

INTERNATIONAL STANDARD

ISO/IEC
11518-6

First edition
1996-06-15

Information technology — High-Performance Parallel Interface —

Part 6:

Physical Switch Control (HIPPI-SC)

Technologies de l'information — Interface parallèle à haute performance —

Partie 6: Commande de commutation physique (HIPPI-SC)



Reference number
ISO/IEC 11518-6:1996(E)

Contents

	Page
Foreword.....	iii
Introduction.....	iv
1 Scope.....	1
2 Normative references.....	1
3 Definitions and conventions.....	1
3.1 Definitions.....	1
3.2 Editorial conventions.....	2
4 CCI and I-Field formats.....	2
4.1 Format.....	2
4.2 Source routeing.....	3
4.3 Logical address.....	3
4.4 Reserved addresses.....	3
5 Switch behaviour.....	3
5.1 Use of INTERCONNECT signals.....	3
5.2 CLOCK signal.....	5
5.3 Connection request successful.....	5
5.4 Breaking a connection.....	5
5.5 Connection request unsuccessful.....	5
 Annexes	
A Routing with the CCI and I-Field.....	6
A.1 General example.....	6
A.2 Source routeing.....	6
A.3 Using a destination address.....	7
A.4 Connections between fabrics.....	9
B Implementation observations.....	10
B.1 Special consideration for Camp-on.....	10
B.2 Interpretation of the W bit and INTERCONNECT signals.....	10
C Bibliography.....	11
 Figures	
1 HIPPI documents.....	iv
2 CCI and I-Field format.....	4
3 I-Field with source routeing, D = 0, and 16 by 16 switch.....	4
4 I-Field with source routeing, D = 1, and 32 by 32 switch.....	4
5 I-Field with destination address and D = 0.....	4
6 I-Field with destination address and D = 1.....	4

© ISO/IEC 1996

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the publisher.

ISO/IEC Copyright Office • Case postale 56 • CH-1211 Genève 20 • Switzerland

Printed in Switzerland

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75% of the national bodies casting a vote.

International Standard ISO/IEC 11518-6 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 25, *Interconnection of information technology equipment*.

ISO/IEC 11518 will consist of the following parts, under the general title *Information technology – High-Performance Parallel Interface*:

- Part 1: *Mechanical, electrical, and signalling protocol specification (HIPPI-PH)*
- Part 2: *Framing Protocol (HIPPI-FP)*
- Part 3: *Encapsulation of ISO/IEC 8802-2 (IEEE Std 802.2) Logical Link Control Protocol Data Units (HIPPI-LE)*
- Part 4: *Mapping of HIPPI to IPI device generic command sets (HIPPI-IPI)*
- Part 5: *Memory Interface (HIPPI-MI)*
- Part 6: *Physical Switch Control (HIPPI-SC)*

Annexes A to C of this part of ISO/IEC 11518 are for information only.

Introduction

This part of ISO/IEC 11518 defines the control for HIPPI physical layer switches. HIPPI by itself is an efficient simplex high-performance point-to-point interface. The physical switch control allows the interconnection of multiple HIPPI based equipments with HIPPI physical layer switches.

Characteristics of this HIPPI physical switch control protocol include

- Support for both source routing and destination addresses.
- I-Fields and CCIs can span multiple physical layer switches within a fabric.
- When a Destination end-point receives a packet, it can easily manipulate the I-Field received to return a reply packet to the Source.
- Support for physical layer switches with differing numbers of ports, all within the same fabric.

Figure 1 shows the relationship of this part of ISO/IEC 11518 (in the solid rectangle) to the other entities shown. HIPPI-SC may be considered a HIPPI component which interprets the signalling information provided to HIPPI-PH in certain switched HIPPI configurations.




Figure 1 – HIPPI documents

Information technology – High-Performance Parallel Interface –

Part 6: Physical Switch Control (HIPPI-SC)

1 Scope

This part of ISO/IEC 11518 provides switch control for physical layer switches using the High-Performance Parallel Interface (HIPPI), a high-performance point-to-point interface between data-processing equipment. This part of ISO/IEC 11518 does not protect against errors introduced by intermediate devices interconnecting multiple HIPPI-PHs.

The purpose of this part of ISO/IEC 11518 is to facilitate the development and use of the HIPPI in computer systems by providing common physical switch control. It provides switch control structures for physical layer switches interconnecting computers, high-performance display systems, and high-performance, intelligent block-transfer peripherals. This part of ISO/IEC 11518 also applies to point-to-point HIPPI topologies.

