Cahners
MICROPROCESSOR REPORT
www.MPRonline.com
THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

AGERE'S PIPELINED DREAM CHIP

PayloadPlus Targets Routers at up to OC48c Speeds

By Kevin Krewell {6/12/00-02}

On January 20 of this year, Agere became the Networking Processor business unit of Lucent Microelectronics, joining the ranks of a half-dozen other network-processor start-ups that have been absorbed into larger corporations. The two-year-old Agere was founded to build a business around Vic Bennett's patent (U.S. No. 5,813,001) for "efficient search of a knowledge base to determine whether an object matches any plurality of knowledge base entries," which allows fast, efficient, and deterministic pattern searches. That search becomes a parallel tree search through the bit pattern to then determine a course of action—in this case, making a routing decision on a data packet. The strength of the algorithm is that it can search the bit stream quickly and can pack the decision tree into a minimal amount of memory. Armed with this algorithm, Agere believes it can classify single-cell (ATM) packets up to and including layer 4 (the transport layer of the OSI model protocol stack) at 2.4Gb/s wire speeds (OC48c, a SONET standard for optical-fiber backbones).

Lucent acquired the young company largely on the basis of its expertise in combining IP (Internet protocol) with fast ATM. Lucent Microelectronics had the expertise in ATM and configurable network components but lacked a programmable solution and IP experience; it got that with Agere. With the addition of the PayloadPlus processor, Lucent can now offer complete OC12 and OC48c line-card solutions that are programmable and flexible enough to support multiple network- and transport-layer protocols. Agere says support for MPLS, IPv6, VoIP, ATM AAL2, or the next new protocol is only a matter of dropping new software onto its hardware. Agere will operate as a wholly owned subsidiary, and its 120 employees will join the engineers at Lucent Microelectronics.

The PayloadPlus Chip-Set Breakdown

The Agere chip set, as Figure 1 shows, consists of three chips: fast-pattern processor (FPP), routing-switch processor (RSP), and Agere system interface (ASI). The first two chips perform wire-speed processing; the ASI and another microprocessor perform packet-speed (back-end) processing. The ASI provides a PCI interface to a microprocessor, such as PowerPC, to execute the routing protocol. The three chips in the Agere set use a total of only about four million transistors. Agere is promising a lot of magic from these modest chips, based mostly on its patented data-matching/pattern-matching logic and a collection of VLIW processors to power the routing engine.

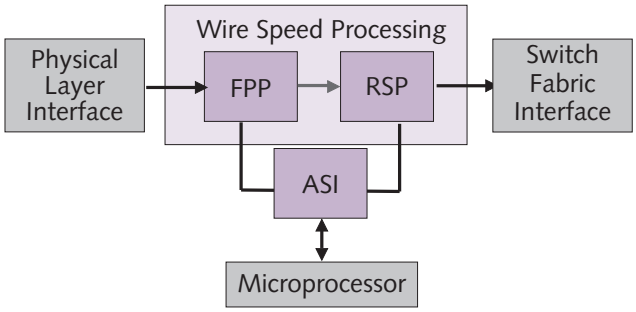

Figure 1. Agere's router solution consists of three chips: the FPP, the RSP, and the ASI. The FPP and RSP process packets at gigabit wire speeds, while the ASI connects to a microprocessor on the line card and processes statistics.

One of the keys to the remarkably low transistor count is the fact that the required data and program memories are located off chip. The chip set therefore requires two 64-bit-wide DRAM interfaces for packet data and seven 64-bit-wide synchronous SRAM interfaces for program and routing control. These nine 64-bit memory channels take up a large number of pins, which caused Agere to break the design into three chips rather than two or one. The three-chip set uses two 655-pin BGA packages (FPP and RSP) and one 448-pin BGA (ASI). The low integration of the chip set is a disadvantage, because the chip set, memory, and support logic take up roughly 20 square inches of board space, and twice that for a full-duplex OC48c design, such as that shown in Figure 2. The chip set dissipates 10W at 133MHz in 0.18-micron CMOS.

To bring the concept to life in the shortest possible time, Agere implemented a proof-of-concept vehicle set in Xilinx FPGAs. This vehicle allowed potential customers to evaluate the technology and lowered the design risk for both customers and Agere. The FPGA implementation runs at 33MHz and can support an OC12 (622Mb/s) SONET line. The FPGA implementation is used in Agere's $75,000 development system, which connects to an OC12 fiber for evaluation.

