An Onboard Processor and Adaptive Scanning Controller for the Second-Generation Precipitation Radar

Mark A. Fischman, *Member, IEEE*, Andrew C. Berkun, William Chun, Eastwood Im, *Senior Member, IEEE*, and Raymond J. Andraka, *Member, IEEE*

***Abstract*—Technology for the 14- and 35-GHz Second-Generation Precipitation Radar (PR-2) is currently being developed by the National Aeronautics and Space Administration Jet Propulsion Laboratory to support the development of future spaceborne radar missions. PR-2 will rely on high-performance onboard processing techniques in order to improve the observation capabilities (swath width, spatial resolution, and precision) of a low-earth orbiting rainfall radar. Using field-programmable gate arrays (FPGAs), we have developed a prototype spaceborne processor and controller module that will support advanced capabilities in the PR-2 such as autotargeting of rain and compression of rainfall science data. In this paper, we describe the new technology components designed for the onboard processor, including an FPGA-based 40×10^9 op/s pulse-compression receiver/filter with a range sidelobe performance of -72 dB, and an adaptive scanning controller which yields a six-fold increase in the number of radar looks over areas of precipitation.**

***Index Terms*—Fault tolerance, field-programmable gate arrays (FPGAs), Global Precipitation Measurement (GPM) Mission, pulse compression, rain, real-time processing and control, spaceborne radar, Tropical Rainfall Measuring Mission (TRMM).**

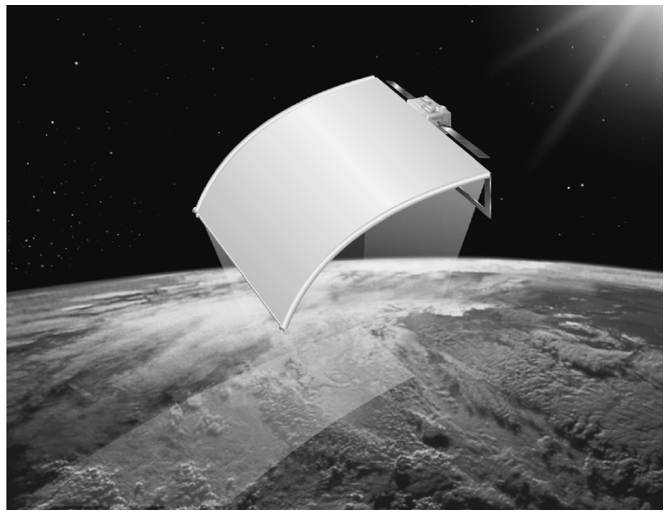

Fig. 1. Conceptual drawing of the spaceborne Second-Generation Precipitation Radar (PR-2), with an innovative 5.3 × 5.3 m reflector antenna based upon deployable membrane technology.

I. Introduction

ADVANCED radar technology has long been sought in the development of an earth-observing satellite system for global measurements of precipitation. The key capability of a spaceborne rain radar is that it can recover the vertical profile and structure of precipitation to infer how liquid water and latent heat are transported in the atmosphere. These data on water mass and heat transport, which until recently have not been well characterized, are essential for monitoring weather dynamics and improving the accuracy of weather forecasting models [1].

The Ku-band (13.8 GHz) Precipitation Radar (PR) onboard the Tropical Rainfall Measuring Mission (TRMM) satellite, launched in 1997, was the first earth-orbiting sensor to retrieve detailed rainfall profiles over tropical and mid-latitude regions [2]. With the success of TRMM, a follow-on instrument called the Dual-Frequency Precipitation Radar (DPR), consisting of one Ku-band radar and one Ka-band (35.6 GHz) radar with coaligned antenna beams for improved rainfall detection, is currently being developed as part of the upcoming Global Precipitation Measurement (GPM) Mission [3]. DPR inherits a conventional slotted-waveguide phased array antenna technology from TRMM which practically limits the aperture size to $\sim$2 m, due to the bulk and mass of the antenna hardware. This upper bound on the antenna aperture limits the horizontal resolution to 5 km and, to a lesser extent, the usable cross-track swath to approximately $\pm 17^\circ$ from nadir (250-km maximum swath coverage in low-earth orbit for the Ku-band radar).

To attain the spatial resolution, swath width, and measurement precision that will be needed to observe rain down to the scale of convective cells using an earth-orbiting spaceborne radar, the National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory (JPL) is leading the development of some breakthrough technologies for a Second-Generation Precipitation Radar (PR-2), with dual-frequency, dual-polarization, and wide-swath scanning capabilities [4]–[6]. Fig. 1 illustrates the concept for the spaceborne PR-2 instrument, with its lightweight antenna deployed in a low-earth orbit. The PR-2 science observations will consist of vertical structures of rain reflectivity at both 13.6 and 35 GHz and at both copolarization and cross-polarization. Such a comprehensive set of measurements from a global satellite sensor is intended

Manuscript received March 1, 2004; revised September 10, 2004. This work was supported by the National Aeronautics and Space Administration.

M. A. Fischman, A. C. Berkun, W. Chun, and E. Im are with the Radar Science and Engineering Section, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109-8099 USA (e-mail: mafisch@mail1.jpl.nasa.gov).

R. J. Andraka is with the Andraka Consulting Group, Inc., North Kingstown, RI 02852 USA (e-mail: ray@andraka.com).

Digital Object Identifier 10.1109/TGRS.2005.844670

TABLE I
SYSTEM PARAMETERS FOR THE SPACEBORNE PR-2

Description	Parameter values	
Spacecraft altitude	400 km	
Operating frequencies	Ku-band = 13.6 GHz	Ka-band = 35 GHz
Polarization	HH, HV	
Effective antenna aperture	5.3 m (Ku-band)	2.1 m (Ka-band)
Antenna gain	55 dBi	
Cross-track scan angle	±32°	
Peak transmit power	200 W (Ku-band)	50 W (Ka-band)
Chirp bandwidth B	4 MHz	
Chirp pulsewidth τ'	50 μs	
A/D sampling rate	20 MHz	
A/D bit resolution	12 bits	
Slant altitude	8 km (locator sweep)	12 km (high-resolution sweep)
Vertical range sampling (nadir)	30 m	
Post-averaged vertical resolution	250 m (high-resolution sweep)	
Pulse repetition frequency (PRF)	4.72–6.52 KHz	
Dwell time	0.26 s	
Horizontal resolution (nadir)	2 km	
Ground swath	500 km	
Noise-equivalent Z_{eq} (single pulse)	10.0 dBZ (Ku-band)	9.8 dBZ (Ka-band)
Effective number of independent samples	~190 (high-resolution sweep)	

to provide information for detailed studies of the microphysics and dynamics of rainfall, snowfall, and dense or precipitating clouds in various climatic regimes and geolocations.

The PR-2 operational geometry consists of a downward looking radar which scans its dual-frequency antenna beams across the flight track, with each scan beginning at nadir and ending at look angles of $\pm 32^\circ$ from nadir. The scanning antenna consists of a 5.3 × 5.3 m cylindrical/parabolic membrane reflector antenna, and a 14- and 35-GHz dual-frequency, electronically steerable feed array. The feed is configured such that the aperture at 35 GHz is under illuminated to provide matched beams at the two frequencies. The RF electronics subsystem consists of: a digital chirp generator used to synthesize a linear frequency-modulated chirp waveform; compact, solid-state transmit/receive (T/R) modules for the dual-frequency electronic beam steering; and upconverters and downconverters for the receive channels. The chirp waveform is generated at an IF frequency and upconverted to both 14 and 35 GHz. The signals are amplified to the desired radiated powers using linear power amplifiers in the solid-state T/R modules, and then sent to the dual-frequency feeds and antenna. There are four receiver channels—two for 14 GHz (H- and V-pol) and two for 35 GHz (H- and V-pol). These measurements are acquired at a horizontal resolution of 2 km and an intrinsic vertical resolution of ~40 m. A summary of the system parameters for the spaceborne PR-2 are listed in Table I.

The PR-2 project also involves the development of a high-performance onboard digital data processor and timing controller, implemented in field-programmable gate arrays (FPGAs). The specially developed electronics for processing and control include an advanced feature for adaptive steering of the coaligned Ku- and Ka-band radar beams. Adaptive scanning complements the 5.3 × 5.3 m aperture antenna technology to make it possible to collect 3-D rainfall profiles with a greater-than-2 times improvement in horizontal resolution (2 km) and in swath width (500 km), compared to the TRMM PR or the GPM DPR.