2 Normative references

The following standards contain provisions which, through reference in this text, constitute provisions of this part of ISO/IEC 11518. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this part of ISO/IEC 11518 are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below. Members of IEC and ISO maintain registers of currently valid International Standards.

ISO/IEC 11518-1:1995, *Information technology, High-Performance Parallel Interface – Part 1: Mechanical, electrical, and signalling protocol specification (HIPPI-PH)*.

ISO/IEC 11518-3:1996, *Information technology, High-Performance Parallel Interface – Part 3: Encapsulation of ISO/IEC 8802-2 (IEEE Std 802.2) Logical Link Control Protocol Data Units (HIPPI-LE)*.

3 Definitions and conventions

3.1 Definitions

For the purposes of this part of ISO/IEC 11518, the following definitions apply.

3.1.1 connection: Condition of the HIPPI-PH when data transfers from a Source end-point to a Destination end-point are possible.

3.1.2 connection control information (CCI): A parameter sent as part of the sequence of operations establishing a connection from a Source to a Destination.

3.1.3 end-point: The equipment at either end of the fabric for a particular connection.

3.1.4 Destination: The equipment at the end of the interface that receives the data.

3.1.5 Destination end-point: The equipment at the end of the fabric that receives the data.

3.1.6 fabric: A group of one or more physical layer switches that can be traversed with one I-Field.

3.1.7 I-Field: A 32-bit field that is sent as part of the sequence of the physical layer operations establishing a connection from a Source to a Destination.

3.1.8 interface: The set of protocols and control signals used to connect a Source and Destination, as defined by HIPPI-PH. Within a fabric, an interface connects an end-point to a switch or a switch to a neighbouring switch.

3.1.9 Logical Address: An address stored in an I-Field that uniquely identifies a Destination end-point or set of end-points.

3.1.10 optional: Features that are not required by this part of ISO/IEC 11518. However, if any optional feature defined by this part of ISO/IEC 11518 is implemented, it shall be implemented according to this part of ISO/IEC 11518.

3.1.11 packet: A data set, as defined by HIPPI-PH, sent from Source to Destination. A packet is composed of one or more bursts.

3.1.12 physical layer switch: A device which allows a single HIPPI physical layer interface to switch between multiple HIPPI physical layer interfaces without involving protocols above the HIPPI Mechanical, Electrical, and Signalling Protocol Specification (HIPPI physical layer).

3.1.13 Source: The equipment at the end of the interface that transmits the data.

3.1.14 Source Address: An address stored in an I-field that uniquely identifies a Source end-point or set of end-points.

3.1.15 Source end-point: The equipment at the end of the fabric that transmits the data.

3.1.16 source routing: A means of packet routing whereby the Source end-point specifies the action of each switch on the way to the Destination.

3.2 Editorial conventions

In this part of ISO/IEC 11518, certain terms that are proper names of signals, state mnemonics, or similar terms are printed in upper case to avoid possible confusion with other uses of the same words (e.g., REQUEST). Any lower case uses of these words have the normal technical English meaning.

A number of conditions, sequences, parameters, events, states, or similar terms are printed with the first letter of each word in upper case and the rest lower case (e.g., Source). Any lower case uses of these words have the normal technical English meaning.

4 CCI and I-Field formats

4.1 Format

The connection control information (CCI) shall be used for controlling HIPPI physical layer switches. Within ISO/IEC 11518-1 (HIPPI-PH) the CCI is used as the I-Field, and is asserted on the HIPPI-PH Data Bus during a connection sequence. The format of the CCI (I-Field) is shown in figure 2. Examples of CCI and I-Field usage for routing are contained in annex A.

L = Locally Administered (bit 31) = 0 designates that the I-Field is defined by this part of ISO/IEC 11518. L = 1 designates that the rest of the I-Field, bits 30 - 0, are locally administered and are not defined by this part of ISO/IEC 11518.

VU = Vendor_Unique (bits 30,29). The contents of the Vendor_Unique bits are not defined in this part of ISO/IEC 11518. Switches shall pass these bits unmodified to the Destination.

NOTE 1 - These bits are available for providing signals to Destinations. Such signals can be used to modify the Destination's behaviour or supply it with additional information on the purpose of the attempted connection.

