

Design and Demonstration of an Advanced On-Board Processor for the Second-Generation Precipitation Radar

Mark A. Fischman, Andrew C. Berkun, Frank T. Cheng, William W. Chun, Eastwood Im
Jet Propulsion Laboratory, California Institute of Technology, Radar Science and Engineering Section
M/S 300-241, 4800 Oak Grove Dr, Pasadena, CA 91109-8099, USA
Tel: (818) 354-3862, Fax: (818) 393-6943, e-mail: mafisch@mail1.jpl.nasa.gov

Ray Andraka
Andraka Consulting Group, Inc.
North Kingstown, RI 02852-1666
ray@andraka.com

Abstract—The Next-Generation Precipitation Radar (PR-2) prototyped by NASA/JPL will depend heavily on high-performance digital processing to collect meaningful rain echo data. Using field-programmable gate arrays (FPGAs), we have developed for the PR-2 a pulse-compression processor and adaptive timing controller that will enable full on-board processing capabilities in a 13 and 36 GHz spaceborne radar. This paper describes some of the new technologies for the on-board processor, including a 40×10^9 op/s bit-serial filter attaining -60 dB range sidelobe performance, and an adaptive scanning control and timing unit (CTU) which yields a 7-fold increase in the radar's dwell time over areas of precipitation.

I. INTRODUCTION

Global three-dimensional rainfall data acquired by a spaceborne precipitation radar provide crucial information about the hydrological cycle. The imaging of liquid water mass in three-dimensions (in altitude and horizontally) is also important for understanding the global climate system because such images reveal how latent heat is transferred through the atmosphere. The Ku-band (13 GHz) precipitation radar onboard the Tropical Rainfall Measuring Mission (TRMM) satellite, launched in 1997, was the first Earth-orbiting sensor to retrieve 3-D rainfall data over tropical and mid-latitude regions [1]. With the success of TRMM, a new class of spaceborne precipitation radars are being developed for follow-on missions with improved sensitivity, spatial resolution, and swath coverage.

NASA/JPL is currently prototyping a Next-Generation Precipitation Radar (PR-2) that offers the advanced capabilities [2] needed to succeed TRMM. The PR-2 project involves the parallel development of several breakthrough technologies [3], including: 1) a large aperture inflatable/rigidizable membrane reflector antenna, 2) compact, solid-state T/R modules for an electronically steerable array at 13 and 36 GHz, and 3) on-board digital data processing and timing control, implemented in field-programmable gate arrays (FPGAs). Fig. 1 illustrates the concept for the spaceborne PR-2 instrument, with its lightweight antenna deployed in a low-Earth orbit. Given the large 5.3×5.3 m antenna aperture and the specially developed on-board processing electronics for adaptive steering of the radar beam, the PR-2 is capable of collecting 3-D rainfall data at a high spatial resolution (2 km footprint) and over a wide cross-track swath (500 km).

The on-board processor, which is the focus of this paper, must be designed to meet the challenging requirements for measuring precipitation from space. Sea surface returns can be as much as 55 dB stronger than the rain return power to become a major source of clutter [4]. Also, for conventional scanning radars, there is a clas-

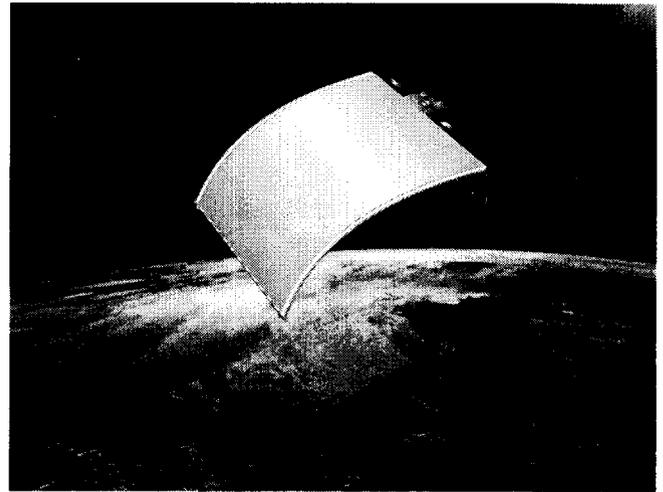


Fig. 1. Conceptual drawing of the spaceborne Second-Generation Precipitation Radar (PR-2), with an innovative 5.3×5.3 m reflector antenna based upon inflatable/rigidizable technology.

sic tradeoff between swath width capabilities and the available dwell time over each footprint [5]. In the following sections, we describe the new features of our prototype data processor and control and timing unit (CTU) that solve these system design challenges.

