



Remote Memory Access Protocol Target Node Intellectual Property

Goddard Space Flight Center, Greenbelt, Maryland

The MagnetoSpheric Multiscale (MMS) mission had a requirement to use the Remote Memory Access Protocol (RMAP) over its SpaceWire network. At the time, no known intellectual property (IP) cores were available for purchase. Additionally, MMS preferred to implement the RMAP functionality with control over the low-level details of the design. For example, not all the RMAP standard functionality was needed, and it was desired to implement only the portions of the RMAP protocol that were needed. RMAP functionality had been previously implemented in commercial off-the-shelf (COTS) products, but the IP core was not available for purchase.

The RMAP Target IP core is a VHDL (VHSIC Hardware Description Lan-

guage) description of a digital logic design suitable for implementation in an FPGA (field-programmable gate array) or ASIC (application-specific integrated circuit) that parses SpaceWire packets that conform to the RMAP standard. The RMAP packet protocol allows a network host to access and control a target device using address mapping. This capability allows SpaceWire devices to be managed in a standardized way that simplifies the hardware design of the device, as well as the development of the software that controls the device.

The RMAP Target IP core has some features that are unique and not specified in the RMAP standard. One such feature is the ability to automatically abort transactions if the back-end logic

does not respond to read/write requests within a predefined time. When a request times out, the RMAP Target IP core automatically retracts the request and returns a command response with an appropriate status in the response packet's header. Another such feature is the ability to control the SpaceWire node or router using RMAP transactions in the extended address range. This allows the SpaceWire network host to manage the SpaceWire network elements using RMAP packets, which reduces the number of protocols that the network host needs to support.

This work was done by Omar Haddad of Goddard Space Flight Center. Further information is contained in a TSP (see page 1). GSC-16467-1

Soft Decision Analyzer

An in-depth insight of a communication system's receiver performance is provided in a variety of operating conditions.

Lyndon B. Johnson Space Center, Houston, Texas

The Soft Decision Analyzer (SDA) is an instrument that combines hardware, firmware, and software to perform real-time closed-loop end-to-end statistical analysis of single- or dual-channel serial digital RF communications systems operating in very low signal-to-noise conditions. As an innovation, the unique SDA capabilities allow it to perform analysis of situations where the receiving communication system slips bits due to low signal-to-noise conditions or experiences constellation rotations resulting in channel polarity in versions or channel assignment swaps. SDA's closed-loop detection allows it to instrument a live system and correlate observations with frame, codeword, and packet losses, as well as Quality of Service (QoS) and Quality of Experience (QoE) events. The SDA's abilities are not confined to performing analysis in low signal-to-noise conditions. Its analysis provides in-

depth insight of a communication system's receiver performance in a variety of operating conditions.

The SDA incorporates two techniques for identifying slips. The first is an examination of content of the received data stream's relation to the transmitted data content and the second is a direct examination of the receiver's recovered clock signals relative to a reference. Both techniques provide benefits in different ways and allow the communication engineer evaluating test results increased confidence and understanding of receiver performance. Direct examination of data contents is performed by two different data techniques, power correlation or a modified Massey correlation, and can be applied to soft decision data widths 1 to 12 bits wide over a correlation depth ranging from 16 to 512 samples. The SDA detects receiver bit slips within a ± 4 bits window and can handle

systems with up to four quadrants (QPSK, SQPSK, and BPSK systems). The SDA continuously monitors correlation results to characterize slips and quadrant change and is capable of performing analysis even when the receiver under test is subjected to conditions where its performance degrades to high error rates (30 percent or beyond). The design incorporates a number of features, such as watchdog triggers that permit the SDA system to recover from large receiver upsets automatically and continue accumulating performance analysis unaided by operator intervention. This accommodates tests that can last in the order of days in order to gain statistical confidence in results and is also useful for capturing snapshots of rare events.