In this paper, we will describe the development of the PR-2's onboard processor/controller system. The processor must meet the challenging requirements for making improved precipitation profiling measurements from space—with better than 60-dB clutter suppression capability [7] for capturing both rain and sea surface returns, with an efficient sampling method to collect multiple radar looks within each rain bin, and with four receive data channels for acquiring both co- and cross-polarized measurements at 14 and 35 GHz. At an analog-to-digital (A/D) converter sampling resolution of 12 bits and 20 MSPS per channel, the PR-2 must support a raw input data rate of 960 Mbit/s. This high input rate drives the radar hardware design toward an onboard processing approach. A primary benefit of onboard processing is that the volume of rainfall data collected over wide swaths can be reduced significantly before it is downlinked to the ground. Also, onboard processing becomes desirable for gaining greater instrument autonomy for advanced radar techniques, such as wide-swath adaptive scanning and pulse compression, that improve the ground swath coverage and rain measurement sensitivity. The design of our FPGA-based data processor for pulse compression and multilook averaging of the received radar echoes is described in Section II. In Section III, a unique timing solution for adaptive scanning of rain targets is presented along with the design of the timing control electronics. The fabrication and testing of the prototype spaceborne PR-2 processor electronics are described in Section IV, followed by a summary in Section V of the technology maturity for an advanced rain measuring mission.

II. PULSE COMPRESSION PROCESSOR

The PR-2 uses a chirp radar technique, which has advantages over the short-pulse radar systems aboard TRMM and GPM that typically sample at a vertical range resolution of 250 m ($\tau = 1.67$ μs transmit pulse). In the PR-2 chirp radar, signal power is gained through the linear amplification and transmission of a longer pulse length ($\tau' = 50$ μs), and range resolution is gained through the chirp bandwidth ($B = 4$ MHz), resulting in an improved intrinsic vertical resolution of $c/2B \approx 40$ m. With the finer vertical sampling, a number of independent radar looks can be averaged together in power within each 250-m vertical resolution volume to reduce Rayleigh fading noise. An increase in the number of samples by a factor of $250\,\text{m}/40\,\text{m} \approx 6$ improves the effective signal-to-noise ratio of the rain echo measurement by as much as $\sqrt{6} \approx 4$ dB. A further benefit of the chirp radar is that, compared to a traditional short-pulse radar, the required instantaneous peak transmit power is reduced significantly by the chirp compression ratio $B\tau' = 200$. (For the PR-2, a transmit power on the order of 50–200 W is sufficient to detect rainfall at a sensitivity of ~10 dBZ.) These moderate transmit power requirements for the chirp radar allow simpler, solid-state transmitter hardware to be integrated into the phased array antenna feeds.

The cost of the chirp radar is that digital pulse compression processing after the RF downconversion and A/D conversion stages is numerically intensive. For time-domain pulse compression using matched finite impulse respone (FIR) filter banks, the number of complex filter coefficients is driven by the 200-fold compression ratio. (A 256-tap, 51.2-μs digital filter design was chosen in the PR-2 design to meet this compression ratio requirement with slight oversampling of the chirp reference function.) Amplitude weighting (windowing) of the reference function must also be appropriately chosen to achieve the −60 dB or better range sidelobes needed to mask out sea surface clutter from rain echoes. The choice of window becomes a tradeoff between range resolution and range sidelobe levels [8]. Our design

Fig. 2. PR-2 airborne prototype radar. (a) Electronic processing and control console aboard the NASA P-3 Orion aircraft. (b) Dual-frequency (14 and 35 GHz) antenna feed structure with mechanically scanning reflector (shown at an antenna range facility). (c) P-3 aircraft deployed for the Wakasa Bay/AMSR-E validation experiment, January–February 2003.

simulations of various window tapering functions showed that an optimal trade could be met for examining rain near the ocean surface by using the square of a $\beta = 6$ Kaiser window [9]. Pulse compression for this Kaiser window (to be shown in Section IV) resulted in a 3-dB range width of 470 ns for a 50-μs chirp pulse. The near surface sidelobes reached -55 dB after only 1.7 μs from the peak return and rapidly fell below -70 dB thereafter.

The dynamic range variation of 60 dB between sea-surface and rain returns also drives the A/D converter resolution in the PR-2. A/D converter dynamic range above the quantization noise floor increases with bit resolution b as [9]

$$DR \approx 6b - 1.25\,[\text{dB}]. \tag{1}$$

Therefore the choice of 12-bit conversion resolution yields a sufficient dynamic range margin of 71 dB in the preprocessed echoes. After pulse-compression (coherent averaging by the time-bandwidth factor of 200-times in power, or 23 dB), a net ideal dynamic range of 94 dB is achieved from the quantization noise floor to the maximum unsaturated return.

Taking into account the complex I/Q arithmetic for the matched filter, the requirement for four separate receive channels (Ku- and Ka-band, horizontal and vertical polarization), and assuming an above-Nyquist complex sampling rate of 5 MHz, the peak processing rate of the PR-2 is

$$\begin{aligned} &256 \text{ filter taps} \times 5 \text{ MHz baseband sampling} \\ &\quad \times 4 \text{ real filters per channel} \\ &\quad \times 2 \text{ math ops. per tap (multiply and add)} \\ &\quad \times 4 \text{ radar channels} = 40.96 \times 10^9 \text{ op/s}. \end{aligned} \tag{2}$$

Traditional microprocessor or digital signal processor (DSP) technology would not be a practical choice for meeting the 40 billion op/s performance requirement, as a large number of processing chips would be needed on board. As an alternative, the recent emergence of high-speed radiation-hardened FPGAs with million-gate densities makes it possible to support this processing rate—while maintaining a low chip count—through the massive parallelization of math operations.

The PR-2 pulse compression data processor concept has been implemented onto two FPGA chips using a distributed arithmetic technique [10]. We have tested the first version of the processor core aboard the airborne prototype of the PR-2 [6], which has successfully collected dual-frequency rainfall data in several airborne field campaigns. These missions include: the 4th Convection and Moisture Experiment (CAMEX-4) [11]; the Wakasa Bay/AMSR-E (Advanced Microwave Scanning Radiometer-EOS) validation experiment [12]; and a joint radiometer-radar technology demonstration with the 10.7-GHz Lightweight Rainfall Radiometer (LRR-X) instrument [13]. Fig. 2 shows the PR-2 radar hardware installed onto the NASA P-3 Orion aircraft as it is deployed for the Wakasa Bay experiment. The pulse-compressed and processed science data from the CAMEX-4 and Wakasa Bay flights are shown in Fig. 3. The multiparameter Ku- and Ka-band rain profiles in these figures include: copolarized radar reflectivity data; vertical Doppler rainfall velocity profiles; and linear depolarization ratios, which accentuate scattering by water-ice particles at the melting layer. As a next step in the PR-2 processor technology development, the data processor design has been modified to interface with a more sophisticated timing controller designed for satellite-based measurements, where a large number of echoes-in-flight must be carefully tracked and averaged together over time.

The signal chain for the PR-2 data processor is illustrated in Fig. 4. The onboard processor is designed around a bit-serial finite impulse response (FIR) filter implemented in a pair of Xilinx Virtex-1000 FPGAs. Data enter the FPGAs as four channels of 12-bit digitized data at a 20-MHz sampling rate. The 4-MHz chirp data are centered around a 5-MHz offset video carrier. The remainder of the bandwidth is noise which has leaked through the front end filters.