II. PULSE COMPRESSION DATA PROCESSOR

The PR-2 uses a chirp radar technique, which has advantages over short pulse weather radar systems like TRMM that typically sample at a vertical range resolution of 250 m. With a chirp rain radar, the number of independent looks within the 250 m cell can be increased by modulating the chirp over several megahertz bandwidth. For example, the PR-2 uses a chirp bandwidth of $B = 4$ MHz—or equivalently a range resolution of $c/2B = 37.5$ m—to increase the number of samples within the range bin by a factor of $N = 6.7$. After these samples are averaged in power to reduce Rayleigh fading noise, the effective signal-to-noise ratio of the radar's rain echo measurement can be improved by as much as $\sqrt{N} = 4.1$ dB. A further benefit of chirp radar is that, because signal power is gained through the use of a long transmit pulse length ($\tau = 51 \mu\text{s}$ for the spaceborne PR-2), the requirements for instantaneous transmit power are relatively low (~ 200 W). This moderate transmit power requirement allows simpler, solid-state transmitter hardware to be integrated into the phased array antenna feeds.

The cost of the chirp radar is that the pulse compression process-

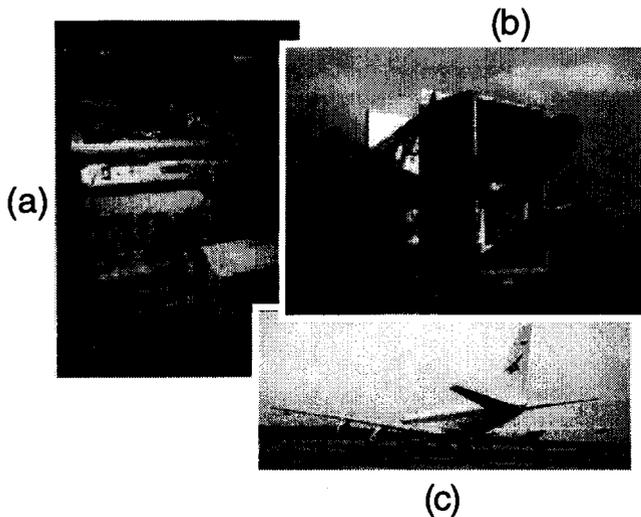


Fig. 2. The PR-2 airborne prototype radar: (a) electronic processing and control console, (b) dual-frequency (13 and 36 GHz) antenna feed structure with mechanically scanning reflector, (c) NASA Dryden DC-8 aircraft used for deployment in the 4th Convection and Moisture Experiment (CAMEX-4).

ing in the receiver section is numerically intensive. To attain sufficiently low range sidelobes of -60 dB for masking out sea surface clutter, a large matched filter bank is needed with 200 or more coefficients and with time-domain weighting of the chirp reference pulse [6]. Taking into account the complex I/Q arithmetic for the matched filter and the requirement for 4 separate receive channels (Ku and Ka-band, H and V-pol), the PR-2 must be able to process incoming data at 40 billion operations per second.

