Technical
Report

Research Triangle Park

INTRODUCTION TO IBM LANSTREAMER ADAPTERS

Ray Zeisz

TR 29.1797 **November 1993**

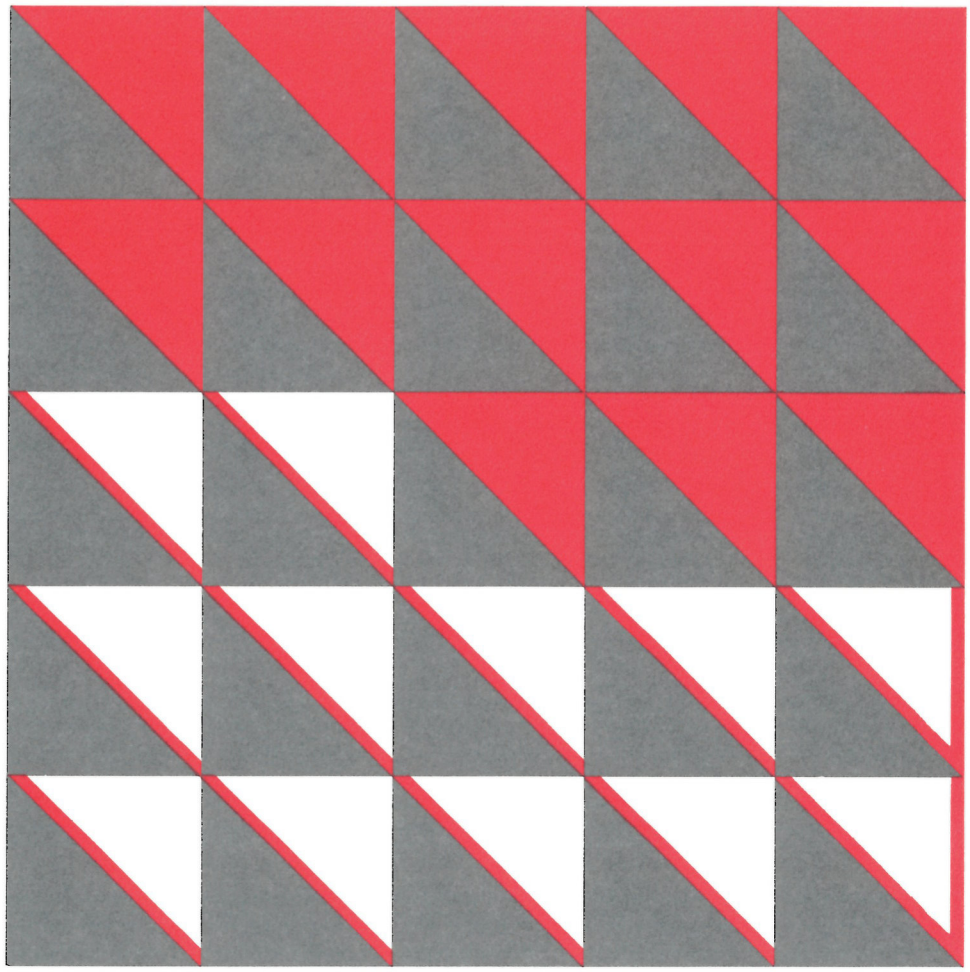

IBM Internal Use Only

INTRODUCTION TO IBM LANSTREAMER ADAPTERS

TR 29.1797

Ray Zeisz

LAN Adapter and New Technology Development
Research Triangle Park, NC

IBM Internal Use Only

ABSTRACT

In March 1993, IBM announced the next generation of Local Area Network adapters, targeted for applications requiring high throughput and low CPU utilization. Called LANStreamer, this technology is the industry's best performing Token-Ring implementation. LANStreamer is IBM's first LAN adapter that streams frames directly between the LAN and system memory. Previous adapters required frames to be temporarily buffered in RAM on the adapter. This and several other functional enhancements are described in this report. The *Streamer family* is IBM's strategic solution for customers demanding high performance at competitive prices for all major system bus platforms. IBM has also announced the EtherStreamer adapter for Ethernet networks. The EtherStreamer provides performance comparable to the Token-Ring LANStreamer, but for 802.3/Ethernet LANs.

