Token Ring Switching

A Technology White Paper

Author, Martin Taylor

December 1994

Inside Front Cover

Blank

The Growing Demand for Bandwidth

The deployment of Personal Computers and Local Area Networks has had a profound impact on the nature of computer applications used in the enterprise. New applications that take advantage of PC and LAN technology are faster to develop and deliver higher levels of productivity than ever before. Downsizing, client-server computing, e-mail, groupware and conferencing applications are all on a steep upward growth curve.

The new LAN-based applications have one thing in common: they are demanding on LAN bandwidth. While 4 Mbps Token Ring may have delivered ample capacity for hundreds of users running mainframe terminal sessions, it can become heavily stressed when new client/server, workflow and groupware applications are introduced.

Token Ring has, in fact, stood up well to the demands that are increasingly being placed on LANs. With the ability to run at 16 Mbps wire speed, and with the token-passing protocol which gives users guaranteed response times at loadings up to and over 95% of bandwidth, Token Ring has proven itself to be a truly industrial-strength LAN technology. By contrast, Ethernet has not fared so well. At loadings above about 40%, collisions eat up all the remaining bandwidth causing users to see widely varying response times. And with its lower wire speed of 10 Mbps, Ethernet's effective carrying capacity is limited to about 4 Mbps, or a quarter that of Token Ring.

As the new networked applications have begun to push shared medium LANs to the limit, so the networking vendors have responded with innovative solutions to the bandwidth problem. One such solution is LAN switching.

Switching is a technique for interconnecting multiple LAN segments simply and efficiently. It allows large LANs to be constructed from many LAN segments. Users on each segment are only contending for bandwidth with the other users on the same segment, so the smaller the segments the greater the average bandwidth available to each user. The switching device has much higher overall capacity than the individual segments, so the aggregate capacity of the LAN is greatly increased.

Building LAN Infrastructures: Bridges and Routers

Large LANs have always consisted of multiple segments, because there are practical limits on the number of stations that can be attached to a single segment. With Token Ring, the maximum permitted is 255 stations on one ring, but in practice few users load rings up with more than about 100 stations.

The traditional method for interconnecting Token Rings is the source routing bridge. This simple 2-port device lends itself to an architecture where workgroup rings are linked by bridges to a backbone ring. However, the historical performance limitations of bridges encouraged network administrators to locate their servers in such a way as to minimize the traffic across the bridges. Typically, this meant attaching each server directly to the ring which connects its main group of users.

Network administrators have found that locating servers on workgroup rings creates a number of problems. The network tends to be inflexible, increasing the burden of dealing with moves and changes, and the servers are more difficult to administer, back-up and secure. However, attempting to move the servers to the backbone often results in performance problems. Users who have enjoyed the high speed response from a locally-attached server are dismayed by the fall-off in performance they see when they have to cross a bridge, and to make matters worse the aggregate traffic loading on the backbone ring can quickly grow to overload levels.

Many network administrators have addressed the problem of aggregate backbone loading by moving to a collapsed backbone architecture based around a large router. The argument here is that the router has a much greater internal capacity than a 16 Mbps Token Ring, so that it can be used to interconnect a large number of rings without running out of bandwidth.

Routers and Network Performance

While the collapsed backbone architecture clearly deals with the aggregate bandwidth issue, large routers do not always deliver the performance benefits that might be expected. The principle reason for this is the delay, or latency, which the router introduces in moving a packet from one ring to another.

All communications protocols are sensitive to latency, though some are much more sensitive than others. The NetWare protocol, IPX, is one of the most sensitive because every packet sent requires an acknowledgement before the next can be transmitted. By holding up both data packets and acknowledgements, routers increase the total round trip time and reduce the rate at which packets can be transferred. The result is a substantial reduction in throughput as seen by the network user. With NetWare, placing a single router between a client and a server will reduce the file transfer throughput between the two by as much as 30-40%. And this degradation in throughput is cumulative, so a transit through two routers will result in a 50-60% slowdown.

