Order this document
by AN1299/D

AN1299

ATM Switch with Shared Memory – A Simple Model

Prepared by: Lakshmi Mandyam

INTRODUCTION

Asynchronous Transfer Mode (ATM) Telecommunications are the mainstay of today's communication systems. ATM transmits data across large networks using high speed transfer of small data payloads to deliver high throughput. The data payloads in ATM are called cells and consist of 53 bytes (the Ethernet data payload can be anywhere from 64 bytes to 1500 bytes and up to 8000 bytes in SONET). The 53 bytes consist of 5 bytes of header information and 48 bytes of data. The header information instructs the switch where to route the data. The various elements of the communication system, such as the routers and switches, should have throughputs of the order of Gigabits per second.

Think of the data cell being transferred as a passenger traveling by air, and think of the ATM switch as an airport. See Figure 1. Let us consider Joe Bit traveling from Austin, Texas to San Jose, California by Air Moto Flight # BY36. He checks in at the ticketing counter and presents his ticket. The airline representative looks at the ticket and determines which flight he is taking and that his flight will be leaving from gate 32K. There are many lounge areas in the airport, but the representative suggests waiting in the closest lounge to the gate which is called the NetRAM #3 lounge. See Figure 2.

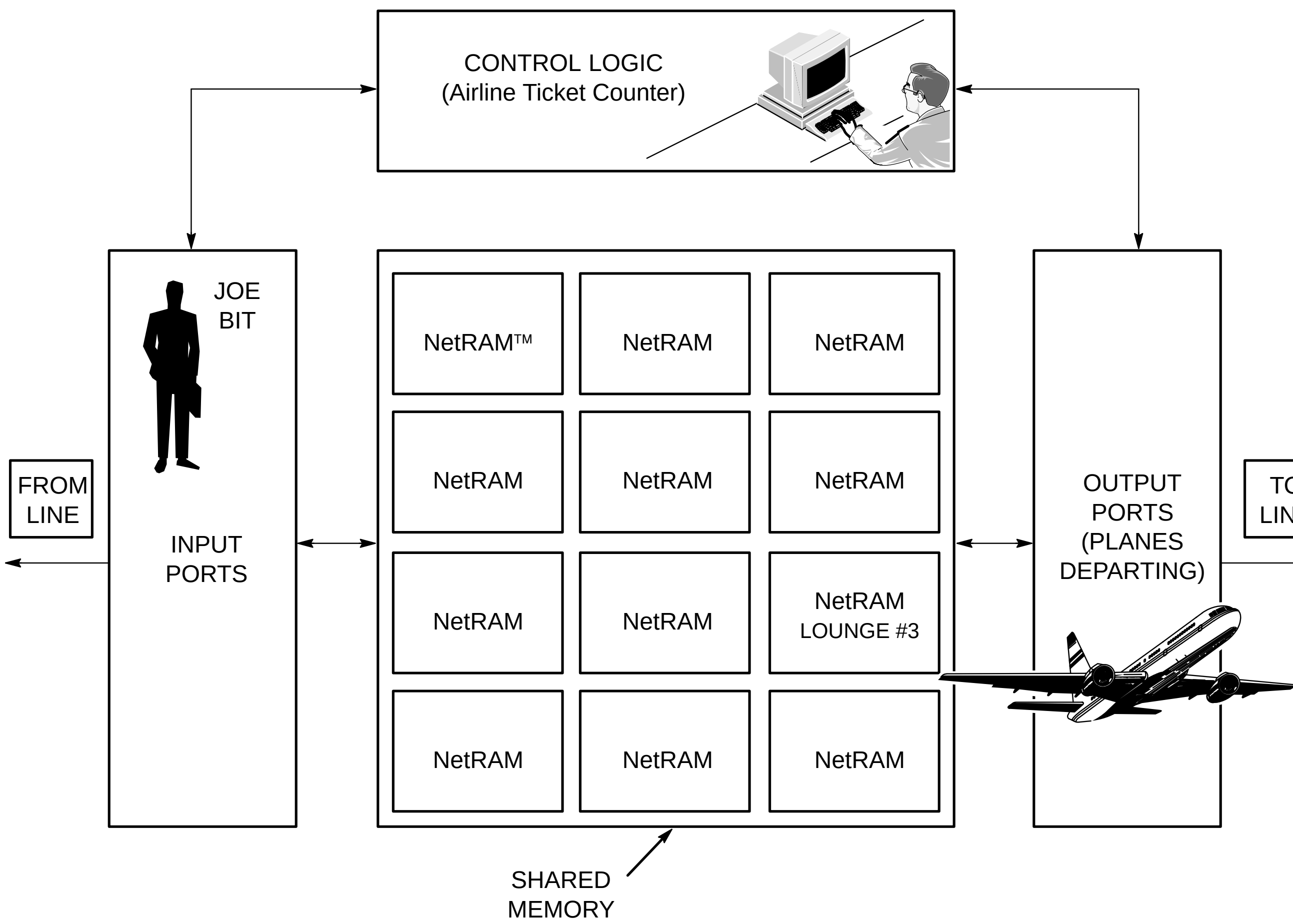

Figure 1. ATM Switch with Airport Analogy

NetRAM is a trademark of Freescale Semiconductor, Inc.

