

Special-purpose SRAMs smooth the ride

DATA AND TELEPHONE NETWORKS ARE EXPERIENCING MORE TRAFFIC OF MULTIPLE TYPES AND AT HIGHER SPEEDS WITH VARYING DELIVERY EXPECTATIONS. ADVANCED SRAM-BASED MEMORIES PLAY A BIG ROLE IN RESPONDING TO PERFORMANCE AND FLEXIBILITY NEEDS, AND APPLICATIONS EXTEND BEYOND THE NETWORK.



If a single CAM doesn't give you enough density, try a module instead (courtesy NetLogic Microsystems).

THE OVERUSED DESCRIPTION OF THE INTERNET AS “the information superhighway” is both unrealistic in its expectations and too narrow in its scope. The phrase conjures up images of data rapidly flowing along a wide path that originates and ends at each computer node. Anyone who has received a file-transfer-protocol (FTP) or Web-site time-out error, heard a “circuits-are-busy”

recording after dialing the modem, or suffered through slow data transfers knows the reality behind the hype.

True, today's communications network contains high-speed fiber-optic backbones, but the last-mile connection to most homes and businesses is twisted-pair wire. Most computers still access the Internet through lowly analog modems. And the computer isn't the only data portal. Televisions and telephones provide other methods of exchanging information. (However, we won't dwell on the average *quality* of such information exchanges.) Coaxial cable and various wireless protocols serve as alternative delivery methods.

To some extent, today's worldwide information network *does* resemble a roadway. The data packets are its vehicles; glass, copper, and air are its pavement.

However, any reasonably efficient roadway contains several other important components. Ramps and intersections link roadway segments with various surfaces, speed limits, and lane numbers. Stoplights and street signs bring order to what would otherwise quickly become chaotic traffic snarls, and, along with maps, direct vehicles along the fastest and shortest routes. Exchange hubs, such as train stations and ports, enable vehicles and their cargo to move among roadways and other trans-

<i>At a glance</i>	94
<i>Standardization: a mixed blessing</i>	95
<i>Memory, logic, or a little of both?</i>	96
<i>Embed it</i>	98
<i>For more information</i>	100

portation systems, such as railways and waterways.

MULTIPORT RAMS

Content-addressable memories (CAMs) and multiport RAMs (Table 1), along with their corresponding logic engines, are the important components in the data-communications and telecommunications analogies to ramps, intersections, lights, signs, maps, and exchanges. Like FIFO buffers, multiport RAMs link data channels with different speeds, bus widths, and protocols. However, unlike FIFOs, both ports' peripherals have fully random access to any location in the memory array. This capability offers numerous benefits.

If, for example, the peripherals represent two CPUs or a CPU and DSP, the processors can decouple their activities, simultaneously accessing data in different areas of the memory and boosting overall system performance. Also, remember that the roadway system contains commuter lanes and toll ways, where travelers benefit from a reduced commuting time in exchange for the inconvenience of car-pooling or shelling out a little money.

Communications networks contain their own examples of services that favor certain travelers. Bulk file transfers, such as FTP sessions, and e-mail messages normally don't have stringent packet-to-packet latency or time-to-completion requirements. However, some customers are willing to pay a little more money for improved average service quality, even if they use the same networking equipment and access methods as their peers. Latency becomes even more critical for telephone service and

AT A GLANCE

▶ Content-addressable memories (CAMs) and multiport RAMs are key contributors to a fast, reliable networking infrastructure.

▶ Multiport RAMs are becoming more proprietary as they migrate to lower voltages, higher densities, and higher speeds.

▶ CAMs have undergone tremendous architectural innovation over the last decade in an attempt to keep pace with exploding network growth and transformation.

▶ As stand-alone chips grow, smaller density versions migrate onto ASICs and PLDs.

for streaming multimedia traffic that runs over the data network.

Because a multiport RAM is a random-access device, one port can retrieve packets in a priority-based order that differs from the sequence in which the other port captured them. The more rigid sequential-access protocol of a FIFO buffer forces one port to access data only in the order that the other port receives it. CAMs play a corroborating role in implementing quality of service.

