Cut-through Token Ring Switching

A Detailed Study of Cut-through Switching on Token Ring Networks

Author, Simon C. Gawne

May 1995

Executive Summary

The emergence of Token Ring switching as a viable LAN internetworking technology is set to revolutionize the way organisations build enterprise networks. Used as a collapsed backbone, Token Ring switches provide a very high speed, low latency connection between multiple ring segments, which dramatically improves the performance of the whole network and allows network services to be centralized. This addendum to Madge's Token Ring Switching White Paper is aimed at technical personnel who require more detailed information about Token Ring switching. It contains detailed technical information about the operation of cut-through switching in Token Ring networks, and covers various issues including the effect of token rotation time on the switch latency and the use of token priority to improve switch network access.

Contents

Overview of Token Ring Switching

Traditional LAN Internetworking Devices

Traditional routers and bridges are known as *store-and-forward* devices, which means a frame has to be completely received from the network before it can be processed to determine its destination. Receiving the whole frame introduces a delay, called *latency*, which can have a serious impact on network performance. A formal definition of latency can be found in appendix A.

The total latency introduced by a store-and-forward router or bridge depends on a number of factors, including:

- the time taken to receive the whole frame
- the time taken to process the frame to determine its destination
- any time involved in copying data across the backplane between ports
- the transmit set-up time on the destination port

The latency of a store-and-forward device varies with frame size. The *minimum* time to receive a 4Kbyte frame from the network is 2,100 microseconds (see appendix B). Hence, the minimum latency which can be achieved in a store-and-forward device carrying standard 4Kbyte data frames is 2,100 microseconds.

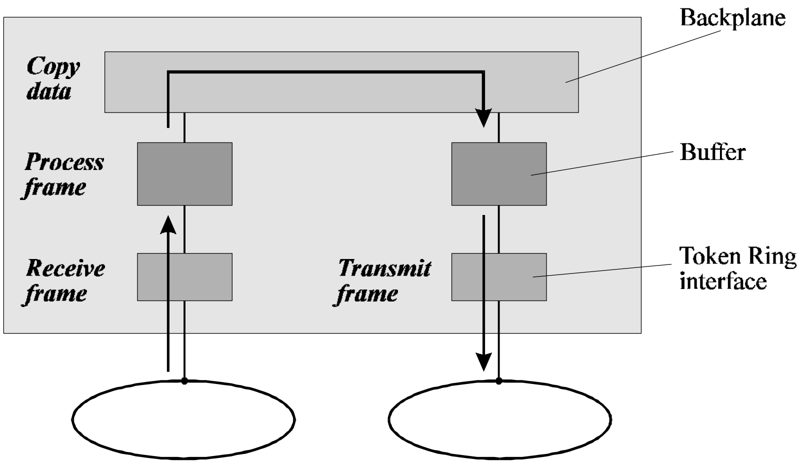

Figure 1: Processes involved in moving data between rings in a store-and-forward device

It should be noted that some vendors claim very low latencies of under 300 microseconds in their store-and-forward devices. However, these times *only* include the time to process the frame, and ignore the time to receive the frame from the network. Once this time is taken into account, the total latency can be well above 2,500 microseconds.

Accessing a file or application stored on a file server may involve many thousands of data frames, so the latency introduced by the internetworking device can have a significant impact on network performance.

Cut-through Switching

Token Ring switching provides a very high speed method of interconnecting multiple ring segments which is both simpler and more cost effective than traditional LAN internetworking devices, such as store-and-forward routers.

Instead of reading the entire frame into memory before making a decision about where to forward it, a switch will take action as soon as the first 20-30 bytes of the frame have been received. Information in the frame header is analyzed and the destination port deduced almost instantaneously. A connection is then made between the input and output ports, and the frame starts transmitting onto the destination ring while it is still being received on the source ring. This high speed frame switching technique is known as *cut-through* or *on-the-fly* switching.

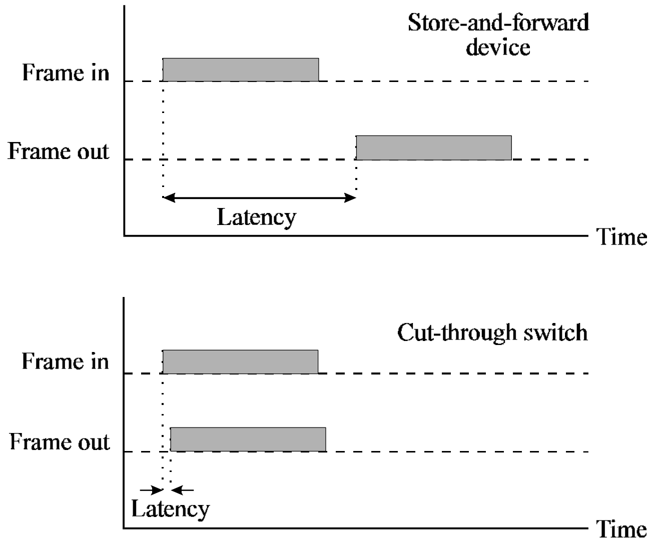

Figure 2: Comparison of latencies for store-and-forward devices and cut-through switches

The total delay or latency introduced by the switch can be as low as 30 microseconds, or 1% of an equivalent store-and-forward device. Dramatically reducing latency in this way allows clients on one ring to communicate with servers on another ring with almost the same performance as if they were both attached to the same ring.

Store-and-forward Switches

LAN internetworking devices have started to appear on the market which use the same layer 2 forwarding techniques as cut-through switches, but are based on store-and-forward designs. These devices are often known as store-and-forward or "buffered" switches.

Architecturally, these devices are multiport bridges and suffer from the same high latency as traditional routers and bridges. When used in real Token Ring networks, these devices deliver only very limited performance improvements and users will see a degradation in performance when communicating with servers located on different rings.