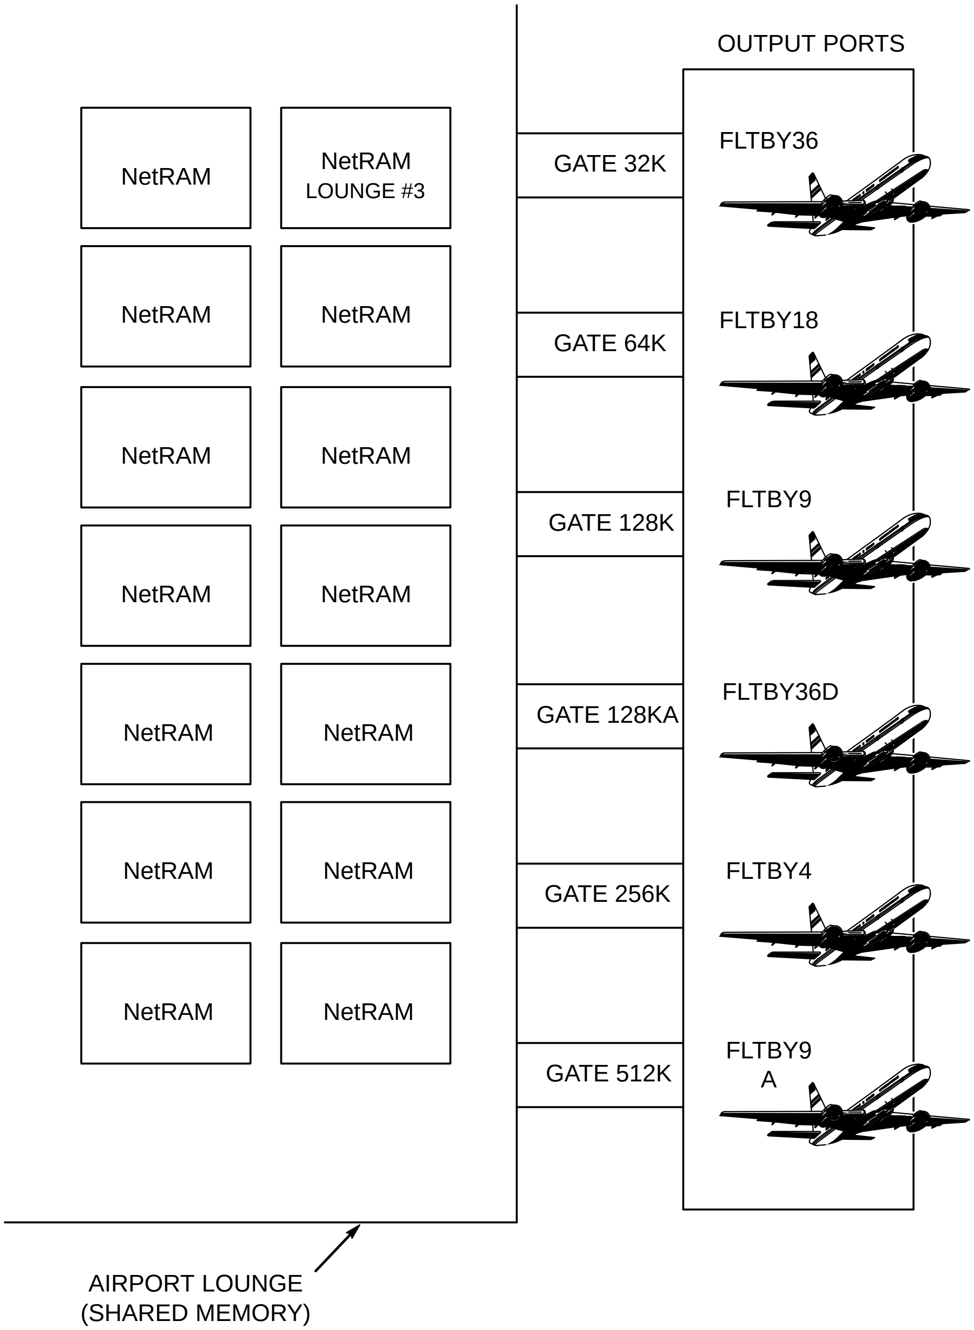

Figure 2. Shared Memory and Output Ports (Airport Analogy)

Joe Bit gets directions to NetRAM #3 lounge and proceeds there to wait for his flight. Similarly, in an ATM scenario, Joe Bit is the data payload and his ticket is the header information. The airline representative at the ticket counter is the control logic of the switch. Gate 32K is the output port and the NetRAM #3 lounge is the memory where the data (Joe Bit) will wait to be read to an output port. In addition, the directions to NetRAM #3 lounge are the write address for the memory location. When the plane is ready to be boarded, the representatives at the gate will call for passengers to board the plane and, at that time, Joe Bit will board the flight. The call for boarding can be considered to be a read address for the memory. Just as in airports where waiting passengers are accommodated in an airport lounge, cells can be buffered in a memory to wait to be read to an output port.