The first stage of processing is composed of a complex demodulator, which translates the 20-MHz digitized samples into 5-MHz complex baseband data. These complex data represent the vertical range samples of the rainfall profile taken at 30-m intervals. Functionally, the complex demodulator consists of a mixer (multiplication by a 5-MHz complex exponential), a

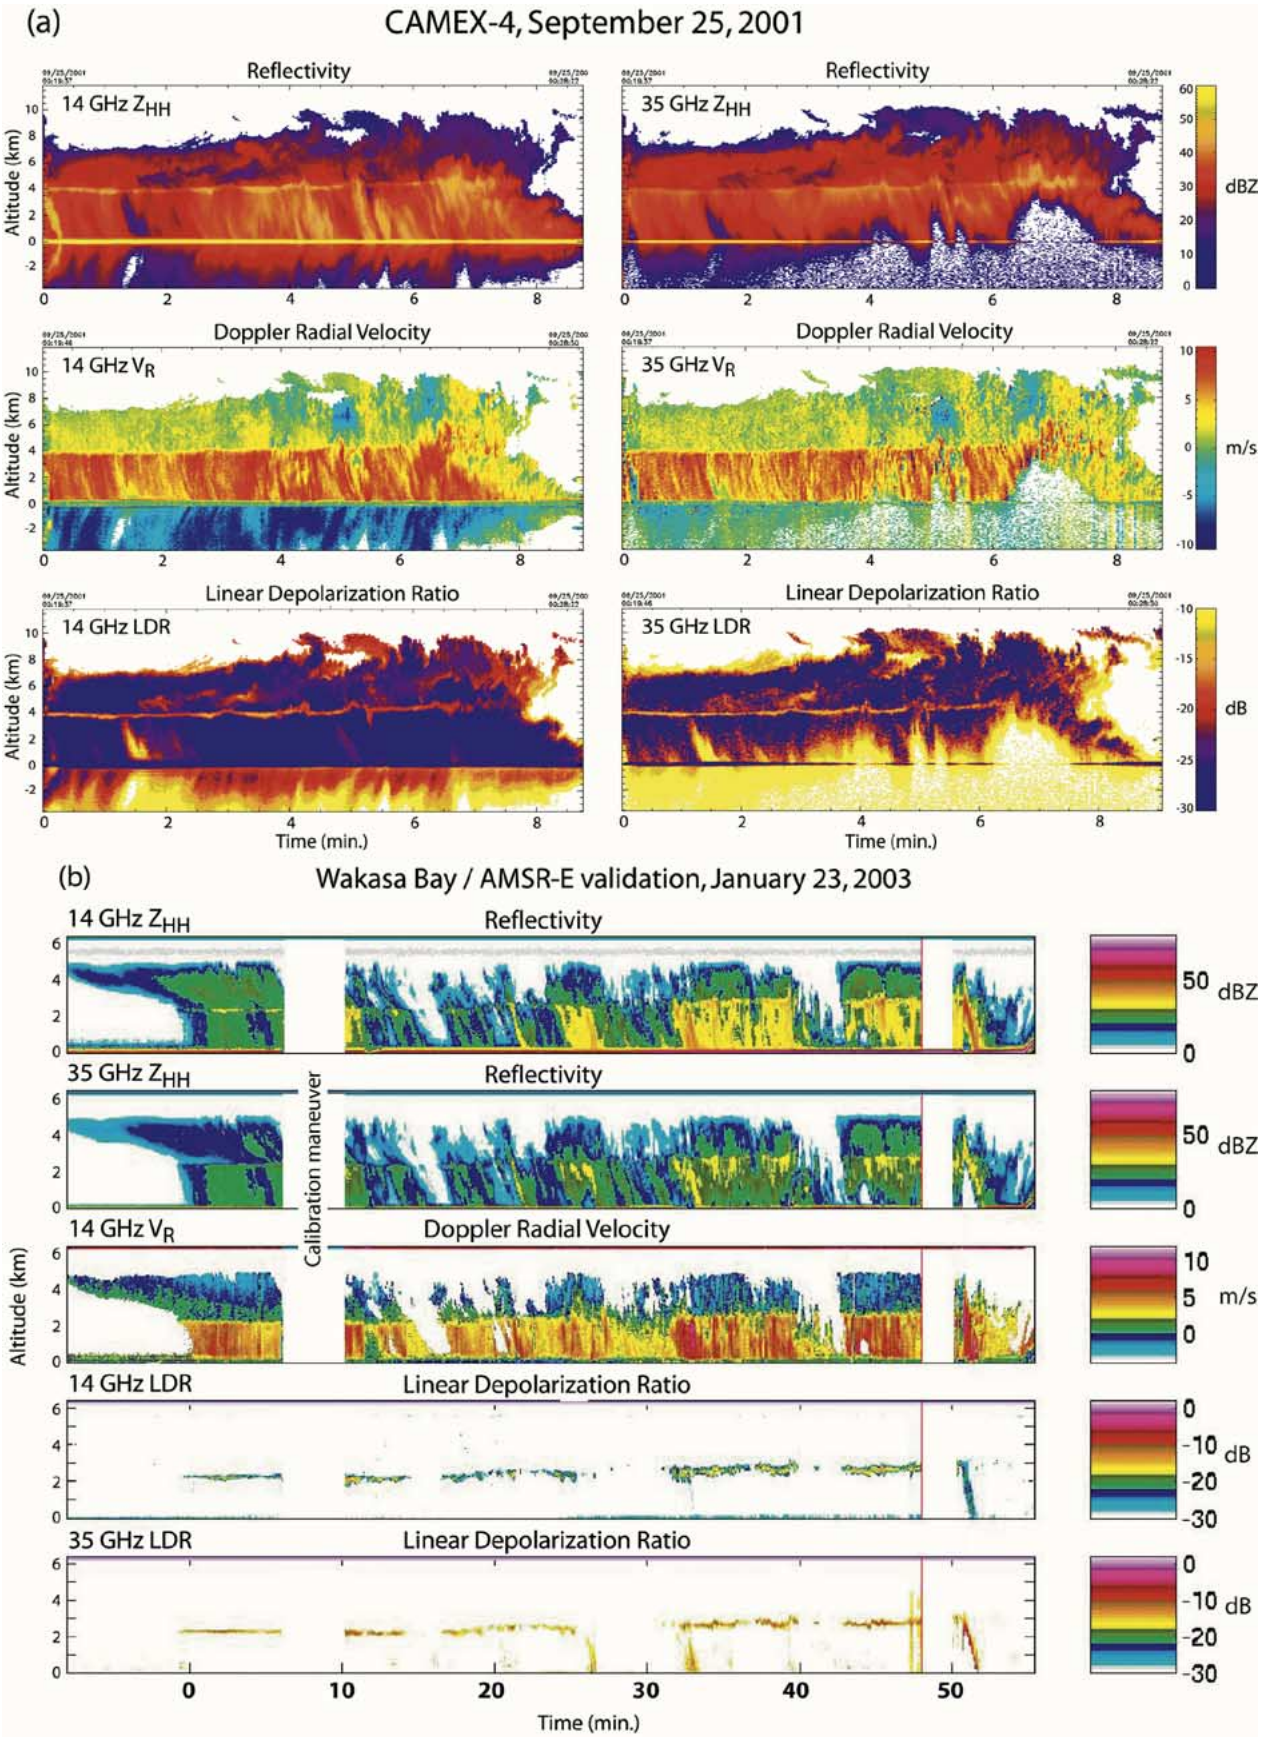

Fig. 3. Multiparameter dual-frequency rainfall data collected by the airborne version of the PR-2 FPGA processors during (a) CAMEX-4 while profiling tropical cyclone Humberto. (b) The Wakasa Bay/AMSR-E validation experiment during an hour-long flight line over the Pacific Ocean, east of the main island of Japan. Graphs show the along-track (nadir) radar observations for copolarized rain reflectivity data at 14 and 35 GHz (Z_{HH}), vertical Doppler velocity measurements (V_R), and linear depolarization ratios (LDRs).