Cut-through in Real Token Ring Environments

Large enterprise Token Ring networks generally support a diverse set of applications and protocols, including TCP/IP, IPX, NetBIOS and SNA. These protocols typically generate two types of frames:

- *small frames:* generally used for data acknowledgements and protocol-specific functions
- *large frames:* used for carrying actual data across the network

Many protocols are able to make use of Token Ring's support for large frames to reduce the number of frames involved in a specific data exchange (16Mbps Token Ring can support up to 17Kbyte frames). For instance, Novell NetWare can support 4Kbyte data frames, and NetBIOS can support even larger data frames. In contrast, Ethernet has a maximum frame size of only 1.5Kbytes. Consequently, the relationship between latency and frame size is particularly important in Token Ring environments.

Cut-through switching delivers major performance benefits over store-and-forward devices for large frames. This is because the time to receive a frame from the network increases with the size of the frame. The latency of a store-and-forward LAN internetworking device, which has to receive the whole frame from the network before processing it, increases with frame size: a store-and-forward device will take significantly longer to move large data frames between rings than small frames. In contrast, because a cut-through switch can start streaming data onto the destination ring almost instantaneously, it can forward both large and small frames with the same low latency.

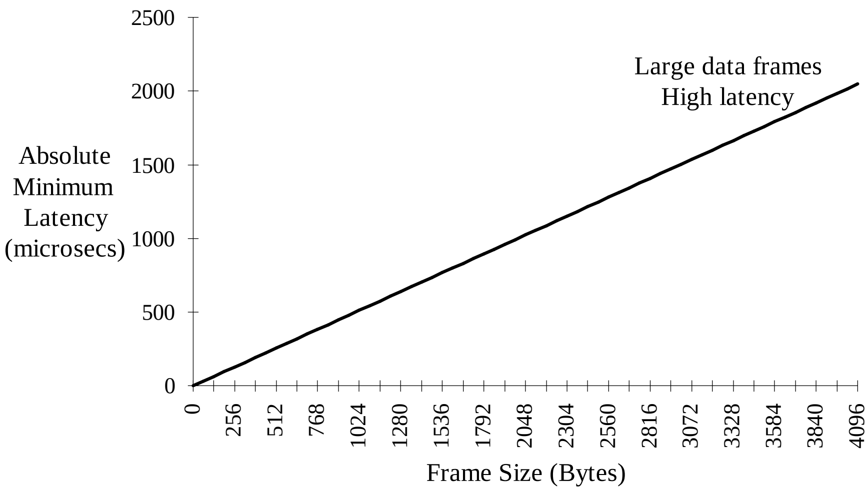

Figure 3: Variation of absolute minimum latency with frame size for a store-and-forward LAN internetworking device

The graph below shows the distribution of frames on a typical Token Ring network. As can be seen, the majority of frames are small frames although there is a local peak around 4Kbytes for large data frames.

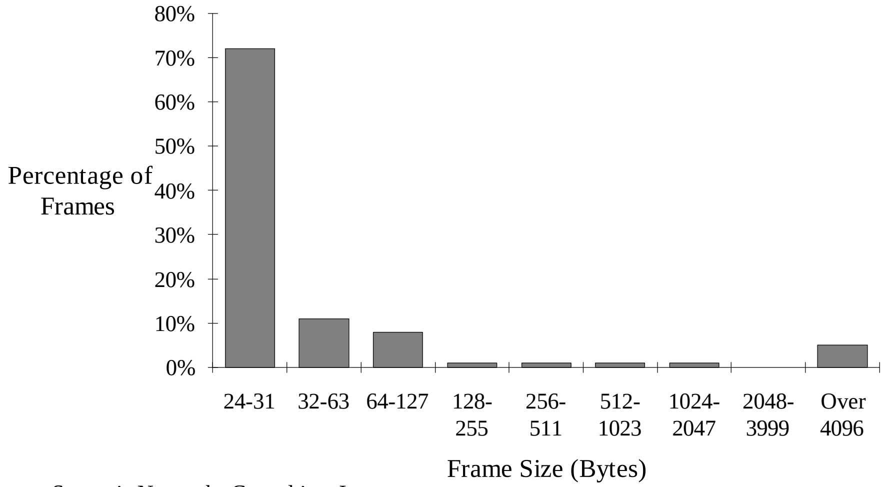

Source: Strategic Networks Consulting, Inc.

Figure 4: Typical frame size distribution on Token Ring

From figure 4, it is possible to determine the proportion of network bandwidth used by the different types of frames. As can be seen below, over 80% of the available bandwidth is used by large data frames i.e. the vast majority of the data on the network is carried by large frame. Consequently, the latency of the LAN internetworking devices has a major influence on the overall performance of the network: high latency store-and-forward switches, bridges and routers significantly degrade network response times. Cut-through switching delivers major benefits to these real-life Token Ring networks.

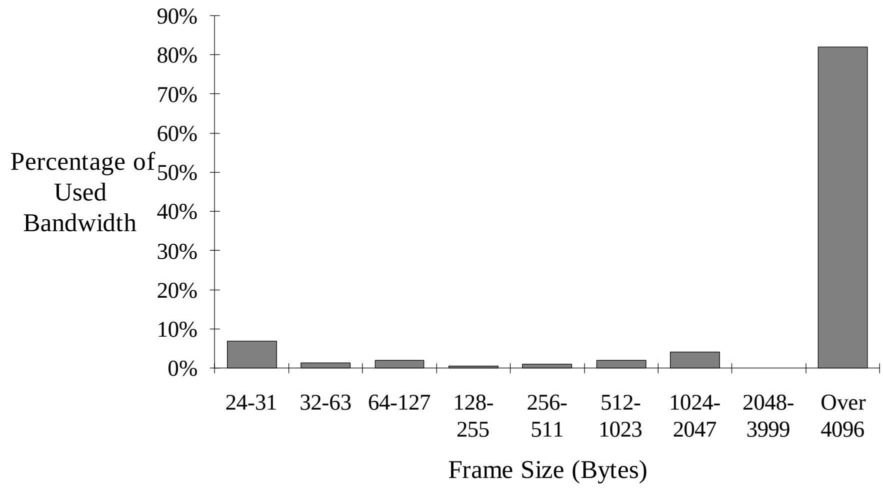

Figure 5: Percentage of bandwidth used by different sizes of frames

Performance Comparisons

The very low latency delivered by cut-through switches is of particular importance in environments where time-sensitive protocols and applications are used.