Latency in routers is a result of the time taken to perform the various steps involved in the routing process. First, a packet has to be received in its entirety from the input port into buffer memory. Then the network address is decoded and the routing decision is made. The Token Ring packet header must be replaced with a new header containing the destination MAC address of the next router or the destination station as appropriate. The packet may then have to be moved across the backplane of the router to another interface card, where finally it can be transmitted onto the output ring.

Latency is not the only performance issue for routers in Token Ring environments. Routers tend to be optimized for Ethernet and TCP/IP environments, where throughputs of 200,000 packets and more are now being achieved through the use of silicon routing accelerators. In a Token Ring environment, where protocols such as IPX and the non-routable LLC and NetBIOS tend to dominate, very much lower levels of performance are being achieved.

Token Ring Switching

Switching offers an alternative means of interconnecting multiple Token Rings which is both simpler and more efficient than routing. In a switch, packets are transferred from one ring to another with negligible delay, because the packets are neither buffered nor processed within the switch.

Instead of reading packets in their entirety into buffer memory before making a decision about where to forward the packet, a switch will take action as soon as the first 20-30 bytes of the packet have been received. Information in the packet header is analyzed almost instantly, and the required destination port is deduced. At this point, a connection is effectively made between the input port and the output port, and the packet immediately starts transmitting onto the destination ring. This technique is sometimes known as "cut-through" or "on the fly" switching.

The total time that a packet is held up within a switch is as little as 30 microseconds, which compares with the 500-4,000 microseconds of delay that routers introduce, depending on packet size.

By virtually eliminating latency, Token Ring switching allows clients on one ring to communicate with servers on another ring with the same performance as if they were both attached to the same ring.

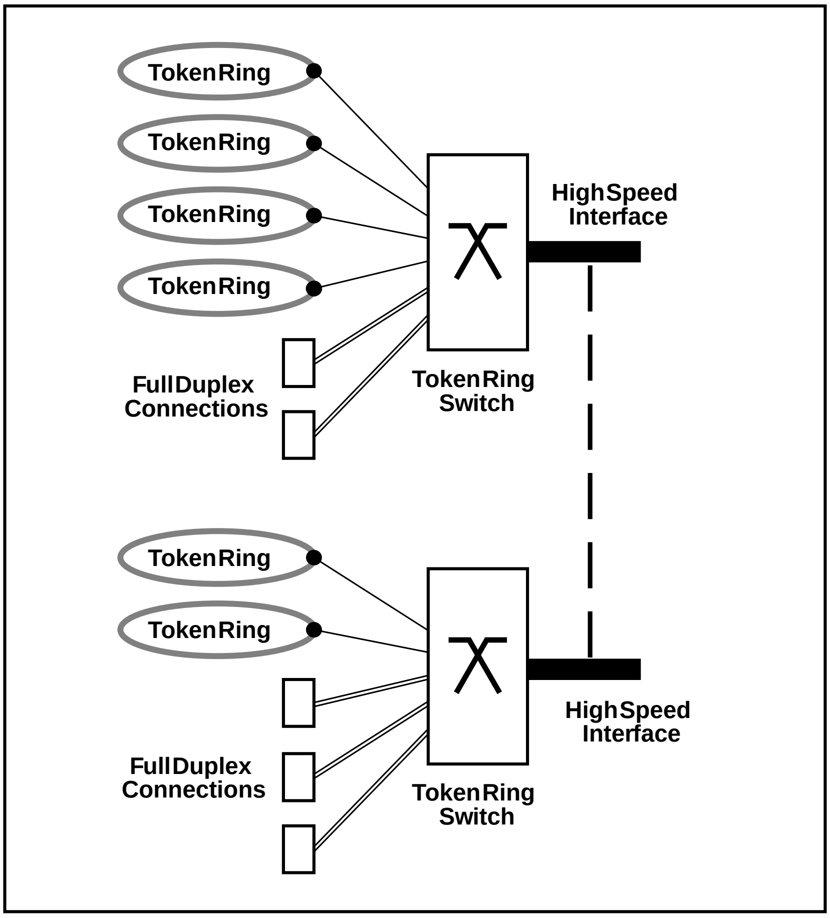

Figure 1: Madge concept of switched Token Ring

Contention and Buffering

Although switching involves the transfer of packets from one ring to another without buffering packets in memory, there are circumstances where buffering is needed. A switch cannot transmit a packet onto a ring if the ring is busy, and in this case the packet must be buffered until the destination ring has gone quiet and the token can be grabbed. Likewise, if packets arrive at two input ports of the switch simultaneously for onward transmission to the same ring, the switch must buffer one of the packets until it has finished forwarding the other.

