Alliance Detours Into Routers

SRAM-Rich Network Processor Is Departure for Memory Vendor

by Tom R. Halfhill

When the only tool you've got is a hammer, every job looks like a nail. So perhaps it's no surprise that a memory company's first processor would devote fewer than 1% of its 61 million transistors to logic.

Alliance Semiconductor's new IPRP-V4 (Internet Protocol Routing Processor) architecture certainly defies classification. Is it embedded-memory logic or embedded-logic memory? Either way, it's designed to solve a growing problem: managing and searching IP forwarding tables to keep up with high-end routers and fiber-optic backbones. The architecture is a significant departure for Alliance, which was founded in 1985 but until now has dealt exclusively in DRAM, SRAM, and flash memory.

IPRP-V4 chips won't replace a router's general-purpose CPU. Instead, they are application-specific coprocessors that would most likely displace the ASICs in core routers—the high-speed routers that act as traffic cops on the fiber backbones of enterprise networks and the Internet. Some members of the IPRP-V4 family might also be suitable for the smaller edge routers that handle packet-switching chores around the fringes of the network.

Finding the Way Without Triple-A

IPRP-V4 chips don't actually create a routing table, transfer data packets, or perform any other high-level tasks that normally fall to a router's CPU. Instead, they do only three things really fast: store the routing table, update the table, and look up entries in the table.

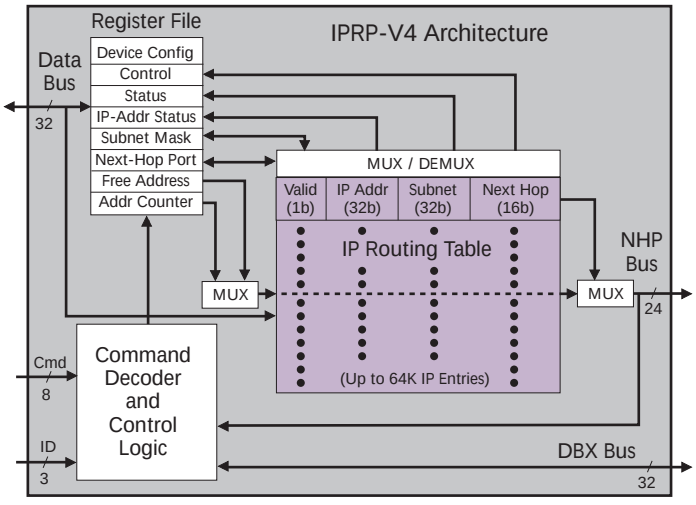

Figure 1. Alliance's IPRP-V4 processors are highly optimized database engines for storing, managing, and searching routing tables.

Every IP packet that traverses the Internet has a source address, a destination address, and a data payload. Routers guide those packets to their intended destinations. To do this, a router maintains a database of destination addresses and periodically exchanges updates with other routers. Each destination address has an associated outgoing-port address to the router that's the next hop across the network.

In an IPRP-V4 chip, this routing table has 81 bits per entry: 32 bits for the destination address, 32 bits for the IP subnet mask, 16 bits for the next-hop port address, and a 1-bit flag that indicates the validity of the entry. The faster the chip can update this table and look up entries, the faster the router can send packets on their journey.

As Figure 1 shows, IPRP-V4 chips are specialized database engines. To minimize off-chip references, they store the routing table in internal SRAM (up to 648 KBytes), which includes some very fast 22-transistor memory cells. These cells are similar to content-addressable memories (CAMs), which typically have 10-transistor cells coupled to external decoding logic. Alliance's 22-transistor cells are based on standard 6-transistor SRAMs, but they have built-in logic for prioritizing and sorting IP addresses. Alliance has applied for a patent on this design.

An IPRP-V4 chip is smart memory, not a true processor. As Table 1 shows, the architecture has an instruction set, but it's more like an I/O command set. There are 27 instructions, most of which allow a router's CPU to read and write to the chip's registers and routing table. The most important instructions are FLPM, which finds destination addresses in the table, and AUTO, which automatically updates the table. All of the instructions are eight bits long, and they execute in a two-stage pipeline.