W = Double-wide (bit 28) = 0 designates that the Source is using the 800 Mbit/s data rate option (DATA BUS is 32 bits wide as defined in HIPPI-PH); the switch shall connect through Cable-A. W = 1 designates that the Source is using the 1 600 Mbit/s data rate option (DATA BUS is 64 bits wide); the switch shall connect through both Cable-A and Cable-B.

NOTE 2 - The W bit is used in conjunction with the INTERCONNECT signals on Cable-A and Cable-B. The INTERCONNECT signals, as defined in HIPPI-PH, tell a switch or end-point that the cable is physically attached to an active HIPPI port. The W bit is used to tell the switch, or Destination end-point, whether or not Cable-B is being used in particular connection.

D = Direction (bit 27) = 0, designates that the right-hand end (least significant bits) of the Routing Control field shall be the current sub-field. D = 1 designates that the left-hand end (most significant bits) of the Routing Control field shall be the current sub-field.

NOTE 3 - When a reverse path exists, a Destination end-point may return a reply to a received packet by simply using the same I-Field that it received with the D bit complemented. For this to work correctly with source routing (PS = 00) then the return path must be symmetrical with the forward path.

PS = Path Selection (bits 26,25). Used to select either (1) a source route, i.e., a specific route through the switches, with output port numbers specified for each switch, or (2) to specify the Logical Address.

00 = source routing: Source selects the route through the switches.

01 = Logical address: Switches select the first route from a list of possible routes.

10 = reserved

11 = logical address: Switches select a route.

C = Camp-on (bit 24) = 0 specifies that the switch shall reply with a connection reject sequence if unable to complete the connection. C = 1 specifies that the switch shall attempt to establish a connection until either the connection is completed or the Source aborts the connection request.

Camp-on is used by the Source to tell a switch to wait for the selected path (or paths) to become available, i.e., the switch should not generate a rejected connection sequence because the selected path is busy. The algorithm used by a switch to select among multiple Sources camped-on to a single Destination is implementation-specific and is not specified in this part of ISO/IEC 11518.

NOTE 4 - A HIPPI rejected connection has a different set of meanings depending on whether or not the Camp-on feature is being used. See clause B.1 for details.

4.2 Source routing

When PS = 00, i.e., source routing, the Routing Control field shall be split into multiple sub-fields, with the size of each sub-field dependent upon the size of the switch that is using it. The number of bits in the sub-field is described as $\lceil \log_2 N \rceil$ where N is the switch size. For example, a 16 by 16 switch would use a four-bit sub-field.

When D = 0, a switch shall use the current sub-field (right most bits of the routing control field) to select the switch output port. The switch shall right shift the routing control field by the number of bits in the sub-field, and shall insert the switch input port number in the left most bits of the routing control field. See figure 3.

When D = 1, the same actions occur except that the current sub-field shall be at the left end of the Routing Control field. The Routing Control field shall be shifted left, and the input port number shall be inserted at the right end of the Routing Control field. See figure 4.

A switch shall not alter the I-Field except when PS=00, and then, only the Routing Control field shall be modified.

4.3 Logical address

When PS = 01 or 11, i.e., logical address, the Routing Control field shall be split into two 12-bit fields. One 12-bit field specifies the address of the Destination end-point(s), the other specifies the address of the Source end-point. When the direction D bit = 0, the right-hand 12 bits shall specify the Destination end-points Logical Address and the left-hand 12 bits shall specify the Source end-points Logical Address (see figure 5). When D = 1, the opposite is true (see figure 6).

4.4 Reserved Addresses

Part of the range of logical addresses are reserved to designate the addresses of network services whose location in the network may vary. All others are available for assignment to specific Destinations.

1111 110x xxxx Reserved for local use
1111 111x xxxx Globally assigned

Globally assigned addresses assigned at the time this part of ISO/IEC 11518 was approved include the following (shown in binary notation).

NOTE - Later registrations will be added as an amendment to this part of ISO/IEC 11518.