Thus a Token Ring switch must be equipped with sufficient buffer memory to deal with these circumstances without dropping packets. This is true for any kind of LAN switch.

Switching by Source Routing

The standard method for interconnecting Token Rings is source routing. With this technique, clients first establish a route to a server using a route discovery process, and then insert information that defines this route in each packet that they send.

Source routing operates at Layer 2, the Data Link layer, and is therefore applicable to all upper layer protocols, whether they include Layer 3 (Network layer) addressing or not. A Token Ring switch can make use of source routing information to make forwarding decisions on each packet that it receives. Because the source routing information explicitly identifies the ring to which the packet should next be passed, the switch can make very rapid forwarding decisions with minimal processing.

There is an alternative to source routing as the basis of Token Ring switching, known as transparent switching. Ethernet does not support source routing, and so Ethernet switches use the transparent switching technique.

With transparent switching, forwarding decisions are made on the MAC address information in the header of each packet. The switch has to learn which MAC addresses are attached to each of its ports and maintain tables of this information. As each packet enters the switch, the learning tables must be updated, and the destination MAC address must be looked up to determine which is the correct destination port.

Switching based on source routing has a number of advantages over the transparent technique. Less work is required for each packet, and therefore less processing power is required, leading to lower costs. All Token Ring applications are compatible with source routing, whereas some will not work with transparent switching. And source routing makes it easier to implement virtual LANs, a subject that will be covered later in this paper.

Full Duplex Token Ring

Token Ring switching provides a means to interconnect multiple Token Rings with very high performance. A switch can act as a collapsed backbone, connecting workgroup rings to other rings which support centrally-based servers.

In almost all LAN environments, many users access a few servers, and the server connection can therefore easily become a bottleneck. Switching helps with this problem, as it enables the server rings to be segmented so as to reduce the number of machines on each ring.

The most heavily loaded servers may well justify having a private ring all to themselves. In this special case, the server is the only device connected to a port on the switch. Instead of running the full token-passing protocol of Token Ring between these two devices, it is possible to dispense with the token and operate this link as a full duplex serial link running at 16 Mbps in both directions at once.

Full duplex Token Ring provides double the bandwidth to any station that handles concurrent bidirectional traffic, as it offers a clear 16 Mbps channel in each direction. Heavily-loaded servers can take advantage of this, since their operating systems will deal with multiple concurrent read and write operations. Video-equipped workstations can also benefit from full duplex Token Ring.

Unlike Ethernet, Token Ring hardware is inherently capable of operating in full duplex mode. To provide full duplex operation, the software controlling the Token Ring chipset simply needs to be changed. Madge's Smart 16/4 Token Ring adapters are designed for software upgradability, and all current adapters will support full duplex mode driver software.

The IEEE 802.5 committee which sets the standards for Token Ring set up a project to define a standard for full duplex Token Ring in March 1994.

Virtual LANs

LAN switches direct traffic so that packets are only sent to the segments where they are needed. This is fine for individually addressed packets, but what about broadcast packets? It turns out that there are additional benefits to be gained from switching if the forwarding of broadcast traffic is controlled in the switch. With the appropriate intelligence applied to the filtering of broadcast packets, switches can help network administrators with moves and changes by allowing the creation of "virtual LANs".

A virtual LAN is a collection of LAN segments connected together by switches where all broadcast traffic originating from any of the segments is seen by all stations on the other segments. By either blocking or enabling the flow of broadcast traffic between designated ports on the switch, virtual LANs can be defined to include or exclude specified LAN segments attached to the switch.