Cypress Semiconductor and Integrated Device Technology (IDT) are the two largest multiport-RAM suppliers. As with many other memory types, the most notable multiport-RAM advancements over the last several years have involved increased density, faster burst and sustained bandwidth, and wider buses. Both companies first use their leading-edge processes for higher volume, commodity SRAMs and later use these processors

to manufacture FIFOs and multiport RAMs. This delayed transition partially explains why most of today's multiport RAMs are asynchronous-interface, 5V memories.

The rest of the picture concerns devices that multiport RAMs interact with. Many common encryption/decryption chips (for virtual-private-networking) and DSPs, such as Texas Instrument's (www.ti.com) 24-bit line, contain asynchronous buses. Asynchronous multiport RAMs offer access times as low as 12 nsec with semaphore, arbitration, and master/slave support for multiprocessor configurations. Synchronous interfaces prevail for devices in which the interfacing logic is a μ P or a multicore ASIC. But with an external register and a bit of logic, you can emulate a synchronous interface on an asynchronous RAM.

Today's pipelined synchronous multiport RAMs run at 100 MHz with buses that are as much as 32 or 36 (parity) bits wide and integrated counters that simplify burst transfers. Lower latency, lower burst bandwidth, flow-through versions are also available, and some devices even allow you to select between registered and pipelined outputs within the same device on a port-by-port basis. Multiple-byte enables accommodate different-sized peripheral buses, and the ports can operate at different frequencies.

Cypress' new Flex36 devices offer a "bus-funnel" feature that integrates the buffering necessary to translate the internal $\times 36$ -bit bus port to external $\times 8/\times 9$ - or $\times 16/\times 18$ -bit alternatives (Figure 1). Most multiport RAMs are two-port offerings, but IDT also sells 8- through 32-kbit four-port devices. IDT also offers 64-

TABLE 1—CAM AND MULTIPORT-RAM SUPPLIERS AND OFFERINGS

Vendor	Circle No.	Multiport RAMs	Reduced-function multiport RAMs	CAMs	Reduced-function CAMs	Multiport-RAM cores	CAM cores and custom CAMs	CAM-interface chips
Altera	301					X	X	
Cypress Semiconductor	302	X				X		
Integrated Device Technology	303	X	X					
Kawasaki LSI	304			X			X	X
Lara Technology	305			X			X	
Lucent Technologies	306					X	X	
Mosaid Technologies	307			X			X	
Motorola	308		X		X			
Music Semiconductor	309			X			X	X
NetLogic Microsystems	310			X			X	
UTMC Microelectronics	311							X

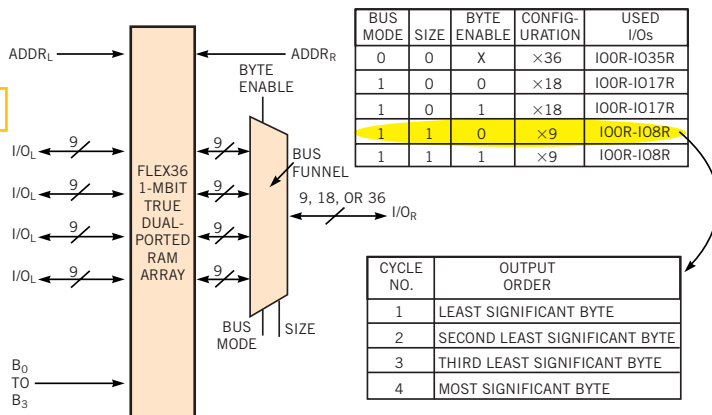
and 128-kbit serial-access memories/RAMs with a FIFOlike synchronous port on one side and a random-access asynchronous port on the other.

