ADVANTAGES OF CAM IN ASIC-BASED NETWORK ADDRESS PROCESSING

INTRODUCTION

Network hardware and software designers are squeezing more functionality into less space, while simultaneously striving to improve performance. In many instances, ASICs and full custom silicon solutions are used in an attempt to achieve both higher density and higher performance. ASICs and full custom solutions coexist with MUSIC CAM solutions to provide the network hardware designer with a compact, high performance solution and a shorter design cycle. The network software designer benefits from the reduced processing requirements of the end system, the reduction of database management tasks, and the increased system scalability.

MUSIC Semiconductors provides a wide range of solutions, from discrete content-addressable memory (CAM) devices and application specific standard products (ASSPs), to fully integrated customer-specific solutions. MUSIC Semiconductors has a strong support infrastructure of Data Sheets, Application Briefs and Application Notes, Application Engineers, and a Sales Team consisting of Customer Service, Area Sales Managers, Field Application Engineers, Manufacturers Representatives, and Distributors. MUSIC also has toll-free telephone numbers for Technical Support and Sales. MUSIC's web site is updated regularly with the latest versions of data sheets and applications information. MUSIC Semiconductors contact information is given at the end of this document.

ASIC WITH INTERNAL SRAM OR CAM

An ASIC containing internal SRAM is not a low-cost approach to a search engine. Commodity (external) SRAM is low-cost because it is sold in large volume and requires less die area than CAM. Internal SRAM in a logic process consumes far more die area than it would in a commodity SRAM process. In the case of an ASIC, internal CAMs have a clear performance advantage over internal SRAM, but internal SRAM no longer has a size or price advantage over CAM. MUSIC Semiconductors offers ASIC solutions with internal CAM. The major design reasons for choosing this approach are board density and simplicity of design. MUSIC designers can integrate customer's logic with MUSIC CAM, providing a low-risk single chip solution. This document also explores the advantages of discrete CAM solutions.

LARGE TABLE ISSUES

With tables too large to be integrated completely into ASIC, the choice is between external CAM and external SRAM with an SRAM controller in the ASIC. From a cost standpoint, some CAM devices are, for a 2K table, about the same cost as an SRAM-based hash engine. The reason for this is that although SRAM is cheaper per bit, the smallest SRAM you can buy is about four times too big for a 2K hash table. As CAM price drops with volume, and SRAM price increases from the current low, the crossover point will be at a higher number of entries. In addition, the CAM-based search engine design is inherently scalable, which is not true of SRAM-based hash or binary search engines. The same design–board, ASIC, etc., could therefore accommodate various amounts of CAM, handling a few hundred entries to several thousand. Also, the CAM's cost per port, ease of use, and time to market advantages need to be considered.

Some designers assume that CAMs may be too expensive for their system without checking the price and considering the performance, hidden costs, and time-to-market issues of algorithm-based solutions. MUSIC shipped more than one million CAMs in 1996, many of them used for very cost sensitive applications.

ASIC USING EXTERNAL SRAM OR CAM

As a solution, many bridge and switch designers use an ASIC addressing external SRAM. However, an ASIC controlling a CAM not only gives better search performance, but better system performance as well, with a much simpler ASIC design. In fact, MUSIC ASSP devices eliminate the need for a CAM control ASIC, as the control as well as address parsing function is integral to the ASSP. MUSIC offers such devices for Token Ring, Ethernet, and FDDI.

Often, an entire incoming Ethernet frame is copied to system memory before the frame is parsed, and the source and destination addresses are extracted. This methodology is being replaced by "snooping" at the port; parsing and processing the source and destination addresses as they are received. This improves system performance by reducing system resource requirements, because the system does not have to store the frame. It is this "snooping," or filtering of addresses at the incoming port, that will be discussed in this document. Both the SRAM and CAM approaches are discussed.

CAMS IMPROVE SYSTEM PERFORMANCE

MUSIC CAMs are parallel processors used to store network addresses. Some of MUSIC CAMs have up to 107 variations of seven basic instructions available to aid the system in network address processing and database management. Processors are chosen for ease of use and flexibility in system design rather than cost, in contrast to how memory devices are chosen. CAM can be controlled with a CPLD, FPGA, or ASIC state machine, freeing the host processor for other system, routing, switching, or bridging tasks. This approach has the added benefit of reduced system software and the associated development and maintenance costs, and improved time-to-market of the end system. CAM search time is completely deterministic—a result is available in one search cycle every time, easing the integration on the hardware level.

The amount of CAM necessary is usually much less in a CAM lookup table than the amount of SRAM necessary in a RAM lookup table. Why? Searching the CAM table is faster and the table maintenance overhead is faster and easier. This means that only the current and recently active entries need to be stored. The prime reason large tables in SRAM-based systems are necessary is because it is more difficult and slower to add new entries in a RAM-based system. A large table is needed to ensure that few Destination Address entries are missed, thereby minimizing broadcasting, which would slow the system. MUSIC CAMs minimize learn time of new entries and simplify database management, therefore fewer entries are missed inherently, which means less CAM memory is needed. With less CAM needed, the cost of CAM becomes a less important factor. It is possible that the amount of CAM memory required can be as little as 10 to 20 percent of what would be necessary in a RAM-based system.