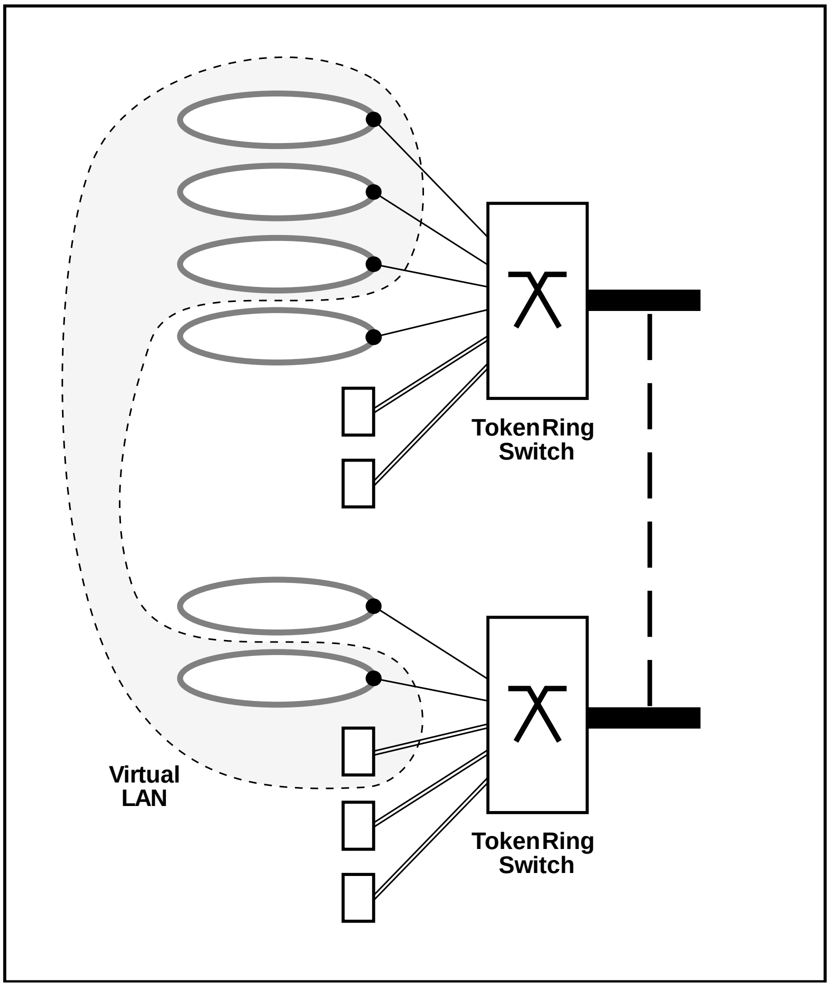

Figure 2: Virtual LANs in a switched Token Ring environment

All the stations on a virtual LAN are able to see all broadcast packets, including service advertisement, address resolution and route discovery packets, that originate within the virtual LAN. Likewise, stations cannot see any broadcast packets that originate on segments which are defined as belonging to other virtual LANs. The result is that stations can only make connections to other stations, servers or gateways that are part of the same virtual LAN.

With virtual LANs, network administrators are able to define a logical network structure that is quite different from the physical layout of the network. A finance department workgroup ring may, for example, be defined as belonging to the same virtual LAN as another centrally-located ring that connects the finance department servers. If the department moves to a new physical location, the virtual LAN may simply be redefined to ensure that the users can continue to access their server ring, while other departments are excluded from accessing it.

Virtual LANs provide the additional benefit that they confine the propagation of broadcast traffic within the set of rings that must receive the broadcasts. This ensures that broadcast traffic occupies only a small proportion of the bandwidth on each segment, and overcomes any concern about "broadcast storms" in a source routed environment.

Some existing Ethernet switch products do allow virtual LANs to be defined, but the obstacle to wider acceptance of this concept is the sheer complexity of controlling broadcast traffic in a LAN installation with multiple switches. In a transparent switching scheme, broadcast packets moving between switches contain no information as to where they originated apart from source MAC addresses. Determining whether to forward these packets to any particular port of the switch requires large look-up tables which are very complex to administer.

With a LAN switching scheme based on source routing, this is no longer an issue, as every broadcast packet contains within it the source routing information that identifies the ring on which the packet originated. The switch can make use of this information to ensure that the broadcast is only forwarded to the segments that are part of the same virtual LAN. This dramatically simplifies the definition of the filtering tables for broadcast packets, and makes virtual LANs spanning multiple switches a realistic proposition.

