



SigmaRAM™ Targets High Speed Networking Applications

Introduction

SigmaRAM is a family of SRAM products jointly defined by the [SigmaRAM Consortium](#). The consortium is a collaborative group of leading SRAM manufacturers who have joined together in establishing the industry's first open standard for networking SRAMs. Consisting of [Alliance Semiconductor](#), [GSI Technology](#), [Integrated Silicon Solutions Inc.](#), [Mitsubishi Electric Corporation](#), [Sony Electronics Inc.](#) and [Toshiba Corporation](#), the organization was formed with the goal of creating an open, [JEDEC](#)-standard memory product family targeting the networking and telecommunication markets

History of Networking Requirements

Telecom applications have become the backbone of modern civilization. Today, we depend on modern packet switched networks to convey a variety of data from voice to video. With the advent of low latency packet switching technology, what used to be different data types are being merged into one form—digital data—and transmitted over a single medium. The combination of these requirements has exposed an ever increasing need for more bandwidth. Higher bandwidth supports time-sensitive transmission of voice and video, and allows for more channels over a single physical medium.

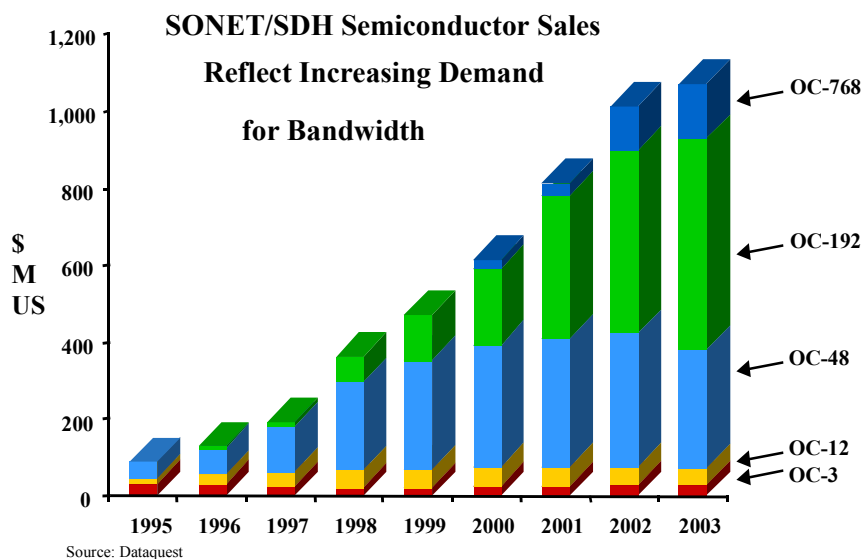


Figure 1: SONET/SDH Semiconductor Sales

Increasing the bandwidth of data communications typically requires an increase in the complexity of network components. Over the past 10 years we have seen multiple revolutions in communications technology and component complexity. Simple LANs to enterprise class switching were originally implemented using single bus, single processor, stored program control architectures. Packet speeds were slow enough that off-the-shelf processors were more than adequate to handle the load. While all such systems require some form of memory storage, the storage elements used in every day PC and server applications were sufficient.

Early communication systems greatly leveraged the use of memory elements designed for computer applications. Already widely in use and relatively inexpensive, these memories were readily adapted to new uses. Today, access products such as edge routers, DSLAMs, and cable boxes operate at speeds much higher than early circuits. It is not uncommon, for example, that an edge router will service a multiple collection of Gigabit Ethernet and OC3 ATM ports, as well as legacy T1 and T3 connections. With enhanced service expectations, multiple protocol translations, and new QoS requirements, the burden of switching and routing elements has greatly increased. Old architectures have given way to a new breed of packet processors.