Latency Sensitive Protocols

A good example of a protocol which is sensitive to latency is the Novell NetWare Core Protocol. During a typical data exchange between a client and a file server, each data frame sent across the network is acknowledged by the end station before the next frame is sent. Hence, when the file server sends the first block of data, it waits for the client to acknowledge that it has received it before transmitting the second block. This type of data exchange is often called a "ping-pong" protocol.

The longer it takes for the end-station to acknowledge the receipt of each frame, the longer the complete data transfer takes. If a store-and-forward internetworking device is placed between the client and server, the increased delay decreases the total number of protocol exchanges in a give time period and reduces the total throughput by as much as 35-50%.

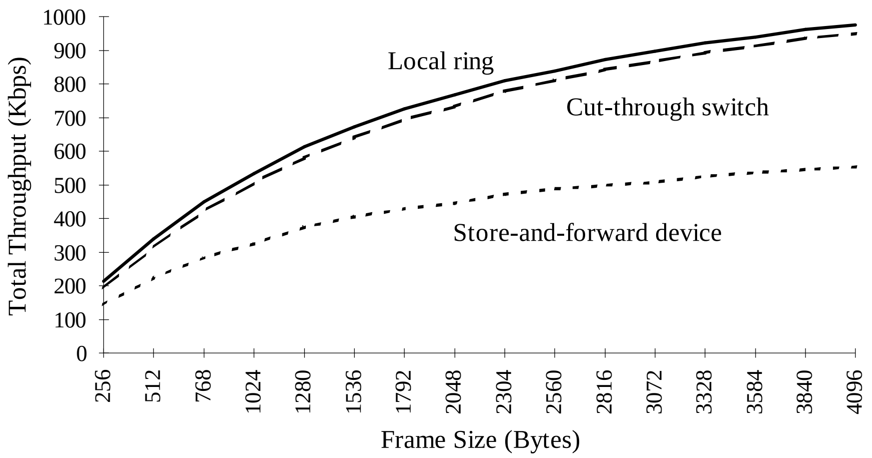

Figure 6: Effect of latency on Novell IPX when communicating across a cut-through switch and a store-and-forward device

In contrast, with a cut-through switch the added delay is insignificant and there is very little degradation in network performance. Consequently, cut-through switching enables network servers to be moved from individual workgroup rings to a central location on the network without affecting network performance.

"Windowed" Protocols

Windowed protocols, such as NetBIOS, are designed to reduce the effect of network latency on application performance. This is achieved by allowing several data frames to be transmitted in series before an acknowledgement is required from the receiving end-station. The number of frames which can be transmitted in series is called the "window size".

However, even with a windowed protocol like NetBIOS, introducing a store-and-forward device can reduce the total network throughput by as much as 15-20%. Cut-through switching is still the best LAN internetworking technology where windowed protocols are used.

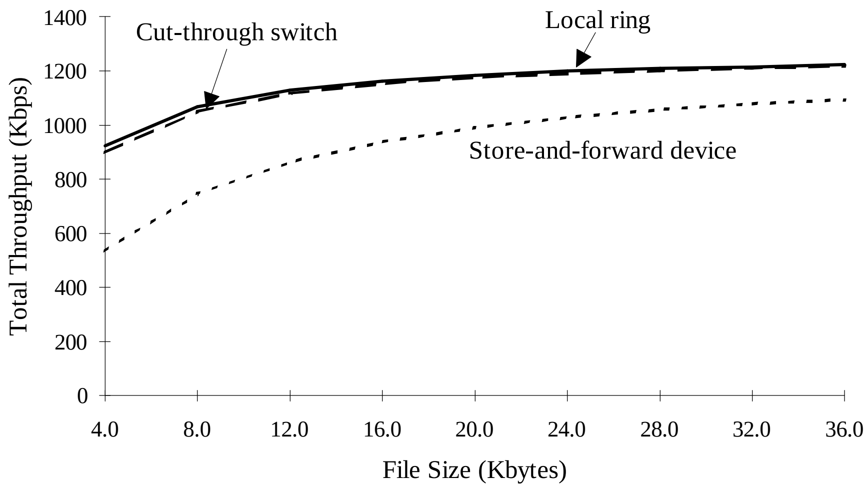

Figure 7: Effect of latency on throughput for windowing protocols across cut-through switches and store-and-forward devices

Multimedia Applications

Other examples of applications which are very sensitive to network latency include video and voice applications. Any large or random delays introduced by LAN internetworking equipment can seriously impact the quality and performance of such applications. Because a cut-through switch streams data frames between rings with very low latency, there are no variable delays introduced between the data frames, and high quality video applications can be supported.

Propagation of Corrupt Frames

Bad or corrupt frames, which contain invalid data bits, occur on networks from time to time. They are caused by a variety of factors, including electromagnetic interference along the network cable from external radiation sources. One of the main criticisms levelled at cut-through switches, particularly in Ethernet networks, is that they propagate corrupt frames. This is because a cut-through switch starts transmitting a frame on the destination port before it has completed receiving it and had chance to check its integrity.