This paper focuses primarily on the IBM Token-Ring LANStreamer Micro Channel adapter; however, most of the concepts presented apply to the entire *Streamer Family* of adapters.

ITIRC KEYWORDS

- LANStreamer
- Token-Ring
- LAN Adapter
- EtherStreamer

CONTENTS

Abstract iii
ITIRC Keywords iii

Introduction 1

LANStreamer Features 3
LANStreamer Bus Interface Enhancements 3
Busmasters and Bus Slaves 3
Performance 6
32-bit Addressing 7
Data Integrity Enhancements 7
Summary 8
LANStreamer On-board Microprocessor 8
LANStreamer Microcode Details 9
Other Functions 10
Multiple Group Addresses 10
Priority Transmit Channel 11
Support for UTP 11
Multiple LAN Adapters 11
Multiple Individual Addresses 11
Enhanced Bridge Support 11
LANStreamer Adapter Hardware Overview 12

References 15
Acknowledgements 15

INTRODUCTION

In March 1993, IBM announced its strategic solution for customers demanding high performance LAN access at competitive prices. The *Streamer family* of adapters is ideally suited for high performance workstations, servers and bridges alike. *Data Communications Magazine* found LANStreamer to be the industry's best performing LAN adapter.

The **IBM Token-Ring LANStreamer MC 32** provides many new functions such as 32-bit busmastering with optional data streaming, on-board support for unshielded twisted pair (UTP) cabling, two prioritized transmit channels, and several other powerful enhancements.

This paper focuses primarily on the IBM Token-Ring LANStreamer Micro Channel adapter; however, most of the concepts presented apply to the entire *Streamer Family* of adapters, including the IBM EtherStreamer for Ethernet LANs.

LANSTREAMER FEATURES

The following sections provide detailed descriptions of the features of LANStreamer including the bus interface, microcode, priority transmit queues, multiple group addressing, unshielded twisted pair cabling and enhancements for improved bridging.

LANSTREAMER BUS INTERFACE ENHANCEMENTS

LANStreamer's most impressive feature and major difference over previous LAN adapters is it's pipelining busmaster capability which provides fast and efficient data transfer. This section is dedicated to clearly defining what busmastering is and how LANStreamer takes advantage of it.

Busmasters and Bus Slaves

A Busmaster is a device that is capable of gaining control of the system bus and transferring data. All PCs have at least one busmaster: the system CPU. The CPU is clearly able to move data between various locations on the bus. Recently there has been a trend to make other devices in addition to the system CPU capable of moving data on the bus. Today, many systems provide busmaster video controllers and disk-drive controllers.

IBM has developed the next generation of LAN adapters as busmasters. Busmaster adapters free the system CPU from having to move data between system memory and the network; now the adapter takes care of this tedious and time consuming task. This differs from the traditional Bus Slave adapters which rely solely on the CPU to transfer data between adapter and system memory. Bus Slave adapters are commonly called **Shared RAM** adapters since they contain a buffer RAM that is shared by the adapter's on-board microprocessor and the system CPU. Figure 1 on page 4 shows the data path for a shared RAM Token-Ring adapter. Note that the CPU is directly involved in this data path.