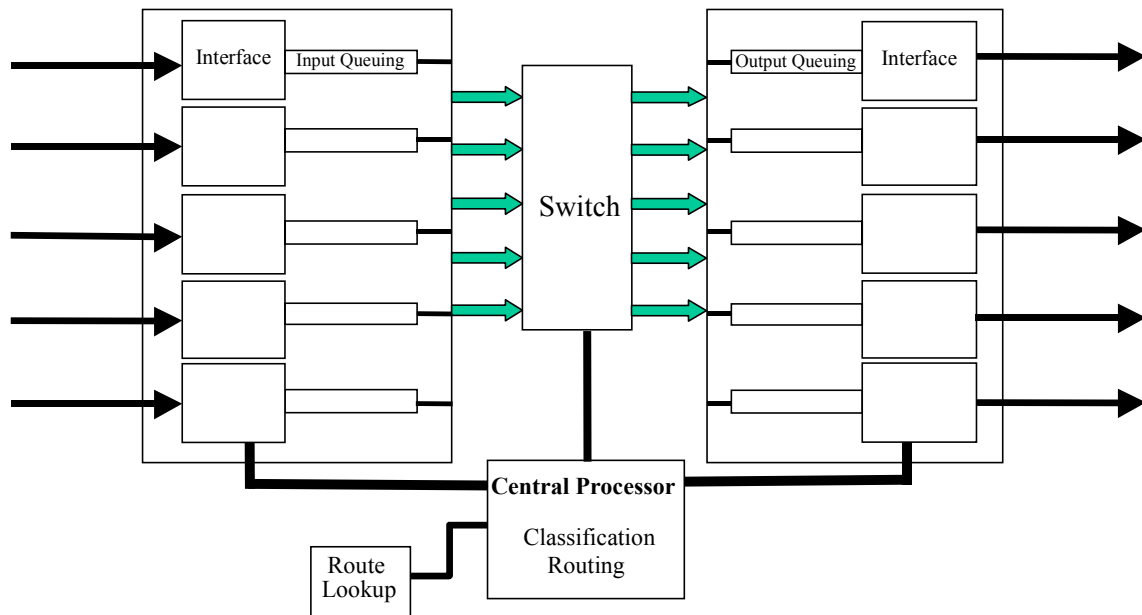


Figure 2: Old Router Architecture

2nd Generation Packet Processing

In the second generation of packet processing, functions were segmented into separate chips. ASICs were developed to optimize each function. Chipsets worked together to service packet processing functions, such as framing, classification, encryption, queuing, load balancing, address translation and routing. While each of these functions has its own special needs, they all share a common denominator in that they require multiple memory

accesses. In an ATM switch, for example, there is a need to translate Virtual Packet Interface/Virtual Channel Interface (VPI/VCI) addresses. Another memory intensive function is classification. Classification encompasses the basic function of identifying packets, headers and fields and comparing them to a fixed database. Some typical classification functions are address forwarding and routing. High-speed classification requires read-modify-write operations. In this case, the memory must be able to change from a read state to a write state in a single cycle, without adding wait states to the processor. Legacy memory components, such as synchronous burst SRAMs, simply cannot do this. However, until now they have been used in this application because nothing better existed.