Token Ring Switching and ATM

Token Ring switches need a high speed interconnect so that large volumes of traffic can be carried between switches. Both FDDI and ATM are seen as appropriate technologies for the switch interconnect.

FDDI offers 100 Mbps capacity between switches, and will support the connection of servers at 100 Mbps, thus eliminating another potential source of network bottlenecks. ATM at 155 Mbps will also provide ample capacity for inter-switch traffic. But when ATM is used together with Token Ring switching, it offers far more than just a means to interconnect switches.

When integrated with a LAN environment, ATM operates in a mode known as "LAN emulation". The idea of this is to make an ATM network behave like a LAN, even though traffic is being carried on point-to-point connections across the network. With LAN emulation, additional services within the ATM network support LAN-like functions such as address resolution and broadcast packets.

ATM is capable of emulating Token Ring LANs. Token Ring packets can be carried across an ATM network just as if they were being carried on a ring - even though in reality they go point-to-point. And Token Ring packets can be switched into an ATM network emulating a Token Ring.

The emulation of a LAN segment on ATM is defined purely by software, and an emulated segment is, in effect, a virtual LAN. Furthermore, it is possible to define multiple virtual LANs within an ATM network. When switched Token Rings are connected to an ATM network, the connections between physical rings attached to Token Ring switches and emulated rings within the ATM network are defined just by software. In fact the whole concept of virtual LANs described above in the context of switched Token Ring, can be extended to the combined Token Ring and ATM network.

This idea of virtual LANs embracing both the Token Ring and ATM domains is extremely powerful and flexible. It provides network administrators with complete freedom to configure and re-configure the linkages between rings and servers in any size of network, just through actions taken at the network management console. And it allows large networks to be built entirely on the basis of switching, largely eliminating the need for slow and inefficient routing.

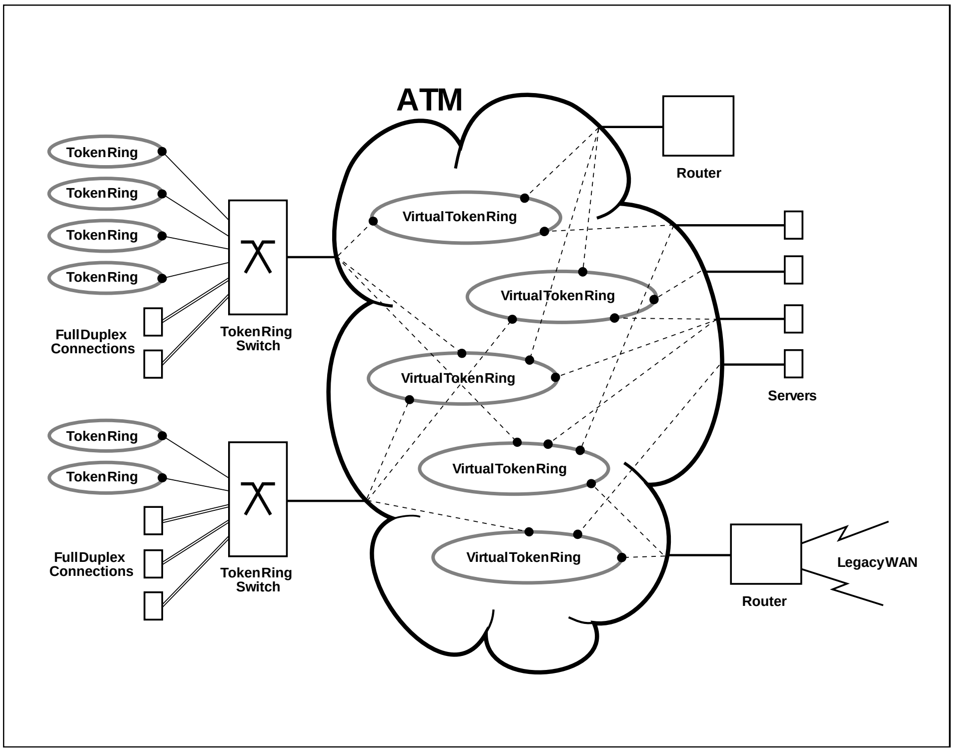

Figure 3: Token Ring and ATM

All LAN Switches Are Not Equal

As so often happens with technology, labels have a way of becoming loosely defined. This is likely to be the case with "switch".