The advantage of having a large airport terminal with many airport lounges is that it can handle heavy passenger traffic efficiently and conveniently. Similarly, the advantage of an ATM switch with a large memory buffer is it can handle enormous amounts of data without loosing cells.

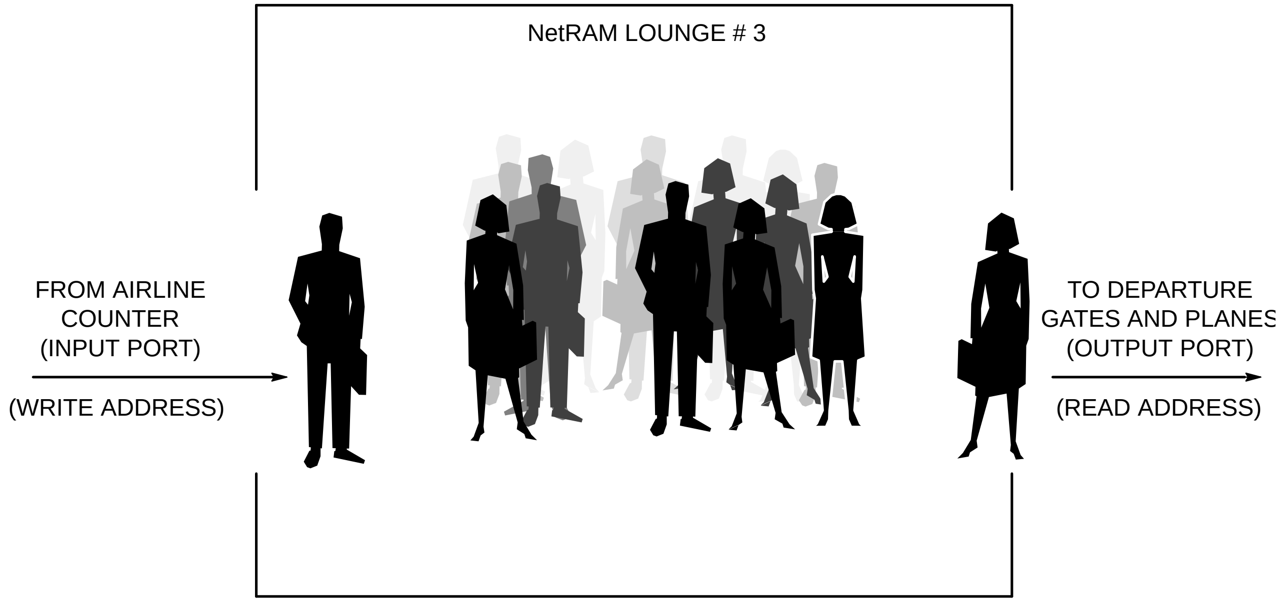

Figure 3. Airport Lounge (Separate I/O, Dual Address NetRAM)

Usually airport lounges have multiple entrances so people can enter the lounge without delay. People can arrive to wait in the lounge for their flight's boarding call and, at the same time, people who hear their boarding call can proceed to their departure gate. See Figure 3. Similarly, the NetRAM is a dual address device where you can perform reads and writes at the same time. Therefore, in an ATM scenario, data cells arriving from the transmission line to the input port can be written to memory, and at the same time, data cells can be read from memory and sent to the output ports.

