Incorporating FDDI MAC Level Bridging

National Semiconductor
Application Note 740
David Brief
February 1991

INTRODUCTION

A natural application for FDDI is as the backbone LAN in large installations with several FDDI networks and many segmented lower speed LANs. The interconnection of these networks is typically accomplished with MAC Layer Bridges and Network Layer Routers.

This Application Note focuses on the unique aspects of MAC layer bridging for FDDI and suggests ways of incorporating these functions into products using the National Semiconductor DP83200 FDDI chipset. The DP83200 chipset provides many capabilities that ease the job of incorporating MAC Layer bridging functions into products. It is likely that MAC level bridging capabilities will be resident in products together with Network Layer Routing Functions, FDDI Concentrator Functions and/or End Station capabilities.

TABLE OF CONTENTS

1.0 MAC LEVEL BRIDGING CONSIDERATIONS ON FDDI
- 1.1 Bridging Protocols
- 1.2 Filtering and Forwarding
- 1.3 Setting and Interpreting the Control Indicators
- 1.4 Stripping Transmitted Frames
- 1.5 Bridgeable Frame

2.0 DP83200 CHIPSET BRIDGING FEATURES
- 2.1 Source Address Transparency
- 2.2 Stripping Mechanism
- 2.3 FCS Transparency
- 2.4 External Matching Interface
- 2.5 Data Structures

3.0 IMPLEMENTING TRANSPARENT BRIDGING
- 3.1 Bridge Block Diagram
- 3.2 Address Matching Alternatives
- 3.3 Developing the Addressing Topology
- 3.4 Void Stripping and Learning
- 3.5 Special Control Indicator Processing

4.0 TRANSPARENT BRIDGING IMPLICATIONS FOR END STATIONS
- 4.1 Indicator Setting
- 4.2 Interpretation of Control Indicators

5.0 IMPLEMENTING SOURCE ROUTING BRIDGING
- 5.1 Address Filter
- 5.2 Discovery Process
- 5.3 Forwarding
- 5.4 End to End FCS Checking

6.0 SOURCE ROUTING BRIDGING IMPLICATIONS FOR END STATIONS

1.0 MAC LEVEL BRIDGING CONSIDERATIONS ON FDDI

The presence of bridges in an FDDI network should enhance the connectivity and performance of the overall network. This should in turn be reflected in the performance seen by end stations.

MAC layer bridging introduces several interesting quirks into the FDDI protocols. These are largely because many of the MAC layer bridging protocols were developed in environments where frame stripping was implicit (In the Ethernet bus topology, packets just disappear and in Token ring there initially was no early token release) and control indicators were not used or were not present.

There are also several options for implementing MAC layer Bridging. In this section, the issues involved with MAC Layer Bridging are discussed in general terms. In subsequent sections, the implications of these topics when using Transparent and Source Routing Bridging is discussed in further detail.

A single bridging method is not defined in FDDI since it is possible to use the same bridging protocols as in the IEEE802 LANs. However, in the areas that relate to interoperability and operation of the FDDI protocols, compromises have been reached on how to run these bridging protocols in FDDI. These compromises are captured in the ANSI X3T9.5 MAC-2 Draft Standard.

1.1 Bridging Protocols

There are two categories of MAC level bridging that are typically performed. They are named according to how much knowledge, responsibility and control an end station (e.g., workstation) has over the routing of the frame through the network. In transparent bridging, the end station has little if any control over the routing of the frame to its destination and is free from all responsibility and work of determining the route to the end station. In Source Routing bridging, as its name implies, the source is responsible for routing a frame to its ultimate destination. To end stations, the bridging function that is accomplished somewhere in the path between the two protocol endpoints appears as part of the network. This bridging function may be accomplished by one or more stations that implement the bridging function. These stations are generically referred to as bridges even though those stations do more than just the bridging function. For example, every station that implements a bridging function also acts as an addressable end station.

With Transparent bridging, the bridges are responsible for maintaining all of the routing information. The most popular, but not the only form of transparent bridging uses what is commonly known as the Spanning Tree protocol for determining how to route frames to their destination. The goal of the Spanning Tree protocol is to create a single path to connect all stations in the connected network at any given time. The Spanning Tree protocol allows extensions for multiple paths between segmented LANs. These paths can be used as backup links or alternatively to provide a form of load sharing. The load sharing is possible across different dialogues, but not across a single dialogue.

Transparent bridging protocols reduce the complexity of the end station while increasing the complexity required by the bridge function to retain the illusion of end station transparency.