High speed, radiation-hardened FPGA chips with million gate densities have recently emerged which can support the high throughput requirements for the PR-2. In the last year, our team partner, Andraka Consulting, Inc., has implemented our concept for a new pulse compression data processor onto two FPGAs using a distributed arithmetic technique [7]. We have successfully tested the first version of the processor core aboard an airborne prototype of the PR-2, deployed during the 4th Convection and Moisture Experiment (CAMEX-4) in the summer of 2001. Fig. 2 shows the PR-2 radar hardware installed onto the NASA Dryden DC-8 aircraft for this experiment. Fig. 3 shows the pulse-compressed and processed science data from a CAMEX-4 flight over Hurricane Humberto in September, 2001. The multi-parameter Ku and Ka-band rain profiles in the figure are as follows: (a)–(b) co-polarized, effective radar reflectivity data taken over a time scale of 8 minutes (the bright line near the bottom of each graph represents the ocean surface return at 0 km altitude); (c)–(d) linear depolarization ratios, which accentuate scattering by ice particles in the melting layer; and (e)–(f) vertical Doppler rainfall velocity measurements. As a next step toward a spaceborne PR-2 mission, we are redesigning our data processor to interface with a more sophisticated timing controller for satellite-based measurements, where a number of echoes-in-flight (EIF) must be tracked and averaged together over time.

A diagram of the signal chain for the PR-2 data processor is illustrated in Fig. 4. The on-board processor is designed around a bit serial finite impulse response (FIR) filter implemented in a pair of Xilinx Virtex-1000 FPGAs. Data enters the FPGAs as four channels of 12-bit digitized data at a 20 MHz sampling rate. The 4 MHz chirp

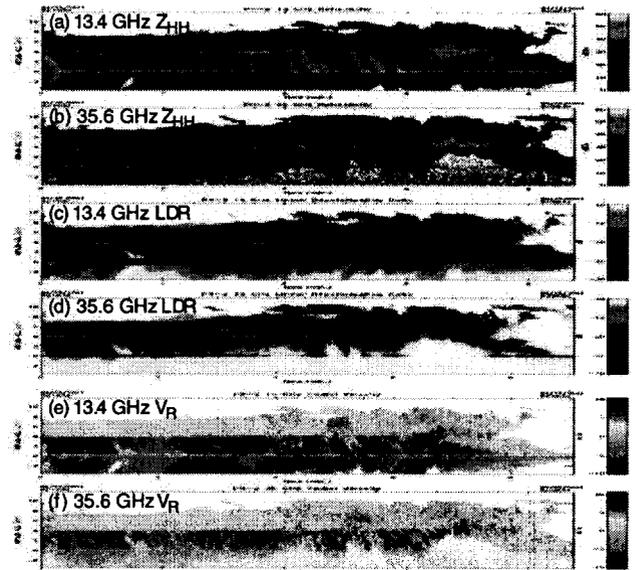


Fig. 3. Multi-parameter, dual-frequency (Ku and Ka-band) rainfall data from Hurricane Humberto, taken with the PR-2 Data Processor during CAMEX-4: (a)–(b) H-pol rain reflectivity data. (c)–(d) Linear depolarization ratios. (e)–(f) Vertical Doppler measurements.

data is 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 a saturation detector followed by an anti-aliasing, symmetric FIR filter, with 64 taps and 16-bit coefficients, which removes the out-of-band noise and the alias centered at 15 MHz. The output of this filter is mixed with a 5 MHz sinusoid to shift the center of the signal spectrum down to 0 Hz. The data is then downsampled by a factor of 4 with no loss of information due to the anti-aliasing filter.

The 5 MHz complex baseband data is stored in a first-in/first-out memory so that pulses containing saturated samples can be eliminated. Data is then fed into a 256-tap, non-symmetric FIR filter which contains the complex conjugate of the expected return. This is where the bulk of the FPGA's processing resources are used, amounting to 20 billion real multiply and add operations per second. The output of this 256-tap filter is the compressed radar return. Complex data is converted to power and averaged over adjacent range bins and multiple pulses. Averaging provides two orders of magnitude data compression. Further range averaging is applied for adaptive scanning, described in the next section, to identify the beams most likely to contain rain.

A number of DSP applications make this processing density possible. Bit-serial filtering makes best use of the speed capabilities of the FPGA. The FIR filter stage uses a lookup table to combine 4 stages of multiplication and addition 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 real estate.