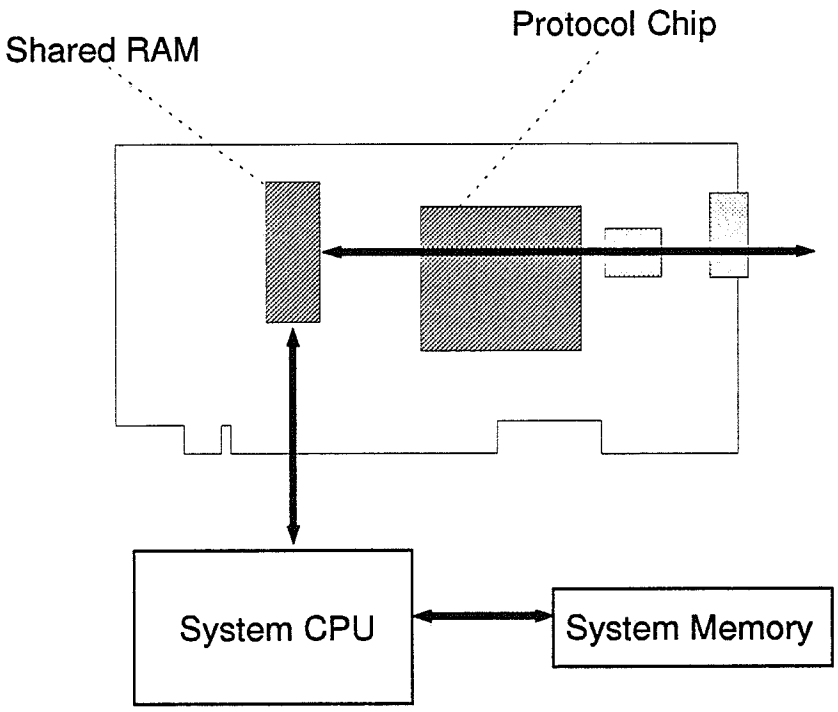

Figure 1. Shared RAM Adapter. *The system CPU is involved in the transfer of each and every byte of every frame.*

Figure 2 on page 5 depicts how the CPU is removed from the "critical path" of a data transfer for a **Store and Forward Busmaster** adapter. This is how the IBM Token-Ring Busmaster Server Adapter/A and the IBM Token-Ring EISA Busmaster Adapter are designed. Notice that while this design provides off-loading of the system CPU it still requires a large amount of buffer RAM on the adapter. This type of adapter requires the CPU only to initiate the data transfer. However, it is still not the optimal design, since frames are still completely buffered on the adapter.