With Source Routing, the job of the bridges is simplified, but the end stations must determine and include explicit routing information in every frame that they transmit. Since a station typically communicates with a small number of stations, and the routing to these stations is relatively static, this does not typically represent a large overhead for end stations.

The IEEE802.1d committee reached a compromise in order to allow Source Routing and Transparent bridging to work in the same extended LAN. The committee has specified a Source Routing-Transparent (SRT) bridge. With this compromise, end stations that participate in source routing protocols also can communicate with stations using transparent bridging provided an SRT bridge is present. For End Stations capable of source routing, this effectively makes all stations connected through the Spanning Tree appear as if they are in the local Source Routing Domain. There is little affect on bridges and end-stations that perform only transparent bridging. For more information on SRT bridging see the recent IEEE802.1d document.

1.2 Filtering and Forwarding

The decision to forward a frame from one port of a bridge to another is based on the addressing information contained in the frame. The address is used as an input to an address filter which contains the necessary information to determine if a frame should be copied and forwarded or rejected and filtered.

Transparent bridging uses the Destination Address field of an FDDI frame to determine if a frame should be forwarded. In a multi-port bridge, the Address filter might also indicate to which port to forward the frame. The set of addresses for which frames are forwarded is configured in the address filter by a learning process. In the future, information contained in the SMT Management Information Base might be used to load this Address filter when all information necessary to develop and predict ring addressing topologies is present through standard management services.

Transparent Bridging might also be implemented with a partial filter where some or all of the frames are copied. Further address filtering then indicates whether or not the frame should be forwarded.

In Source Routing Bridging, the decision to copy a frame for forwarding is based on its 16-bit ring address being listed in the routing information field contained within the INFO field of the FDDI frame. The presence of this routing field is indicated when the most significant bit of the SA field is 1. End stations participating in the Source Routing protocol are only required to recognize their own address as the destination and are not required to process the routing information field for purposes of forwarding.

1.3 Setting and Interpreting the Control Indicators

The meaning of the Address Recognized and Frame Copied control indicators (A_Ind and C_Ind) is impacted in an extended LAN. After a long and sometimes heated debate over the meaning of this frame status information, the FDDI Committee reached agreement on a compromise solution.

The A Indicator, when set, indicates that a precise address match occurred. The C Indicator, when set, indicates that the frame was copied successfully by a station and that either the frame will be delivered to the correct protocol endpoint or an error will be indicated back to the sending station.

Source routing bridges set both the A and C Indicators if it recognizes and copies the frame for forwarding just as if its address had appeared in the DA field. End stations can therefore assume that if transmitted frames return with both the A and C Indicator set, that the frame is being forwarded to its ultimate destination. If a transmitted frame returns with the A Indicator set but the C Indicator as reset, then the station might assume that buffer congestion has occurred somewhere between the source and the destination buffer. The end station could use this as an indication to stop transmitting any more frames because there is likely buffer congestion somewhere in the path and additional frames that are transmitted would probably not be copied anyway, and would waste system and ring bandwidth. The exact interpretation and the recovery implied is dependent on the protocols used. Protocols that indicate back pressure in an extended LAN are still largely proprietary.

The interpretation of the control indicators when Transparent bridging is present is slightly more complicated. A transparent bridge will only set the A_Ind if the frame is explicitly addressed to the bridge as before. However, for frames that are copied with the intent to forward the C_Ind may optionally be set. This capability is indicated in the FDDI SMT MIB.

If a transmitted frame returns with A_Ind = R and C_Ind = S, an end station can then assume that the frame has been copied by a bridge that intends to forward the frame to its ultimate destination. If a transmitted frame returns with A_Ind = R and C_Ind = R, an end station can either assume that the frame was not forwarded by a transparent bridge if it knows that only bridges that implement the option are in the route to the ultimate destination. This can be used to stop transmission of additional frames until the buffer congestion condition subsides. If bridges that do not implement the C setting option are in the route to the ultimate destination, then the station can not assume anything about the forwarding of the frame (as in Ethernet).

If a transmitted frame returns with A_Ind = S, an end station can assume that the frame was addressed to a station on the local ring. The interpretation of the C Indicator depends on the option implemented by the addressed station. An optional status clearing capability may be implemented in stations. In this case, a bridge that has not yet learned of a station on the ring can copy a frame for forwarding, before the destination receives the frame. When a frame is received by a status clearing station with the received A_Ind = R and C_Ind = S, a station implementing the option will change the C_Ind to reset if it cannot copy the frame. The DP83266 MACSI™ and DP83261 BMAC™ devices implement this option. This capability preserves the meaning of the A_Ind =S, C_Ind = R combination on a local ring, and the transmitting station then knows that the destination did not copy the frame. In this way the presence of transparent bridges implementing the first option will not destroy any status used in optimizing the performance on a local ring.