Switching is closely related to bridging, in that both provide a connection between LAN segments operating at Layer 2. The distinguishing feature of switching is that packets are not buffered in the device, they are forwarded on direct connections between input and output ports. Bridges, by contrast, operate in store and forward mode, buffering the entire packet in memory before sending it to the destination ring. Therefore bridges, like routers, introduce substantial latency, while switches do not.

Devices that claim to be Token Ring switches have already started to appear on the market. Some of these products actually operate in store and forward mode, and should therefore really be described as multiport bridges. These devices will not provide the performance enhancing benefits of true switching.

Some switch vendors have claimed that Token Ring cannot be switched by cut-through techniques because the switch has to wait for the token to arrive at the destination port before it can start transmitting. However, the average wait before the switch acquires the token on a typical ring is equivalent to only a few bytes of extra latency. Switch latencies at least as low as the best cut-through Ethernet switches are certainly achievable.

Madge Switched Token Ring Technology

Madge’s Token Ring switch technology is designed to achieve minimal latency and maximum utility. It will use source routing and provide cut-through packet transfers. Token Ring switching, like Madge's other Token Ring technologies, is meant to be cost effective while making no compromises in performance or quality. To that end, Madge will employ technologies it has pioneered in order to leap beyond the architectural limitations of off-the-shelf components.

The principle challenge in building a Token Ring switch is in moving data between the network wire and the switching fabric with sufficiently low latency. To meet this challenge, Madge has developed new Token Ring protocol handler silicon, known as the RingRunner™ chip. This chip is to be made available to other vendors who wish to build Token Ring switches or other devices such as Token Ring to ATM converters, under a licensing agreement with Texas Instruments.

Like bridges and routers, Token Ring switches need to interconnect with other switches or to backbone media. Thus, Madge’s Token Ring switches are to be built with a high-speed interface to other switches and to servers via FDDI or ATM.

Madge's Token Ring switches will be based upon four core technologies - the Token Ring protocol handler, the associated Token Ring MAC processing software, a hardware frame filter, and a non-blocking switch fabric.

RingRunner - Token Ring Protocol Handler

Designed and developed by Madge, the RingRunner converts between the bit stream on the Token Ring wire and the byte stream for input/output to the switching fabric. It handles all packet sizes at wire speed, and is capable of passing over 48,000 packets per second, with a total latency of only a few bytes. The chip also handles the Token Ring protocol functions including destination address matching, source route filtering, ARI/FCI bit insertion, CRC generation and checking, and token generation. RingRunner has been optimized for bridge, router, switch and Token Ring to ATM conversion applications, and it is suitable for both source routing and transparent switching.

Token Ring MAC Software

The MAC frame processing software is designed for use with the RingRunner chip. It handles the Token Ring MAC frames which provide the Layer 2 management functions in Token Ring. The software is written in C for portability, and is fully compliant with IEEE 802.5 Draft 5. It has been tested on TI-based adapter cards and on Intel processors, and it can be easily modified to accommodate future high speed Token Ring developments.

Hardware Frame Filter

Incoming packets are routed within the switch using specially designed hardware. This subsystem routes packets based on destination address information, source route information, broadcast address information, and specific routing filters.