Historically, a surprisingly high degree of function, timing, package, and pinout compatibility existed between Cypress and IDT (see sidebar “Standardization: a mixed blessing” and Reference 1). This de facto compatibility is breaking down, however, with both vendors’ latest high-density $\times 8/\times 9$ - and $\times 16/\times 18$ -bit offerings. Cypress uses a 100-pin TQFP for both its 5 and 3.3V versions, and IDT moves to a 128-pin TQFP for 3.3V. Pinout and package incompatibility also exists at the $\times 36$ -bit bus width.

IDT claims that the larger TQFP allows the company to add power and ground pins necessary for upcoming 125- and 133-MHz versions. Cypress hints that a large die and high power consumption forced IDT’s move to a larger TQFP. Both of these future speed targets are “magic numbers”: 125 MHz is one-eighth the bit rate of Gigabit Ethernet, and 133 MHz is twice the speed of the 66-MHz PCI buses that now appear in communications equipment.

These performance targets are at the edge of the toggle rates you can achieve with low-voltage TTL, and companies will in the not-too-distant future move to HSTL (high-speed transceiver logic) and SSTL (series-stub-terminated-transceiver logic). Even more radical architectures are now in the planning stages. These architectures include dual-port RAMs with one port offering a zero-la-

Figure 1



Bus funneling is an alternative approach to byte enables in delivering bus-width flexibility (courtesy Cypress Semiconductor).

tency bus for a μP interface and the other supporting double-data-rate (DDR) transfers for a DSP. Multiport-RAM packaging alternatives also include DIP, PLCC, PQFP, and BGA.

COST REDUCTIONS

Cypress has aggressively moved its higher density dual-port RAMs to an eight-transistor (8T) bit structure, whereas IDT continues to use the six-transistor, two-resistor (6T2R) approach. The 6T2R approach traditionally delivers smaller die for a memory array of a certain density, but those differences disappear as processes near 0.2 μm . The 8T structures have lower stand-by power consumption, which suits applications in which you want to increase system reliability by eliminating heat-

dissipating fans (Reference 2). However, both dual-port structures are still larger than the six-transistor and four-transistor, two-resistor (4T2R) cells that commodity SRAMs use, which need not support multisourced read and write accesses (Figure 2).

Motorola has harnessed 4T2R cells as the foundation of its dual-port product line. The 1- and 4-Mbit separate-I/O NetRAMs offer distinct read and write data buses that share a common address bus. Motorola promotes its 1-Mbit dual-I/O NetRAMs as primary alternatives to true multiport RAMs; 4-Mbit versions will follow. NetRAMs lack asynchronous ports; both ports must run at the same speed, or, at a minimum, you must derive them from the same clock. Also, the parts are slower than Cypress’ and IDT’s fastest

STANDARDIZATION: A MIXED BLESSING

At a recent Joint Electron Device Engineering Council (www.jedec.com) meeting, NetLogic Microsystems kicked off an industry-standardization group for content-addressable memories (CAMs). As competitive and fast-moving as the CAM industry is, I’m skeptical about whether this effort will—or even should—succeed. Granted, standardization not only increases the number of compatible product sources, but also drives down prices.

Unfortunately, the slow com-

mittee-driven standardization process opposes the frenetic pace of the networking business, and standardization implies no application- or customer-specific features. Several CAM vendors acknowledge that a significant percentage of their business comes from custom devices for specific clients.

For the vendors, standardization is a mixed blessing. On one hand, none of them, particularly the market leaders, are enthused about the concept of less proprietary devices for which price is

increasingly the sole differentiator. On the other hand, they would like to expand the application and customer base for their CAMs to include other silicon-search-engine functions, such as cache translation, look-aside buffers, data compression and encryption, audio- and image-pattern recognition, and database mining.

The CAM suppliers admit that the complex, proprietary nature of their devices makes the devices intimidating to understand and to implement and

somewhat limits their success. Application expansion will occur only if the new applications can employ devices with features that have been fine-tuned over many years of networking usage. Multiport-RAM manufacturers also target applications beyond the network, such as redundant-array-of-inexpensive-disk controllers. Official standardization and lower prices, along with other enhancements, might help these applications to emerge.

devices. But if none of these criteria are important to you, you should consider NetRAMs for their higher density and lower cost potential.