(Corrupt frames should not be confused with the standard MAC beacon and soft error frames found on Token Ring. These frames are part of the IEEE 802.5 error recovery mechanisms and are an integral part of Token Ring. Switches, like layer 2 bridges and layer 3 routers, do not propagate MAC frames between rings.)

This section examines the mechanisms built into Token Ring for handling corrupt frames and discusses the effect of corrupt frame propagation in a switched Token Ring environment.

CRC Checking

A Cyclic Redundancy Check (CRC) is a number located at the end of a frame which is used to check the integrity of the data in the frame. If the CRC does not match the data in the frame, the frame is corrupt. The receiving end-station uses this information to determine whether to discard the frame.

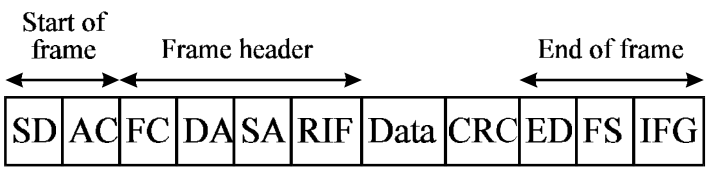

SD = Starting Delimiter (1 byte)
AC = Access Control (1 byte)
FC = Frame Control (1 byte)
DA = Destination Address (6 bytes)
SA = Source Address (6 bytes)
RIF = Routing Information Field (0-32 bytes)
CRC = Cyclical Redundancy Check (4 bytes)
ED = Ending Delimiter (1 byte)
FS = Frame Status (1 byte)
IFG = Interframe Gap

Figure 8: Token Ring frame format

The use of the CRC prevents corrupt data being received from the network into the end-station's memory.

CRCs and Switching

Because store-and-forward switches, bridges and routers receive a complete frame from the network before processing it, they have the opportunity to check the frame's CRC. If the frame is corrupt (invalid CRC), the store-and-forward device can prevent further propagation across the LAN by discarding the frame.

Although cut-through switches will propagate bad frames between rings, the effect on the network and the attached end-stations is negligible because they are extremely rare on Token Ring networks and the end-station discards them anyway. To understand this in more detail, consider a frame being transmitted from client A on ring 1 to server B on ring 2 across a LAN internetworking device. Immediately after transmission the frame is corrupted. The events that follow depend on the type of internetworking device used:

- Store-and-forward switch, bridge or router connecting ring 1 to ring 2: the frame will be analysed by the internetworking device and the invalid CRC identified. The frame is discarded. After some time period (dependent on the protocol used), station A will detect that it has not received a response from B and will resend the frame. This time there are no problems with the frame and it will be sent to ring 2. The communication process will then continue as normal (see figure 6).
- Cut-through switch connecting ring 1 to ring 2: the frame will be forwarded by the switch to ring 2, where the adapter card in end-station B receives it, identifies it has an invalid CRC and discards it (see figure 7). No data is corrupted in the end-station's memory. Next, either end-station B will immediately request a re-transmission from A or, after some time period, end-station A will detect that it has not received a response from B and will resend the frame. Either way, communication automatically recovers.

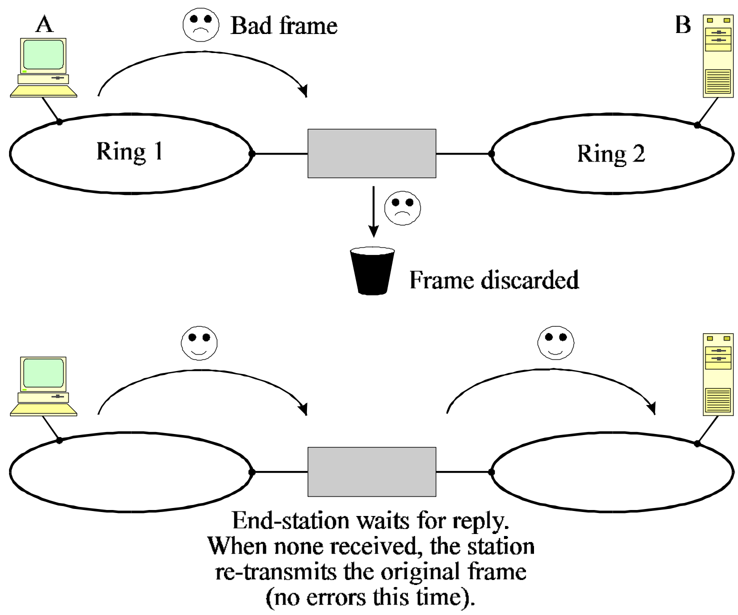

Figure 9: Error control in store-and-forward environments

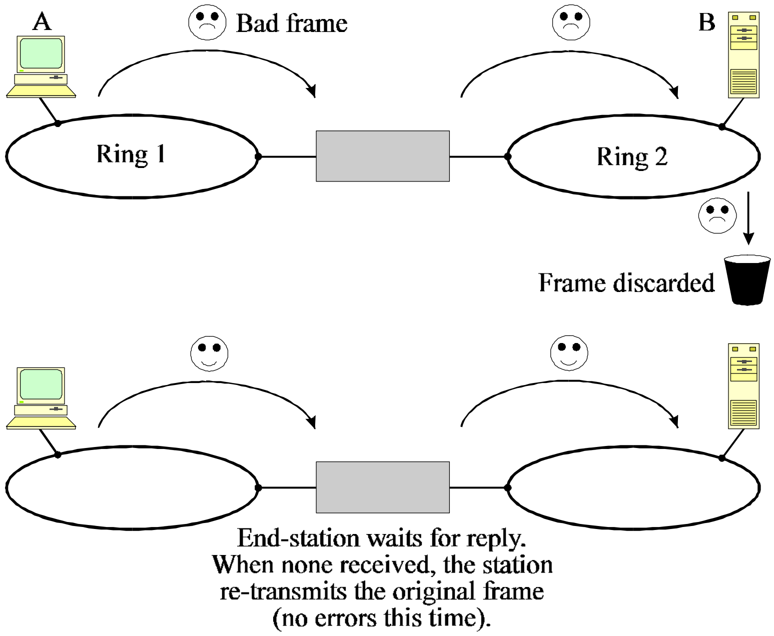

Figure 10: Error control in cut-through environments

Consequently, even though cut-through switches propagate bad frames, no data is corrupted in the end station's memory and the communication process recovers in an identical way to a store-and-forward environment.