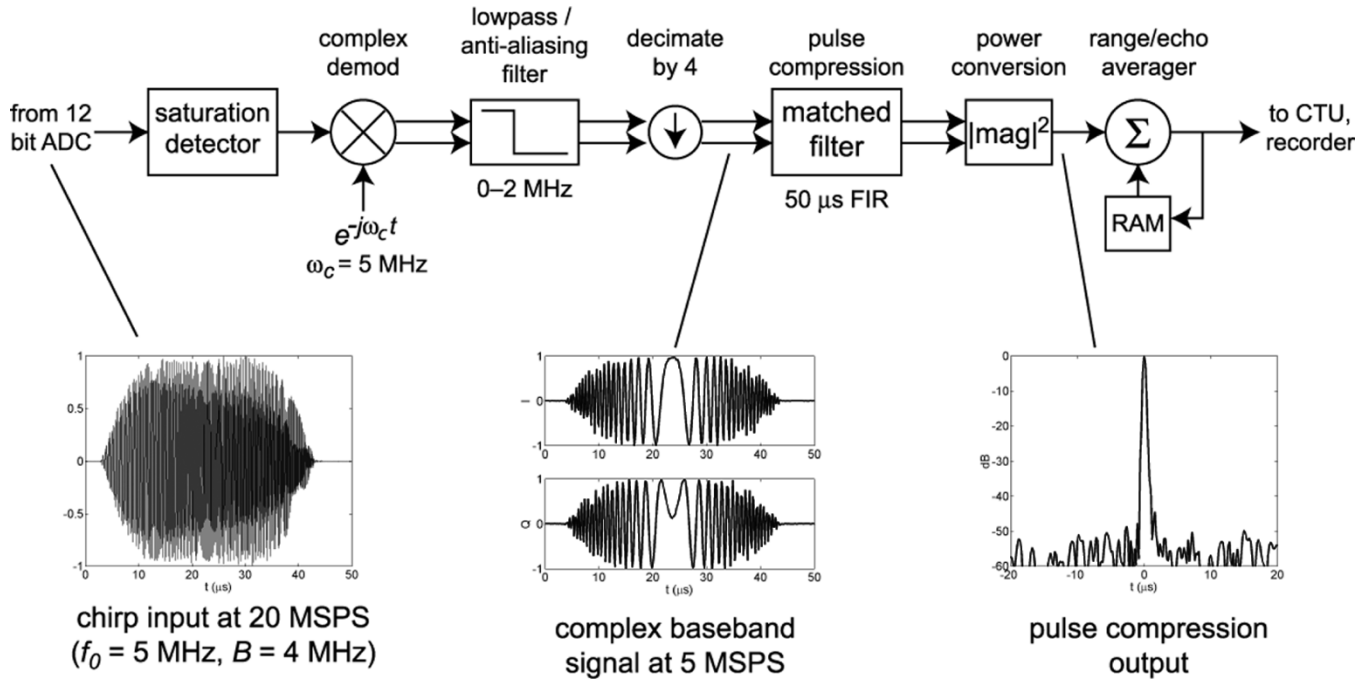

Fig. 4. Functional block diagram of the digital data processor for the PR-2.

low-pass antialiasing filter, and a decimator which reduces the output data rate by a factor of 4. There is no loss of information in the decimate-by-4 operation, given that the signal of interest is limited to less than 5-MHz bandwidth and that the aliased frequency components have been removed. The complex demodulator is implemented with a polyphase demodulator, which combines the mixing, filtering and decimation operations in one stage [10], [14]. A 64–tap, 16-bit polyphase FIR filter is included in the digital hardware implementation of this demodulator.

The 5-MHz complex baseband data are stored in a first-in/first-out memory so that digitized echoes containing saturated samples (digital samples at the positive or negative rails of the A/D converter) can be detected and eliminated before range compression. (Saturated digital samples occur when there is an unusually high backscattered speckle from the sea surface). Range line data which is saturation-free is fed into a 256-tap, nonsymmetric FIR filter containing the complex conjugate of the expected return. This is where the bulk of the FPGA processing resources are used, amounting to 20 billion real multiplication and addition operations per second per chip. The output of this 256-tap filter is the fixed-point compressed radar return. Complex data are converted to power and coherently averaged over multiple range samples and multiple independent pulses to improve the signal power estimate and provide onboard data volume reduction. The adaptive scanning technique, described in the next section of this paper, can be used to identify the radar beams most likely to contain rain. It will be shown that with adaptive scanning, the averaging of multiple radar looks in range and in time can provide over two orders of magnitude science data volume reduction.

A number of digital signal processing applications make this processing density possible. Bit-serial filtering makes the best use of the speed capabilities of the FPGA. The FIR filter stage uses a lookup table to combine four stages of multiplication and addition operations into a single configurable logic block. The video filter, demodulator, and downsampler are actually four parallel filter stages which compute only the samples which are kept after decimation. Complex-valued filtering is implemented by running delay stages and multipliers at double speed and switching between real and imaginary components of each received waveform, saving half of the FPGA's logic resources.

III. Adaptively Scanning Radar

A. Timing Solution

A unique feature of precipitation is that it is often sparsely distributed over the earth's surface. Even for weather in tropical regions, which accounts for more than half of the total global rainfall, precipitation occurs over only 4% of the surface area [7]. This statistic can be exploited in the PR-2 by using an adaptive scanning technique [15]—that is, the radar can use preliminary "quick-scan" data to electronically steer the radar beam to only those areas which contain precipitation, and to ignore areas that are precipitation-free.

Adaptive scanning becomes all the more necessary as the observation requirements for future spaceborne rainfall measuring missions move toward higher horizontal resolution and larger swath widths—the two opposing requirements that compete against the available dwell time. Compared to TRMM's PR observations or to the DPR observations planned for the GPM mission, the PR-2 represents a 12-fold increase in the number of radar beam positions per swath area (a six-fold increase due to horizontal resolution improvements in two dimensions and a two-fold increase in swath width). Given the limited dwell time of a 7-km/s low-earth orbiting satellite, the timing sequence for interlacing the transmit and receive beams becomes critical. The radar must be able, in real time, to selectively dwell on only those beam locations that have rainfall and then generate a transmit/receive timing solution that can capture 60 or

TABLE II
ADAPTIVE SCANNING PARAMETERS

	Locator sweep	High-resolution sweep
Vertical resolution	2 km	250 m
Number of vertical samples per resolution volume	67	8
Number of cross-track beam locations	250	24
Number of pulses per beam location	2	32
PRF	5.54-6.52 KHz	4.72-5.56 KHz
Available sweep time	90 ms	170 ms
Net number of samples averaged	67×2 = 134 times	8×32 = 256 times

more independent radar looks needed [16] to accurately estimate the reflectivity of each rain bin. Conventional cross-track scanning radars, such as the TRMM PR, scan in a predetermined, right-to-left sequence at a constant pulse repetition frequency (PRF). However, a constant PRF scanning method does not make the most efficient use of the satellite's available dwell time over the rain target, due to the resultant "dead time" or gaps between the transmitted and received pulses. A more intricate transmit and receive timing scheme with variable PRF has been proposed for the GPM DPR design to increase the number of radar looks within the cross-track sweep period as the beam is steered inclusively, from right to left [17]. This scheme does not accommodate adaptive scanning, though, where the beams must be steerable over an arbitrary sequence of positions.

A unique, variable PRF timing solution has been developed for the electronically scanning PR-2 radar that can autotarget rain having any possible spatial distribution. The entire scan sequence is 260 ms long (to provide an along-track sampling interval of approximately 2 km), and consists of the following steps: 1) a 90 ms *locator sweep* (quick-scan) that performs a coarse measurement of the range profile return power over a 0–8-km altitude and over all cross-track beam locations (∼250 beams); 2) a *bubble-sort* algorithm which ranks the 250 beam locations in order of highest to lowest radar return powers; and 3) a 170 ms *high-resolution sweep* in which a fraction of the cross-track beam positions with the strongest rain echo profiles are resampled over a larger altitude (12-km was used for the technology demonstration) and over a greater number of radar looks. Table II lists the parameter values used to execute the locator and high-resolution sweep sequences.

The input to the PR-2 steering algorithm is a sequence of beam locations to be observed for either locator or high-resolution sweeps. The timing algorithm then can be posed as a counting problem: In what order should the beams be counted, and how should the PRF be varied, to gather the largest possible number of independent radar looks from the rain scene? Simultaneously, how can it be guaranteed that there are no overlaps between the radar's transmit and receive intervals? Our timing solution is based on several rules to meet these requirements:

1) The receive window length (the slant altitude range for the rain scene) is held constant throughout the sweep.
2) The beam sequence always moves from nadir to alternating left and right sides of the cross-track scan—that is, from the shortest to longest range delays to the target—so that the average PRI (pulse repetition interval) is monotonically increasing.
3) The number of echoes-in-flight (EIFs) is determined by the largest integer number of pulses that can fit within the shortest (first) range delay of the scan sequence

$$EIF = floor\left(\frac{2H}{c \cdot PRI}\right) \tag{3}$$

where the $floor$ function is defined here as rounding the argument down to the nearest integer, and H is the spacecraft altitude; this EIF value is held constant throughout the scan.
4) The beam scanning order is counted in an alternating pattern within pairs of adjacent beams (e.g., a beam number sequence $\{1,2,1,2,1,2,1,2,\ldots,3,4,3,4,3,4,3,4,\ldots\}$). This ensures first that the satellite, traveling at ∼7 km/s, has time to move laterally between pulses by one-half of the 5-m aperture length so that multiple radar looks of the same rain bin are statistically independent. Second, the adjacent beam sequence ensures that each series of pulses has essentially the same round-trip range delay time (and therefore equal PRF over each beam position pair) so that Rule 2 is never violated.

Using the above rules for transmit/receive timing and beam steering, the PR-2 can focus its beam resources mainly on those areas of the swath that contain rain. This, in turn, buys the capability to collect rainfall data over a larger set of beam locations per unit time. For the PR-2, 250 unique beam locations are effectively surveyed within the 260 ms cross-track scan interval, yielding a spatial sampling rate of 960 beams per second. This represents an order of magnitude improvement in spatial sampling compared to the TRMM PR, which collects rainfall data from 49 beam locations within a 565 ms interval (87 beams per second). The effective improvement in spatial sampling is what enables the PR-2 to make concurrent gains in both swath width and horizontal resolution for the rain measurements.