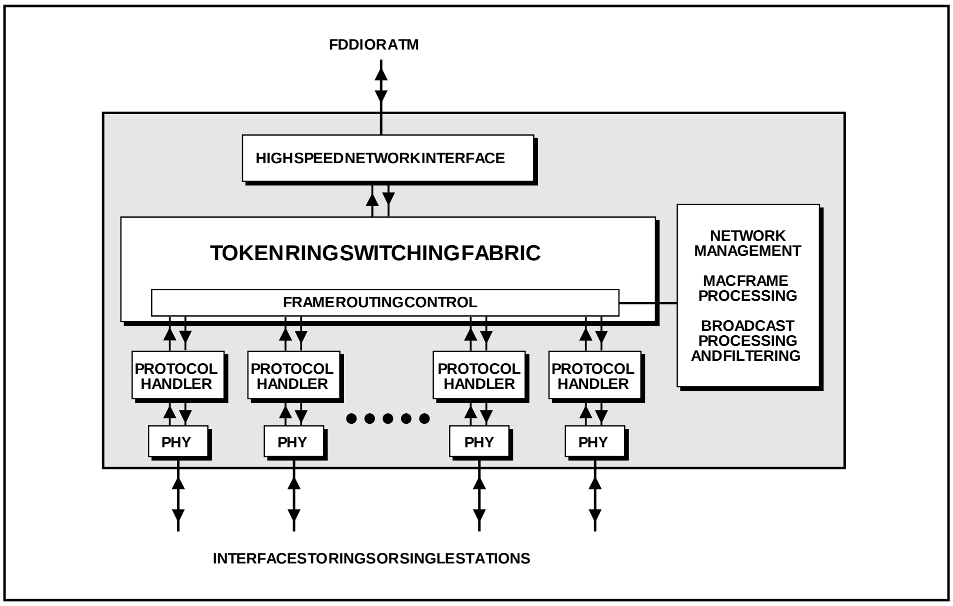

Figure 4: Madge Token Ring Switching Architecture

Frame Switch Fabric

Frame switching is carried out in a non-blocking fabric based on a multiport memory architecture. It is capable of queuing full length Token Ring frames, and provides output buffering so that incoming packet traffic will not be blocked if an output port is busy or if another station is transmitting on the destination ring.

Price/Performance

Token Ring switching is an inherently low cost technology by comparison with collapsed backbone routers. The packet forwarding decisions made by switches are simple, and packets are not changed as they pass through the switch. The movement of packets from input to output ports is carried out in dedicated silicon on parallel paths, and no processor is involved in the data transfer. The result is that very high throughputs and packet handling rates can be achieved with relatively little processing power.

Token Ring switching will provide higher performance than backbone routers with lower latency, at a price that will be well under half the cost of a typical router-based solution.