Errors on Ethernet and Token Ring

The propagation of corrupt frames by cut-through switches has been a particular problem on Ethernet networks. This is because Ethernet is a collision-based technology: if two stations try to transmit data onto the network at the same time, a collision error is generated which causes both stations to stop transmitting and wait for a random length of time before trying again. Because of the random wait time built into this process, it is not possible to determine precisely which stations will have access to the network at any given time. This issue becomes critical in a heavily loaded network and is responsible for Ethernet's low 4 to 5 Mbps maximum throughput in such networks.

Problems arise in busy cut-through switching environments with large Ethernet segments. Consider a station transmitting data onto an Ethernet segment which is received by an Ethernet cut-through switch. Once the switch has determined the destination of the frame, the switch will start streaming data onto the destination Ethernet segment. If another station on the source segment starts transmitting, then a collision error is generated and both stations stop transmitting. The cut-through switch stops receiving data and an aborted frame, often called a *collision fragment*, is generated on the destination segment. The busier the network, the more fragments which are generated. These fragments can seriously disrupt the network and reduce the available bandwidth for normal data traffic.

The situation on Token Ring is significantly different. Token Ring is based on a token passing technology which precisely defines how and when end-station gain access to the network. At the heart of Token Ring is the token: a three byte frame which is passed from station to station around the ring. Before any end station can transmit data onto the network, it must take possession of the token. Consequently, it is not possible for multiple end-stations to transmit data onto the same ring at the same time, and therefore Token Ring networks do not suffer from the frame fragmentation problem described above.

In addition, because Token Ring has very sophisticated PHY and MAC management functions, and generally operates across high grades of cable, such as Category 5 UTP and STP, the number of network errors is extremely low. The IEEE 802.5 Token Ring standard specifies a *worst case* allowable station frame error rate of 10^{-5} or one frame in 100,000. Actual Token Ring stations have much lower error rates. As a result, corrupt frames appear very infrequently on Token Ring networks and cut-through switching technology is extremely well suited to Token Ring environments.

Hybrid Schemes

Some vendors have proposed hybrid schemes in which the cut-through switch monitors the level of corrupt frames being received on each port, and if the level exceeds a user-defined threshold the switch automatically changes to slower store-and-forward mode, allowing frames to be discarded.

Madge believes such schemes are not necessary on Token Ring for the reasons described above. In particular, because the error rate on Token Ring is so insignificant, these hybrid switches will very rarely enter store-and-forward mode on the majority of Token Rings and users will never see this feature in use.

In addition, these schemes can cause problems if particular types of network error do occur. For example, in the event that a faulty adapter card attaches to the network and starts sending bad frames, the hybrid switch will detect invalid frames on one port and will put the switch into store-and-forward mode. As a result, all frames coming from that ring will be slowed down, all users will see a degradation in network performance, and the organisation loses the benefits which made it choose cut-through switching in the first place.

Summary

Although cut-through switches propagate corrupt frames between rings, the effect on the network and the attached end-stations is negligible. Because Token Ring operates across high grades of cable and has very sophisticated PHY and MAC functions, the frame error rate on Token Ring networks is extremely low and errors rarely occur. In the unlikely event that a frame is corrupted, the receiving adapter card will identify that it has an invalid CRC and will discard the frame. None of the data in the host system is affected, and the communication protocol will automatically recover in the same way to a store-and-forward environment.

Token Rotation

In Token Ring networks, because a station has to wait for the token before it can transmit data, the time it takes for a token to travel round the ring strongly influences the rate at which stations can transmit data. This wait time is particularly important for cut-through Token Ring switches, since data cannot be streamed onto the destination port until the switch captures the token.

Token Rotation and Token Wait Times

The time taken for the token to travel completely round a ring depends on two factors: the number of stations on the ring, and the total cable length. Appendix C describes in detail how to calculate the token rotation time. The average time a station must wait before it can claim the token is half the token rotation time. The graphs below illustrate the effect of ring size (number of stations) and average lobe length on the average token wait time. As can be seen, the typical token wait time varies between 10 and 40 microseconds.