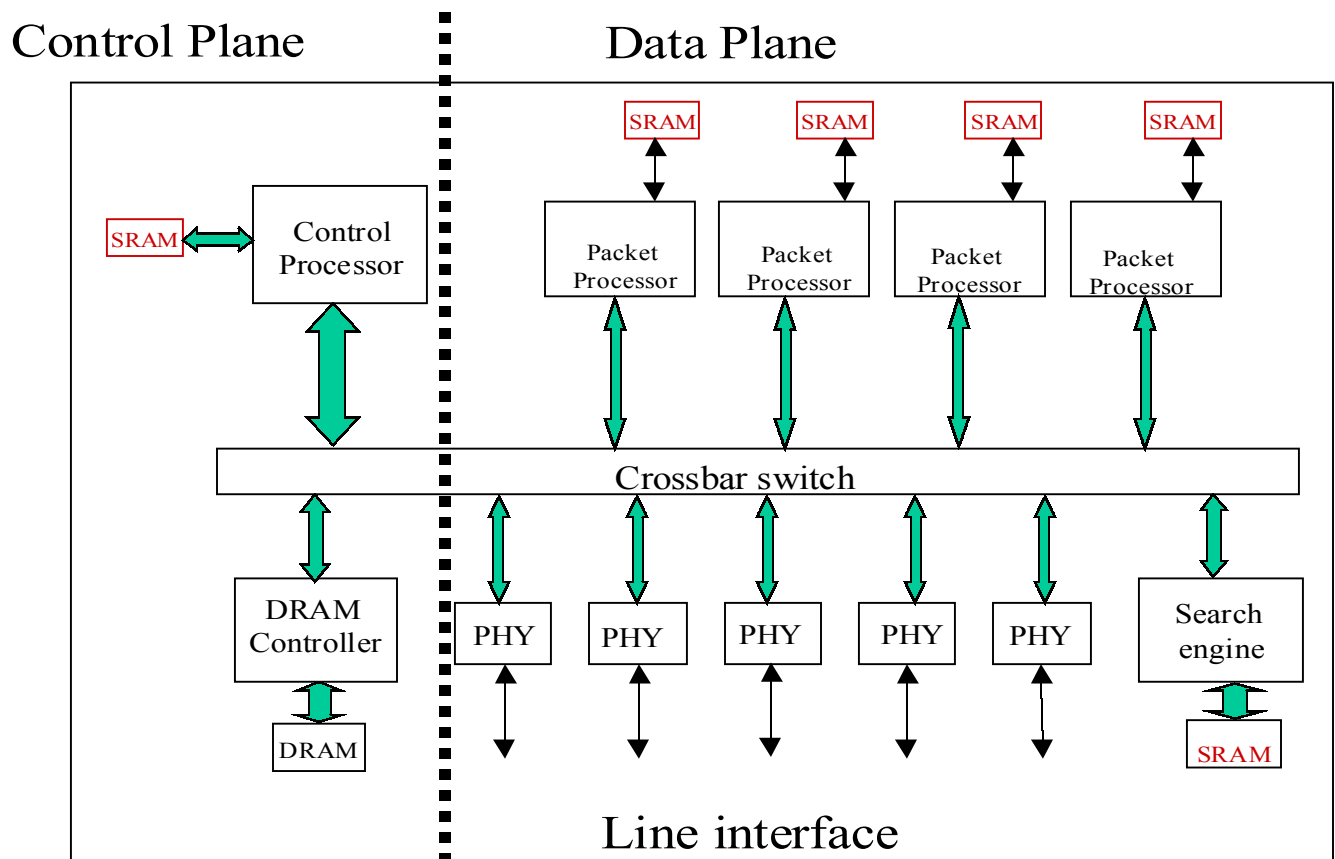


Figure 3: New Router Architecture

New System Requirements

Third generation systems require the raw throughput of ASIC designs coupled with flexibility in delivering timely market solutions. This has led to the advent of the network processor. As a specialized device, a network processor is optimized to perform programmable packet processing functions. New requirements that drive the 3rd generation systems have also spawned a new generation of memory devices. Unlike the memory devices of previous generations, which were designed for computer (i.e., cache) applications, these new devices were developed specifically for networking applications.

Distributed architectures place a premium on fast, low latency memory systems able to turn around quickly from reading to writing and vice versa. Indeed, one of the salient requirements, next to raw speed, is the ability to avoid the bus turnaround penalty inherent in all synchronous memories designed for cache. Conversely, the fundamental requirement in a cache memory is to be able to read data very quickly. Burst protocols were developed specifically to avoid the requirement of reasserting the address before each read cycle. Because a typical cache read restored several “lines” of data, the burst protocol was an optimal way to replace multiple lines of memory. Write operations were hidden in the overhead of a cache miss.¹

Memory access in third generation systems typically fall into three categories: packet buffering, table storage, and fast lookup. The buffer function occurs at the front end of processing as packets are received. If the rate of data streaming exceeds packet processing, speed buffering may be required. Multiple packets can also be placed into temporary storage to concatenate payloads. This is especially true when processing small cells, as in ATM. Not all buffering is sequential. In IP transport, each packet can travel over separate paths. This means that packets sent in order can be received out of order. To reorder packets requires random retrieval from memory. A common bus memory architecture with no read/write penalty is ideal for this application.

Once the data stream is parsed, each individual frame is processed. In many cases, this will require configuration parameters from a table. These memory requests are usually a single random access. The number of reads will greatly exceed the number of writes. This favors a common bus memory architecture². Because processing occurs at wire speed, the memory access must be fast enough to keep up. Burst protocols are not useful because, while a burst increases data bandwidth, it tends to decrease address bandwidth.³ Therefore, a fast Single Data Rate (SDR) device is preferred over slower Double Data Rate (DDR), as DDR devices all operate in Burst mode only.