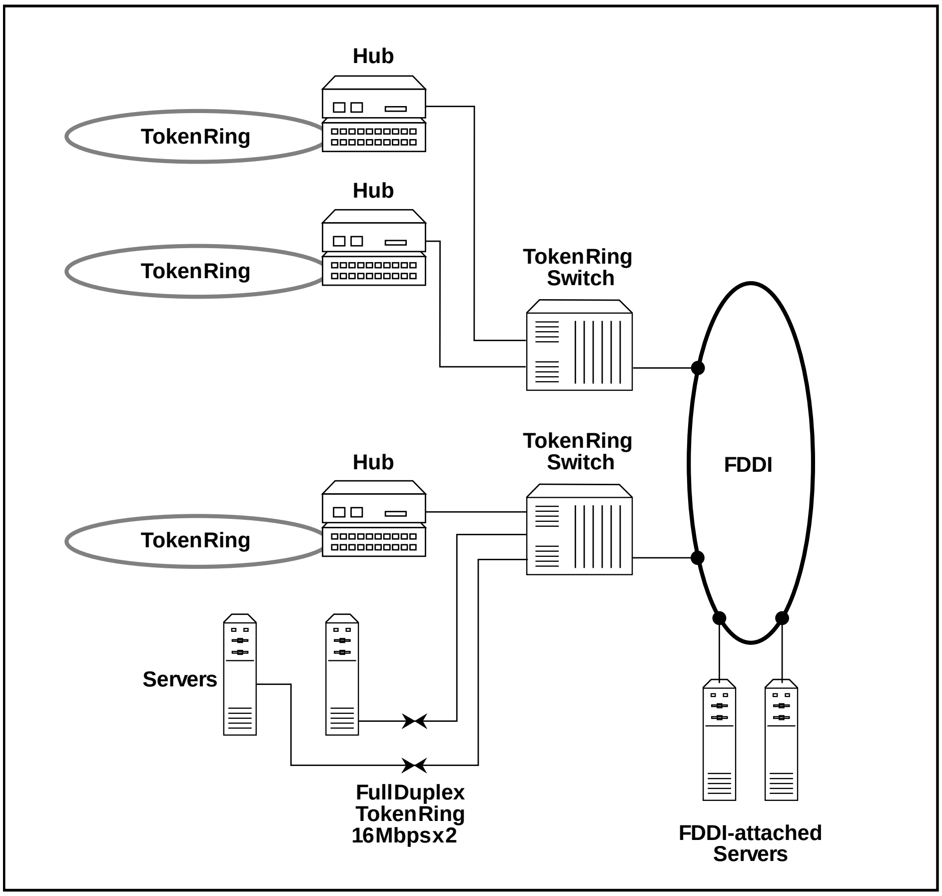

Figure 5: Token Ring Switches With FDDI Interconnection

Compatibility With Existing Token Ring Infrastructure

Madge's design objectives in developing Token Ring switching technology have been defined entirely by detailed customer feedback. The result is a solution that enhances the performance and utility of an existing Token Ring infrastructure with the very minimum of change.

Madge's Token Ring switch will comply with IEEE 802.5 standards for Token Ring, and will be fully compatible with IBM Token Ring networks. Existing Token Ring hubs, wiring concentrators, lobe cabling and adapter cards can all be left in place - only the trunk cabling needs to be re-routed to connect into ports on the switch.

The standards for full duplex Token Ring are only now emerging, but any changes to the switching technology needed to maintain compliance with the standards can be made entirely in software. Users of Madge's Smart 16/4 Token Ring adapters have excellent investment protection, in that the adapter cards they are purchasing today will be software upgradable to operate in full duplex mode in the future. To take advantage of full duplex mode, users only have to connect their servers directly to ports on the switch as if it were a conventional hub, and then simply install the full duplex driver software on the server.

Path to the Future

Token Ring switching offers a new way for network administrators to interconnect multiple Token Ring segments, with greater performance, and at lower cost.

The very low latency achievable with Token Ring switching makes it possible for network administrators to locate servers and other network resources wherever is most convenient within the local network infrastructure, without worrying about the performance implications of the number of hops between workgroups and servers. And by supporting the definition of virtual LANs, switching allows logical connections to be set up between workgroup rings and servers throughout the infrastructure. These logical connections offer a considerable degree of network security, while at the same time confining broadcast traffic and protecting against broadcast storms.

The low latency of switching provides an additional benefit: it prepares the LAN infrastructure for video. The store and forward operation of bridges and routers plays havoc with picture quality because of the variable delays introduced to video packet streams. A switch-based infrastructure introduces no such delays.

Switching provides a highly affordable incremental upgrade that greatly enhances the value of existing investments in Token Ring. More than this, switching provides a portal to the future. By opening up the Token Ring infrastructure to future integration with ATM, switching delivers the best from today's technology while offering a stepping stone to tomorrow's.

Inside Back Cover

Blank

Amsterdam
Atlanta
Bangkok
Berlin
Boston
Brussels
Cape Town
Chicago
Cologne
Copenhagen
Dallas
Denver
Detroit
Frankfurt
Hong Kong
Jakarta
Johannesburg
Kuala Lumpur
London
Los Angeles
Madrid
Minneapolis
Munich
Nashville
New York
Oslo
Paris
Philadelphia
San Jose
Seattle
Singapore
Stockholm
Sydney
Tokyo
Toronto
Vancouver
Washington DC

Madge Europe
Madge Networks Ltd
Loudwater
High Wycombe
Bucks HP10 9QZ
England
Tel: +44 1628 858000
Fax: +44 1628 858011

Madge Americas
Madge Networks Inc
2310 North First Street
San Jose
Calif. 95131-1011
United States
Tel: +1 408 955 0700
Fax: +1 408 955 0970

Madge Asia
Madge International Ltd
64-01 Central Plaza
18 Harbour Road
Wanchai
Hong Kong
Tel: +852 2593 9888
Fax: +852 2519 8022

Madge Japan
Madge Japan KK
Believe Mita
43-16 Shiba 3-chome
Minato-ku, Tokyo 105
Japan
Tel: +81 3 5232 3281
Fax: +81 3 5232 3208

Madge Networks reserves the right to change specifications without notice.

980-276-01 11/94