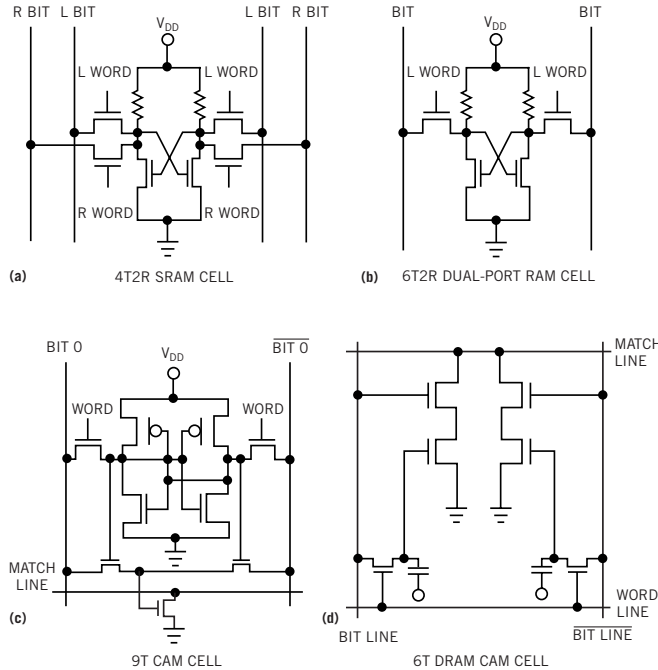
To the outside world, a NetRAM looks

like a standard synchronous dual-port RAM. All of its addresses, data, and control signals are latched on the rising clock edge. Internally, however, the memories read or write one port's information on

the rising clock edge and the other port's information on the falling clock edge. Motorola reports that it hopes, by early next year, to offer devices for sampling that eliminate the port-to-port clock dependency, but the internal data phase-shifting that enables the simpler cell structure may always keep NetRAMs at a speed grade or two lower than an 8T or 6T2R, dual-port equivalent.

IDT has its own lower cost product line, a bank-switchable dual-port memory (Figure 3). In addition to simplifying the cell structure, IDT eliminates the port-contention logic that a true multi-port contains. (This logic handles cases in which both ports simultaneously attempt to read and write or write the same array location.) Instead, these 64k×16- and 32k×16-bit asynchronous devices in 5 and 3.3V versions subdivide into four banks. One of the port's controls manages each bank. Bank management takes

A standard four-transistor, two-resistor or six-transistor SRAM cell (a) is smaller and less complex than cells that dual-port RAMs (b) and CAMs (c) use; a DRAM-based CAM (d) provides another potential way to reduce cost (a through c courtesy Motorola; d courtesy Mosaid).



MEMORY, LOGIC, OR A LITTLE OF BOTH?

Over time, suppliers have added increasing amounts of logic intelligence to their content-addressable memories (CAMs)—either as separate chips or integrated alongside the memory array. Music Semiconductor has led the charge, reflecting its self-described transition from “the CAM company” to “the packet-accelerator company.” Music supplies companion chips for its LANCAMs, including the Token Ring interface, the Fibre Distributed Data Interface, and both 10- and 10/100-Mbit Ethernet filters.

Music also offers several CAM-based-routing coprocessor families. The MUAA has a 32-bit synchronous I/O bus and contains 2048 80-bit entries (with 4096- and 8192-entry versions in the works). With aged and learned entry queues and auto-

matic aging, MUAA supports 48 ports of 100-Mbit Ethernet, four ports of Gigabit Ethernet, or 12.5 million Layer 4 look-ups/sec. You can program the CAM-versus-associated-RAM partitioning ratio from 32-to-48 to 80-to-0.