Slip and quadrant performance are displayed in real time in addition to being logged for later analysis. The SDA pro-

vides methods of evaluating the significance of BER (bit error rate) test results as well as techniques to characterize imbalances due to inter-symbol interference or inter-channel interference. Techniques employed in the SDA provide a means to statistically characterize receiver performance efficiently using a minimum of accumulated test results with a definable level

of error. These advanced techniques have broad application to other fields that rely on evaluation of binomial experiments (pass/fail, true/false, go/no-go), and allow evaluators to extract more information from fewer trials.

This work was done by Glen Steele, Chatwin Lansdowne, Joan Zucha, and Adam Schlesinger of Johnson Space Center. For fur-

ther information, contact the JSC Innovation Partnerships Office at (281) 483-3809.

This invention is owned by NASA, and a patent application has been filed. Inquiries concerning nonexclusive or exclusive license for its commercial development should be addressed to the Patent Counsel, Johnson Space Center, (281) 483-1003. Refer to MSC-24798-1.

Distributed Prognostics and Health Management With a Wireless Network Architecture

Distributed architectures prevent total system failure during emergencies, allowing parts of the system to continue to function, and making overall system recovery faster.

Ames Research Center, Moffett Field, California

A heterogeneous set of system components monitored by a varied suite of sensors and a particle-filtering (PF) framework, with the power and the flexibility to adapt to the different diagnostic and prognostic needs, has been developed. Both the diagnostic and prognostic tasks are formulated as a particle-filtering problem in order to explicitly represent and manage uncertainties in state estimation and remaining life estimation. Current state-of-the-art prognostic health management (PHM) systems are mostly centralized in nature, where all the processing is reliant on a single processor. This can lead to a loss in functionality in case of a crash of the central processor or monitor. Furthermore, with increases in the volume of sensor data as well as the complexity of algorithms, traditional centralized systems become — for a number of reasons — somewhat ungainly for successful deployment, and efficient distributed architectures can be more beneficial.

The distributed health management architecture is comprised of a network of smart sensor devices. These devices monitor the health of various subsystems or modules. They perform diagnostics operations and trigger prognostics operations based on user-defined thresholds and rules. The sensor devices, called computing elements (CEs), consist of a sensor, or set of sensors, and a communication device (i.e., a wireless transceiver beside an embedded processing element). The CE runs in either a diagnostic or prognostic operating mode. The diagnostic mode is the default mode where a CE monitors a given subsystem or component through a low-weight diagnostic algorithm. If a CE detects a critical condition during monitoring, it raises a flag. Depending on availability of resources, a networked local cluster of CEs is formed that then carries out prognostics and fault mitigation by efficient distribution of the tasks. It should be

noted that the CEs are expected not to suspend their previous tasks in the prognostic mode. When the prognostics task is over, and after appropriate actions have been taken, all CEs return to their original default configuration.

Wireless technology-based implementation would ensure more flexibility in terms of sensor placement. It would also allow more sensors to be deployed because the overhead related to weights of wired systems is not present. Distributed architectures are furthermore generally robust with regard to recovery from node failures.

This work was done by Kai Goebel of Ames Research Center, and Sankalita Saha and Bhaskar Sha of Mission Critical Technologies, Inc.

Inquiries concerning rights for the commercial use of this invention should be addressed to the Ames Technology Partnerships Division at 1-855-NASA-BIZ (1-855-6272-249). Refer to ARC-16450-1.

Minimal Power Latch for Single-Slope ADCs

A CMOS implementation for remote sensing applications results in further reduction of power consumption and noise.

NASA's Jet Propulsion Laboratory, Pasadena, California

Column-parallel analog-to-digital converters (ADCs) for imagers involve simultaneous operation of many ADCs. Single-slope ADCs are well adapted to this use because of their simplicity. Each ADC contains a comparator, comparing its input signal level to an increasing reference signal (ramp). When the ramp is

equal to the input, the comparator triggers a latch that captures an encoded counter value (code). Knowing the captured code, the ramp value and hence the input signal are determined. In a column-parallel ADC, each column contains only the comparator and the latches; the ramp and code generation

are shared.

In conventional latch or flip-flop circuits, there is an input stage that tracks the input signal, and this stage consumes switching current every time the input changes. With many columns, many bits, and high code rates, this switching current can be substantial. It