In an airport lounge, people can enter and leave the gate simultaneously. The average number of people who go through the lounge is very high for a given time period. We can think of this as the throughput of the lounge. For the NetRAM, approximately 36 bits can be written and 36 bits can be read from the NetRAM every 16 ns in the 32K x 36 version. This means that a total of 72 bits can pass through the NetRAM every 16 ns or 4.5 Gigabits (Gbits) every second. The NetRAM then has a throughput of approximately 4.5 Gbits/sec. This is higher than a conventional single address SRAM. The throughput mentioned above is just for a single NetRAM. Typically switches have NetRAMs connected in parallel. For example, if four NetRAMs are connected in parallel, the number of bits jumps to 144 bits every 16 ns or 18 Gbits/second. If sixteen NetRams are connected in parallel (sixteen wide), the throughput jumps to 72 Gbits/second.

The dual address feature enables the user to enhance system performance by using the simultaneous read/write feature. For example, when the data comes into the switch it can be written to the memory while data is being read to the output ports. Therefore, the system control does not have to wait until the read to the output port is finished before writing the data from the input port. This greatly enhances the throughput of the switch.

The ATM platform has the ability to handle traffic of different priorities. Let us go back to the airport analogy to better understand the concepts. Jane Byte, who is traveling on the same flight as Joe Bit, has a first class ticket. She checks in ten minutes after Joe Bit and is also sent to NetRAM lounge #3. Usually the first boarding call is for first class passengers. Jane Byte will board the plane before Joe Bit, even though she arrived at the lounge after him, because she has a first class ticket which has a higher priority than Joe's economy ticket. Similarly in ATM, data with a higher priority will be output from the switch first. The pass–through feature of the NetRAM enables the user to handle high priority traffic by allowing the data to go directly from the input to the output

In an airport, planes of different passenger capacities can be handled. For example, because there is a higher demand for flights from Austin to Dallas, there are frequent flights between the two destinations so passengers can get to their final destination faster. However, there is not much demand for flights from Austin to Brownsville; therefore, there are fewer flights, so it may take longer to reach a destination. Similarly, in ATM, data payloads of different sizes can be handled, and in some cases, there are dedicated connections depending on how often a connection is used.

Imagine a large airport such as the Dallas/Fort Worth Airport. This airport has hundreds of departure gates with numerous paths between the gates such as airport monorails, walkways, and people movers. Imagine you are looking down at the airport with the roof removed. The connections are so numerous and complicated, they look like the threads in a woven fabric. Similarly, in a shared memory in an ATM switch, the connections between the different memories are so numerous that it looks like a woven fabric. This is what is called "Switch Fabric."

ATM is one of the state–of–the–art protocols used in telecommunications today. The key to this technology is speed of transmission and therefore the throughput of the switch. The NetRAMs system performance enhancing features are:

- Each NetRAM has a throughput of 4.5 Gbits/second. It can be used in parallel with other NetRAMs which can multiply the throughput. For example, sixteen NetRAMs together have a bandwidth of 72 Gbits/second.
- The NetRAM's external clock runs at 66 MHZ which is at the PCI bus rate.
- The Dual addressing scheme enables the user to perform simultaneous Read/Writes in the same clock cycle. This provides performance equivalent to 133 MHZ of a conventional Single Address RAM.
- The price will be at a better competitive price point than to other dual ports of this depth and performance.

Thus, the NetRAM can greatly enhance system price/performance.

How to Reach Us:

Home Page:
www.freescale.com

E-mail:
support@freescale.com

USA/Europe or Locations Not Listed:
Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
+1-800-521-6274 or +1-480-768-2130
support@freescale.com

Europe, Middle East, and Africa:
Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
support@freescale.com

Japan:
Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku,
Tokyo 153-0064
Japan
0120 191014 or +81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:
Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:
Freescale Semiconductor Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
1-800-441-2447 or 303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com