The MUAC coprocessor also includes a 32-bit synchronous I/O bus, but the internal data width is 64 bits. The company currently offers 4096- and 8192-entry versions. MUAC supports as many as 42 ports of 100-Mbit Ethernet, 78 ports of OC-3c, 19 ports of OC-12, or 28.5 million packets/sec of Internet Protocol Version 4 (IPV4) Layer 3 classless interdomain routing (CIDR). The device contains 32-bit ternary compare instructions, and the index that the match function generates can access an external RAM containing port-mapping and other associated data.

Music claims that its Epoch is a multimedia-ready, quality-of-service-aware, integrated on-chip switch that processes layers 3 and 4 of the IP stack. Epoch delivers as many as 1.4 million packets, or “flow classifications,” per sec, supports as many as 16 ports, and handles IPV4 (including multicast) and Internetwork Protocol Exchange in its hardware. Target applications include digital-subscriber-line-access multiplexers, remote-access servers, work-group switches and routers, WAN access-switch and edge routers, and LAN private-branch-exchange core switches. Epoch interfaces control not only the processor but also a variety of external memories, including, not surprisingly, MUAC CAMs.

NetLogic Microsystems is also heading down the value-added-

logic-plus-memory path. The company organizes its CIDR processor as 32,768 40-bit entries. It sustains look-ups at a 50- or 66 million-packet/sec rate with table updates at clock speed and a one-clock latency on the optional match flag.

Unlike a hashing algorithm, the CIDR processor internally manages its table, requiring no processor-controlled entry resorting.

Input buses include a 32-bit destination address, a 5-bit prefix length, 8 bits of associated data, and a 12-bit instruction. Outputs include a 16-bit longest-prefix-match address, a 5-bit prefix length, 8 bits of associated data, a 3-bit status-flag set, an exact-match flag, and 15 bits of multi-chip cascade control. NetLogic hopes to make the CIDR processor available for sampling in October.

the form of either software semaphores or hardware bank-select signals, and interrupts and mailboxes support interport communication.

For lower densities, you should also consider integrating multiport RAM onto your ASIC or PLD (see sidebar “Embed it”). Higher performance, lower power consumption, reduced board

space, and eliminated packaging expenses are all potential benefits of this approach. Unlike DRAM, SRAMs are process-compatible with standard logic, so you don’t have to grapple with a short list of foundry sources or trade off logic switching speed. Your actual cost savings and performance, however, depend on how the logic vendor imple-

ments its on-chip RAM.

If multiport RAMs are the on- and off-ramps of the network, the CAM is its traffic cop and navigator. Think of the CAM as a search engine—a sort of “silicon Yahoo.” Instead of your giving the CAM an address and the CAM’s returning data to you, you provide the CAM with the data that you are searching for, and it returns

EMBED IT

At some point, integrating memory on ASICs and PLDs begins to make sense. At *what* point that occurs, however, is a tough question to answer. What benefit will the application derive from the integration? How much do alternative stand-alone memories cost? Is the memory technology compatible with a standard logic process? If not, what trade-offs do you need to make? How much memory density can you squeeze onto a piece of silicon?


Shortcomings to the contrary, plenty of chips embed standard SRAM, ROM, EPROM, EEPROM, flash memory, FIFO buffers, and even DRAM. Embedded multiport-RAM and content-addressable-memory (CAM) arrays are also now emerging, and they will invariably become more popular. Multiports on cell-based ASICs are the most widespread example. Lucent Technologies’ memory compiler generates cores with any combination of as many as five read and write ports. Cell complexity grows as the number of ports increases, because, unlike Motorola’s NetRAMs, these devices are true multiport structures. Simultaneous-port-access collision detection and avoidance are the responsibility of external logic; the core lacks this support. Chip Express (www.chipexpress.com) also offers embedded RAM blocks on its gate-array products.

All FPGA vendors offer embedded dual-port RAMs for at least some portion of their product lines. Depending on how the vendor implements its on-chip arrays, however, the

amount of memory you can embed and the speed at which it runs vary (Reference A). Are the arrays small, distributed, and look-up-table-based? Are they, therefore, highly flexible but slow when you chain them together to form high densities? Or are they large and discrete with higher performance but less adaptability? Does the FPGA choose an approach midway between these extremes, or does it combine both extremes? Does the memory core support dual address and data ports, running at distinct clock frequencies, or does the vendor emulate this capability by doubling the memory array and adding extra logic control (analogous to the technique Motorola uses in its NetRAMs)?