The implications of the compromise that was reached for transparent bridging is that end stations wishing to take advantage of the control indicators are required to keep status information about each station they are communicating with in order to determine how to interpret the control indicators.

1.4 Stripping Transmitted Frames

In FDDI, the MAC is responsible for stripping every frame that it transmits into the ring. In the case of both Source Routing and Transparent bridging, the SA and DA fields of the frame contain the original source and ultimate destination stations for the frame. Therefore, when a bridge forwards a frame onto an FDDI ring, it is not possible to use the SA field to recognize frames to be stripped. For this reason, a bridge is required to implement an alternative stripping mechanism. Other stations may also find it useful to implement an alternative stripping mechanism.

The MAC-2 standard does not specify a single stripping mechanism, but rather suggests examples of a number of different stripping mechanisms that can be implemented within the bridge station. Any method which is interoperable with the rest of the FDDI Standards and meets several general criteria may be employed.

One property of the ring that is exploited to implement alternate stripping mechanisms is the property that all frames transmitted will return to the station that transmitted them before any frame transmitted by any other station is received. This allows the station to use a special frame to mark the end of one or more frames that are transmitted during a service opportunity (while holding the token). If this method is used, the MAC-2 Draft standard suggests the use of a special type of Void frame that contains this stations DA and SA, this is called a My__Void Frame. This is differentiated from regular Void frames by the non-Null DA. It is differentiated from Void frames from other stations (Other__Voids) by the presence of this stations SA field. By transmitting a My__Void frame before a token is issued, and stripping until it returns, all frames transmitted by this station are removed from the ring in the forward mode of operation.

Errors in this process, such as the creation on corruption of a My__Void frame, result in either Understripping, where not enough frames stripped, or Overstripping, where additional frames are stripped unnecessarily.

Understripping is undesirable because it causes duplicate packets to be delivered to stations on the ring. Any algorithm that stops stripping upon receipt of a valid token cannot have a probability of understripping less than that of a token being erroneously created from a frame. This requires the conversion of the FC field to that of a token and the conversion of two consecutive data or Idle symbols to T symbols. Thus, a limiting factor to the probability of understripping is approximately:

$$1.25 \times 10^{-22}$$

In order to prevent understripping more than once in 100 years, the probability of understripping should be:

$$1.25 \times 10^{-22} \leq \text{probability (understripping)} \leq 1.3 \times 10^{-14}$$

Overstripping is also detrimental to ring performance. It causes frame loss that would otherwise not have occurred. Assuming any other factors which might cause frame loss (e.g., receiver congestion) are eliminated, the frame loss rate is bounded by the link bit error rate.

In order that overstripping not cause an excessive increase in total frame loss rate, it should not cause the loss rate to increase more than 10% above the loss rate due to link bit errors alone. Therefore:

$$\frac{\text{probability (frame loss due to overstripping)}}{\text{probability (frame loss due to link errors)}} < 0.1$$

Several mechanisms exist that might be used to meet these frame stripping requirements. A combination of mechanisms may be required to achieve adequate robustness in the stripping algorithm. Some examples, but not all, of mechanisms that might be used as components of stripping algorithms include:

(a) Transmitting one or more My__Void frames before issuing the token to mark the end of a burst of transmitted frames, and stripping until a My__Void frame is received. (The MACSI and BMAC devices transmit two My__Void frames when the stripping option is invoked)

(b) Stop stripping upon receipt of valid Tokens, other Void frames or MAC frames (e.g., Beacon or Claim frames). (The MACSI and BMAC devices implement all of these options)

(c) Counting the number of frames transmitted and the number of frames stripped. (The MACSI and BMAC devices do not implement this option)

To ensure interoperability of different stripping methods, and to minimize under stripping or over stripping, the bridging appendix in the MAC-2 Draft Standard recommends that a bridge:

(a) Transmit at least one My__Void frame immediately before issuing a token.

(b) Stop stripping when it receives its own My__Void frame or any valid Void frame with SA other than its individual address.

In addition, an alternative stripping method is required to stop stripping when it receives a valid token or clears Ring__Operational.

Some stations, including stations using the DP83266 MACSI device, are implemented in such a way that a shared resource (e.g., bus or memory bandwidth) is used at some point in a frame. To allow design of these stations without undesirable performance implications, the MAC-2 Draft standard requires stripping to begin no later than the seventh symbol after the end of the SA field.