1111 1110 0000 Messages pertaining to switch configuration, including HIPPI-LE Address Resolution requests as described in RFC 1374 "IP and ARP on HIPPI." [1]

1111 1110 0001	All IP protocol traffic conventionally directed to the ISO 8802-1 broadcast address as described in RFC 1042 "Standard for IP transmission over 802 networks." [2]
1111 1110 0010	RFC 1112 Host extensions for IP multicasting Class D addresses not assigned below. [3]
1111 1110 0011	RFC 1131 OSPF specification All Routers (Class D address 224.0.0.5) [4]
1111 1110 0100	RFC 1131 OSPF specification All Designated Routers (Class D address 224.0.0.6) [4]
1111 1110 1000	ISO/IEC 9542:1988 CLNP ES-IS all ES's [5]
1111 1110 1001	ISO/IEC 9542:1988 CLNP ES-IS all IS's [5]
1111 1110 1010	ISO/IEC 10589:1992 IS-IS all Level 1 IS's [6]
1111 1110 1011	ISO/IEC 10589:1992 IS-IS all Level 2 IS's [6]
1111 1110 1100	ISO 8802-1 MAC Bridging flooding [7]
1111 1110 1101	ISO 8802-1 MAC Bridging Spanning Tree Protocol [7]
1111 1111 1111	Unknown or unassigned address. This value should never be used to address a Destination or Destinations. It can be used to indicate that the Source is unaware of its Source Address in the CCI, or to signify an unknown Logical Address in higher layer protocols.

The protocols used to access these services and the means whereby these services keep track of their configuration of the network is outside the scope of this part of ISO/IEC 11518.

5 Switch behaviour

A HIPPI physical switch has input ports (attachments to HIPPI Sources) and output ports (attachments to HIPPI Destinations). This clause defines how an HIPPI physical switch behaves with regards to the states of the HIPPI-PH control signals on the input and output ports for a particular connection operation.

5.1 Use of INTERCONNECT signals

As defined in HIPPI-PH, each switch input port and output port shall generate an INTERCONNECT signal when that port is "on-line"; i.e., powered on and enabled for HIPPI connections. Each switch input and output port shall




Figure 2 – CCI and I-Field format




Figure 3 – I-Field with source routing, D = 0, and 16 by 16 switch




Figure 4 – I-Field with source routing, D = 1, and 32 by 32 switch




Figure 5 – I-Field with logical addressing and D = 0




Figure 6 – I-Field with logical addressing and $D = 1$

monitor the received INTERCONNECT signal and shall use this signal to validate all other HIPPI control signals.

NOTE - A switch port may deassert the INTERCONNECT signal when that port is disabled for maintenance or diagnostics.

5.2 CLOCK signal

The HIPPI CLOCK signal generated by the switch output port (a HIPPI Source) shall be continuous and shall conform to the HIPPI-PH specification at all times.

5.3 Connection request successful

Once a connection is completed the switch shall be transparent, with the exception of switch induced latency, to the HIPPI signal sequences.

NOTE - The switch acts as a repeater under the constraints imposed by subclause 7.9 of ISO/IEC 11518-1, and can change the number of idle words between bursts or packets.

5.4 Breaking a connection

Either the Source end-point or the Destination end-point may break a connection.

5.4.1 Source deasserts REQUEST

When the Source end-point deasserts the REQUEST signal, the switch shall break the connection. The switch shall not wait for the Destination end-point to deassert the CONNECT signal to break the connection. The Source end-point will see the CONNECT signal go false regardless of Destination end-point actions.

NOTE - This immediate disconnection frees the associated input port for the next connection-request. This maximizes the efficiency of the HIPPI attachment by allowing the Source end-point to make a new connection without waiting for the propagation delay and the Destination end-point disconnect turn-around time.

5.4.2 Destination deasserts CONNECT

When the Destination end-point deasserts the CONNECT signal, the connection through the switch shall be broken. The switch shall not wait for the Source end-point to deassert the REQUEST signal to break the connection. The Destination end-point will see the REQUEST signal go false regardless of Source end-point actions.

5.4.3 INTERCONNECT false

If the INTERCONNECT signal, received by either the switch input port or the switch output port, goes false during any stage of a connection then the connection through the switch shall be broken.

5.5 Connection request unsuccessful

Connection requests can be unsuccessful due to

- Unavailable Destination end-points
- Unavailable fabric resources
- Errors

5.5.1 Down-stream connection reject

A rejected connection sequence can be initiated by either a down-stream switch or the Destination end-point. A rejected connection sequence shall be propagated through the switch without change. When the switch detects the CONNECT signal is false, at the end of the sequence, the connection through the switch shall be broken.