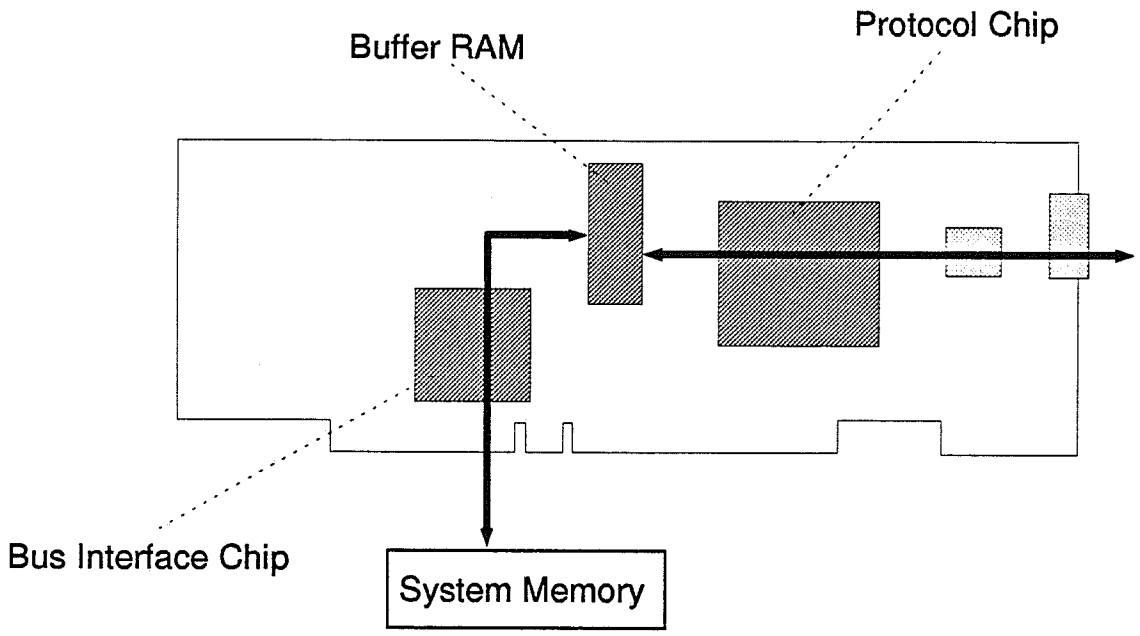

Figure 2. Store and Forward Busmaster Adapter. *The system CPU is only used to initiate data transfers. The adapter must buffer the entire frame.*

Figure 3 on page 6 demonstrates LANStreamer's design. This design eliminates the buffer RAM and replaces it with a small queue, commonly called a **FIFO** (First In First Out). This design allows frames to be moved directly between the system memory and the LAN without being entirely stored on the adapter. The importance of this feature will become clear in the following sections. The small FIFO is placed on the adapter to allow for momentary loss of the system bus by the adapter when other adapters in the system are transferring data.

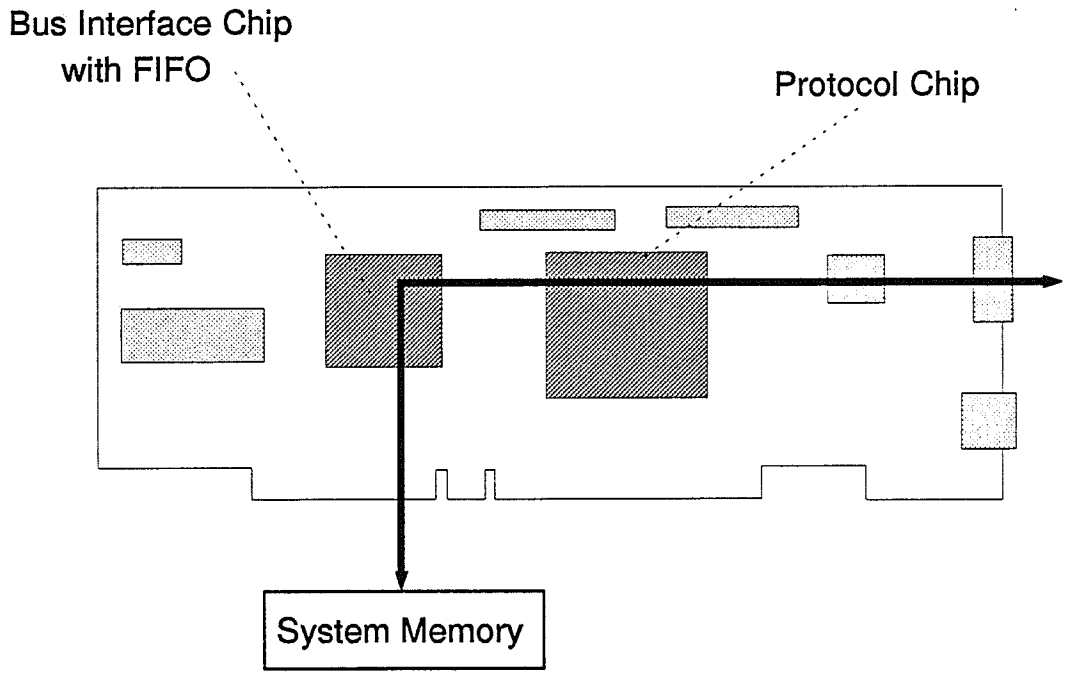

Figure 3. LANStreamer Adapter. *The adapter has no buffer memory. Frames are transferred directly between the LAN and system memory.*

Performance

The two most important factors in a LAN adapter's performance are **throughput** and **latency**. Throughput refers to how much data may pass through the adapter in given amount of time. A common analogy is that of a pipe carrying water. The wider the pipe is, the larger the throughput capacity. Latency refers to the amount of time it takes to move data between the system memory and the LAN. In the pipe analogy, the latency would be determined by the length of the pipe. As can be seen from the pipe analogy, it is possible to increase the throughput of an adapter without decreasing the latency. However, for applications requiring high-bandwidth, such as multi-media, low adapter latency is crucial. Excessive delay in video and audio transmissions causes a noticeable lag that becomes frustrating to the user.