In the MACSI device, the copy decision can begin as early as the 4th byte of the INFO field. Since all fragments should be less than this size (they should have at most 6 symbols in the INFO field), fragments will not result in any wasted bus and memory bandwidth. Longer fragments caused by stripping errors with more than 4 bytes in the INFO field might be copied by the MACSI device if an address match occurs. The termination status reported with every copied frame by the MACSI device indicates these frames as being stripped frames.

For a frame whose SA is the MAC's individual address, it is only necessary to strip based on the recognition of the SA field.

1.5 Bridgeable Frame

The FDDI frame format defines the first byte of every frame transmitted as the Frame Control Field. The Frame Control Field is used to determine the Protocol Endpoint such as MAC, SMT, LLC Asynch, LLC Synch, Implementer, Reserved. Currently, only LLC Asynch frames are bridgeable. LLC Synch, Implementer and Reserved frames are beyond the scope of the standard. MAC and SMT frames must never be forwarded from one ring to another (this would cause serious problems for MAC and SMT protocols). Only LLC Asynch Frames are intended to be bridged.

The use of restricted dialogues that traverse more than one ring are also beyond the scope of the Standard.

Any LLC Asynch frame with at least 4 data bytes after the SA is considered bridgeable. In a worst case condition, many short frames with minimal preamble can be received. Using long addresses with 4 bytes of Info and B bytes of preamble, this corresponds to over 400,000 frames per second!

Large frames might also present difficulties for bridges. In Basic mode FDDI has a maximum frame size of 4500 bytes and in Hybrid mode FDDI has a maximum frame size of 8600 bytes. Without use of a segmentation/reassembly protocol, it is not possible to bridge frames that are larger than the size permitted on a LAN. For example an 8600 byte from a ring operating in Hybrid mode could not be forwarded into a ring operating in Basic mode. However, a 4500 byte frame could be forwarded in either direction. Similar problems appear when forwarding frames between Ethernet and FDDI.

2.0 DP83200 CHIPSET BRIDGING FEATURES

The DP83200 FDDI chipset incorporates several features to ease implementing bridging functions. These capabilities make possible simple and efficient bridge designs. The features are covered below followed by a description of how they are used in Source Routing and Transparent Bridges and End Stations.

2.1 Source Address Transparency

Normally the Source Address is transmitted from the MAC parameter RAM as added protection since the SA is used to strip frames. The SA transparency capability allows the SA to be transmitted from the Data Stream as opposed to from the parameter RAM. This is important since bridges work with the original source address on frames that are to be forwarded. This is different from the bridge station's address.

In the MACSI device SA transparency is a channel parameter. Every frame transmitted on a channel either uses SA transparency or does not depending on the current programming of the channel configuration register. This register can be modified reliably between requests, not on a per frame basis. An implementation using the MACSI device might transmit all of the bridged traffic on one channel and all of the local traffic on the other. (The BMAC device can handle changes on a per frame basis.)

The BMAC device also supports the capability to provide separate transparency control over the routing information indicator bit of the SA. For historical reasons the BMAC device signal is called SAIGT since its routing information indicator falls in the same relative location as the individual/group within the DA field.

2.2 Stripping Mechanism

In FDDI, every station needs to strip every frame it transmits. Typically this is accomplished by stripping based only on the Source Address. However, in bridging applications where the SA is not necessarily of this station, this station needs to either watch out for all of its frames (and use CAM technology to assist the strip decision) or use some other mechanism. The MAC-2 Draft Standard states that the strip points of all frames is before the fourth byte of the INFO field. That implies that any fragment with more than three bytes of INFO is an illegal fragment.

The National stripping mechanism accomplishes the stripping by sending out two My__Voids before the token and stripping everything until the first My__Void returns. The second My__Void is stripped on the basis of its SA, as in the normal stripping mechanism.

2.2.1 Algorithm

- Stripping of frames with SA received = MSA or MLA still occurs
- If the stripping option is requested by any Frame of a Service Opportunity (i.e., while a token is held)
  - at the end of the Service Opportunity (before the token is issued) two My__Void frames are transmitted.
  - Stripping continues until:
    - a valid My__Void frame is received (normal case)
    - a valid Other__Void frame is received
    - a valid token is received
    - a MAC frame (other than My__Claim) is received

      Stripping does not stop on receipt of a My__Claim in order to allow removal of all My__Claim fragments that would otherwise be present after the Claim token process is won by this station.
    - a MAC Reset occurs
- Void Frames

  Three types of Void frames are used in this algorithm, the Void frame, the My__Void frame and the Other__Void frame. During normal operation Void frames are transmitted when L__Max expires and when frames are aborted. My__Void frames are transmitted at the end of a service opportunity when stripping has been requested for any frames of that service opportunity. Other Void frames are Void frames transmitted by some other station on the ring.