The physical layer (PHY) input to the FPP is configurable to support the ATM-standard 32-bit Utopia interface or the PMC-Sierra POS-PHY interface. ATM cell reassembly (SAR function) and layer 2 classification are performed on the FPP before the packet is stored in the PC133 SDRAM. The FPP performs layer 3 and above processing, using its fast pattern- matching logic. The packet is read out of the SDRAM and sent to the RSP for routing, with the conclusion of header processing appended to the packet in the Destination ID (DID) bits. The bandwidth provided by the 64-bit-wide PC133 memory allows one write and one read of the packet from memory at OC48c data rates.

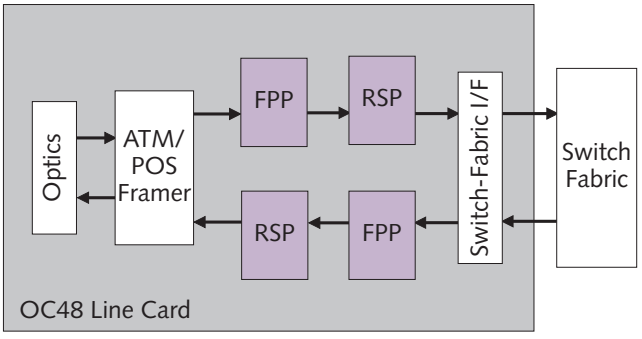

Figure 2. A full-duplex OC48c line card would require two sets of chips. For an OC12 design, it is possible to multiplex one chip set to process both directions. The Lucent switch fabric shown here can be PI-X, for OC48c, or Lucent Atlanta, for OC12.

Agere's Pipelined Approach

To classify packets, Agere pipelines the processing of the wire-speed data stream for low-latency scanning of the layer 3 and layer 4 headers. The FPP, shown in Figure 3, assembles the packet data and uses the patented pattern-matching logic to drill into the protocol headers.

The RSP, shown in Figure 4, uses information discovered by the FPP and appended to the packet to determine the packet priority and apply the appropriate traffic management. The RSP has three VLIW engines to perform traffic management, traffic shaping, and stream editing. It is the pipeline of all these 133MHz VLIW engines that allows Agere to maintain such high data throughput.

The ASI is part glue chip and part statistics gatherer, as well as a part-time coprocessor for the FPP. Its primary role is to connect the FPP and RSP to a 66MHz, 32-bit PCI bus, which provides a standard interface to the line-card microprocessor. The microprocessor could conceivably be any general-purpose processor, but Agere has so far used only x86 and PowerPC processors. The microprocessor must run a real-time operating system, such as VxWorks, and it is responsible for bookkeeping tasks like initializing the chip set and maintaining routing tables.

Using the 8-bit packet over the SONET/SDH (POS) interface from the RSP, the ASI takes control packets containing routing information culled by the RSP and passes that information to the microprocessor to update routing tables. The microprocessor can use the ASI to inject packet data into the FPP through another 8-bit POS interface. The ASI, shown in Figure 5, also provides statistical counters that can be polled by the microprocessor.

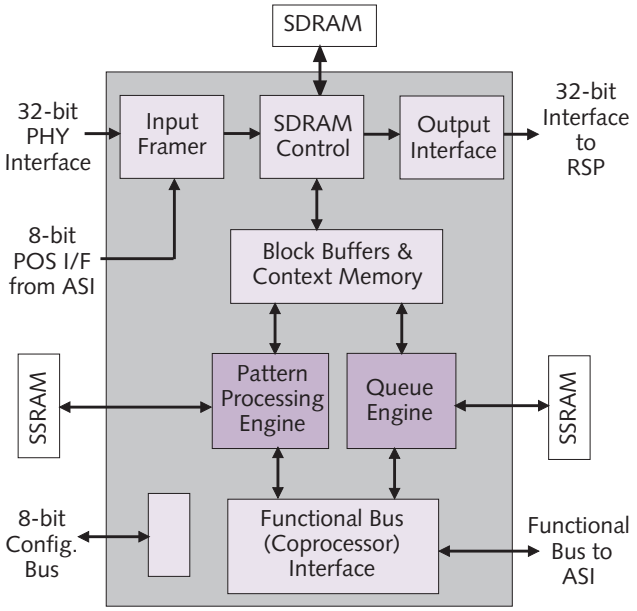

Figure 3. The FPP is the brains of the chip set. The Pattern Processing Engine searches through the reassembled packets for layer 3 and above routing information.