LANStreamer has made considerable improvement in latency over previous LAN adapters. In addition, the latency on LANStreamer is independent of frame length, while shared RAM adapters have a latency that is directly proportional to frame length. Studies have found that a shared RAM adapter has a latency of over 2000 microseconds for a 4096 byte frame; while LANStreamer can begin to pass the same frame to the network in less than 30 microseconds. In fact, LANStreamer is able to pass a frame of any size in about 30 microseconds.

Throughput of LANStreamer has been maximized by the use of a 32-bit data path. Most adapters on the market pass only 8 or 16 bits of data during each bus cycle. LANStreamer is capable of moving 32 bits of data each cycle, doubling the throughput over previous adapters. In some newer systems, LANStreamer is capable of **data streaming**, whereby the duration of bus cycles is cut in half, thus doubling the throughput *again*. In data streaming mode, the first bus cycle explicitly supplies the address for the transfer, but each additional cycle has an implied address of the next successive location. Data streaming is thus very useful when large contiguous blocks of memory are moved from one location to another. Data streaming provides peak data transfer rates of 40 million bytes per second on newer Micro Channel computers. This allows for sustained throughput of 16 Mbps without frame loss. In addition, LANStreamer is capable of detecting when it is installed in a computer capable of data streaming, and will automatically take advantage of this feature.

These enhancements make LANStreamer up to four times faster than the fastest 16-bit busmaster adapters currently available from other vendors.

32-bit Addressing

Not only does LANStreamer support 32-bit data moves, but it is also capable of 32-bit addressing. This allows over 4 billion bytes of system memory to be directly addressed by LANStreamer. Several competitors have implemented 24-bit addressing with 32-bit data transfers and proclaimed the adapter as being "32-bits". LANStreamer is one of the few adapters with 32-bit addressing as well as 32-bit data transfers, making it ideal for systems with more than 16 Megabytes of system memory.

Data Integrity Enhancements

LANStreamer performs **parity checking** on the data it transmits. This provides added robustness since the adapter can now detect and report storage errors that occur in system memory or on the system bus. Parity checking is also independently performed on the address bus.

An additional feature called **selected feedback monitoring**, allows the adapter to detect and report conditions in which the adapter accesses an unimplemented memory location. This will not occur under normal conditions; however, under unusual circumstances such as a program error or electromagnetic interference, the adapter may attempt such an access. LANStreamer is ready for these conditions.

These features allow for added robustness especially valuable for systems running mission critical applications and real-time control systems.

Summary

LANStreamer's single greatest asset is its FIFO-based busmastering capability. IBM has never before produced such an adapter. Studies have shown that while the IBM Token-Ring Busmaster Server Adapter/A can pass 3,000 frames per second, a LANStreamer is able to pass 48,000 frames in the same amount of time. IBM has reached the pinnacle in LAN adapter performance with LANStreamer.

LANSTREAMER ON-BOARD MICROPROCESSOR

Every IBM Token-Ring adapter has been designed with a 16-bit microprocessor on-board. This processor is responsible for executing the IEEE 802.5 Token-Ring protocol as well as performing requests from the device driver. On shared RAM adapters, the IEEE 802.2 Logical Link Control (LLC) processing may also be performed by the adapter's micro-processor. However, in the LANStreamer, the LLC processing is performed by the system CPU, which is typically faster than the Token-Ring adapter's microprocessor.[1] The primary function of LLC is *assured delivery*, whereby the adapter guarantees frames are received by the destination using automatic retry and sequenced responses.