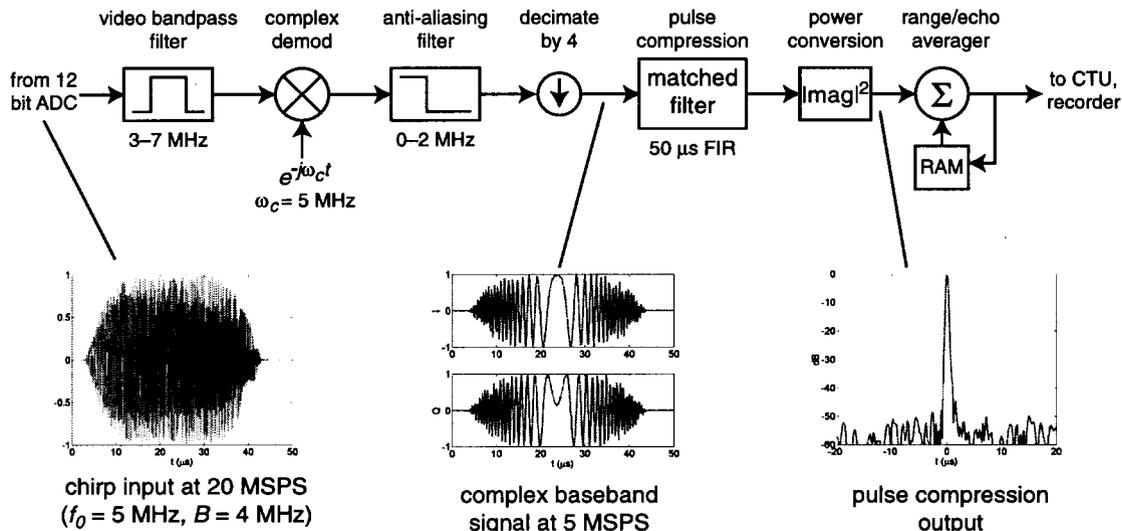


Fig. 4. Block diagram of the digital data processor signal flow for the PR-2.

III. CONTROL AND TIMING UNIT

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 [4]. This statistic can be exploited in the PR-2 by using an adaptive scanning technique [5]—that is, the radar can use its own 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. With adaptive scanning, the dwell time for the rain target can be increased 7 times beyond that of a conventional cross-track scanning radar without sacrificing swath width.

We have developed a specialized control and timing unit with an adaptive scanning algorithm onto a Virtex-1000 FPGA. This CTU generates the transmit and receive timing for an entire 300 ms sweep cycle, consisting of: 1) a locator sweep that performs a rough measurement of echo return power over all cross-track beam locations (248 beams) and over an 8 km altitude range; 2) a bubble-sort algorithm which ranks the 248 beam locations from highest to lowest return powers; and 3) a high-resolution sweep, which takes additional radar looks of the top 24 ranked beams over a longer 12 km altitude range. The timing solution for the CTU must be generated on-the-fly with minimum dead-time between EIFs, while also avoiding collisions between transmit and receive echoes.

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 to using state machines is that they can respond immediately and in-parallel to a large number of timing interrupts. Fig. 5 shows a diagram of the key logic modules developed for the CTU. The sweep engine module includes an EIF counter array which keeps track of the location of up to 32 radar echoes and provides feedback to guarantee that transmit and receive pulses are interleaved. A bubble-sort machine takes echo 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 then sends beam steering data from an SRAM lookup table to 2,300 phase-

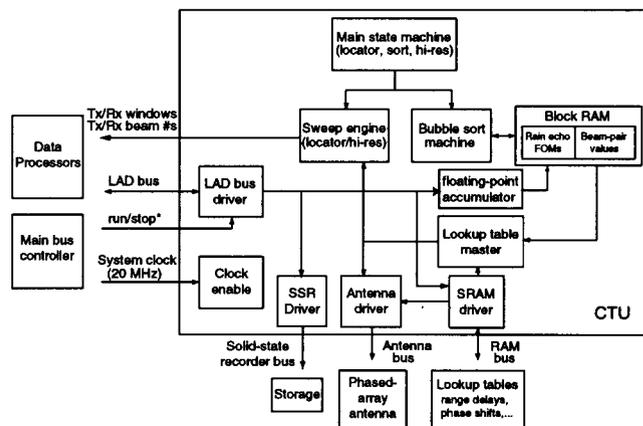


Fig. 5. Verilog module design and chip-to-chip interface for the Control and Timing Unit FPGA.

shifters in the PR-2's antenna array to define the present locations of transmit and receive beams. All of the digital circuitry for the CTU has been designed at a register transfer level using the Verilog hardware description language. Currently this design has a total equivalent gate count of 79,000 gates. To guarantee that setup and hold timing requirements are met, we implemented the CTU logic as a synchronous design (with all registers tied to a common 20 MHz clock) and ran it through a static timing analysis.

