Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

STORAGE PROCESSOR LEVERAGES LEON

Network RAID Controller Based on Free SPARC V8-Compatible Core

By Tom R. Halfhill {5/2/05-02}

How to beat the high cost of living: design a new chip around bits and pieces of LEON, a freely licensable SPARC-compatible processor core from the European Space Agency. The finished RAID controller chip is now solving space problems of a different sort by bringing affordable network-attached storage (NAS) to home and small-business users.

Left on the ground this time are intellectual-property (IP) vendors such as ARC International, ARM, MIPS Technologies, and Tensilica. One of them almost certainly would have scored this design win if a free alternative to their licensable processor cores were not available. However, designing a processor using a free core involves trade-offs that will keep IP vendors from losing significant business for the near future.

The new LEON-derived chip is the IT3107 network storage processor from Infrant Technologies, a four-year-old privately held company in Fremont, California. ("Infrant" stands for "infrastructure new technology.") This is the second RAID controller Infrant has designed using parts of LEON1, a 32-bit processor core adhering to the SPARC V8 instruction-set architecture. The European Space Agency freely distributes a synthesizable VHDL model of LEON1 under a GNU license. (See *MPR 1/16/01-01*, "A GNU SPARC?")

Infrant chose to start with LEON for two reasons. First, it believes the SPARC-compatible architecture is more suitable than other licensable architectures for what amounts to an embedded server. According to Infrant, even the MIPS architecture isn't as good for this application as SPARC is. The company's second reason for choosing LEON was, frankly, to avoid paying a commercial IP vendor hundreds of thousands of dollars in upfront licensing fees and chip royalties.

In addition to offering the IT3107 as a standard part, Infrant also sells complete IT3107-based system boards and RAID software to OEM companies making NAS products. The bundled software includes Infrant's own RAIDiator embedded operating system and an easy management program that runs inside a web browser. Infrant's customers can produce their own RAID subsystems by simply wrapping the system board in a plastic case and adding Serial ATA (SATA) disk drives. This is the kind of turnkey product kit many of today's OEMs demand.

To bring this product concept to market more quickly, Infrant also sells its own RAID boxes at retail. Its newest product is called ExpandaNAS, because users can start with only one disk drive and add more later. Intended for home and small-business users needing large amounts of plug-and-play disk storage with RAID protection, ExpandaNAS has internal connectors for up to four SATA hard drives. Using the largest SATA drives currently available (400GB), users can stuff up to 1.6TB of storage into a box retailing for less than $500 (without drives). Although the ExpandaNAS product would seem to compete with the IT3107-based NAS products from Infrant's own customers, the company says its boxes have limited retail distribution and haven't upset any customers.

Designing a Processor With LEON

To create the IT3107, Infrant started with the LEON1 integer unit, which implements all the SPARC V8 integer instructions and the usual SPARC register windows. Normally, storage controllers don't need floating-point math, so the IT3107 has no FPU. (LEON1 doesn't include an FPU, either, but it has an

interface for the Sun MicroSPARC FPU core, available under a Sun community source license.) As a result, Infrant's core is a basic 32-bit RISC processor with a conventional five-stage uniscalar pipeline. Fabricated in 0.15-micron CMOS, the IT3107 runs at 280MHz.

What distinguishes the IT3107 from other LEON-based chips is the application-specific IP wrapped around the processor core—or rather, the processor cores, because the IT3107 is an asymmetric dual-core chip. One core is the CPU, which runs the RAIDiator operating system and other embedded software, assisted by a hardware RAID engine. The second core powers the DMA controller, which Infrant calls the DataJunction. The DMA controller connects the RAID engine to a 64-bit DDR-1 memory interface running at 140MHz (effectively 280MHz), providing 2.2GB/s of bandwidth to PC2700 DRAM.

In addition, the IT3107 has a four-channel 1.5Gb/s SATA ×4 controller, a Gigabit Ethernet controller with TCP/IP offload hardware, a NAND flash-memory controller for diskless system boots, a 32-bit/33MHz PCI host controller, and a miscellaneous serial controller, which supports a UART, Two-Wire interface, LED interface, and general-purpose I/O (GPIO). There's also a JTAG boundary-scan interface for debugging. The physical link layer (PHY) for SATA is built in, but the Ethernet controller requires an external PHY. The whole chip design fits into a 449-pin PBGA package requiring a 1.5V core supply, 3.3V I/O, and 2.5V DRAM. Figure 1 shows a block diagram of the IT3107.