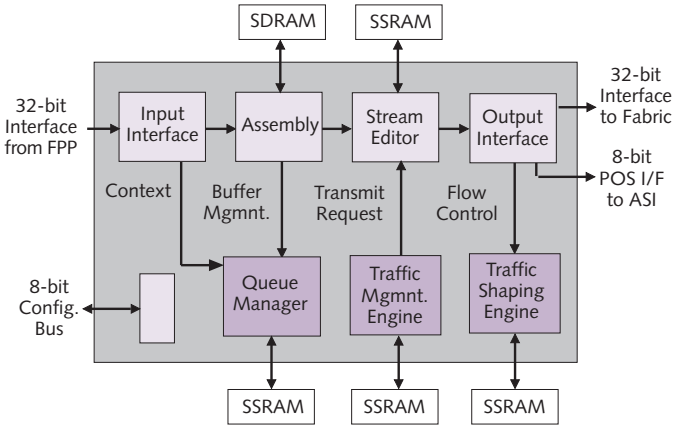

Figure 4. The RSP block diagram shows the five synchronous SRAM memory channels that are required to keep the VLIW processors humming along.

Software: Where the Rubber Meets the Road

One of the challenges of a unique architecture is finding someone who can program it. People who understand the nitty-gritty of protocol processing are a rare breed and must be used effectively. The software tool chain must therefore provide high-level support for maximum productivity. Most vendors with RISC-derived network processors use tools adapted from existing instruction-set standards. This is the advantage of using existing processor cores, such as MIPS and StrongArm, where tools are readily available and only minor modifications are required.

Agere has taken a different, higher-level approach that it calls functional programming. The FPP is programmed in a proprietary language often used to classify protocol information and is designed to perform that task with great efficiency. Agere created a C-like scripting language to allow its customers to develop custom routing protocols. If writing code for all those processors is too much work, Agere is wisely offering standard management algorithms as part of the development package.

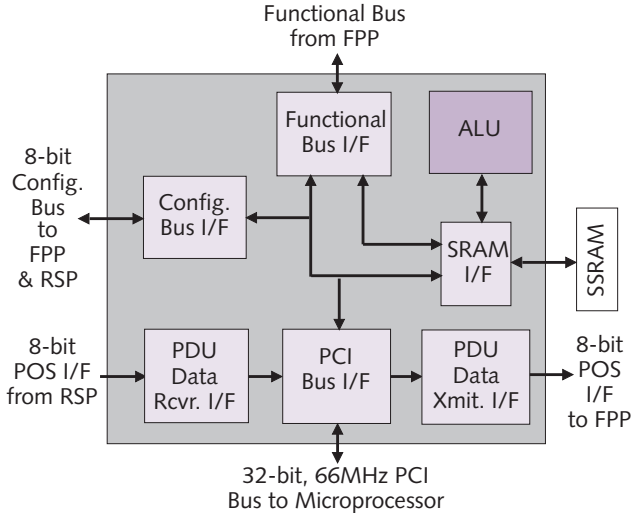

Figure 5. The ASI is the glue that ties the FPP and RSP to the line-card microprocessor. The ASI is also used as a coprocessor to the FPP, through the Functional Bus, to gather statistical information on packet traffic.

For More Information

Agere's Web site at *www.agere.com* contains a number of white papers and some product information. Detailed product information is available under NDA.

Everyone Loves Network Processors

The patented fast-pattern-matching logic is the heart of the FPP and allows packets to be classified in a deterministic manner, at wire speeds. The advantage of the FPP over the competition is that it can process OSI layer 2, 3, and 4 protocols at line speeds up to OC48c without a custom ASIC or a content-addressable memory (CAM). Processing the lower protocol layers is not particularly new or particularly hard at slower data rates, but many of the RISC solutions will find it difficult to scale to OC48c and OC192 (9.8Gb/s) data rates. Agere is promising that its chip-set architecture is scalable to OC192.

Today, Agere's competitors are mostly vendors of custom ASICs and a few vendors of network-processors, such as C-Port. But the market for edge routers is attracting a lot of vendors. Most network processors have taken a core processor and added coprocessors that work in parallel. That approach relies on the natural parallelism of packet processing. It's also easier to create a processing core and then step multiple cores onto a die. The key issue is keeping those cores busy. C-Port (now a part of Motorola) appears to be closest to Agere on target market and processing power. Intel's IXP1200 (see *MPR 9/13/99-01*, "Intel Network Processor Targets Routers") addresses the low end of this application with the StrongArm-based design it acquired from DEC. Sitera (now a part of Vitesse) appears to be taking a similar philosophy, but it is using a fast MIPS core instead of StrongArm.

At this time, there seems to be no shortage of specialized network processors and niche designs in which to use them. Each new vendor has found customers willing to invest in its version of network-acceleration silicon—Agere is engaged with about 20 customers, including a Lucent design. In many cases (Agere, C-Port, Sitre), the next logical step was acquisition by a larger, more established company. So the present state of affairs is telling prospective network-processor designers this: If you build it, someone, somewhere, will design it in, or, at the very least, buy your company. ◇