For a long time, Lattice Semiconductor (www.latticesemi.com) with its ispLSI6192DM was the only CPLD manufacturer to offer on-chip dual-port RAM. Cypress Semiconductor now joins the party with its Delta39K family. Each 128-macrocell logic block contains two 8-kbit single-port SRAM arrays. An additional 4-kbit dual-port or FIFO block resides outside each logic block and directly connects to the hierarchical routing channels. The vendor’s 39K100, for example, contains 240 kbits of memory. The memory block includes all required dual-port logic, including independently running port clocks and limited arbitration support, thereby not using available general-purpose PLD resources. Cypress also adds

Figure A



You can expand the width and depth of Altera’s 32×32-bit CAM core, which supports both encoded (a) and unencoded (b) match-output schemes.

multiblock cascading support and dedicated flag-generation logic when you configure the memory block as a FIFO buffer.

What about CAMs? Several ASIC suppliers, including IBM (www.ibm.com), Kawasaki LSI, LSI Logic (www.lsilogic.com), and Lucent Technologies, offer flexible CAM cores. Don’t be surprised if a few of today’s discrete CAM suppliers in the near future add core capability to their product portfolios. Lucent’s compiler can automatically generate 2×1- to 1204×72-bit CAMs. Each CAM entry also includes a valid bit. With the match-register and invalidate function, you can clear any portion of the entries in one cycle, including the entire array. Lucent notes that matching is a power-intensive operation that becomes critical when you integrate the CAM array alongside other fast-toggling logic. Lucent’s 10-transistor binary cells minimize standby power consumption. Kawasaki LSI claims to offer a low-power seek approach that uses NAND-based match lines

instead of the more common NOR variant.

Finally, consider CAM integration on PLDs. Currently, the only example of this approach is Altera’s Apex 20KE family, although more alternatives will surely follow. The lowest level Apex ternary-CAM building block is a 4-nsec-access 32×32-bit array that Altera constructs from an embedded array block. The CAM supports both encoded and unencoded outputs, as well as limited forms of both depth and unencoded width expansion that the Quartus megafunction design tool automatically generates (Figure A). Array initialization is a two-cycle operation involving first true and then complement forms of the bit sequence, and a third optional write cycle handles don’t-care bits.

REFERENCE

A. Dipert, Brian, “Counting on gate counts? Don’t count on it,” *EDN*, Aug 3, 1998, pg 52.

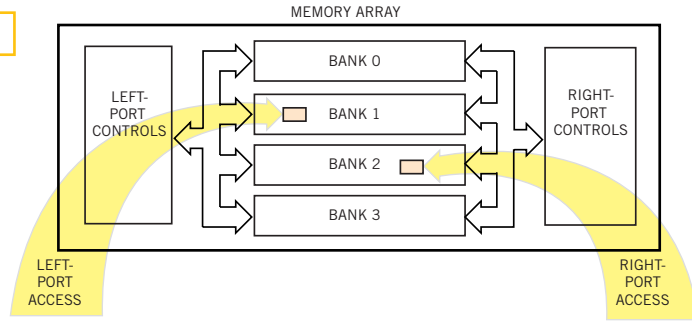
an address and other information related to the first match it finds.

The earliest CAMs were basic in their operation. These CAMs began with the early 1990s' LANCAM family from today's CAM-industry leader, Music Semiconductor. Data to match (the comparand), comparison instructions, and match results all traveled across a single bus with separate pins for control and match-flag functions. CAM array densities extended only to tens of kilobits, and only one global mask register existed to disregard specific bit values during match operations. These devices' close-to-100-nsec seek times were adequate for the data rates and network loads of the time.