IV. SPACEBORNE PROCESSOR TESTING AND ASSEMBLY

A. Test Fixture for Electronic Boards

A PC-based test fixture has been developed for data acquisition and control of the prototype PR-2 boards. Two National Instruments Data I/O cards installed in the computer drive up to 64 digital ports going to a main bus controller chip (a radiation-hardened Actel FPGA), which is interfaced to the Virtex FPGAs through a local bus. We have also written a graphical user interface (GUI) in the LabView visual language for the various controls and indicators needed for benchtop operation of the prototype boards. These GUI

functions include:

- Generating and loading matched-filter coefficients for the Data Processor's pulse compression module
- Generating and loading into CTU memory a range delay lookup table, based on the spacecraft's altitude
- Setting the CTU's "zero-page registers," which define basic parameters such as the transmit and receive window lengths, number of independent radar looks per beam, number of beam positions in the high-resolution sweep, and so on
- Manual read/write control to Data Processor and CTU memory
- High level controls for configuring, resetting, and enabling or disabling the FPGAs

To facilitate testing of the complicated multiple beam timing scheme, an Echo Delay Simulator (EDS) test fixture is currently being developed. The EDS substitutes for the chirp generator, upconverter, radar antenna, rain target, receiver, and A/D converter in the end-to-end radar system. The EDS will be implemented on a commercial FPGA card manufactured by Annapolis Micro Systems, Inc.

The EDS waits for a transmit command and a cross-track beam location value from the Control and Timing Unit. This beam location number is then used to index a range-delay lookup table, which holds the radar pulse's round trip delay time to the rain target and back. The delay value is loaded into one of 32 independent count-down timers that represent at any moment the status of up to 32 separate echoes-in-flight. When a count-down timer reaches zero—and provided that the antenna bus port from the CTU is set to receive at the same cross-track beam location—the echo delay simulator will send the digitized range return data back to the Data Processor's ADC port. We have written a MATLAB software tool that can generate the rain profiles for theoretical weather systems and convert these profiles into packets of digitized range samples for the EDS.

Because the EDS only sends out its packet of data when the correct receive beam location is commanded at the correct time by the CTU, the reception of data is proof of the CTU's timing solution. Using the MATLAB tool, we can test the instrument's auto targeting capability for different weather systems in order to provide a robust test environment for the spaceborne radar.

B. Range Sidelobe Performance

A hardware description language simulation of the Data Processor FPGA was run to test the processor's ability to remove sea surface clutter from a rain image. The chirp radar returns for light rain targets at 1, 3, and 5 km altitudes above the ocean were calculated and then input to the FPGA firmware under test. The resultant pulse-compressed, power detected output, plotted in Fig. 6, shows that the rain targets are in fact detectable even for rain echoes that are 60 dB weaker than the sea surface return. Note also from the graph that the range sidelobes from the ocean reflection (the detected power between the ocean and the 1 km rain scatterer) extend 72 dB below the ocean return. This range sidelobe performance in the digital processor enables the PR-2 to measure rainfall with a high level of sensitivity, even when these observations are made over the ocean.

C. Adaptive Scan Timing Efficiency

The CTU was operated over an entire adaptive scan cycle to measure the timing efficiency e of the radar—that is, the percentage of the total sweep time dedicated to transmitting a pulse or receiving an echo:

$$e = T_{on}/T_{net}, \quad (1)$$

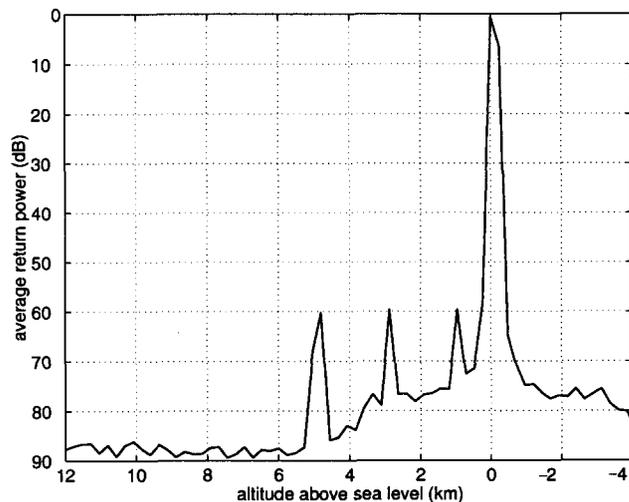


Fig. 6. Hardware description language simulation of the spaceborne Data Processor FPGA for the case where weak rain scatterers are observed over the ocean.

rain profile	total scan time (ms)	timing efficiency (%)
worst case	271.4	84.7
random	259.9	88.5
best case	248.7	92.5

TABLE I
TIMING EFFICIENCY RESULTS FOR THE ADAPTIVE SCAN ALGORITHM.

where T_{on} is the transmit time plus receive time and T_{net} is the total time for a scan. For the spaceborne PR-2,

$$T_{on} = b_l N_l (\tau + t_{rx(l)} + t_s) + b_h N_h (\tau + t_{rx(h)} + t_s) \quad (2)$$

where b is the number of beam locations, N is the number of independent looks per beam, t_{rx} is the receive window time length, t_s is the transient switching time for the radar's transmit/receive circulators, and subscripts l and h denote the parameters for locator and high-resolution sweeps, respectively. Based on the sweep parameter values discussed in Section III, and setting the number of independent looks to $N_l = 2$ and $N_h = 32$, the T_{on} time calculated in (2) is 230.0 ms.

The test results for the total sweep cycle time T_{net} and timing efficiency e are listed in Table I for several types of rain profiles: 1) a worst case profile where rain cells lie only at nadir and at the extreme edges of the swath; 2) a moderate profile where the rain has a random, uniform distribution over the swath; and 3) a best case profile where the rain cells are clustered only around nadir. Efficiencies vary depending on the minimum range delay available in a scan and on how many pulse repetition intervals can fit within that minimum delay. The results indicate that our timing solution for the radar's adaptive scan will range from 85–93%, depending on how the rain is distributed.

To test the speed of the bubble-sort processor, several random rain profiles were generated in software for the full 248 beam scan and were input to the CTU Verilog simulator. The sorting machine had an average run time of 1.9 ± 0.1 ms, which shows that ranking of rain locations can be completed within a small fraction (less than

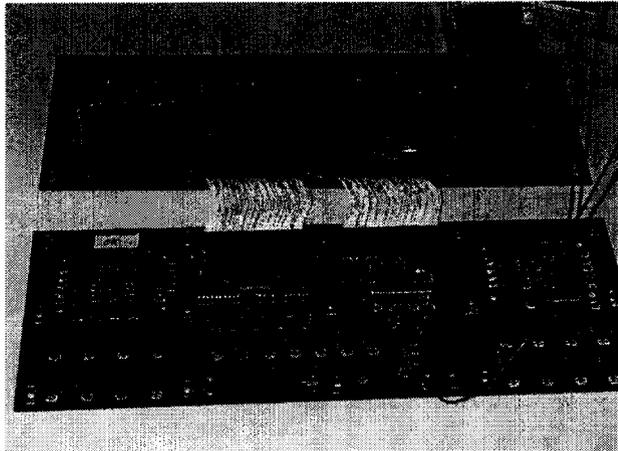


Fig. 7. Mounted printed circuit boards for the CTU and bus controller (top) and for the Data Processor FPGAs (bottom).