A cartoon illustration of the steps involved in the adaptive scan is shown in Fig. 5. In Fig. 5(a), the radar first builds up echoes-in-flight at beam positions closest to nadir by transmitting at the minimum PRI, determined by

$$PRI_{(\min)} = 2\tau' + \frac{2r_{\max}}{c} \tag{4}$$

where r_{max} is the slant altitude (either 8 or 12 km, for locator or high-resolution sweeps). Fig. 5(b) illustrates the time sequence for the radar beam positions: in accordance with Rule 2, the beam cross-track location is gradually steered further from nadir as time progresses. After the initial EIF pulses have been built up, the space between the satellite and ground is filled with both transmit and receive pulses. The scanning algorithm then enters a closed-loop mode in which the next transmit pulse follows the falling edge of the most recent receive window. Because the receive window is held constant to prevent collisions [Rule 1, Fig. 5(c)], the maximum altitude of the rain image tapers off slightly at the edge of the swath (a 13% decrease at a 30° look angle).

B. Control and Timing Unit

The timing algorithm discussed above can be realized in digital circuitry to make a real-time beam controller for the PR-2.

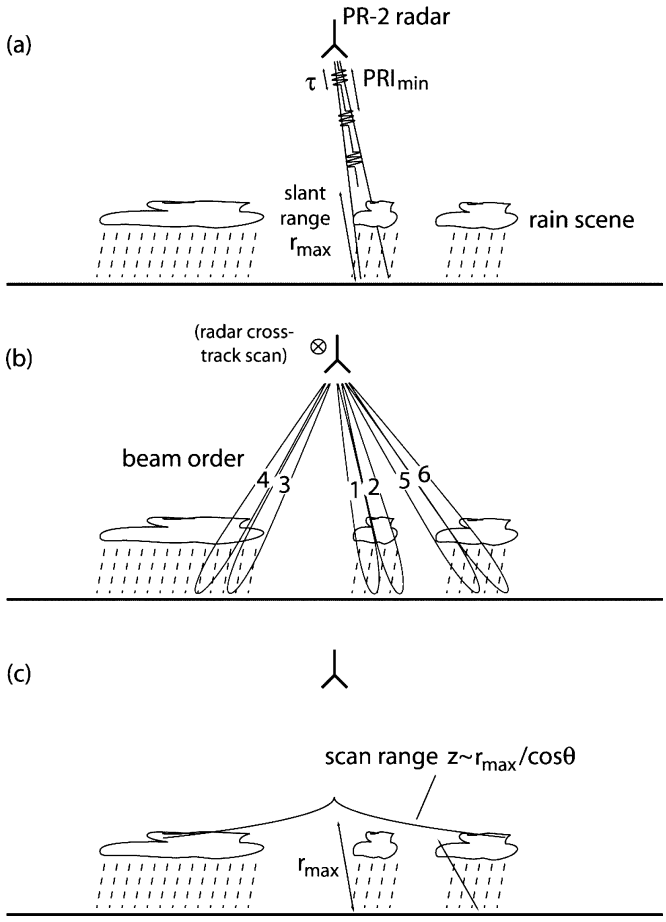

Fig. 5. Illustration of steps involved in the adaptive scanning sequence. (a) The initial buildup of echoes-in-flight (transmit only). (b) Beam steering sequence, from nadir to alternating left and right sides of the cross-track swath. (c) The tapered altitude that results from using a constant range swath (receive window length) throughout the scan.

We have designed and implemented a specialized adaptive scanning control and timing unit (CTU) onto one Virtex-1000 FPGA chip which scans over the full 500-km, 250-beam swath. Because of the critical timing constraints of the adaptive scan, we chose to implement the CTU with a network of custom-made state machines in the FPGA rather than with a microprocessor core. The advantage of state machines is that they can respond immediately and in-parallel to the large number of timing interrupts (up to 15–20) associated with the EIF pulses. Fig. 6 shows the system architecture of the CTU. The controller comprises several submodules, including the following:

- a sweep engine module/EIF counter array which keeps track of the location of each radar echo and provides feedback to guarantee that transmit and receive pulses are interlaced;
- a bubble-sort machine which averages the radar return values from the locator sweep and executes a sorting algorithm of order $O(N^2)$ in the block RAM of the FPGA;
- an antenna driver module that sends beam steering data from an SRAM lookup table to 2300 phase-shifters in the PR-2's antenna array to define the present locations of transmit and receive beams;
- drivers for a solid-state recorder (to store science data) and for a local bidirectional bus (to communicate with other chips onboard).

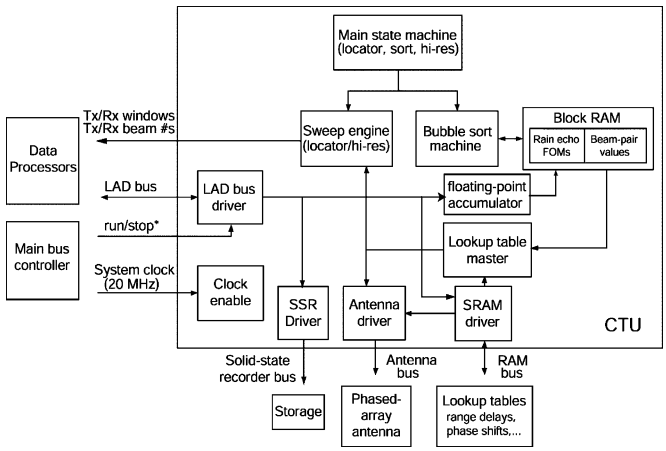

Fig. 6. Module design and chip-to-chip interface for the control and timing unit FPGA.

TABLE III
TIMING EFFICIENCY RESULTS IN THE ADAPTIVE SCAN ALGORITHM

rain profile	total scan time (ms)	timing efficiency (%)
worst case	277.5	82.9
random	260.0	88.5
best case	245.3	93.7

All of the digital circuitry for the CTU has been designed at the register transfer level using the Verilog hardware description language. This design synthesizes to an equivalent gate count of ~79 000 logic gates.

The CTU was tested over entire adaptive scan cycles to measure the *timing efficiency*—the percentage of the total sweep time dedicated to transmitting a pulse or receiving an echo. Verilog simulations were run with several types of rain profiles to test varying levels of difficulty for the autotargeting algorithm: 1) a *best case* profile where all heavy rain is located near the nadir direction (i.e., very little range delay variation); 2) a *worst case* profile in which there is a small amount of rain at nadir and the rest of the rain is at the extreme left and right edges of the swath (a large range delay variation between beams); and 3) a *random case* where the rainfall is uniformly distributed over the swath. Table III lists the timing efficiency results for the full 250-beam scan. The results indicate that the timing solution for the adaptive scan yields a high efficiency, from 83% to 94%, depending on how the rain is distributed.

The CTU tests described above were run with about 10% of the beams (24 out of 250) selected for the high-resolution sweep. It should be pointed out that this represents just one of many possible settings for the adaptive scan. Numeric registers were designed into the FPGA so that the timing solution parameters could easily be changed through software. Given a high-resolution sweep time of 170 ms and a PRI of 200 μs or more, there is time to sample approximately 800 echoes in high-resolution mode. The number of independent echoes was therefore set to 32 so that the product $(24 \text{ beams} \times 32 \text{ looks} = 768 \text{ echoes})$ can fit within the time limit. For comparison, consider a conventional scanning radar with the same total available time of 260 ms per cross-track scan. The number of independent echoes that can be measured is then: total scan time$/$(total number of beams $\times$ PRI)≈ 5 echoes. It follows that in our adaptive scan configuration tested with

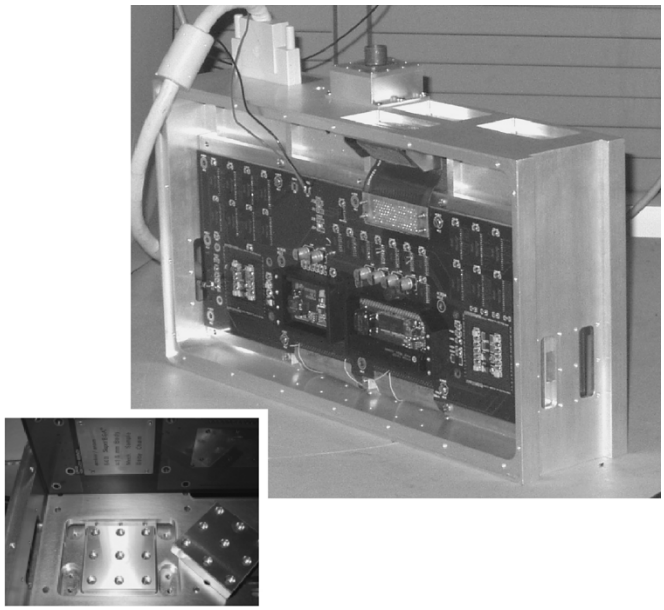

Fig. 7. Space-grade chassis with assembled printed circuit boards for the PR-2 processor/controller. Machined heat sink blocks for conductive cooling of the FPGA ball-grid array packages [inset].