5.5.2 Switch-generated connection reject

The switch input port logic shall initiate a rejected connection sequence to the Source end-point, and shall not complete the connection through the switch, in the following situations

- When the selected output port(s), or Destination end-point, does not exist or is unavailable. For example, when the INTERCONNECT signal received by the selected output port is false.
- When the selected output port(s) is(are) busy, i.e., connected to some other port, and Camp-on = 0.
- When a data parity error is detected on the I-Field while connecting through this switch. The switch is not obligated to check the I-Field parity after the REQUEST signal has been propagated through the switch.
- When the I-Field Path Selection (bits 26,25) specifies an addressing mode not supported or disabled.
- When the I-Field Double-wide (bit 28) selection conflicts with the switch capabilities or the end-point capabilities. Clause B.2 discusses such conflicts.
- When the I-Field Locally Administered (bit 31) selection is not supported by the switch.

5.5.3 Connection request contention

When two or more input ports vie for the same output port only one input port can "win". The winner shall be connected through the switch. The loser(s) shall either wait for the output port to become available (i.e., Camp-on = 1), or the input port(s) shall generate a rejected connection sequence (i.e., Camp-on = 0).

Annex A (informative)

Routing with the CCI and I-Field

A.1 General example

The CCI and I-Field are used to control HIPPI physical layer switches, supporting the interconnection of many HIPPI devices. Figure A.1 is an example of a small general configuration that will be used to describe the operation of the CCI and I-Field as specified in clause 4. Three hosts, A, B, and C, are shown, but there will probably be many more hosts connected in an actual configuration. The switching fabric is the interconnection mechanism, in this example the four switches and the interconnecting HIPPI links.




Figure A.1 – Physical layer switch example

Two types of operation are specified: (1) source routing, and (2) logical address. The direction of interpretation of the Routing Control field is also under user control, allowing a Destination end-point to return a reply by simply using the same I-Field that was received with the direction bit complemented.

A.2 Source routing

With source routing, the Source end-point specifies the action of each switch on the way to the Destination. For example, to go from host-A to host-B, host-A could specify 2,6,9. This would result in switch-1 selecting output port-2, switch-2 selecting output port-6, and switch-3 selecting output port-9. Alternatively, host-A could have specified 7,5,9 to go through switch-4 instead of switch-2.

For connection or packet routing, this involves an end-point (the "originator" or "Source") having to know the physical route to a particular Destination before a connection can be established. When using multiple switches (where the switches have no intelligence with regards to network routing), the Source has to establish the entire physical route for a given connection. Source routing can often facilitate low latency connections

because the switches have no burden of decision making during the connection process. However, source routing can be unattractive for a configuration with the following features: (1) large and/or dynamic configuration (the hosts must keep track of the interconnection configuration), (2) "blocking" configuration (there is a path to the Destination available, but it is not the one specified, and the one specified is unavailable), and (3) no mechanism that allows a node to "discover" the addresses and associated routes for other nodes

Source routing can be very valuable for diagnostics. By specifying a particular path, a diagnostic program can determine if that path is operable. Source routing can also be attractive when certain paths are more desirable to use than other paths based on criteria that the switches are unaware of, e.g., usage costs associated with particular links.

When using source routing with the CCI and I-Field as specified in clause 4, PS = 00 and the Routing Control contains the route information. The eight high-order bits of the I-Field will be called Ctl. If the switches are 16 x 16 switches using 4-bit addresses, and the direction bit (D) = 0, then to go from host-A to host-B through switch-2, host-A would set the I-Field = Ctl,x,x,x,9,6,2, (in hexadecimal). When received at switch-1, switch-1 picks off the low order four bits (the low order because D=0, four bits because the switch size = 16 x 16) and connects input port-1 to output port-2. It also shifts the Routing Control field of the I-Field to the right and inserts switch-1's input port number in the left end of the Routing Control field. When received at switch-2, the I-Field will be Ctl,1,x,x,x,9,6. Likewise, when it reaches switch-3 it will be Ctl,3,1,x,x,x,9. When it reaches host-B, it will be Ctl,8,3,1,x,x,x.

In summary:

Ctl = L = 0 (not a locally administered I-Field)
 VU = xx (Vendor Unique bits ignored)
 W = x (this example is not concerned with width)
 D = 0 (direction bit)
 PS = 00 (source routing)
 C = x (this example is not concerned with Camp-on)

At A and input to switch-1: I-Field = Ctl, x, x, x, 9, 6, 2
 At input to switch-2: I-Field = Ctl, 1, x, x, x, 9, 6
 At input to switch-3: I-Field = Ctl, 3, 1, x, x, x, 9
 At host-B: I-Field = Ctl, 8, 3, 1, x, x, x

