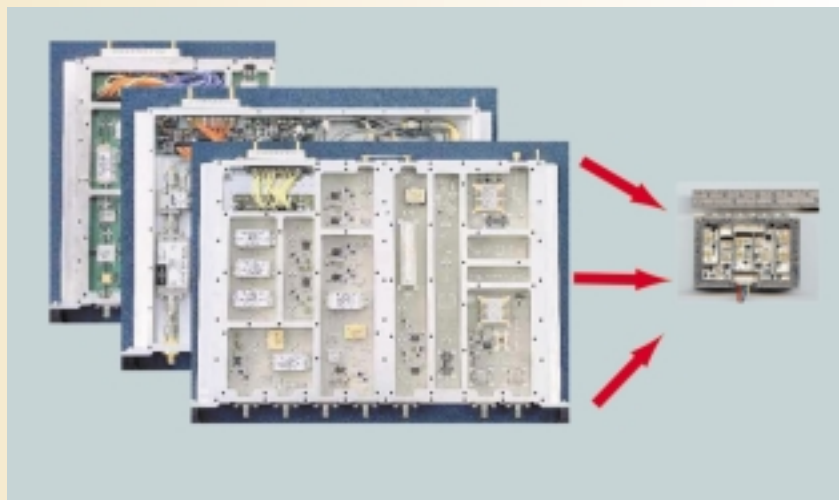


Figure 4. Analog RX/EX comparison of older generation microstrip (left) to miniSAR LTCC (right).

The image rejection of the demodulator is >60 dBc — 20 dB better than the older generation SAR analog quadrature demodulator. Furthermore, the performance is scalable within the FPGA. In other words, better image rejection can be achieved by simply using more FPGA resources. For FPGAs, more gates equates to better performance. This is one of the fundamental benefits of reconfigurable logic for high-end radar applications where it is desired or necessary for system performance to ride the technology curve (also known as Moore's Law). Fortunately, FPGA speed and gate density improve every year.

Another DRX element, one with analog ancestry, is the variable decimation range plane filter. This filter is based on multiple panes that are basically high-performance MAC cells. The MACs are efficiently implemented in modern

FPGAs where high-performance multipliers and associated high-speed RAM are embedded in the FPGA fabric. Two features of the filter allow its decimation factor, and hence its bandwidth, to be changed on the fly: its polyphase structure, and the programmable delay of the connections between its panes. A two-stage implementation of the pane filter (See Figure 3) allows selection of one of 38 IF bandwidths, whereas the former analog receiver switched between one of three discrete filters. The additional IF bandwidths are paramount to the proper selection of radar pulse width as a function of radar operating range, which in turn maximizes the system SNR. The added feature of programmable filter coefficients adds further to the flexibility of the DRX.

An additional filter, termed the APF allows pulses or azimuth samples, which

are vectors of pre-filtered samples from the ADC, to be filtered in the pulse-to-pulse or radar Doppler domain. The APF is extremely valuable to systems such as miniSAR where Doppler bandwidths as illuminated by the radar antenna exceed the Doppler bandwidth of the processed image. The APF can significantly reduce the real-time processor load for fine-resolution imaging scenarios. The primary logic resource required by the APF is fast, dual-port RAM, which is abundant in modern FPGAs.

Also, in the DRX, there is a feature to facilitate module- and system-level testing. A built-in data generator, which is basically a dual accumulator direct digital synthesizer, allows contrived IF signals to be generated without the analog receiver connected to the DRX. The phase, frequency, and chirp rate of the data generator synthesizer are fully programmable, allowing simulation of point and pseudo-distributed targets at various ranges. Pulse-to-pulse phase and frequency modulation are also available to simulate radar Doppler effects due to motion. A linear-feedback shift register noise generator is implemented in the data generator to emulate a random noise floor.

With all functions active, the DRX processor, as implemented in two six-million gate FPGAs, demonstrates a raw DSP throughput of 32 GMACs per second (32 billion multiply-accumulate operations per second).

Attacking Size and Weight

This article has described ways in which the FPGA has enabled SAR miniaturization. The other half of the challenge is to reduce the size and weight

of what's left. In terms of the remaining analog portion of the radar RX/EX, the recent size reduction has been phenomenal. The use of LTCC multilayer RF circuit technology, in conjunction with the analog subsystem simplification as enabled by the modern FPGA, has resulted in a 40 times reduction in the physical size of this component over microstrip modules (See Figure 4).

Besides the REA (shown in Figure 1), miniSAR requires a physically separate antenna and gimbal assembly. The miniSAR antenna and gimbal assembly demonstrates a 4 times reduction in weight and a 10 times reduction in swept volume over the older-generation SAR antenna and gimbal assembly.

Eventually, the gimbal, transmitter, and antenna functions will be integrated into a highly compact, lightweight active phased array electronically steerable antenna. This, in combination with further reductions in the size and weight of the REA will realize the ultimate miniSAR goal of a sub-20 pound SAR.

WD&D

About the authors:

Dale Dubbert's areas of specialty include RF and microwave design, signal analysis, SAR system design, and digital signal processing. He received his MS degree in electrical engineering in 1986 from Kansas State University.

Armin Doerry has worked in numerous aspects of synthetic aperture radar design since 1987. He has author/coauthor of over 70 papers and technical reports, and holds 4 patents. He received his Ph.D. degree (1995) from the University of New Mexico, Albuquerque, New Mexico. Dubbert and Doerry are distinguished members of the Technical Staff at Sandia National Labs.

George Sloan is a Principal Member of the Technical Staff and an adjunct instructor of Electrical Engineering at the University of New Mexico. Mr. Sloan received his MSEE degree from Brigham Young University in 1987. Since then he has worked in numerous aspects of radar design and development, and is currently the project leader for Sandia's SAR miniaturization program.

Ken Karnofsky is the Marketing Director of DSP and Communications for The MathWorks. He holds a degree in Systems Engineering from the University of Pennsylvania. Mathworks is the principle collaborator on this project.