LANStreamer adapters have no buffer RAM on-board. This is drastically different from the shared RAM adapters which typically have 64 KB of RAM for buffers on-board.

In the store-and-forward cases, the data is moved from the system memory to the adapter. In a bus slave or shared RAM adapter, the system processor would perform this move of data. In a busmaster, the adapter would perform the data move, independent of the system CPU. Next, the adapter determines that the frame may be transmitted, and the frame is then copied from the adapter memory to the network.

In the FIFO case, the frame is copied from the system memory directly to the network. Since the LANStreamer has no RAM on-board and since frames never wholly reside on the adapter, the adapter's on-board microprocessor is not able to perform LLC processing of frames. This is not a drawback as one might suspect at first glance. Since most of today's personal computer systems are based on 80386 or newer technology, the system CPU is better suited for LLC processing anyway. By moving the LLC processing from the on-board processor to the system CPU, still further performance improvements are realized. However, since there is no buffer RAM on the adapter, slightly more system memory is required by the device driver.

---

1 Historical Note: Originally, the LLC processing was performed on the adapter only because older system CPU's such as the 8088 and 8086 were not fast enough to perform the LLC processing.

The Media Access Control (MAC) processing is always performed by the adapter's on-board processor. The MAC processing is minimal as compared to the LLC processing. Having the on-board microprocessor dedicated to performing the MAC protocol also provides robustness to the network as the on-board processor is never busy performing another task. This guarantees that an adapter cannot adversely affect the ring due to lack of MAC processing.

LANStreamer Microcode Details

The adapter's on-board processor executes a program from a Read Only Memory (ROM). This program is commonly called the Token-Ring microcode. This microcode performs three basic functions:

- Diagnostics
- 802.5 Protocol Processing (Ring Task)
- System Interface Functions

The Diagnostics are performed by the on-board processor each time the system is booted and the adapter is initialized by the system CPU. The LANStreamer has been equipped with a completely new set of diagnostics that test all major components on the adapter.

Among other operations tested, the diagnostics verify that the adapter is able to send and receive data. A diagnostic has been added which allows the adapter to send and receive frames through the system bus memory. This comprehensive test ensures that the adapter has a functional data path for both transmitting and receiving.

The Ring Task is IBM's fully standard implementation of the IEEE 802.5 Token-Ring Media Access Control (MAC) protocol. The Ring Task is responsible for handling all events that occur on the Token-Ring.

In particular, some of the operations the Ring Task is responsible for include:

- Neighbor Notification
- Active Monitor Functions
- Standby Monitor Functions
- Beaconing and Resolution of Beaconing
- Inserting the Adapter Into the Network
- Fault Detection

The System Interface provides a method for the device driver to communicate to the adapter. This allows the device driver to request the adapter microcode to perform such commands as Open Adapter or Close Adapter. For example, the microcode will only insert the adapter into the network when it has been instructed to do so via an Open Adapter

command from the system CPU to the System Interface. The System Interface is also used to inform the system CPU of conditions that are present on the network such as beaconing.

OTHER FUNCTIONS

In addition to all of the enhancements discussed so far, there are several more subtle functions waiting to be used by state-of-the-art applications. The following sections briefly discuss these new features.

Multiple Group Addresses

Group addresses allow multiple adapters to receive the same frame. This is useful when a server needs to communicate the same data to multiple stations at the same time. Previous Token-Ring adapters allow for only a single group address. LANStreamer now allows 256 unique group addresses to be address matched. This provides complete filtering of multicast frames in hardware. Most of today's applications broadcast frames to all adapters on the network; forcing the device driver running in the system CPU to decide to keep or discard the frame based on the encapsulated data.

In a multicast environment, each application could be assigned a group address and then only stations needing to copy frames would copy them. This is drastically different from today's network designs whereby applications such as TCP/IP use up to 40% of every system CPU on the network (even CPUs not running TCP/IP) since they transmit quite often to the All Stations Address.[2]

Group addressing is supported by all major LAN types including Ethernet and FDDI. This makes multiple group addressing especially attractive for heterogeneous networks.

