Proceedings of the 39th Midwest Symposium on Circuits and Systems, Ames, Iowa, August 18-21, 1996, pages 73-76.

A CONTROLLER CHIP FOR A SCALEABLE ATM SWITCH NODE

Paul Shipley, Michael Weeks, Magdy Bayoumi

The Center for Advanced Computer Studies,
University of Southwestern Louisiana,
Lafayette, Louisiana 70504
e-mail: mcw4900@usl.edu

ABSTRACT

This paper presents the architecture of the control chip for an Asynchronous Transfer Mode (ATM) chipset, used in a distributing banyan network. This chip, with 4 switch chips, forms a scaleable 16x16 switching element that runs at 155 MHz. It is larger and much more efficient than previous switching nodes. The architecture uses shared multibuffering, since it is less bandwidth limited than a shared buffering switch. Such criteria make the switch more scaleable, although the control logic increases in complexity. Very high performance is required of the chip and thus a number of special circuits have been created to achieve this performance. The chip has been prototyped in 1.0 micron CMOS using a mixture of static and dynamic logic.

1. INTRODUCTION

An Asynchronous Transfer Mode (ATM) network transfers data quickly, due to several reasons. First, the destination does most of the error checking, which makes the switching elements faster and simpler to design. Second, ATM networks typically use lightwave systems (fiber optics), which have very low error rates. Third, the small, fixed cell size reduces the network latency, and makes the switching more flexible. A cell never needs to wait longer than 3 microseconds at a transfer rate of 155 Mb/s. Because of these advantages, researchers have concluded that ATM will dominate telecommunications traffic over the next ten years [6].

Each ATM cell contains 5 bytes of header followed by 48 bytes of payload. This fixed cell size of 53 bytes meets the conflicting needs of the ability to transfer large data files, such as still images, and the speed to transfer continuously sampled data, like a phone conversation.

This paper presents the control chip for an ATM chipset. The first section that follows introduces the distributed banyan network. The second section discusses the switching node. The third section examines the architecture of the control chip in more detail. The control chip implementation is discussed in the fourth section. The final section presents a summary of the design.

2. DISTRIBUTED BANYAN NETWORK

The distributed banyan network architectures are designed to allow the efficient construction of very large ATM switches. These architectures consist of a base fabric of switching nodes, interconnected using a distributed banyan packet switch [1]. The distributing banyan network is composed of two types of layers, distribution layers and routing layers. Distribution layers send packets to the emptiest available chip in the next layer without regard to the routing information. Routing layers switch according to the routing bits of the packets. These layers are alternated with a routing layer following after one or more distribution layers. Identical switching nodes can be used in both layers by adding a modest amount of extra control logic and a pin to indicate which type of layer it is in.

Proceedings of the 39th Midwest Symposium on Circuits and Systems, Ames, Iowa, August 18-21, 1996, pages 73-76.

This control chip described in this paper. allows for a 16x16 switching element [2]. Previous designs have had a 4x4 switching node size. The larger switching node achieves higher efficiency, better performance, and lower cost than previous designs.

3. ARCHITECTURE OF THE SWITCHING NODE

The control chip and 4 switch chips form a 16x16 switching element, as shown in figure 1. Each bit-sliced switch chip receives 1/4 of all the packets. For example, the first switch chip gets the first 1/4 of each packet. The control chip gets the packet headers from the switch chips, as well as the input flow control. This information allows the control chip to reconstruct the sequencing of the packets and to adjust the internal flow of packets according to back-pressure. The control chip outputs flow control data, address, and destination routing information to the switch chips. The control chip determines where incoming packets should be stored as well as when and where each of the packets stored in the switch chips should be sent out.