CAMs found their first widespread usage in Layer 2 routers. Before 1993, only three classes of IP (Internet Protocol)-address allocation existed (Figure 4), so the number of mask registers could be small and still deliver fast searches with minimal mask-change overhead. As network administrators added active nodes, the router's software updated the CAM table with the IP addresses and corresponding ports. When a packet came through with an appropriate destination address, the CAM output the port value.

Times have changed. Routers now handle a number of Layer 3 and 4 functions, such as network-address-translation (NAT), firewall, URL, applet blocking, and other security protocols. They

Figure 3



Bank-switchable dual-port memories have more features but a lower cost than true dual-port alternatives (courtesy Integrated Device Technology).

also handle protocol translation; quality of service; and additional data types, such as telephony and streaming audio and video. These added responsibilities mean that the CAM array contains more types of data, which increases both the array size and the number of required mask registers. They also mean that seek performance must accelerate, because routing algorithms access the CAM several times for each packet.

Binary-CAM vendors are now offering samples of chips with densities of 1 Mbit and beyond and fast seek times even with 66-MHz clock rates. Although performance, power, cost, and board-space advantages favor a single-chip device, minimal- or no-glue depth expansion supports multichip modules.

CAMs also incorporate varying de-

grees of automatic array updating. Previously, to add data to a CAM, you first had to find an unused location within the array (stored data can also consist of valid and aging bits) and then manually store the information in that location. Some newer CAM architectures automatically store data when you issue the Learn command; they also occasionally support an Unlearn command. This ability eliminates the time-consuming free-location search. A few CAMs even automatically output the data's location in one cycle as part of the Learn command.

If the CAM array is full, does the CAM automatically discard old data and replace it with the new information? This cache-controller-like capability might be valuable in some cases, but it is expensive to implement. Therefore, CAM manu-

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, circle the appropriate numbers on the Information Retrieval Service card or use EDN's InfoAccess service. When you contact any of the following manufacturers directly, please let them know you read about their products in EDN.

Altera Corp
1-408-544-7000
www.altera.com
Circle No. 301

Kawasaki LSI USA
1-408-570-0555
www.klsi.com
Circle No. 304

Mosaid Technologies Inc
1-613-599-9539
www.mosaid.com
Circle No. 307

NetLogic Microsystems Inc
1-650-961-6676
www.netlogicmicro.com
Circle No. 310

Cypress Semiconductor
1-408-943-2600
www.cypress.com
Circle No. 302

Lara Technology Inc
1-408-894-1821
www.laratech.com
Circle No. 305

Motorola Inc
1-800-521-6274
www.motorola.com/sps
Circle No. 308

UTMC Microelectronics Systems
1-719-594-8000
www.utmc.com
Circle No. 311

Integrated Device Technology Inc
1-408-727-6116
www.idt.com
Circle No. 303

Lucent Technologies
1-610-712-4331
www.lucent.com
Circle No. 306

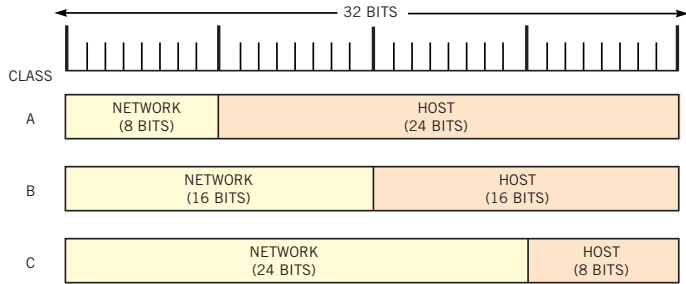
Music Semiconductor
1-908-979-1010
www.music-ic.com
Circle No. 309

SUPER CIRCLE NUMBER
For more information on the products available from all of the vendors listed in this box, circle one number on the reader service card. Circle No. 312

facturers do not currently offer it. Because the Learn command is costly both in the required on-chip logic and the necessary test time for each device during manufacturing, lower cost variants often eliminate it. CAMs alert system logic to full and almost-full (with user-configurable threshold) status via status registers, output flags, or both.