Functionally, IPRP-V4 chips are slave coprocessors that rely heavily on the router's CPU. They cannot fetch their own instructions; the CPU invokes commands over an 8-bit command bus. They have no branches, so the CPU must manage all flow control. And the CPU gets the result of their operations—the next-hop port address for an IP packet—by invoking instructions that return the matching address over the 24-bit NHP (next-hop port) bus.

An Elephant's Memory for Routers

Alliance is offering nine members of the IPRP-V4 family to accommodate a variety of routers. There are three speed grades (66 MHz, 50 MHz, and 33 MHz), with three amounts of on-chip memory for the routing tables (648K, 324K, and 162K, which store 64K, 32K, or 16K table entries, respectively). The chips are pin compatible, providing an easy upgrade path for router vendors.

Alliance says the chips are fast enough to keep up with multiported gigabit routers at wire speeds. At 66 MHz, an IPRP-V4 chip can feed 28 routing ports at 1 Gbit/s per port, or 10 optical-carrier OC-48 ports at 2.4 Gbits/s per port, or three OC-192 ports at 9.6 Gbits/s per port.

To put those figures into perspective: today's fastest Internet backbones are OC-3 lines that operate at a relatively slow 155 Mbits/s. The next-generation Internet 2 will probably use OC-12 (622 Mbits/s) and OC-48. Those are also the fastest lines supported by Cisco's most powerful 12000-series routers. And the highest-density IPRP-V4 chips have enough memory to store the entire routing table of a major Internet hub such as MAE West—about 61,000 entries.

Future routers that need even more capacity can gluelessly cascade up to six of the 66-MHz chips, or eight of the 33/50-MHz chips, on the 32-bit database extension (DBX) bus. When linked in this manner, the chips can search a single routing table as large as 512K entries, with no loss of performance. To a router, the multichip forwarding table looks like a single database. The chips can execute parallel searches across the database and return an IP address in a single cycle—the same throughput as lookups on a single chip.

The DBX bus is limited by capacitive loading, however. Higher-frequency chips increase the loading, so fewer of them can be linked together without losing performance.

Worse Than ZIP Codes

As the name implies, Alliance designed IPRP-V4 for today's IPv4 standard, which has 32-bit IP addresses. One problem the architecture must deal with is the complex addressing format of this standard.

A shortage of IP addresses caused by the exploding growth of the Web forced the Internet Engineering Task Force to create the classless interdomain routing (CIDR) protocol. The "cider" protocol substitutes variable-length network prefixes for some fixed-length IP addresses. A network prefix can point to a gateway (or more than one gateway) that leads to numerous IP addresses. But a prefix is shorter than an IP address, in the same way that the mailing label on a magazine addressed to a company is shorter than a mailing label addressed to a person at that company.

One consequence is that a routing table may contain more than one match for a single destination. There may be an entry for a specific node (e.g., an individual PC) as well as a CIDR entry for the network to which the node is attached. The router prefers the most specific address, which is the longest entry. Some routers have to sort the entries by length and re-sort them whenever the table needs updating. IPRP-V4 chips can use the AUTO instruction, which updates the entries without flushing and re-sorting the table.

Similarly, an IPRP-V4 chip can browse its table for the longest prefix match by using the FLPM instruction, which returns a result for the first match in two clock cycles. After that, the two-stage pipeline is primed, and subsequent searches return a match every cycle.

Price & Availability

Alliance (San Jose, Calif.) is currently sampling the ASN64132 chip, which runs at 33, 50, or 66 MHz and has enough SRAM for 64K routing-table entries. The company will sample the ASN32132 (32K entries) and the ASN16132 (16K entries) in 3Q99. All three versions will enter production in 3Q99. Prices for the 66-MHz parts will be $150 (ASN64132), $75 (ASN32132), and $50 (ASN16132) in 10,000-unit quantities. The 50-MHz chips cost 10% less, and the 33-MHz chips cost 20% less. For more information, go to *www.alsc.com*.