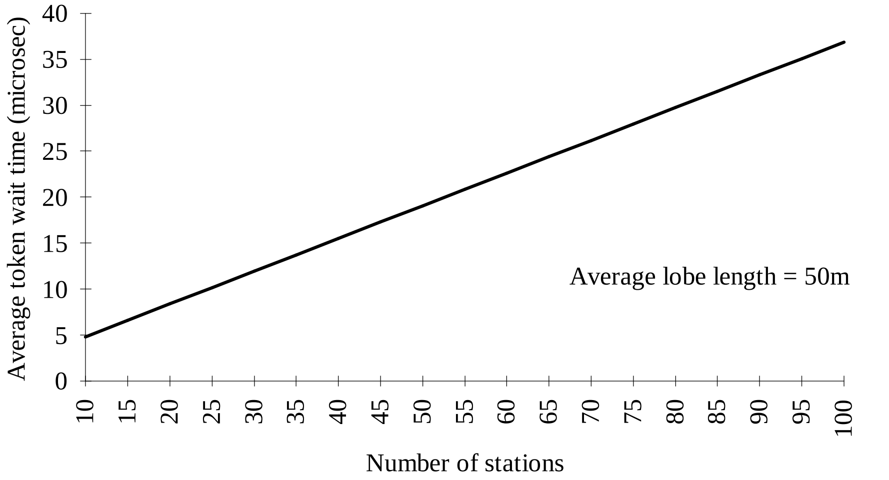

Figure 11: Relationship between token wait time and number of stations per ring

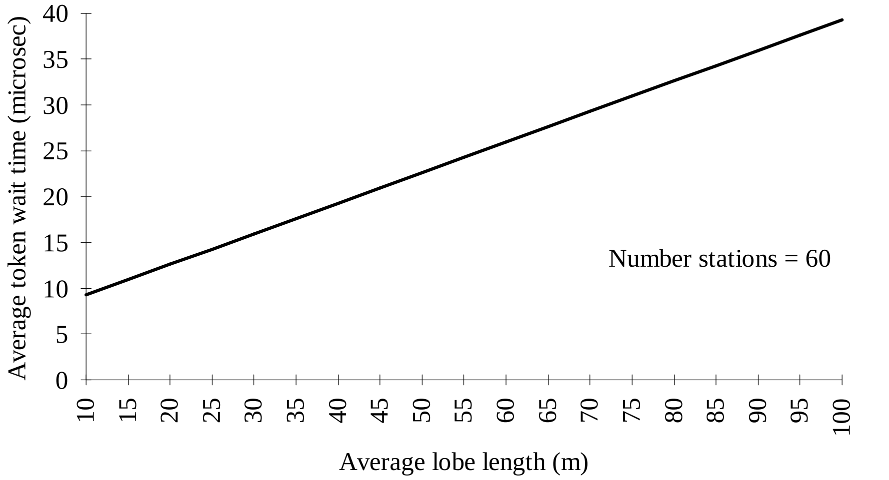

Figure 12: Relationship between token wait time and average lobe length

Token Rotation and Cut-through Switches

It has been argued that it is not possible to perform cut-through switching on Token Ring because the switch has to wait for the token and therefore buffer (store in memory) the frame. As can be seen, the average wait time for the token on a typical workgroup ring is around 20 microseconds. Hence, the average latency of a cut-through switch is:

time to process frame	30 microseconds
average token wait time	20 microseconds
total delay	50 microseconds

For a store-and-forward switch, bridge or router, the latency for 4Kbyte data frames exceeds 2,000 microseconds. Cut-through switches still achieve far lower latencies compared to store-and-forward devices even when the token rotation time is taken into account.

Token Priority

A criticism often levelled at Ethernet switches is that they do not perform well in heavily loaded network environments where there is a lot of traffic passing between clients on one segment and servers on other segments across the switch. This is because the switch port has the same access rights to the available bandwidth as all the clients on the workgroup segment but has a disproportionate amount of data to send.

For example, consider 30 clients located on one segment logging into a server across a switch. Each client will request data from the server and the switch will carry this data from the server segment to the client segment. At any instant, the switch may have, say, 10 frames queued ready for transmission onto the workgroup segment (and there will be 10 clients waiting for the data). However, because the switch has the same access rights as any of the normal end-stations, if the network is busy and the remaining 20 stations also want to transmit data, the switch will receive only 1/21 of the network access time and it will take considerable time for it to transmit all its queued frames.

Token Ring is able to avoid this problem by the use of a mechanism called *token priority*, which is built into the IEEE 802.5 Token Ring standard and supported by all token ring adapters.

Priority Values

Token priority can take a value between 0 and 7. Priority 0 is the lowest priority and is used by network stations for normal data transmissions; priority 7 is the highest priority and is reserved for the low level MAC frames which are used to keep the ring operating. Priority level 4 has been reserved by IEEE 802.5 for internetworking devices, such as Token Ring switches, which require greater access rights to the network than normal end-stations.

Figure 8 on page 13 shows the format of a standard Token Ring frame. Near the start of the frame is an 8 bit field called the Access Control (AC) field which is used to control the priority of the frame.

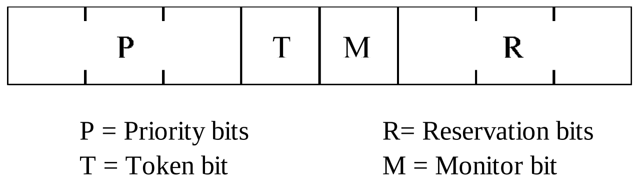

Figure 13: Access Control (AC) field with priority and reservation bits

There are two priority values in the AC field. The *Priority Bits* (P-field) indicate the priority of a token and therefore directly controls when stations are able to use the token. The *Reservation Bits* (R-field) allow stations with high priority to request that the next token be issued at the requested priority.

Operation of Token Priority

To understand how token priority works on a real Token Ring network, consider a ring with 4 stations on it as shown in figure 14.

Station A starts to transmit a normal data frame onto the ring with a P-field set to 0 and an R-field set to 0. The adjacent station B has frame queued at priority level 1. The frame travels from A to B, B checks the R-field and sees that it is lower than its own priority. It therefore sets the R-field to 1 to indicate that it would like the next available token.