  The My__Void Frame provides a convenient marker to delimit stripping. The My__Void frame is used to distinguish it from the Void frame.

  Void frames are used within a service opportunity in order to make sure that valid frames are sent every F__MAX and to limit the maximum preamble size. This typically occurs after a frame is aborted or after the token has been captured when a frame is not ready to be transmitted.
  - My__Void frame
    - FC = Void
    - DA = MSA or MLA (MSA used if enabled)
    - SA = MSA or MLA (MSA used if enabled)
    - FCS
    - ED, FS
  - Void frame
    - FC = Void
    - DA = Null
    - SA = MSA or MLA (MSA used if enabled)
    - FCS
    - ED, FS
  - Other__Void frame (Detection Criteria)
    - FC = Void
    - DA = any
    - SA = not ((MSA and MSA enable) or (MLA and MLA enabled))
    - FCS
    - ED, FS

2.2.2 Robustness

- Two My_Void frames are used to improve the robustness over using just one My_Void Frame
  - Overstripping
    - Overstripping will only occur when both My_Void frames are corrupted or lost
      - even in this case, overstripping will be limited by Other_Void frames which could occur before during or after the next station's burst of frames.
    - assume
      - that noise in receiver major course of error this implies that there is no correlation of errors across frames.
      - maximum ring load
      - 100 stations
      - 100 byte frames
      - 1 frame per station per token
      - 10 e^{-10} link error rate
      - with one My_Void frame
        - overstripping occurs once every 13 minutes
      - with two My_Void frames
        - overstripping occurs once every week
- Understripping Probability:
  - this would only occur if
  - some other frame turns into a well formed My_Void or Other_Void frame
  - a My_Void or Other_Void is transmitted from the data interface.

2.2.3 Affects on Synchronous Allocation

The bandwidth used by the two My_Void frames is taken from the synchronous bandwidth allocation and effectively reduces the maximum allocatable synchronous bandwidth by the time it takes to transmit the two My_Void frames. The My_Voids can be sent with the short address to minimize the use of ring bandwidth. With the MACSI and BMAC devices, if short addresses are enabled, the MACSI and BMAC devices will transmit My_Void frames with the short address.

The bandwidth required from the synchronous pool is independent of the number of stations on a ring using the Void Stripping option and is dependent only on the maximum number of My_Void frames that can be transmitted before the token is issued. If all bridges in the ring utilize the MACSI or BMAC device, then only 34 byte times of the synchronous allocations needs to be used.

In a ring where both Synchronous and Asynchronous services are being used, if THT expires while transmitting a frame, two My_Voids will then be transmitted followed by the token. At the next station and all downstream stations THT will have already expired and only synchronous requests may be serviced. In this way the token will go all the way around the ring until it reaches the station after the station where THT expired during the frame transmission. All other stations will have the opportunity to send all of their synchronous traffic. In the case where all of the stations using synchronous bandwidth line up perfectly and use all of the synchronous bandwidth on one token rotation, the recovery required conditions will still not become true at any station provided that the synchronous bandwidth was not overallocated.

2.2.4 Why Not Use a Frame Counter?

Our approach is easier to implement and less likely to overstrip than a stripping mechanism that uses a frame counter, in the presence of line errors.

The first My_Void frame provides equal protection with or without a frame counter. However, a frame counter will miss on a hit in any data frame that causes that frame to be lost. A second Void frame will only be missed on a hit on the Void frame itself. Since the length of the Void frame is significantly less than the total length of the transmitted data frame(s), it is significantly more likely that the counter will miss a data frame due to noise than that a Void frame will be missed due to noise. This implies that a frame counter approach is more likely to cause understripping than an approach using a second My_Void.

For additional protection, stripping is stopped on receipt of any well-formed Void frame from another station, which minimizes interference with other implementations and avoids cascaded overstripping when multiple stations are stripping and multiple My_Void frames from one station are lost. With the frame count approach, if both data frame(s) and the single My_Void frame are lost (much more probable than losing a pair of My_Void frames), overstripping could eat up another station's My_Void frame, as well as its data frame(s). This can cause cascaded overstripping unless additional logic is added to selectively strip frames based on their FC values. Other implementations may use some kind of selective strip filter logic; however, this results in fragments on the ring, whereas our strip mechanism can be used to clean up all garbage except multiple valid MAC or well-formed void frames (the first such frame is stripped), or tokens (which are stripped but regenerated).