If a router distributes its table across multiple IPRP-V4 chips, the FLPM instruction finds the longest prefix match among them and returns the result at the same single-cycle rate. A three-bit device ID indicates which chip's memory contains the match.

The future IPv6 standard will expand an IP address to 128 bits, creating enough capacity to give every light bulb on the planet its own IP address (should it ever become necessary). Alliance is already working on an IPRP-V6 architecture that widens the data bus to 128 bits.

Name	Description	Cycles
Register Read/Write Instructions		
CERR	Clear error	1
LDAC	Load address counter	1
LDCR	Load control register	1
LDNHPR	Load next-hop-port data register	1
LDSMR	Load subnet-mask data register	1
RDAC	Read address counter	1
RDCR	Read control register	1
RDFR	Read free-address register	1
RDIPDR	Read IP-address data register	1
RDNHPDR	Read next-hop-port data register	1
RDSMR	Read subnet-mask data register	1
RDSR	Read status register	1
Database Read/Write Instructions		
CMTYAC	Delete database entry at address counter	1
LDDBAC	Load database entry at address counter	3
LDDBFA	Load database entry at free address	3
LDNHPAC	Load next-hop-port ID field	1
RDDB	Read database entry at address counter	3
RDDBINC	Read database entry, inc address counter	3
Matching Entry Instructions		
AUTO	Auto update	4
CMTYME	Delete matching entry	2
FLPM	Find longest prefix match	1
LNHPME	Load next-hop-port ID at matching entry	2
RNHPME	Read next-hop-port ID at matching entry	2
Special Instructions		
GFA	Get free address	1
NOP	No operation	1
RSFA	Read system free address	1
SRES	Software reset	9

Table 1. The IPRP-V4 instruction set has only 27 operations.

Alliance's foundry, UMC in Taiwan, will manufacture the current-generation chips in a 0.25-micron five-layer-metal process. The largest-capacity chip has a 168-mm^2 die and consumes an estimated 5.7 W at 66 MHz; consumption drops to 1.5 W for the smallest part at 33 MHz. Future generations will move to a 0.18-micron process, which will permit Alliance to boost clock speeds into the 100-MHz range and cram even more memory on a chip.

Undercutting the Competition

Today's routers generally fall into three categories. Older routers, such as the Cisco 7000, cache the routing table in SRAM and use a CPU to search the table. But it's difficult to scale this approach to high-end routers, because CPUs aren't keeping up with the growth in network bandwidth.

Faster routers assist the CPU by embedding the search algorithm in an ASIC, again caching the table in SRAM. A third approach combines an ASIC and a CPU with CAM. Either solution can deliver more performance than a CPU alone, but both are limited by memory bandwidth. Also, the memory is expensive, and ASICs take a long time to develop.

Alliance says it can deliver more performance at a comparable price while shortening development cycles, because IPRP-V4 chips are standard parts, not custom ASICs. At $150 for the fastest, largest-capacity part, the Alliance chips are indeed competitively priced against large ASICs.

In the networking market, however, entire companies are built on proprietary solutions for these problems. These companies can be stubbornly religious about their intellectual property. They are loath to change architectures until convinced that their current strategy has gathered not only the low-hanging fruit but also the high-dangling produce.

In addition, Alliance is a relatively small company ($48 million in revenue in its latest fiscal year), with no track record outside of the memory market. Network processing is attracting some heavyweight competition—including Intel, which will soon announce a StrongArm-based chip that performs some of the same functions as IPRP-V4 chips.

Alliance's ace could be its unusual architecture. By adding a small amount of logic to some fast SRAM, Alliance has created a niche product that can sell for a higher price than SRAM alone. Yet the cost is still low for a chip that's going into an expensive router—low enough, perhaps, to undercut the more processor-centric solutions from Intel and other competitors. Memory vendors are accustomed to slim margins that would traumatize an Intel accountant. If Alliance can score some small design wins and work its way up, the company's detour into network processing won't be a wrong turn. M