A key issue is that MUSIC CAMs provide faster address processing, not just faster address lookup. Although the individual basic lookups are faster, system performance improves due to the considerably reduced incidence of broadcasting necessary when an address lookup fails, which happens every time a new address comes on line or an old one goes off, but is addressed by another source.

WHY IS BROADCASTING REDUCED?

The worst-case broadcast performance of a CAM-based system is once per new address. This is not true of SRAM-based systems using hash or binary search methods. In a hash engine, the hash value is calculated first (which usually happens quickly). Then that memory location is checked to see if it matches the desired address. If it does not, another memory location is calculated and its contents compared. If the desired value is not found, the cycle repeats until found or a search limit has been reached. In either instance, the worst-case performance of a hash engine will broadcast several times before learning the new address.

The problem is worse for a binary search engine, especially if the table is very large. In addition to taking several searches to find the desired address worst-case, the list must be ordered. This poses a problem when learning, because the list needs to be resorted after learning each new address. Again, this process takes time, and the same new address will be broadcast several times before the system can learn and filter it.

For both the hash and binary search methods, address processing lookup tables in SRAM-based systems require other list maintenance routines. For example, a routine to update time stamps of active entries and purge old entries, which usually consumes a large amount of time, may potentially require many tens of microseconds or even milliseconds. The host processor is now tied up for this task, possibly keeping it from doing other routing, bridging, switching, or system functions.

CAMs are self-maintaining, because no sorting or indexing is required. New insertions virtually happen automatically with the “Move to Next Free address” instruction. Time stamp updates and Port ID updates are just as simple, with the “Move to Highest Match” instruction. Purging based on Time Stamp is easily accomplished by searching for the Time Stamp value to be purged, and marking all matching locations “Empty.” In a CAM-based system, Aging/ Purging routines to remove old entries are executed in a few hundred nanoseconds — less than the null packet time of most FAST protocols. All of this overhead can be accomplished with little or no host processor intervention.

WHY FASTER ADDRESS PROCESSING?

Usually the system wants to have address attributes (called Associated data) stored with, or accessible with the network address. These attributes may be history or aging bits, port identifiers, protocol identifiers, static/dynamic entry marker bits, etc. MUSIC CAMs store these in the RAM partition connected directly to the same address location in the CAM array, therefore accessible on the next instruction cycle for read or modify. Algorithm-based lookup approaches, such as RAM tables, generally have difficulty

giving fast access to this associated data. MUSIC's current CAMs provide fast access in 50 to 120 nanoseconds, depending on speed grade, type, and I/O width. Some devices provide access to associated data simultaneously. This can also be considered part of the list maintenance task that MUSIC devices simplify as well.

THE SCALABILITY ADVANTAGE

CAM-based designs are scalable due to cascadability of CAMs. Most algorithm-based systems are not. Once you choose a table size with an SRAM-based system, it tends to be fixed. With CAM, the designer can decide on a small amount of CAM, and plug in more only if required after system performance monitoring (RMON techniques may help this). Conversely, the designer can choose a larger amount of CAM initially and reduce it with a simple parts list deletion rather than a redesign, or offer different lookup table sizes dependent upon the end-user application by just populating or not populating CAM sockets.

PARTITIONABILITY IMPROVES DESIGN FLEXIBILITY

MUSIC CAMs are partitionable into CAM and RAM partitions, or Associative and Associated data fields. The intrinsic auto-lookup gives a consequent inherent translation capability not easily obtained with an algorithm based approach. This means that MUSIC CAM can be used in protocol translations and header processing applications. Ethernet addresses can be easily assigned to groups or VLANs, and vice-versa. With a little extra RAM, Ethernet addresses can be converted to IP addresses, and conversely, IP Addresses converted to Ethernet Addresses (ARP caches). The MU9C1965A, with its 128-bit word width, allows this conversion internal to the CAM. ATM VPI/VCI and UNI/NNI conversions can be made simply and easily, and in a known (fully deterministic) amount of time.

MULTIPLE PROTOCOL ADDRESSES IN A COMMON CAM DATABASE

By clever use of the validity bits, whole sections of the CAM database can be excluded from or included in a comparison operation, allowing one device or cascaded array of devices to be used for multiple functions within one design. For example, one device can be used for basic address filtering for two different networks or workgroups without conflict, or basic address filtering can coexist with header processing or VLAN work. The possibilities are endless.