2.3 FCS Transparency

FCS Transparency determines if the FCS calculated by the CRC Generator in the MAC transmitter will be appended to the frame. This allows diagnostic testing of the CRC checker and generator. This option is also used to perform end to end FCS checking. This could be used in FDDI to FDDI transparent bridges, but is not useful in transparent bridges between FDDI and Ethernet because the FC field and bit ordering of the SA and DA change the value of the FCS.

2.4 External Matching Interface

The MACSI device accepts several inputs from an external address filter in order to cause frames to be copied based on external address matches and also to affect the setting of the control indicators and incrementing of the frame counters.

The External A flag (EA) signal to the MACSI device causes the A Indicator to be set and the frame copied count to be incremented if the frame is successfully copied. The External Copy (ECOPY) signal to the MACSI device causes the frame to be copied when asserted from the beginning of the frame until ECIP is deasserted.

The External M_flag signal to the MACSI device causes the current frame to be stripped. This is useful if an alternative to the resident Void stripping algorithm is used.

An implementation of an address filter would typically use the interface between the MACSI and PLAYER+™ devices. This address filter would detect a starting delimiter and then sequence into the frame.

2.5 Data Structures

The data structures used by the MACSI have been optimized to allow efficient transfer of data. The Output and Input data structures are symmetric and the data and descriptors are segregated.

3.0 IMPLEMENTING TRANSPARENT BRIDGING

Bridging that appears transparent to end stations can be considered the most difficult to implement. In this section topics relating specifically to transparent bridging are discussed. In subsequent sections, new topics relating to end stations and source routing bridges are discussed.

3.1 Bridge Block Diagram

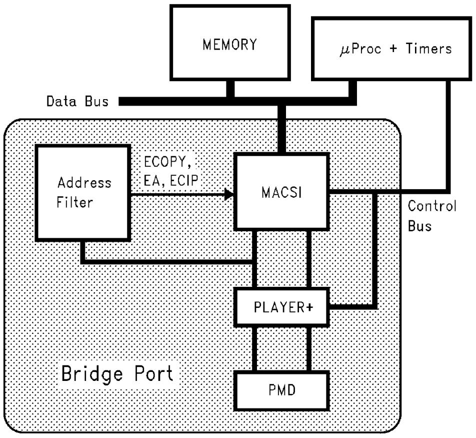

TL/F/11119–1

FIGURE 1. Bridge Block Diagram

The block diagram in *Figure 1* shows a generic architecture for a multi-port bridge, highlighting the requirements for a single port in the bridge.

3.2 Address Matching Alternatives

In any bridging application, additional address recognition logic is required. This may range from totally promiscuous copying with filtering only in Software, to combinations between Hardware and Software, to complete filtering in real time with dedicated Hardware.

The advantage of doing the address matching in real time in dedicated hardware is to be able to set the Address recognized indicator accurately (in explicit bridges) and to set the Frame Copied indicator appropriately (when the intent is to forward the frame).

The address matching logic, the Address Filter, can either be implemented in very wide CAMs (to handle 48-bit addresses) or can be implemented in sophisticated SRAM lookup tables. One way of using an SRAM to get a reasonable address density is to use each successive byte or nibble as an index into the RAM implementing a TRIE structure (M-ary tree) in hardware. Each successive byte or nibble is combined with data from the previous cycle and used together as part of the address for the next cycle.

Bridges that can handle 400,000 frames per second containing over 4K addresses are becoming available.

As an alternative to Forwarding Bridges where for each received frame, the bridge determines which port, if any, to forward it to, there are also Filtering Bridges. Filtering Bridges reject frames with certain addresses only and accept all others. For example, it is reasonable to reject all frames with an address on this ring (in the ring map), and copy all others. This approach is particularly appropriate in designs with two ports since there is only one place that the frame can be forwarded to. In bridges with more than two ports, it is appropriate if the forwarding decision is made using additional software support.

Similar logic is used to provide additional Group Addressing Capabilities and possibly Individual Addressing Aliasing.

However the address filter is implemented, the Ring Engine relies on the External A flag signal (EA) in order to set the transmitted A Indicator. Similarly, the system interface portion of the MACSI device expects the External Copy signal (ECOPY) combined with a latching signal, ECIP, External Compare In Progress. When ECIP is deasserted ECOPY is sampled to determine if the frame should be copied by the MACSI device. The MACSI device will copy the frame to either Channel 1 or Channel 2 depending on its current configuration.