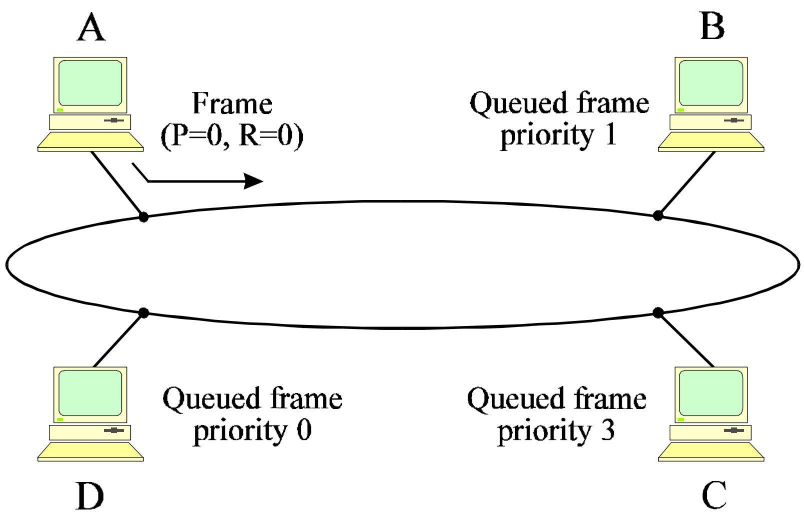

Figure 14: Example ring where stations are using token priority

The frame carries on to station C, which also has a frame queued but this time at an even higher priority level 3. C also checks the R-field, sees that the value of 1 is lower, and raises the R-field to 3. The next station D has a frame queued with priority 0. Since this is less than the frame's R-field, D does nothing.

After travelling round the ring, the frame eventually reaches station A again. A strips the frame from the network and generates a token with a P-field set to 3 (obtained from the R-field). The R-field is reset to 0.

The token travels round the ring to station B. Because the P-field is higher than B's priority value, B cannot claim the token. B again sets the R-field to 1.

The token then travels to station C. The P-field has the same value as C's priority, so C is able to convert the token to a data frame (by resetting the T-bit in the AC-field) and attaching its data. After the frame has travelled around the whole ring back to C, C removes the frame and generates a token with *unchanged* P- and R-fields (set to 3 and 1 respectively).

The only station which can reduce the priority P-field in a token is the station which raised it in the first place. This ensures on a busy network that low priority stations still gain access to network resources (although at a reduced rate) even when there are multiple stations transmitting at higher priority. In the example here, the station which raised the value of the P-field is A. So when the token from C reaches station A, A generates a new token with a P-field of 1 (obtained from the R-field) and an R-field reset to 0. When the new token reaches B, B is able to claim it (its priority is the same as the token's P-field) and transmit data onto the network.

Token Priority and Switches

Token Ring switches can operate at priority level 4, which ensures that they have greater access to network resources. For example, on a very busy ring with many stations and a single switch interface, the switch can have access to network resources 50% of the time. This allows it to efficiently carry data from centralized network services, such as file servers and mainframes, to the ring segments requesting it.

Early Token Release

Early Token Release (ETR) is an option which has been added to the Token Ring standard which further improves the performance of a Token Ring network. With ETR, once a station has transmitted a frame it can generate a free token *before* the preceding frame has travelled round the ring and been stripped off the network by the originating station. ETR increases the available bandwidth for data transmissions and minimizes the dead-time of the network.

ETR can influence the operation of token priority on large rings carrying small frames. This is because the priority of the token (P-field) is determined from the R-field in the preceding frame. If the token is transmitted before the start of the frame has arrived back at the station, there is no way to determine the correct priority and the token is transmitted with a default P-field of 0. With large frames this is not a problem because the frame can travel round the ring and start to be stripped from the network while the station is still transmitting the end of the frame. The effect only occurs if the frame is shorter than the total ring length.

On a typical workgroup ring with 50 stations and an average lobe length of 50m, the largest frame size which does not fully support token priority with ETR is 76 bytes (see appendix C). The vast majority of networks comprise a mixture of large data frames (between 1 and 4 Kbytes) and small data requests-acknowledgements. Consequently, token priority and ETR operate perfectly well together on Token Ring networks.

Summary

Token priority enables important network components, such as Token Ring switches, to have greater access rights to network resources than standard end-stations. This allows them to efficiently carry data from centralized network services, such as file servers and mainframes, to the ring segments requesting it. Token priority is part of the IEEE 802.5 standard and is fully compatible with all Token Ring adapters.

Conclusion

Token Ring switching is revolutionising the way organisations build enterprise networks. With the application of high speed cut-through switching technology, Token Ring switches provide very low latency data transfers between ring segments which not only reduce network response times, but allow organisations to centralize and consolidate network services, and prepare their networks for future LAN applications.

Cut-through switching is particularly important in Token Ring environments because the vast majority of the data is carried by large frames: large frames take significantly longer to traverse a store-and-forward device than a cut-through switch. In addition, because of the operation of token priority, cut-through switches are able to efficiently carry data from centralized network services to the ring segments requesting it.

Although cut-through switches propagate corrupt or invalid frames, the effect on the network and the attached end-stations is insignificant. This is because the frame error rate on Token Ring networks is extremely low, and on the rare occasions when errors do occur the end-stations discard them anyway.

Cut-through Token Ring switching provides an excellent solution to today's network performance problems which fully protects investment made in existing Token Ring equipment and extends the life of Token Ring networks well into the future. With Token Ring switching, organisations have the perfect platform to migrate to ATM as and when required.

Appendix A

Definition of Latency

The term *latency* is used to describe the delay introduced into traffic flow across the network by the individual network components. Latency is particularly important for backbone internetworking devices, since it directly affects the transit times of frames passing between rings and can have a significant effect on the overall performance of the network.