10% autotargeting coverage, the number of independent radar looks improves by a factor of $32/5 = 6.4$ compared to a traditional radar. Because of the built-in flexibility of the CTU software registers, it is easy for the PR-2 operation to be adjusted to make tradeoffs between swath coverage versus number of radar looks, depending on the type of weather being observed and on the requirements of the rain rate retrieval algorithm.

IV. Spaceborne Processor Assembly and Testing

The physical design for the PR-2 processor/controller module consists of a pair of FR-4 (epoxy fiberglass) multilayer printed circuit boards and an aluminum chassis (Fig. 7), both which have been specially designed for a space environment. Key semiconductor parts suitable for space are the Virtex-1000 FPGA, which has passed radiation testing and has flown in NASA's 2003 Mars Exploration Rover mission, and the S-class 15-MHz static RAM chips. Two main aspects of the design that increase the reliability and survivability of the FPGA hardware are: 1) configuration memory scrubbing to recover from radiation-induced logic upsets; and 2) thermal design of the chassis for conductive cooling of the data processor FPGAs. The processor's operation has been successfully confirmed in a series of thermal chamber tests over an ambient temperature range of -20 °C to $+70$ °C.

Fig. 8 shows the experiment setup for the PR-2 processor, along with a PC-based test fixture and National Instruments Data I/O card to control and acquire data from the prototype boards. In the processor design, an Actel 1280 FPGA serves as an interface to the host computer, to a memory manager for the protected memory bank, and to the configuration bus for the Xilinx FPGAs. A graphical user interface (GUI) has also been developed in the LabVIEW visual language to provide the various controls and indicators needed for benchtop operation of the PR-2 processor. With the benchtop test fixture, entire adaptive scan cycles can be run in real time and memory contents examined to verify the radar processor operation.

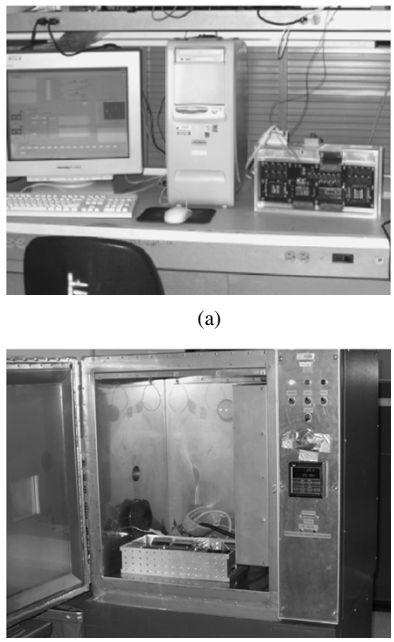

(a)

(b)

Fig. 8. Test setup for the spaceborne PR-2 processor chassis. (a) Benchtop testing with boards interfaced to a personal computer with LabVIEW control software. (b) Processor boards and radar echo signal generator in the thermal chamber.

A. Range Sidelobe Performance

The PR-2 processor boards were tested in the thermal chamber in free-running mode with all four ADC input channels connected to a custom-made signal generator. This signal generator waveform is programmable through read-only memory to simulate the chirp radar response to a rain target at offset video frequencies (after digitization of the received signal but before digital I/Q demodulation). Using the chirp waveform generator, the performance of the data processor's pulse compression filter can be evaluated. Specifically, the most critical figure-of-merit is the range sidelobe level response of the processor, as these sidelobes affect how well sea surface clutter can be suppressed from the rain image.

The signal generator was programmed to simulate the chirp radar return for light rain scatterers at altitudes of 1, 2, and 3 km above the ocean surface to test the linearity of the range compression processor at small signal levels. The return power from the rain cells was set to 60 dB below the ocean return, which corresponds to the rain-to-surface return ratio for a 14-GHz TRMM-class radar viewing a 1-mm/hr rain rate at nadir [7]. The FPGA processors were set to average the radar return samples in range by a factor of 2. At a 5-MHz baseband sampling rate for the processors, this range averaging translates into a vertical sampling interval of 60 m/pixel.

The resultant pulse-compressed, power detected output, plotted in Fig. 9, shows that the rain targets are detectable directly above the ocean for light rain rates of 1 mm/hr. For comparison, the FPGA processor was also run and the output recorded for the case of an echo with no rain over the ocean (to

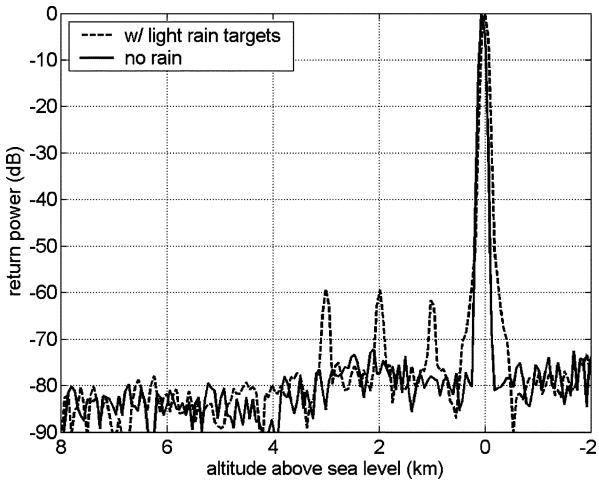

Fig. 9. Pulse-compressed output recorded from the FPGA-based data processor hardware. The range response plots are for simulated light rain targets located at 1, 2, and 3 km above the ocean surface and for an ocean surface return without rain. The low peak range sidelobe performance of −72 dB from the sea surface clutter indicates that a high detection sensitivity for rain can be achieved (down to a 1-mm/hr rainfall rate).

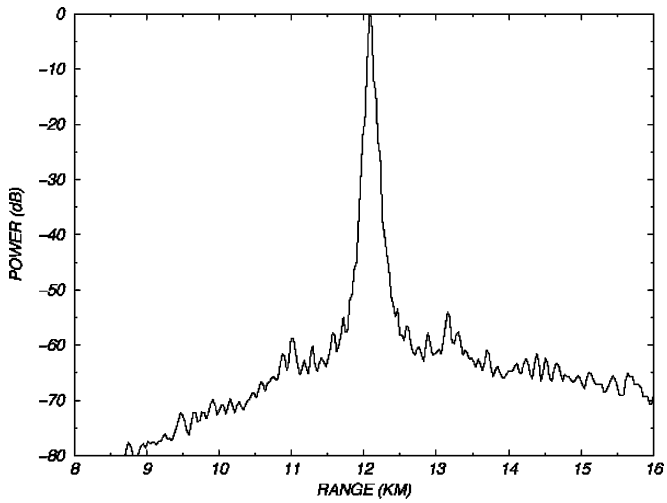

Fig. 10. Pulse-compressed response from the airborne PR-2 prototype processor, measuring a clear-air (rain-free) ocean target ($\tau' = 40$ μs chirp length). The data were acquired during a NASA DC-8 aircraft engineering flight (in preparation for the CAMEX-4 field campaign) in which an end-to-end range sidelobe suppression capability of 60 dB was demonstrated.

emulate a single-point scatter response off of the sea surface), using a sharper amplitude taper on the transmitted chirp pulse. It is observed from the graph that the peak range sidelobe levels from the ocean reflection extend at least 72.3 dB below the ocean return power. This range sidelobe level performance in the FPGA processor enhances the spaceborne PR-2's ability to detect light rainfall in the presence of high-energy clutter from the sea surface. It should also be noted that in the end-to-end radar system, the sea-surface sidelobe performance is limited by phase and amplitude errors and distortion introduced in the RF hardware chain (for example, by microwave power amplifiers and low noise amplifiers). As a reference, Fig. 10 is included to show the end-to-end pulse compression response test results for the airborne PR-2 prototype radar/processor, where it can be seen for a clear-air ocean target that a net sidelobe suppression of 60 dB was achieved in the entire radar system.