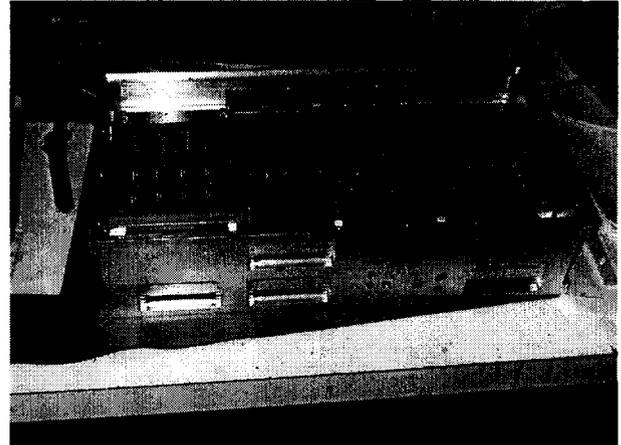


Fig. 9. Photo of boards assembled in the aluminum chassis.

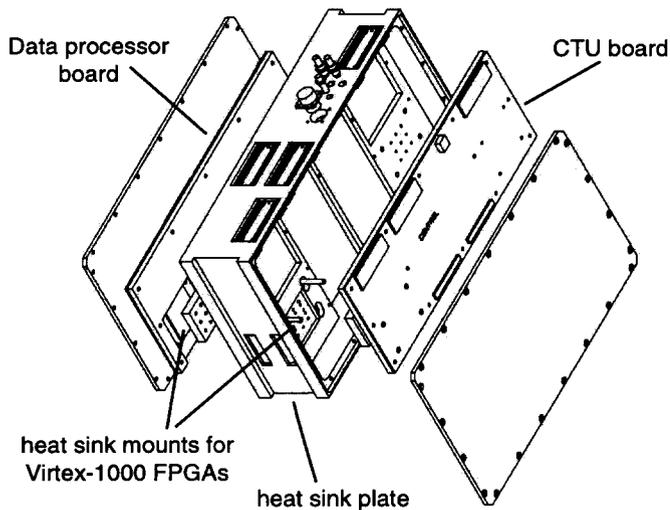


Fig. 8. Mechanical chassis design for the spaceborne PR-2 processor/controller.

1%) of a total sweep cycle.

D. Space Qualified Boards and Chassis

The physical board design for the data processor and CTU (Fig. 7) provides the standard interfaces in a satellite instrument and uses equivalent S class parts. Key parts are the Virtex-1000, which has passed radiation testing and is being evaluated for S class, and the 15 MHz static RAM chips. Key interfaces include a 1553 spacecraft command bus, data output bus for a solid state recorder, telemetry, science ADC, the bus for the electrically steerable antenna, an arbitrary waveform generator for the chirp transmission, and a test port for ground-based operation. The board is also being mounted in a custom-made chassis, shown in Figs. 8–9, which is capable of removing the 20 watts of heat from each FPGA in a vacuum and stabilizing the board under 10 G's of launch acceleration.

E. Error Detection and Correction

One of the challenges of operating in a radiation environment such as space is that any memory cell or register can be corrupted by a single event upset (SEU). The Virtex FPGA devices, which use SRAM-based technology for holding their configuration data, are therefore susceptible to SEUs. Triple redundancy is commonly used to mitigate radiation-induced errors, but that approach was deemed too costly for the PR-2 application. The approach chosen is instead to allow the possibility of corrupted configuration data, but to halt the FPGA operation momentarily and clear out any possible upsets once every 5 minutes. 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 minutes duration.

In the PR2 design, the Xilinx FPGA configuration data and initial memory contents are loaded into a protected memory bank. Each memory word contains 16 bits of raw data and an appended 6 bit Hamming code. A one-time programmable, antifuse-based Actel 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 minutes (500 scrub cycles), the Actel part reads out the memory contents and reconfigures the Xilinx FPGAs and their command tables. In this way, an SEU cannot permanently disable the instrument.

The Error Detection and Correction (EDAC) feature was tested in the lab by writing to the protected memory with one of the data bus bits stuck in a logic-low state. In the first period, the FPGA would not configure with the corrupted data, but on the next cycle the table had been corrected and the Xilinx FPGA initialized successfully. In another test, one bit of memory was shorted to ground and released during the board operation. Again, the FPGA would not load with the corrupted data, but it did load successfully on the following cycle, after protected memory had been scrubbed.