A practical example of how multiple group addresses might be used is that of a stock market. Each stock is assigned a unique group address. For each transaction that takes place, a frame is broadcast with that stock's group address as the destination. Brokers may then select the stocks that they preferred to monitor by having the respective group addresses set in the adapter. This would allow the broker's CPU to be interrupted only by frames for which that broker is interested. The LANStreamer would allow brokers to monitor several hundred stocks simultaneously, without requiring their systems to be interrupted needlessly for transactions which do not interest them.

[2] Frames destined to the All Stations Address are automatically copied by all adapters on the network.

Priority Transmit Channel

LANStreamer adapters have two prioritized transmit channels. This is ideal for workstations which are running multi-media applications. In order for multi-media to be effective, the latency must be kept to a minimum. By queueing multi-media frames on the high priority channel, they will be transmitted before any lower priority frames. In addition, the high priority transmit channel is allowed to request tokens of up to priority six on the Token-Ring. This allows latency sensitive frames to be transmitted in front of frames that are not sensitive to delay.

Support for UTP

In the past, Token-Ring adapters connected via unshielded twisted pair cabling required a *Type 3 Media Filter* to be attached to the adapter. This filter, which improved the signal to noise ratio, had a price of about $50. LANStreamer adapters have removed the need for the Type 3 Filter completely. By adding the filter circuitry and an RJ-45 UTP jack to the adapter, the need for the media filter has been completely eliminated. This allows the adapter to be connected directly to UTP cabling which is prevalent in many establishments.

Multiple LAN Adapters

The LANStreamer technology has removed the concept of Primary and Alternate adapters that existed with shared RAM adapters. Now, up to six LANStreamer adapters may be installed in a single Micro Channel computer. This is very useful for servers which need to service multiple Token-Rings. Connecting a server to multiple Token-Rings may reduce the amount of traffic needing to cross bridges or routers.

Multiple Individual Addresses

LANStreamer adapters are equipped with a feature which allows them to receive frames for multiple individual addresses. These addresses will be verified for uniqueness on the Token-Ring, for security reasons. Up to 32 consecutive locally administered addresses may be used by a single LANStreamer adapter at one time. This could be especially useful for gateways, bridges and routers.

Enhanced Bridge Support

The LANStreamer has been ideally designed for use in bridges. Bridges are devices that connect two or more LAN segments together, thus allowing a network to become very large. The LANStreamer may be used with the ***IBM LANStreamer Token-Ring Bridge Program/DOS*** to provide a throughput of over 15,000 frames per second. This is a huge step up in performance over the older shared RAM based bridges.

Additionally, LANStreamer has been equipped with an enhanced **Address Match Function** that performs **Multi-port**[3] bridge route matching in hardware. This is unlike software matching whereby the system is burdened with frames that are copied, only to be later discarded by the system CPU. LANStreamer can independently provide Single Route and All Route matching for all 4096 possible ring numbers.

LANSTREAMER ADAPTER HARDWARE OVERVIEW

LANStreamer technology is based on two VLSI chips. The Protocol Chip contains the logic needed for the Token-Ring interface while the Bus Interface Chip contains the logic needed to interface with the system bus.

The Protocol Chip is comprised of the following key units:

- A 16-bit Microprocessor
- RAM (for microcode variables)
- Local Bus Control Logic
- Token-Ring Protocol Handler
- Token-Ring Analog Front-End Interface

The Protocol Handler contains the logic that the microcode uses to perform functions such as inserting into the network. It also encompasses the state machines for token operation, delimiter detection and address match algorithms. The analog front-end is the part of the adapter which converts the data into signals which can be placed on a wire.

The Bus Interface Chip contains FIFO's for the transmit and receive channels, an interface so that the system CPU may communicate with the on-board microprocessor, and the logic necessary to perform the complex busmaster operations that the adapter supports. The Bus Interface Chip is LAN protocol independent.