3.3 Developing the Addressing Topology

In order to correctly determine which frames should be copied and forwarded and which frames should be rejected and filtered, it is necessary to develop an accurate addressing topology. Once the addressing topology is established, it can be used to determine the location of every address and the route (one or more) to every address from any other address. In Spanning Tree transparent bridging there is one route between any address pair, and in most implementations just the next hop needs to be known, not the entire route.

The addressing topology can be developed in several ways. One way, popularized by Ethernet networks where there is no alternative, is to effectively learn the topology by watching the traffic. By seeing what frames come from where, it is possible to deduce what side of the bridge or to which port frames are connected.

Once the addressing topology is developed, it is used by the address filter to determine which frames to copy (and which port to forward them to) and which frames to reject.

The addressing topology must be maintained using some kind of aging or refreshing process. Changes in the addressing topology are relatively slow, except in the presence of configuration changes. In a configuration change (such as a wrap or an additional wrap), a small adjustment can have a large affect on routing between address pairs.

3.4 Void Stripping and Learning

When using Void Stripping, it is desirable to inhibit copying of frames transmitted by this station. A learning algorithm that monitors frames that are received to determine their source would get confused if transmitted frames are also received. In order to implement this function, additional logic is necessary between the BMAC and BSI-2™ devices to monitor when a frame is transmitted with the strip option. The MACSI device provides an output pin,(LEARN) that replaces this logic.

In order to inhibit copying, a Stripping Flag is suggested. While this flag is set, no frames should be copied. The STRIPPING flag is set when a STRIP is asserted and TXPASS = 0 and TXRDY = 0. The STRIPPING flag is cleared when one of the following conditions occurs:

a valid My__Void frame is received [normal case]
a valid Other__Void frame is received
a valid token is received
TXRINGOP of the BMAC device is deasserted (this includes the cases where a MACRST occurs and MAC frames are received).

This requires a station to decipher the Void and Token Frame Control Fields and also monitor the BMAC device AFLAG and MFLAG signals to determine the Void frame type. Use of the BMAC device TXRINGOP is also necessary. The plan for the integrated solution is to provide this stripping flag as an output to be used in external copy decisions.

3.5 Special Control Indicator Processing

For Stations implementing transparent bridging special handling of the C Indicator is allowed by the MAC-2 Draft Standard.

The rules for setting the A and C control indicators are as shown below:

```
if the frame is addressed to this station then
    set A Indicator
    if frame copied successfully then set C Indicator
if the frame is not addressed to this station then
    leave A Indicator unaltered
    if A Indicator received as R and frame copied optionally set the C Indicator
```

For frames copied with the intent to forward for which a station address did not match, only the C Indicator should be set, not the A Indicator.

To accomplish this, since the current MACSI and BMAC devices will not set the C Indicator without setting the A Indicator as well, it is necessary to intercept the PHY Request byte stream between the MACSI or BMAC devices and the PLAYER+ device. Fortunately, the difference between the R and S symbol is only a single bit. Thus, when a frame is copied with the intent to forward it, the received A Indicator was not set, and the station's AFLAG is not set, the C Indicator must be changed from an R to an S. This occurs one byte time after EDRCVD is asserted and requires that PRD(0) be changed from a 0 to a 1.

This is only required in bridges implementing this option. In the integrated solution, an extra pin will be provided to help implement this function and remove the requirement for this small amount of external logic.

4.0 TRANSPARENT BRIDGING IMPLICATIONS FOR END STATIONS

Although the goal of transparent bridging is to remove all implication from end stations, there are a few areas that are affected as indicated below.

4.1 Indicator Setting

For End Stations and Bridges which also act as end stations, the MACSI and BMAC devices always implement the optional status clearing function defined in the MAC-2 Draft Standard.

Specifically, for frames received with A__Ind = R and C__Ind = S for which the station recognizes the address and intends to set the A Indicator but the frame was not successfully copied, the A Indicator is transmitted as S and the C Indicator is transmitted as R.

Support by the MACSI and BMAC devices of this function is in line with the desire to support all end station requirements very well and support other station types (bridges, routers, concentrators) with minimal additional logic.

5.0 IMPLEMENTING SOURCE ROUTING BRIDGING

The high level block diagram for a multi-port Source Routing Bridge is remarkably similar to that of the Transparent Bridge. The major differences come in the responsibilities and implementation of the Addressing Filter.