In a NAS subsystem, disk drives connect to the SATA interfaces and transfer data to and from the host system over a network using the chip's Gigabit Ethernet port. The PCI interface supports additional peripherals—such as external disk drives, printers, and connections to wireless networks—but isn't involved in moving data to the disk array, so it's not a bottleneck. Although the IT3107 has some logic to accelerate Triple-DES (3DES) cryptography, that hardware lies fallow for now, owing to the lack of an industry-standard RAID security protocol.

At $120 in 1,000-unit quantities, the IT3107 seems pricey. It's in the same price range as network storage controllers for much higher end RAID applications, such as AMCC's new PowerPC 440SPe. Announced at Fall Processor Forum 2004, the three-way superscalar PowerPC 440SPe runs at more than twice the clock frequency of the IT3107 and has Gigabit Ethernet, three PCI Express interfaces (eight lanes total), PCI-X, a DDR-2 memory controller, a 256KB L2 cache, and many more features. Yet it costs $110 (533MHz) to $128 (667MHz), virtually the same price as the IT3107, albeit in 10,000-unit quantities. (See *MPR 10/25/04-01*, "Embedded CPUs Zoom at FPF.")

Despite its price, the IT3107 is better suited for the lower-end RAID applications it targets. It integrates the SATA controllers and PHYs (the PowerPC 440SPe requires external chips for those functions), dedicates a second 32-bit processor core to DMA, and includes a license for Infrant's RAIDiator embedded software. The PowerPC 440SPe is clearly designed for RAID subsystems in high-performance servers.

RAIDiator Has User-Level Management Tools

Infrant's RAIDiator software is based on Linux and optimized for RAID applications. It supports RAID levels 0, 1, and 5, plus a proprietary RAID level called xRAID, for which Infrant has applied for a U.S. patent. With xRAID, users can add disks without manually reconfiguring the array. RAIDiator also supports hot swapping and hot spares. The latter feature allows users to designate one drive as a backup that automatically joins the RAID if another drive fails.

RAIDiator is much more than an embedded OS. It also provides a user-level configuration and management tool that runs inside a browser window, with step-by-step wizards for common tasks as well as an advanced-user mode. The management program runs on Windows, Macintosh, Linux, and Unix systems; has multiple levels of security (including separate passwords for user accounts); supports multiple clients; and automatically sends email messages to the user's mailbox if the subsystem detects a disk error.

In the past, OEMs would have written their own software such as RAIDiator. Now, many OEMs expect suppliers to provide near-finished designs with system boards and software. Consequently, Infrant employs more than twice as many software programmers as hardware engineers, despite the extra design effort required to create its custom processor core. The IT3107 and ExpandaNAS board provide a

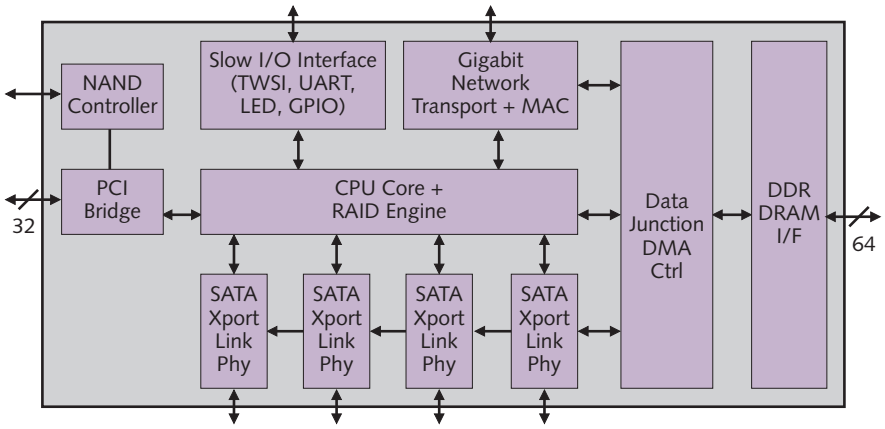

Figure 1. Infrant IT3107 block diagram. This network storage controller has two 32-bit RISC cores derived from the synthesizable LEON1 core freely licensed by the European Space Agency. Additional IP makes it a specialized chip for network-attached storage (NAS) applications. One processor core is the CPU, which runs embedded RAID software, and the other core is in the DataJunction DMA controller.