Figure 4 on page 13 shows a typical LANStreamer adapter with each component labeled and indications of the logical data paths.

[3] A Multi-port bridge is one that is connected to three or more LAN segments at once.

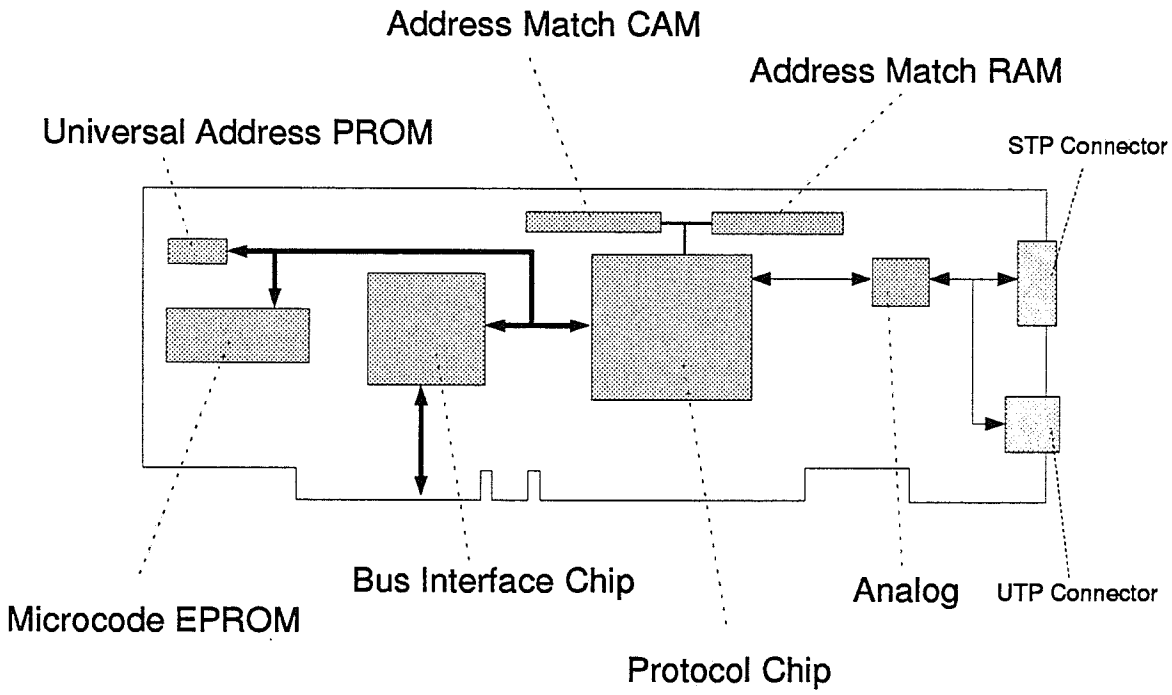

Figure 4. IBM Token-Ring LANStreamer MC 32 Adapter. *Anatomy of an adapter.*

The Microcode EPROM is nonvolatile storage which contains the program used by the on-board microprocessor. The Universal Address PROM contains the unique "Burned In Address". The Address Match RAM and CAM[4] are used for the group address matching and the multi-port bridge Routing Information matching.

4 Content Addressable Memory

REFERENCES

For information on LAN Protocols:

- **IBM LAN Concepts and Products** GG24-3178-02
- **IBM Token Ring Architecture Reference** SC30-3374-02
- **IEEE 802.5 Standard** (Token Ring)
- **IEEE 802.2 Standard** (Logical Link Control)

For information on LANStreamer Performance:

- **Data Communications Magazine** February 1993

ACKNOWLEDGEMENTS

The author wishes to thank the following people for their input and guidance: Ken Christensen, Joe McDonald, Franc Noel, Steve Polge, Norm Strole, Kathy Wilhelm, and Ken Wilson.