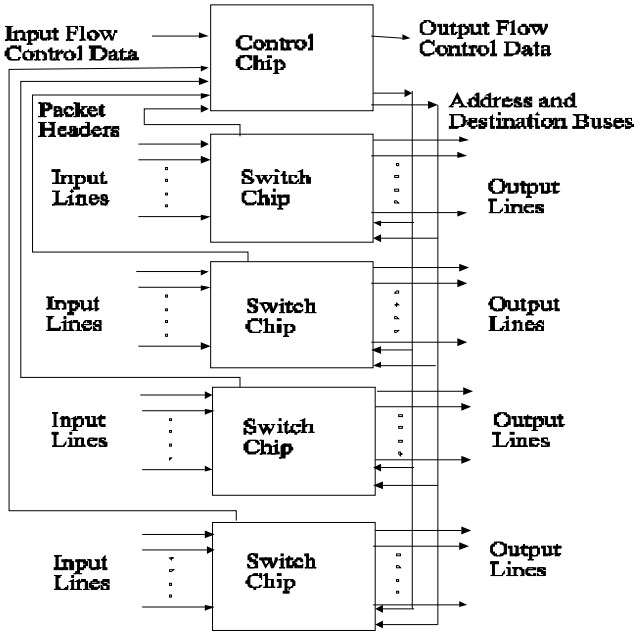

Figure 1. Diagram of the 16x16 Switching Node

4. ARCHITECTURE OF THE CONTROL CHIP

The switch chips store the packets and send them to the next node on the way to their destination. Each node can store 256 packets. The control chips send commands to the switch chips, which store the packets in order of the commands.

The control chip is composed of four components, Fig. 2. First is the flow control unit, which communicates with the next level of control chips. It indicates the capacity of this switch level, from empty to full. The second component, the address file, keeps track of which of the 16 cells are valid. The third component, the comparator, looks for the oldest packet. The fourth part, the register file, makes up the bulk of the control chip. It holds the 256 21-bit headers, and contains over 130,000 transistors.

The control chip decides which packet is the oldest via a timestamp comparator. The 256 timestamp bits are ANDed with each other in parallel, until only 1 of the 256 bits is a logical 1. Next, an encoder translates the position of this "winning" bit to an 8 bit packet address. The control chip sends the address to the switch chips, and the switch chips send out the packet corresponding to this address. Eight bits of destination data are also sent from the control chip to the switch chips. The switch chips route the packet to the next node of the switch.

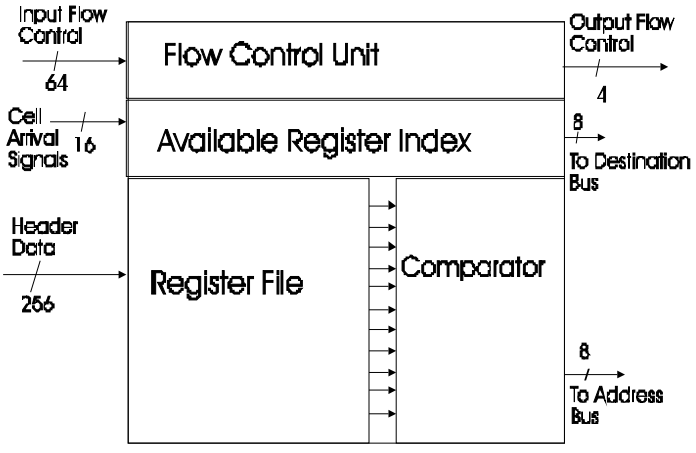

Figure 2. Structure of the Control Chip

Proceedings of the 39th Midwest Symposium on Circuits and Systems, Ames, Iowa, August 18-21, 1996, pages 73-76.

5. IMPLEMENTATION OF THE CONTROL CHIP

Due to the high performance requirements of the control chip, a number of specialized circuits are required for its implementation. These circuits include many components working at 155 MHz, which is four times the speed of the rest of the chipset. Pipelining is used to help achieve these speeds. Special mixed-signal devices are used in critical places to increase speed and save area. The combination of these features allows the chip to meet its performance goals.