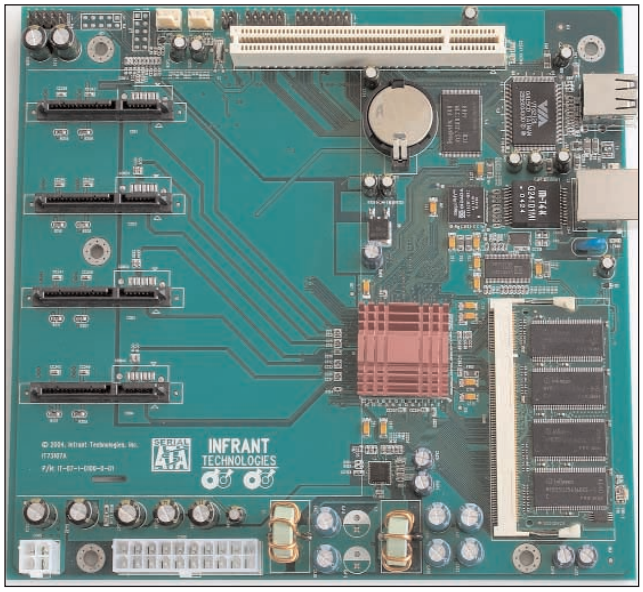

Figure 2. Infrant's ExpandaNAS system board comes populated with an IT3107 processor (beneath a heatsink, slightly lower-right of center), 128MB of SO-DIMM memory (lower right), and the RAIDiator software embedded in on-board flash memory (upper right). A long-life battery (near the PCI slot, at top) powers a real-time clock for RAID event logging and time stamping. Four slots for SATA disk drives are at left. The system board is mostly empty to leave room for hot-swapping the drives. To make a RAID box for NAS, just about all an OEM must add is some industrial design.

turnkey OEM solution. Figure 2 is a photo of Infrant's ExpandaNAS system board as offered to customers.

Weighing the Design Trade-Offs

Licensing a processor core from a commercial IP vendor would have been faster than designing a custom processor core, even after borrowing the LEON1 integer unit. However, because the custom core is a relatively simple RISC design, Infrant believes the time devoted to the project wasn't excessive. Infrant's president and CEO Paul Tien estimates the custom core required about two man-years to design and perhaps ten man-years to debug and verify, with four hardware engineers working full time. (Meanwhile, more than ten programmers were busy writing RAIDiator.)

Had Infrant licensed a suitable processor from a commercial IP vendor, the core would have been ready to use off the shelf. Licensable-IP cores are preverified, so the engineers would have saved time during that part of the project, too, although the whole chip design would still require verification. Infrant says it preferred SPARC to other licensable architectures, but the company could have started with a customizable processor from ARC, MIPS, or Tensilica and added the application-specific features needed. (Tensilica's Xtensa processors even have SPARC-like register windows.) The main trade-off is whether the cost of licensing commercial IP would exceed the revenue lost by reaching the market later.

Ownership is another potential issue. LEON1 is covered by the GNU Lesser General Public License (LPGL), a variation of the GNU General Public License (GPL). Under the terms of the LPGL, users can freely modify the LEON1 source code and combine it with their own code to produce a derivative work. They needn't publish the source code of other IP combined with LEON1, but they are obligated to publish any modifications to the free core. With the IT3107, the demarcation between LEON1 source code and proprietary source code should be clear: LEON1 is in VHDL, and Infrant wrote everything else in Verilog. Infrant is obligated to publicly share only the VHDL of the slightly modified LEON1 integer unit.

Microprocessor Report believes Infrant made a reasonable trade-off by designing a custom processor core—at least on this occasion. The time-to-market disadvantage is less important for a young company trying to establish itself in a new market. Convincing home and small-business users they need a networked RAID subsystem and can manage it themselves will probably take more effort than Infrant's design project did.

However, in a more established market, in which eager customers are lining up at the door of a more established company, licensing a commercial processor core is usually more sensible. Commercial IP cores are expensive, but they are also preverified, have better hardware/software-design tools, run more embedded operating systems, come with expert technical support, and are verified to work with libraries of other IP. If Infrant someday outgrows its homemade processor core, we won't be surprised to see the company open its checkbook and license a processor from a commercial IP vendor. ◇

Price & Availability

Infrant is shipping the IT3107 network storage processor now, as well as the ExpandaNAS system boards and RAIDiator embedded software. In 1,000-unit quantities, the IT3107 costs $120, including a RAIDiator license. A nearly identical chip, the IT3102, supports only two-channel SATA instead of four-channel SATA and costs $80. For more information, visit *www.infrant.com*.

For more information about LEON from the European Space Agency, visit *www.estec.esa.nl/wsmwww/core/soc.html*. For information about LEON2 and LEON3, newer versions of the LEON core, see *www.gaisler.com*. For information about the GNU Lesser General Public License, see *www.gnu.org/copyleft/lesser.html*.