There are two definitions of latency which are applied to LAN internetworking devices. ***Cut-through latency*** is defined as the time difference between the start of the frame arriving at the internetworking device and the commencement of data transmission on the destination LAN segment. It is a true measurement of the transit time across the LAN internetworking device. Cut-through latency is best illustrated in the diagram below, as it appears in store-and-forward devices and cut-through switches:

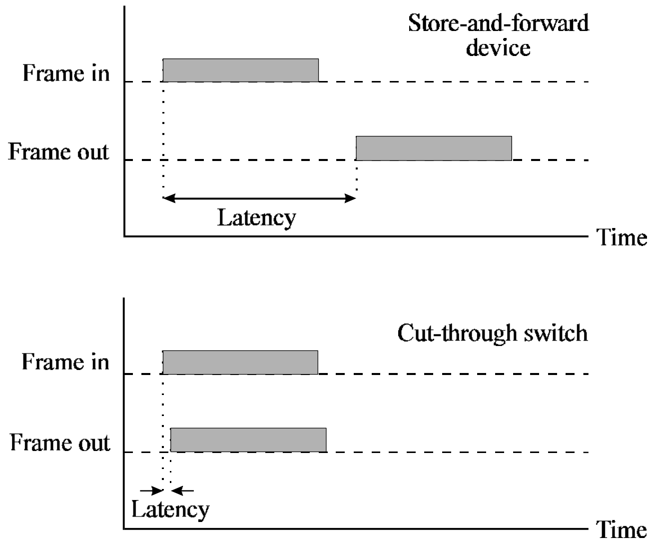

Figure 15: Definition of cut-through latency

As can be seen, because the store-and-forward device has to receive the complete frame before processing it, the minimum latency of a store-and-forward device is equal to the time taken to receive the entire frame from the network. In the case of a standard 4Kbyte data frame, this is 2,100 microseconds. Most store-and-forward devices have latencies well in excess of this.

Store-and-forward latency is defined as the time difference between the end of the frame arriving at the internetworking device and the commencement of data transmission on the destination LAN segment. This is shown in the diagram below.

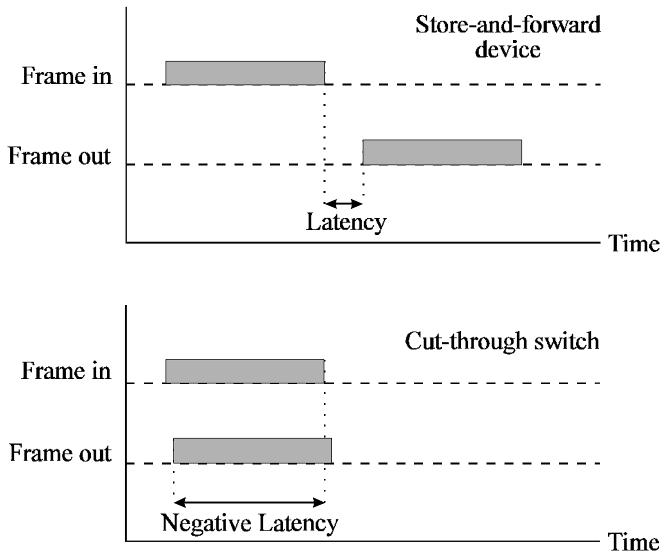

Figure 16: Definition of store-and-forward latency

This definition of latency is really a measurement of how long the store-and-forward internetworking device takes to the process a frame, and does not take into account the time required to receive the frame from the network. Because the frame processing time is often independent of the size of the frame, this definition of latency is particularly useful for store-and-forward device since it typically yields a single number for all frame sizes.

For cut-through switches, the "store-and-forward" latency is actually *negative* since the frame is transmitted onto the destination segment before the end of the frame is received.

It should be noted that some vendors mix the usage of the term latency depending on the type of device being referred to. This is dangerous, since one definition of latency refers to the transit time across the device, whereas the other refers to the frame processing time. For the purposes of this white paper, all references to latency refer to cut-through latency.

Appendix B

Calculating the frame receive time

The time to receive a typical 4Kbyte data frame on a 16Mbps Token Ring is calculated as follows:

$$4\text{Kbytes} = 4096 \times 8 \text{ bits}$$

$$16\text{Mbps} = 16\times10^6 \text{ bits per sec}$$

$$\text{time to receive frame} = \frac{4096 \times 8}{16\times10^6}$$

$$= 0.0021\text{sec or } 2{,}100 \text{ microseconds}$$

This is the length of time a store-and-forward LAN internetworking device takes to receive a data frame from the network. This is the theoretical minimum latency of a store-and-forward internetworking device forwarding 4Kbyte data frames.

Appendix C

Calculating the token rotation time

In order to calculate the total time for a token to travel round a ring, the following factors need to be taken into account:

- each network station delays the token by approximately 2.5 bits (0.156 microsecs)
- there is a 8.9 bit delay (0.556 microsecs) for each 50m of lobe cable or 100m of trunk cable
- the active monitor adds a average delay of 5 bytes or 2.5 microsecs (this includes the elastic buffer)

Hence, the total token rotation time t_r is:

$$t_r = 0.156n + 0.556\Sigma l_{lobe}/50 + 0.556\Sigma l_{trunk}/100 + 2.5$$

where n=number of stations, l_{lobe} = lobe lengths, l_{trunk} = trunk lengths

For a typical Token Ring workgroup (no trunk connections), this can be written more simply as follows:

$$t_r = 0.156n + 0.556nl_{av}/50 + 2.5$$

where l_{av} = average lobe length.

The average time a station has to wait for the token (ignoring Early Token Release) is

$$t_{av} = t_r/2$$

For an average workgroup ring with 50 stations and an average lobe length of 50m, the token rotation time is 30 microseconds i.e. the average wait time is 15 microseconds.

Calculating the ring length

The ring length (in bits) is calculated as follows:

$$t_r = 2.5n + 8.9\Sigma l_{lobe}/50 + 8.9\Sigma l_{trunk}/100 + 40$$

For a typical workgroup ring comprising 50 stations with an average lobe length of 50m, the ring length is

$$t_r = 610 \text{ bits or } 76 \text{ bytes}$$

Amsterdam

Atlanta

Bangkok

Berlin

Boston

Brussels

Cape Town

Chicago

Cologne

Copenhagen

Dallas

Denver

Detroit

Frankfurt

Hamburg

Hong Kong

Jakarta

Johannesburg

Kuala Lumpur

London

Los Angeles

Madrid

Minneapolis

Munich

Nashville

New York

Ontario

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.

980-283-01 7/95