The register file, Fig. 3, stores the header data for each of the packets stored in the switch chips. It is composed of 256 register units, Fig. 4. The register unit has 2 sets of D-latches, one for the routing data and one for the timestamp data. The main building block of the register file is the register unit. Sixteen register units wired together form a register module. The stored data consists of 5 bits of routing data and 16 bits of timestamp data. The routing data is used to determine when to present the timestamp data to the comparator forming a content addressable memory. At the beginning of a new cell cycle, the register will load the routing information in one clock cycle, and load the timestamp data in the following clock cycle. In the routing layers, if the routing information stored by the register unit matches the 5 bits of routing information sent along the routing bus, then the register unit sends the 16 timestamp data bits to register unit in high speed serial. By contrast, for the chips in distribution layers, the timestamp data is always sent regardless of the routing data. Four select lines strobe the multiplexor to send all 16 timestamp bits through the output, one at a time. Register units are used instead of SRAM due to speed concerns. The register units provide higher bandwidth, and are content addressable. The register file takes up most of the space on the control chip, using an area of 6,300x7,500 lambda. A routing filter connects the inputs to the register file. For each of the 16 buses going to the register modules, a multiplexor sends 5 bits of routing information first, followed by 16 bits of timestamp data.

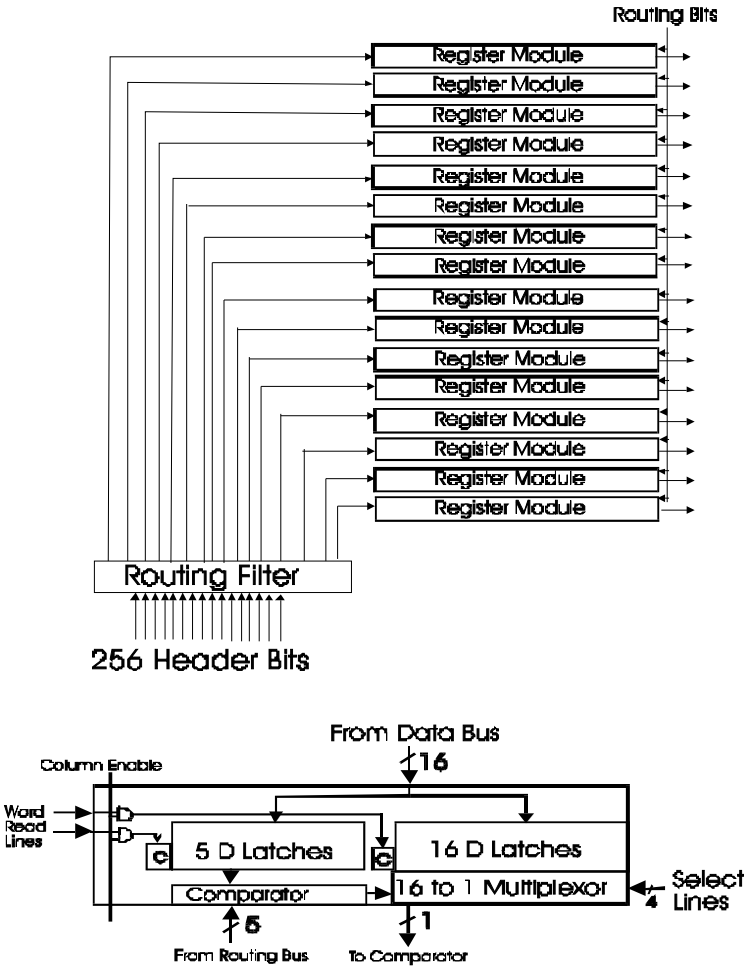

Figure 4. A Register Unit

The comparator is used to determine the oldest timestamp value, as shown in figure 5. Its zero detector determines when all inputs are 0. When this occurs, all of the inputs are replaced with the outputs of the register in the comparator. This allows the comparison to be skipped for that bit position. Figure 6 shows a sub-circuit used both in the zero detector and the encoder. This circuit has two outputs, one corresponding to at least one input being high, and the other to more than one input being high. From these two values three states can be distinguished: no inputs high, exactly one input high, and two or more inputs high. This mixed signal circuit is used because it is both faster than and much smaller than a conventional combination circuit.