Table lookup is the third memory access category. In processing layer 3 (addressing) the processor identifies the packet destination and chooses a route path based on parameters determined by the routing protocol. Memory accesses for address lookup are in the form of random search. CAMs are ideally suited for this type of search. However, CAM searches are typically slow.

New SRAMs Meet The Challenge

Today’s high performance networking systems require a new kind of memory support. Very fast access time, high cycle rates, no bus turnaround penalty, low power, and high density. Meeting the demands of 3rd generation systems requires memory to be able to

¹ This is assuming a copy-back protocol, widely considered as the highest performing way to configure a cache subsystem.

² A common bus memory architecture executing the NBT (No Bus Turnaround) protocol can execute multiple reads or writes without penalty. Separate bus architectures are much less efficient at these operations.

³ Address bandwidth is the rate at which random addresses can be asserted on a system or device.

execute fast, random, multiple reads, change from reads to writes in one clock cycle, and sustain both data and address bandwidth exceeding wire speed. This describes the SigmaRAM family of synchronous SRAMs. Designed specifically for networking applications, the family of devices features higher speed, new packaging, an improved clocking scheme, multimode operation, and low power.

Clock speed is one of the key features of the SigmaRAM family. At a clock rate of 333 MHz, a 72-bit output can deliver up to 24Gb/s throughput. Compare this with 5–7Gb/s, the typical output of SRAMs designed for burst-cache operation⁴. The ability to change from reads to writes in one clock cycle is another key feature supported in the SigmaRAM family. This is a necessity when performing fast lookup at OC192 packet speeds.

SigmaRAM packaging is also designed for networking. While legacy burst RAMs use advanced BGA packaging, the pinout favors the use of multilayer circuit boards requiring buried vias to reach pads deep in the array. SigmaRAMs were designed to minimize this problem by placing all signal pads within the first three outer rows. This maximizes the escape space and eases routing hassles. In addition, the omission of buried vias allows networking boards to utilize less expensive technology.

While legacy SRAM architecture features a common read/write bus, SigmaRAMs include a separate bus option. No single architecture is right for all situations. The separate bus option specifically optimizes back-to-back read-writes. The reason for this is that in the separate bus architecture, reads are implemented on the first clock cycle and writes on the second, alternating in this fashion. The architecture can implement multiple reads or writes, but in that case each alternate cycle is wasted. A separate bus architecture also eliminates any possibility of bus contention⁵, a potential issue in any common bus architecture. When operating at the high data rates SigmaRAMs offer, a special output clock is available allowing system designers to easily synchronize the data. Using a 3 ns cycle time in DDR mode, there is a new data output approximately every 1.5 ns. Output clocks⁶ fine tune the registration of the data valid window.

The SigmaRAM Family

The SigmaRAM family consists of devices featuring both common and separate I/O busses. Each architecture optimizes a specific sequence of read and write commands. In addition, the family includes SDR and DDR operation and supports Late Write (LW) and Double Late Write (DLW) protocols in various combinations. The combination of wide bus widths, from x18 to x72, and high clock rate (up to 333 MHz) gives the SigmaRAM family industry-leading performance. Product availability by company is shown in the

⁴ Cache SRAMs are typically narrower and slower. Metrics assume a x36 burst SRAM at 200 MHz with or without Burst mode.

⁵ Bus contention occurs when the time for one device to release the bus is greater than the time for another to assert the bus. When both devices assert control simultaneously, power is wasted which, if excessive, could affect operation.

⁶ Echo clocks are generated coincident with the data to help register data downstream.

consortium [roadmap](#). Detailed information about SigmaRAM can be found in the [product presentation](#).

Conclusion

Networking components and requirements have followed Moore's Law. The explosion in communication technology, service and bandwidth requirements continue to drive component manufacturers accordingly. Network elements now routinely operating at OC192 rates and above require higher memory throughput. SigmaRAMs combine legacy features designed for computer cache applications with new features targeting networking applications. SigmaRAMs represent a new paradigm in SRAM technology designed to meet these requirements.