If you think a dual-port RAM's 8T or 6T2R cell structure is more complex and expensive than that of a standard SRAM, wait until you see a nine- or 10-transistor binary-CAM bit (Figure 2c). The CAM uses these extra transistors to implement the XNOR compare. XNOR outputs combine in a wired-OR fashion to generate the match bit, and some CAMs also support a multimatch output. CAM bits are expensive real estate, so they find limited use for supplemental functions. More commonly, CAM outputs become index addresses for a lower cost commodity SRAM or DRAM that stores this additional data. Comparand, instruction, result, and control buses are often distinct on newer CAMs; in lower

Figure 4



Rigidly defined IP-address allocations disappeared with the emergence of classless interdomain routing.

cost variants, their widths may be less than the internal data width. In those cases, multiple bus cycles read or write the entire data sequence.

THE MASK

Classless interdomain routing (CIDR) and NAT of nonroutable IP addresses for LANs have extended the life of 32-bit IP Version 4 (IPV4). By dissolving rigid class boundaries, CIDR more efficiently uses the available IPV4 address space. However, this approach increases the corre-

sponding seek burden on CAMs and control logic, because every address bit is potentially significant, and a few global mask registers are insufficient for the task. CIDR instead uses the “longest-prefix-match” technique.

Enter ternary CAMs. In these devices, each data bit contains not only a cell storing the 1- or 0-bit value but also a corresponding mask bit indicating valid or invalid status. This approach instantly halves the maximum number of stored addresses for a given-sized CAM array

and makes the seek circuitry more complex. But ternary CAMs have proved themselves equal to the task of overcoming these drawbacks: Thanks to 0.25- μ m manufacturing processes, vendors have announced 1-Mbit devices with sustained 66-Mbps, 64-bit search performance. Even more complex CAMs and logic to control them are on the way (see sidebar “Memory, logic, or a little of both?”).

Just as with multiport RAMs, lower cost CAM alternatives also exist. Motorola’s 256-kbit and 1-Mbit CAMs are standard SRAMs with on-chip logic that runs a hashing algorithm beginning at the array midpoint and successively zeroing in on the matching data. The device assumes that you order data from lowest to highest value in successive array locations, so each update requires a time-consuming reorder step.

Seek time using hashing or derived al-

gorithms is also nondeterministic because the time it takes to match data depends on the number of successive recursive-search steps the logic must execute until it determines whether a match exists. However, if your design needs moderate performance in which worst-case seek times approach 200 nsec and if table updates are infrequent, these devices may offer a cost-effective approach to implementing a wide, dense CAM function, such as in asynchronous-transfer-mode applications. Kawasaki LSI and UTMC Microelectronics also sell CAM longest-prefix-match hashing-logic engines that interface to discrete or embedded SRAM and DRAM arrays.

Even more radical cost-reduction options may soon appear. Embedded-DRAM pioneer Mosaid Technology hopes to begin sampling a DRAM-based ternary CAM by year’s end. Instead of using a 17-transistor SRAM ternary-CAM

cell, Mosaid will employ a 6T DRAM cell structure (**Figure 2d**). Mosaid claims that, after an initial pipeline latency akin to a standard DRAM random-access time, its DRAM-based CAM will sustain 66 million searches/sec.

Mosaid’s system can initiate CAM refresh when it isn’t otherwise accessing the CAM. Alternatively, the CAM can self-refresh and alert its status to the system via hardware flags. With respect to increasing bandwidth, Mosaid officials comment that DDR techniques, such as those that the company’s DDR synchronous DRAMs pioneered, are an obvious approach. Mosaid also plans to support single-cycle learn-plus-match and glueless depth expansion. □

REFERENCES

1. Dipert, Brian, “SRAMs strive to specialize,” *EDN*, Nov 5, 1998, pg 62.
2. Dipert, Brian, “Embedded memory: the all-purpose core,” *EDN*, March 13, 1998, pg 34.