What is new with Source Routing Bridging is the presence of the Routing Information Field. The presence of the Routing Information Field is denoted when the most significant bit of the SA field is set to One. The Routing Information Field contains a string of 16-bit bridge numbers that the frame is to be routed to. These 16-bit bridge numbers are considered aliases of a certain bridge.

Source Routing Bridges fall under the category of explicit bridges used in the MAC-2 Draft Standard. In explicit bridges, the addresses that the bridge recognizes are considered its aliases, and the A Indicator is set for this class of address matches. (For this reason, in these implementations it is possible to connect the EA and ECOPY signals of the MACSI device together.)

5.1 Address Filter

The Address filter is much simpler in Source Routing bridges than in transparent bridges. The address filter is required to parse the Routing Information Field, when present, and look for this bridge's 16-bit ID number. If the number is present in the Routing Information Field then the frame is copied and the control indicators are set appropriately using the EA, ECOPY, and ECIP inputs to the MACSI device. The frame is then forwarded to the next destination in the list.

5.2 Discovery Process

In order for an end station to determine the route to another end station, the discovery process is necessary. In one variation of the discovery process, a station transmits out several all route frames. Each bridge then adds its bridge number into the Routing Information Field, until the addressed station has been reached. The addressed station then transmits back a frame according to the "best" route. This route is then used for future transmission.

5.3 Forwarding

Once the bridge's, bridge number is detected in the Routing Information Field, the frame can be forwarded to the next bridge number using the appropriate port (in a multi-port bridge).

Since all frames that are transmitted, must be stripped before the 7th symbol of the INFO field, the Void stripping still must be used, because it would not be tolerable to strip based on a bridge number later in the frame.

There are some subtle requirements placed on the transmission of frames in Source Routing Bridges. For example in some frames, the Most Significant Bit of the SA remains unchanged for some frames and is forced to zero in others. The user can meet these requirements using the Void Stripping (VST) and Source Address Transparency (SAT) options of the MACSI's request configuration registers in conjunction with the Bridge Option Select bit (BOSEL) in MCMR2. See the MACSI datasheet for more details.

5.4 End to End FCS Checking

Between FDDI and FDDI rings complete End to End FCS checking is possible. It may also be possible to provide this type of service between IEEE802.5 and FDDI. In any event, the FCST (FCS Transparency) input is used to control this option.

6.0 SOURCE ROUTING BRIDGING IMPLICATIONS FOR END STATIONS

There are a few requirements placed on End Stations that participate in Source Routing protocols. The end station maintains the responsibility for discovering the route to its peers using the All route frame. Once the route has been discovered, it must be used in all future correspondences as part of the Routing Information Field.

An End station has no requirements for any external address matching logic. End stations can use the ability to transmit only the MSB of the SA (the Routing Indication Indicator) from the data stream using the SAT bit to assert SAIGT and configuring BOSEL (BOSEL = 1) to choose VST not to assert STRIP, (VST = 0). The reason the rest of the SA can come from the Ring Engine is to insure proper stripping based only on the transmitted SA. This implies that in End Stations there is no need to use Void stripping.

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

National Semiconductor Corporation
2900 Semiconductor Drive
P.O. Box 58090
Santa Clara, CA 95052-8090
Tel: 1(800) 272-9959
TWX: (910) 339-9240

National Semiconductor GmbH
Livry-Gargan-Str. 10
D-82256 Fürstenfeldbruck
Germany
Tel: (81-41) 35-0
Telex: 527649
Fax: (81-41) 35-1

National Semiconductor Japan Ltd.
Sumitomo Chemical Engineering Center
Bldg. 7F
1-7-1, Nakase, Mihama-Ku
Chiba-City,
Ciba Prefecture 261
Tel: (043) 299-2300
Fax: (043) 299-2500

National Semiconductor Hong Kong Ltd.
13th Floor, Straight Block,
Ocean Centre, 5 Canton Rd.
Tsimshatsui, Kowloon
Hong Kong
Tel: (852) 2737-1600
Fax: (852) 2736-9960

National Semiconductores Do Brazil Ltda.
Rue Deputado Lacorda Franco
120-3A
Sao Paulo-SP
Brazil 05418-000
Tel: (55-11) 212-5066
Telex: 391-1131931 NSBR BR
Fax: (55-11) 212-1181

National Semiconductor (Australia) Pty, Ltd.
Building 16
Business Park Drive
Monash Business Park
Nottinghill, Melbourne
Victoria 3168 Australia
Tel: (3) 558-9999
Fax: (3) 558-9998

National does not assume any responsibility for use of any circuitry described, no circuit patent licenses are implied and National reserves the right at any time without notice to change said circuitry and specifications.