INTERNAL MASK REGISTERS IMPROVE SYSTEM PERFORMANCE FOR IP

The mask registers in MUSIC CAMs and the ease of moving data to masks allow for easy hierarchical, group or subnet lookups – not an easy task for algorithm-based systems. These mask registers are also useful in list maintenance when purging or modifying list content. Only CAMs offer flexible, reconfigurable multiple mask registers.

LAYER 3 (IP) SWITCHING

Implementing Layer 3 Switching using CAMs is fairly simple, but far more complicated for SRAM-based search engines. The reason for this is the search time to find the longest matching address and list maintenance tasks are greatly reduced in the CAM-based approach. CAMs are well-suited to processing IP addresses because of their ability to selectively mask bits during compare operations and to resolve the multiple matches that may result. IP addresses can be written into the CAM in hierarchical order such that the match at the lowest level of IP hierarchy for any given IP address can be found during a single compare operation. Such an approach is particularly attractive for switches and routers that use the CIDR (Classless Internet Domain Routing) protocol.

For Layer 3 switching, both LAN and IP addresses can be written into the same CAM word. Compare operations can search for either address independently of the other by selectively masking bits during compares. Even if you are designing a Layer 2 switch, CAMs provide the flexibility to allow the design to be upgraded to a Layer 3 switch in the future.

SUMMARY

MUSIC CAMs are parallel processors, unlike SRAM devices. The arguments for using CAMs in a network system in the past, using discrete devices, are still valid in systems implementing much higher levels of integration. As networks grow, systems grow to handle the traffic. To compound the network switch, bridge, or router's problem, protocols are getting faster. The CAM technology provides a solution to both problems economically and efficiently, regardless of whether the discrete devices are used, or a full custom or ASIC solution integrating CAM is provided.

MUSIC SEMICONDUCTORS STANDARD PRODUCT OFFERINGS

16-bit Interface, 64-bits Wide

Device P/N	Density	Organization	Supply	Speed Grades
MU9C3480A	16Kb	256x64	5V	50,70,90,120ns
MU9C3480L	16Kb	256x64	3.3V	70, 90,120ns
MU9C5480A	32Kb	512x64	5V	50,70,90,120ns
MU9C5480L	32Kb	512x64	3.3V	70, 90,120ns
MU9C1480A	64Kb	1024x64	5V	50,70,90,120ns
MU9C1480L	64Kb	1024x64	3.3V	70, 90,120ns
MU9C2480A	128Kb	2048x64	5V	50,70,90,120ns
MU9C2480L	128Kb	2048x64	3.3V	70, 90,120ns
MU9C4480A	256Kb	4096x64	5V	50,70,90,120ns
MU9C4480L	256Kb	4096x64	3.3V	70, 90,120ns

32-bit Interface, 32-bits wide

Device P/N	Density	Organization	Supply	Speed Grades
MU9C4320L	128Kb	4096x32	3.3V	90,120ns

32-bit Interface, 64-bits Wide

Device P/N	Density	Organization	Supply	Speed Grades
MU9C1485A	64Kb	1024x64	5V	50, 70, 90, 120ns
MU9C1485L	64Kb	1024x64	3.3V	70, 90, 120ns
MU9C2485A	128Kb	2048x64	5V	50, 70, 90, 120ns
MU9C2485L	128Kb	2048x64	3.3V	70, 90, 120ns
MU9C4485A	256Kb	4096x64	5V	50, 70, 90, 120ns
MU9C4485L	256Kb	4096x64	3.3V	70, 90, 120ns

32-bit Interface, 128-bits Wide

Device P/N	Density	Organization	Supply	Speed Grades
MU9C1965A	128Kb	1024x128	5V	50, 70, 90, 120ns
MU9C1965L	128Kb	1024x128	3.3V	70, 90, 120ns

32/16-bit Interface, 80-bits Wide

Device P/N	Density	Organization	Supply	Clock Speeds
MUAA2K80	160Kb	2048x80	3.3V	20ns
MUAA4K80	320Kb	4096x80	3.3V	20ns
MUAA8K80	640Kb	8192x80	3.3V	20ns

Consult factory for current device and speed grade availability

MUSIC Semiconductors Agent or Distributor:

MUSIC SEMICONDUCTORS

http://www.music-ic.com
email: info@music-ic.com

USA Headquarters
MUSIC Semiconductors
254 B Mountain Avenue
Hackettstown, New Jersey 07840
USA
Tel: 908/979-1010
Fax: 908/979-1035
USA Only: 800/933-1550 Tech. Support
888/226-6874 Product Info.

Asian Headquarters
MUSIC Semiconductors
Special Export Processing Zone 1
Carmelray Industrial Park
Canlubang, Calamba, Laguna
Philippines
Tel: +63 49 549 1480
Fax: +63 49 549 1023
Sales Tel/Fax: +632 723 62 15

European Headquarters
MUSIC Semiconductors
Torenstraat 28
6471 JX Eygelshoven
Netherlands
Tel: +31 45 5462177
Fax: +31 45 5463663