V. CONCLUSION

We have presented an overview of the Next-Generation Precipitation Radar's on-board digital electronics design, which includes a pulse compression processor with very low range sidelobe performance and a new adaptive scanning controller. The on-board processor/controller will serve as a key subsystem for a future satellite precipitation radar mission having dual-frequency (13 and 36 GHz) capabilities.

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.

REFERENCES

- [1] 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 Sensing*, vol. 39, pp. 102–116, Jan. 2001.
- [2] E. Im, S. L. Durden, G. Sadowy, A. Berkun, J. Huang, M. Lou, B. C. Lopez, Y. Rahmat-Samii, and S. Rengarajan, "System concept for the next-generation spaceborne precipitation radars," in *Proc. 2000 Aero. Conf.*, vol. 5, (Big Sky, MT), pp. 151–158, IEEE, Mar. 2000.
- [3] 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 Aero. Conf.*, vol. 4, pp. 4/1811–4/1823, IEEE, 2001.
- [4] R. Meneghini and T. Kozu, *Spaceborne Weather Radar*. Norwood, MA: Artech House, 1990.
- [5] D. Atlas, "Adaptively pointing spaceborne radar for precipitation measurements," *J. Appl. Meteor.*, vol. 21, pp. 429–431, Mar. 1982.
- [6] 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 Sensing*, vol. 32, pp. 211–213, Jan. 1994.
- [7] R. Andraka and A. Berkun, "FPGAs make a radar signal processor on a chip a reality," in *Proc. 33rd Asilomar Conf. Signals, Syst. Comput.*, (Monterey, CA), IEEE, Oct. 1999.



Mark A. Fischman (S'95-M'01) received the B.S. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, in 1989, and the M.S. and Ph.D. degrees in electrical engineering from the University of Michigan, Ann Arbor, MI, in 1996 and 2001, respectively. He worked at the Sony Corporation television development groups, Tokyo, Japan and Pittsburgh, PA, from 1989 to 1995. In 2001, he joined the Radar Science and Engineering Section at the Jet Propulsion Laboratory, Pasadena, CA, where he is currently developing the on-board processing and adaptive scanning control systems for

the Next-Generation Precipitation Radar satellite instrument. His research interests include the development of advanced digital receivers, phased-array and aperture synthesis techniques for microwave radar and radiometry.



Andrew C. Berkun is a Digital Engineer at JPL. 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, Shuttle Imaging Radar-C, JPL Airborne SAR, and GeoSAR, and most recently, the Second-Generation Precipitation Radar breadboard, and the 94 GHz cloud profiling radar for the ESSP CloudSat Mission.

Frank T. Cheng received the B.S. degree in electrical engineering from the National Chiao Tung University in 1978 and the M.S. degree in electrical engineering from the University of Missouri in 1981. He has fifteen years of working experience in the area of parallel data processing for SAR processor and scientific data simulation. He has contributed to the processor algorithm and architecture design for the Magellan SAR data processor and was responsible for the Magellan Engineering SAR processor implementation. He was also responsible for the RADARSAT ScanSAR processor architecture design and the integration and test of the Alaska SAR Processor systems in the Alaska SAR Facility (ASF). He is currently interested in FPGA design for digital signal processors.

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



Dr. Eastwood Im is the Radar Instrument Manager of the NASA ESSP CloudSat Mission. He is also the Supervisor of the Atmospheric Radar Science and Engineering Group at JPL. He has extensive experience in spaceborne meteorological radar science and remote sensing, and advanced radar system studies. Dr. Im 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 the Principal Investigator in three NASA earth science research studies: the IIP dual-frequency wide-swath scanning rain radar antenna study; the CAMEX-4 airborne radars rainfall observation experiment, and the validation of the EOS AQUA AMSR-Es rainfall measurements. He has been a member of 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. Dr. Im is a Senior Member of the IEEE, and a member of American Meteorological Society. He has published over 120 articles in refereed journals and conference proceedings, and 10 technical reports.



Raymond J. Andraka, P.E. 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 5 years with Raytheon Missile Systems designing radar signal processors, and 3 years of signal detection and reconstruction al-

gorithm development for the US Air Force. He also spent 2 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. Ray earned his BSEE from Lehigh University in Jan 1984 and his MSEE from the University of Massachusetts in June 1992.