Proceedings of the 39th Midwest Symposium on Circuits and Systems, Ames, Iowa, August 18-21, 1996, pages 73-76.

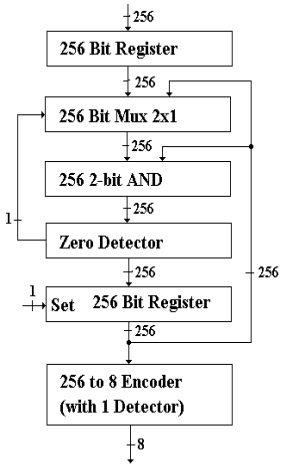

Figure 5. Structure of the Comparator

6. SUMMARY

This paper presents the design of the control chip used in a distributed banyan ATM switch. A switching element consists of one control chip, and four identical switch chips. Simulations have confirmed that the new switching node enabled by the presented control chip has over twice the storage efficiency of previous nodes for the distributing banyan network. This control chip allows for a 16x16 switching element, which is larger than previous designs. The larger size and storage efficiency are achieved with the specialized circuits of the control chip. Although the control chip logic is complex, it allows for a much better switching node than was previously available.

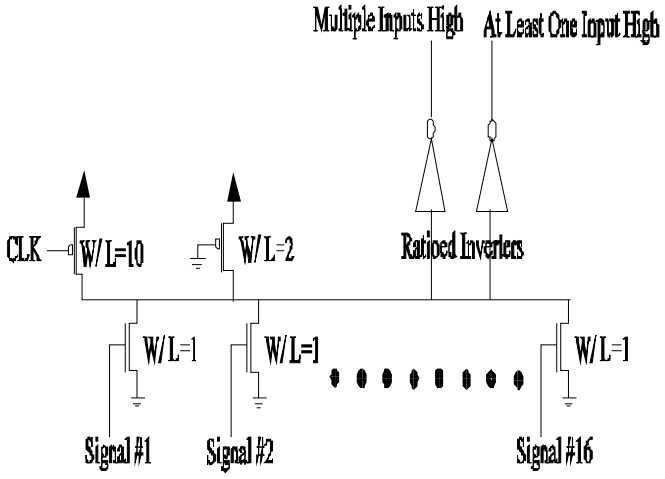

Figure 6. Sub-circuit Determining 0, 1, or More Than 1 High Inputs

ACKNOWLEDGMENT

The authors acknowledge the support of the DAMA Center Research Project, DOE No. 49X-SU016C.

REFERENCES

[1] Shipley, Paul, "VLSI Architectures for a New Class of ATM Switches", *Ph.D. Dissertation*, 1996.

[2] Shipley, Paul; Sayed, Sherif; Bayoumi, Magdy; "A High Speed VLSI Architecture for Scaleable ATM Switches", *Proc. of the IEEE Great Lakes Symposium on VLSI*, March 1996.

[3] Weste, Neil; Eshraghian, Kamran; Principles of CMOS VLSI Design: A Systems Perspective, Addison-Wesley Publishing Company, Reading, Massachusetts, 1993.

[4] Mukherjee, Amar, Introduction to nMOS and CMOS VLSI: Systems Design, Prentice Hall, Englewood Cliffs, New Jersey, 1986.

[5] Lane, James, “ATM Knits Voice, Data on any Net”, *IEEE Spectrum*, February, 1994, pages 42-45.

[6] Miller, Arthur, “From Here to ATM”, *IEEE Spectrum*, June 1994, pages 20-24.

[7] Ruiu, Dragos, “Testing ATM Systems”, *IEEE Spectrum*, June 1994, pages 25-27.

[8] Mirfakhraei, Nade, “Design of a CMOS Buffered Switch for a Gigabit ATM Switching Network”, *IEEE Journal of Solid-State Circuits*, vol. 30, no. 1, pp. 11-18, January 1995.