B. Power and Thermal Characteristics

During the free-run testing of the PR-2 processor, measurements of the chassis load currents and FPGA device temperatures were made under a variety of configuration conditions to calculate the power consumption rating for each chip. Table IV lists the power consumption data and FPGA temperature diode measurements at an ambient temperature $T_A = 25$ °C for the cases where: 1) one data processor chip and the CTU are configured and running at the full throughput rate of 20×10^9 op/s; 2) neither data processor is configured, but the CTU is configured and running full sweep sequences; and 3) none of the Virtex FPGAs are configured.

From Table IV, the data processor power consumption can be inferred as the difference in net chassis powers between conditions (1) and (2) ($P_{\rm DP} = 14.8$ W per FPGA chip). Likewise for the CTU chip, the power consumption is estimated at 1.0 W, i.e., the power difference between conditions (2) and (3). The junction temperature of the data processor FPGA is known from voltage measurements of an internal diode that has a temperature coefficient of -1.38 mV/°C. At full throughput and an ambient temperature of 25 °C, the data processor junction temperature stabilized to $T_J = 35.1$ °C.

TABLE IV
CHASSIS POWER CONSUMPTION AND DATA PROCESSOR FPGA TEMPERATURE DATA AT 25 °C AMBIENT TEMPERATURE

condition	chassis net load	FPGA junction temperature T_J
1) One Data Processor configured and running at full throughput	28.1 W	35.1 °C
2) No Data Processors configured, CTU configured and running	13.3 W	27.9 °C
3) No FPGAs configured (quiescent load)	12.3 W	27.9 °C

The chassis' effectiveness in removing heat from the FPGA chip can be evaluated using the thermal resistance expression for heat transfer

$$T_J = T_A + \theta_{\rm JA} P_{\rm DP} \tag{5}$$

where $\theta_{\rm JA} = \theta_{\rm JC} + \theta_{\rm CA}$ is the junction-to-ambient thermal resistance, and $\theta_{\rm JC}$ and $\theta_{\rm CA}$ are the thermal resistances from junction-to-case and case-to-ambient. (In the space environment, ambient temperature is defined at the heat sink wall of the chassis, which is opposite the connector side.) Solving for $\theta_{\rm JA}$ in (5) yields a thermal resistance of 0.68 °C/W. For comparison, the junction-to-case thermal resistance of the Virtex-1000 package (a 560-pin ball grid array) has a typical quoted rating of $\theta_{\rm JC} = 0.8$ °C/W—a value that is actually slightly higher than the empirical $\theta_{\rm JA}$. This indicates that the conduction of heat away from the FPGA is limited by the device package itself and not by the aluminum chassis, and therefore that the machined heat sink mounts in the chassis design are highly efficient in removing the 15 W of heat away from each data processor chip. Table V summarizes the power and thermal data results for the FPGA processor.

C. Fault Tolerance

One of the challenges of operating in a space radiation environment is that any memory cell or register can be corrupted

TABLE V
FPGA POWER AND THERMAL SUMMARY FOR THE PR-2 PROCESSOR/CONTROLLER BOARDS

Data Processor power consumption/temperature	P_{DP} = 14.8 W per chip T_J = 35.1 °C (25 °C ambient)
Data Processor junction-to-air thermal resistance	θ_{JA} = 0.68 °C/W (mounted in chassis)
CTU power consumption	P_{CTU} = 1.0 W
chassis quiescent power	P_Q = 12.3 W
chassis net power budget	$2P_{DP}+P_{CTU}+P_Q$ = 42.9 W

by a single event upset (SEU). The Virtex FPGA devices, which use SRAM-based technology for holding configuration data, are susceptible to SEUs. Triple module redundancy is commonly used to mitigate radiation-induced errors, but that approach was deemed too costly in terms of circuit power consumption and mass for the PR-2 application. The approach chosen instead is to allow the possibility of corrupted configuration data, but to halt the FPGA operation momentarily and clear out any possible upsets once every 5 min. In this way, the instrument is guaranteed to keep operating for its mission lifetime, with the possibility of a few outages which would not exceed 5-min duration.

The Xilinx FPGA configuration data and initial memory contents are loaded into a protected memory bank by the host computer during the initial power up sequence. Each memory word contains 16 bits of usable configuration data and an appended six-bit Hamming code for error correction. The one-time programmable, antifuse-based Actel 1280 FPGA, which is immune to reconfiguration by SEUs, continuously reads the memory bank one word at a time and corrects any single bit errors as they appear. After 5 min (500 scrub cycles), the Actel part reads out the memory contents and reconfigures all Xilinx FPGAs and their command tables within 75 ms. A relatively small amount of science data is thus traded to ensure that bit errors cannot accumulate or permanently disable the processor.

The reliability of the EDAC feature was tested in the thermal chamber by writing to the protected memory with one of the data bits forced in a logic-low state. In the first period, the Virtex FPGAs would not configure with the corrupted data, but on the next cycle the memory contents were corrected and the FPGAs initialized successfully. This memory scrubbing test was repeated at several ambient temperature settings over the testing range (-20 °C, $+25$ °C, $+40$ °C, $+50$ °C, $+60$ °C, and $+70$ °C). In all cases, the FPGAs successfully recovered from the software-induced bit errors.

V. SUMMARY OF TECHNOLOGY READINESS

The development of the Second-Generation Precipitation Radar onboard processing hardware has led to several important technical breakthroughs to support an advanced global rain measuring mission: 1) the timing solution for autotargeting of rain can be implemented within a single FPGA chip, and with a timing efficiency of 83% to 94% over the total available scan time; 2) pulse-compression of the chirp radar signal can be accomplished with low range sidelobes (-72 dB in digital loopback tests) using a novel FPGA-based design for the receive processor; 3) the power consumption of each data processor FPGA is confirmed at 15 W at the full throughput rate, and the net power budget for the entire processor/controller is 43 W; 4) the conductively cooled design of a space-grade chassis yields a low, 10 °C temperature gradient between the FPGA substrate and the ambient temperature, thus improving the long-term reliability of the device; 5) the digital electronics successfully recovered from software-induced bit flips using a memory scrubbing technique, which points toward a solution to SEUs encountered in space. These results help establish the baseline instrument requirements and reduce the risk for the spaceborne PR-2.

Many of the design aspects of the PR-2 processor/controller have risen in maturity from concept formulation to component or subsystem validation in a relevant environment. Onboard data processing is now at a very mature level. The core components of this design—the functions for pulse-compression and averaging of multiple radar looks—have been proven in laboratory with a target echo simulator as well as in several airborne precipitation experiments. The unique timing solution for the PR-2 began as a theoretical formulation for guaranteeing the autotargeting of rain without overlapping of the transmit and receive pulse sequences, and has now matured into working FPGA firmware for the radar's CTU. This CTU firmware also represents a technical step forward in that it is one of the first high density (100 000+ gate) FPGA-based timing controller designed for a spaceborne sensor. Tests of the CTU demonstrate the efficiency of using dedicated state-machine logic instead of a microprocessor for time-critical control of the radar. Finally, a memory scrubbing approach for correcting radiation upsets in the FPGA configuration data was developed and proven as an attractive alternative to more complicated triple module redundancy techniques.

ACKNOWLEDGMENT

The research described in this paper was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The authors thank S. L. Durden for helping to define and review the adaptive scan system design for the PR-2, and S. Tanelli for preparing the airborne PR-2 data products from the CAMEX-4 and AMSR-E validation experiments.

REFERENCES

[1] G. M. Flaming, "Requirements for global precipitation measurement," in *Proc. IGARSS*, vol. 1, Toronto, ON, Canada, Jun. 2002, pp. 269–271.
[2] T. Kozu, T. Kawanishi, H. Kuroiwa, M. Kojima, K. Oikawa, H. Kumagai, K. Okamoto, M. Okumura, H. Nakatsuka, and K. Nishikawa, "Development of precipitation radar onboard the Tropical Rainfall Measuring Mission (TRMM) satellite," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 1, pp. 102–116, Jan. 2001.
[3] D. F. Everett, S. W. Bidwell, G. M. Flaming, T. B. Rykowski, E. F. Stocker, and L. E. Braatz, "Engineering of the global precipitation measurement system," in *Proc. 2003 Aerospace Conf.*, vol. 1, Big Sky, MT, Mar. 2003, pp. 141–153.
[4] E. Im, S. L. Durden, G. Sadowy, A. Berkun, J. Huang, M. Lou, B. C. Lopez, Y. Rahmat-Sammi, and S. Rengarajan, "System concept for the next-generation spaceborne precipitation radars," in *Proc. 2000 Aerospace Conf.*, vol. 5, Big Sky, MT, Mar. 2000, pp. 151–158.
[5] G. Sadowy, A. Berkun, S. Durden, J. Huang, E. Im, B. Lopez, M. Lou, Y. Rahmat-Samii, F. Y. Liu, and S. Rengarajan, "Technologies for the next generation of spaceborne precipitation radars," in *Proc. 2001 Aerospace Conf.*, vol. 4, Big Sky, MT, Mar. 2001, pp. 1811–1823.
[6] G. A. Sadowy, A. C. Berkun, W. Chun, E. Im, and S. L. Durden, "Development of an advanced airborne precipitation radar," *Microw. J.*, vol. 46, pp. 84–98, Jan. 2003.
[7] R. Meneghini and T. Kozu, *Spaceborne Weather Radar*. Norwood, MA: Artech House, 1990.

[8] A. Tanner, S. L. Durden, R. Denning, E. Im, F. K. Li, W. Ricketts, and W. Wilson, "Pulse compression with very low sidelobes in an airborne rain mapping radar," *IEEE Trans. Geosci. Remote Sens.*, vol. 32, no. 1, pp. 211–213, Jan. 1994.
[9] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
[10] R. Andraka and A. Berkun, "FPGAs make a radar signal processor on a chip a reality," in *Proc. 33rd Asilomar Conf. Signals, Systems Computers*, Monterey, CA, Oct. 1999.
[11] R. Kamineni, T. N. Krishnamurti, R. A. Ferrare, S. Ismail, and E. V. Browell, "Impact of high resolution water vapor cross-sectional data on hurricane forecasting," *Geophys. Res. Lett.*, vol. 30, no. 5, pp. 38-1–38-4, Mar. 2003.
[12] S. Tanelli, J. Meagher, S. L. Durden, and E. Im, "Multiparametric airborne radar observations of the melting layer during the Wakasa Bay experiment," in *Proc. 31st Conf. Radar Meteorology*, Seattle, WA, Aug. 2003, pp. 33–34.
[13] C. Ruf, C. Principe, T. Dod, B. Gosselin, B. Monosmith, S. Musko, S. Rogacki, A. Stewart, and Z. Zhang, "Lightweight rainfall radiometer STAR aircraft sensor," in *Proc. IGARSS*, vol. 2, Toronto, ON, Canada, Jun. 2002, pp. 850–852.
[14] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
[15] D. Atlas, "Adaptively pointing spaceborne radar for precipitation measurements," *J. Appl. Meteorol.*, vol. 21, pp. 429–431, Mar. 1982.
[16] P. Amayenc, J. Testud, and M. Marzoug, "Proposal for a spaceborne dual-beam rain radar with Doppler capability," *J. Atmos. Oceanic Technol.*, vol. 10, pp. 262–276, Jun. 1993.
[17] S. Kobayashi and T. Iguchi, "Variable pulse repetition frequency for the Global Precipitation Measurement Project (GPM)," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 7, pp. 1714–1718, Jul. 2003.

Mark A. Fischman (S'95–M'01) received the B.S. degree in electrical engineering from the California Institute of Technology, Pasadena, in 1989, and the M.S. and Ph.D. degrees in electrical engineering from the University of Michigan, Ann Arbor, in 1996 and 2001, respectively. His thesis research at Michigan involved the development of a 1.4-GHz direct-sampling digital correlation radiometer to facilitate phase-stable, single-chip receiver architectures for Synthetic Thinned Aperture Radiometry (STAR).

He was an Electronics Design Engineer with the Sony Corporation Television Product Design groups, Tokyo, Japan, and Pittsburgh, PA, from 1989 to 1995. In 2001, he joined the Radar Science and Engineering Section, Jet Propulsion Laboratory, Pasadena, where he is currently a Senior Engineer in the Atmospheric Radar Science and Engineering group. At JPL, he is implementing the onboard radar processing hardware systems for several research and flight projects, including the Second-Generation Precipitation Radar, the Lightweight L-band Space Based Radar, and the Wide Swath Ocean Altimeter instrument. He is a Co-Investigator for the Spaceborne Doppler Precipitation Radar Processor project funded under the NASA Advanced Information Systems Technology (AIST) program. His research interests include advanced microwave and digital receiver design for large-aperture phased-arrays, interferometric radar, and synthetic aperture radar, and the development of real-time satellite radar processing techniques for improving spatial coverage and observation times.

Dr. Fischman received a NASA New Technology Report Award and Space Act Board Award in 2004.

Andrew C. Berkun received the B.S. degree from Rensselaer Polytechnic Institute, Troy, NY, and the M.S. degree from the California Institute of Technology, Pasadena, in 1990 and 1997, respectively, both in electrical engineering.

He has been a Hardware Engineer at the Jet Propulsion Laboratory (JPL), Pasadena, for 15 years. He has been actively involved in the digital electronics design, implementation, and testing in many spaceborne and airborne radar programs at JPL, including the Cassini radar for remote sensing of Titan (Cognizant Engineer for timing unit), Shuttle Imaging Radar-C (Cognizant Engineer for digital data handling assembly), JPL Airborne SAR, GeoSAR (Cognizant Engineer for timing unit and ADC/buffers and Systems Engineer), and most recently, the Second-Generation Precipitation Radar spaceborne prototype (Principal Investigator), as well as the 94-GHz cloud profiling radar for the ESSP CloudSat Mission. His interests include real-time processing engines for radar data using FPGA implementations and unusual radar timing solutions.

William Chun, photograph and biography not available at the time of publication.

Eastwood Im (S'81–M'85–SM'01) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Illinois, Urbana, in 1981, 1982, and 1985, respectively.

He has been with the Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, since 1986. He has been involved in several design studies on spaceborne earth remote sensing radars, including precipitation radars, altimeters, and millimeter-wave cloud radars. He is currently the Radar Instrument Manager of the NASA ESSP CloudSat Mission and the Supervisor of the Atmospheric Radar Science and Engineering program at JPL. He has been a member of the Tropical Rainfall Measuring Mission (TRMM) Science Team and a Principal Investigator in TRMM radar rain profiling algorithms and TRMM radar calibration, and the Principal Investigator of the IIP Second-Generation Precipitation Radar (PR-2) technology study. He has also been a Co-Investigator of several other NASA research studies. Currently, he is a member of NASA's Precipitation Measurement Missions Science Team and the Principal Investigator in four NASA earth science research tasks, including the CAMEX-4 airborne radar's rainfall observation experiment, the EOS AQUA AMSR-E's rainfall measurement validation, the IIP dual-frequency wide-swath scanning rain radar antenna study, the IIP NEXRAD-in-Space radar study, and the Titan Cloud/Precipitation Radar and Altimeter instrument design study. He is also the Co-Investigator on the NASA ESE Advanced Information System Technology (AIST program) advanced spaceborne rain radar processor and is the Guest-Investigator on the radar study of Titan's atmospheric methane precipitation in NASA's Cassini Mission. He has been a Member of a science steering group on the NASA Global Precipitation Mission since its inception, focusing particularly in the areas of radar performance and technology development and infusion. His research interests are in the areas of spaceborne meteorological radar science and remote sensing, and advanced radar system and technology studies. In particular, he has developed radar algorithms for the retrievals of precipitation intensity and surface topography, and the two-color laser altimetry algorithms for measuring surface pressure and tectonic motions. He has published over 150 articles in refereed journals and conference proceedings, and 10 technical reports. He was the Associate Editor of the *AMS Journal of Atmospheric and Oceanic Technology* in 1999–2000.

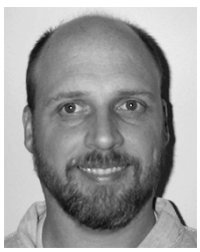

Raymond J. Andraka (S'82–M'83) received the B.S.E.E. degree from Lehigh University, Bethlehem, PA, and the M.S.E.E. degree from the University of Massachusetts, Amherst, in 1984 and 1992, respectively.

He is the President of the Andraka Consulting Group, Inc., a digital hardware design firm he founded in 1994. His company is focused exclusively on high-performance DSP designs using FPGAs. He has applied FPGAs to signal processing applications including radar processors, radar environment simulators, sonar, HDTV, digital radio, spectrum analyzers, image processing, and communications test equipment. Ray's prior signal processor design experience includes five years with Raytheon Missile Systems designing radar signal processors and three years of signal detection and reconstruction algorithm development for the U.S. Air Force. He also spent two years developing image readers and processors for G-Tech, where he set the company time-to-market record for a new product. He has also authored many conference papers dealing with high-performance FPGA design and signal processing applications and is a regular contributor to the comp.arch.fpga, comp.lang.vhdl and comp.dsp newsgroups.