Host-B can send a reply packet back to host-A by complementing the D bit (direction) in the Ctl field. Now the switches will use the left most bits of the Routing Control field as the output port selector, shift the Routing Control field to the left, and insert the input port on the right end. In summary, it would look like:

Ctl = L = 0 (not a locally administered I-Field)
 VU = xx (Vendor Unique bits ignored)
 W = x (this example is not concerned with width)
 D = 1 (direction bit *Note that it changed*)
 PS = 00 (source routing)
 C = x (this example is not concerned with Camp-on)

At B and input to switch-3: I-Field = Ctl, 8, 3, 1, x, x, x
 At input to switch-2: I-Field = Ctl, 3, 1, x, x, x, 9
 At input to switch-1: I-Field = Ctl, 1, x, x, x, 9, 6
 At host-A: I-Field = Ctl, x, x, x, 9, 6, 2

Switches other than the 16 x 16 switches of the example can also be used. In that case the number of bits used by each switch will be determined by the switch size. For example, a 64 x 64 switch will use six bits for selection of its output port, shift off six bits, and add in six bits of input port number.

Mixing switches of different sizes in the same fabric also works. Each switch uses, and shifts, the appropriate number of bits.

This use of the Routeing Control field assumes that all of the cabling is symmetrical. Because the backwards path is not guaranteed to be available, (a) host-B can use the received I-Field as described above with multiple tries or Camp-on, or (b) host-B can use an alternate path which must be determined by other means, e.g., a look-up table.

The reverse path problem is minimized with a system running under the following rules:

A fabric-wide limit for maximum connection time. This will minimize the time it takes for a particular path to become available.

The switches have the capability of allowing connection requests to "camp-on" a particular output port to wait for it to become available.

The output port contention algorithm is fair.

A.3 Using a logical address

Logical addresses specify where the packet is to be delivered, not the route to take to get there. Source end-points use the same Logical Address to reach a particular host, no matter where the Source is located; note that this is not the case with source routing.

With logical addresses the intermediate switches are responsible for picking an appropriate route. The Path Selection bits of the I-Field tell the switches whether to (1) use only the first output port from a list of possible ports, or (2) to select any output port from the list of possible ports. Hence, the switches that use logical addresses must do the port selection as well as the port switching, where source routing switches only need to do the port switching.

A major advantage of using logical addresses is that only the switches need to know the fabric interconnection topology, and the hosts only need to know the logical addresses. Hence, if a link or switch fails, the other

switches can route around it without the hosts having to know about it or do anything special.

Logical addresses are specified in clause 4; PS = 01 or 11 and the Routeing Control field contains the address information. The Routeing Control field contains two 12-bit fields, one for the Source Address and the other for the Logical Address. The direction bit (D) specifies which field is which.

The use of symmetrical cabling is not necessary with Logical addresses, just the need for a path back to the Source. Management is also easier since only the switches need to know how to route among themselves, the hosts do not need to know the fabric interconnection configuration, and there will probably be many fewer switches than hosts. The switches can also use alternate paths to avoid congestion without having to go back to the Source host and try another whole route.

Security is also enhanced by the fact that the switches control the routing and can exercise some control (via their look-up tables) over which paths are legitimate for certain hosts. One feature that is desirable is to have some way for the hosts to "discover" the Logical Address of the hosts, for example, via a name server.

The Logical Address may specify either the end-point of the fabric (i.e., fabric-specific), or the unique address of the host attached to the fabric end-point (i.e., host-specific). The difference between the two is seen when a host is moved from one fabric end-point to another. In this case, the fabric-specific address would change but the host-specific address would not change. There are advantages and disadvantages to both schemes.

It is envisioned that switches can be built to use look-up tables to house routing paths to Destinations. A look-up table can be indexed using a 12-bit Logical Address. The look-up table would be used to hold a possible route (or routes) for each Destination. The management of the tables can be based on either of the addressing schemes: fabric-specific or host-specific.

A.3.1 Fabric-specific addresses

In the following example, the fabric-specific addresses specify fabric end-points. Several schemes can be envisioned for fabric-specific addresses, but in this example the fabric-specific addresses are further broken into a switch number and port number. The fabric-specific address of host-A in figure A.1 is 1-1, host-B is 3-9, and host-C is 4-6. The eight high-order bits of the I-Field are called Ctl.

If the direction bit (D) = 0, then to go from host-A to host-B, host-A would set the I-Field = Ctl, 1-1, 3-9 (in hexadecimal) (3-9 is the Logical Address and 1-1 is the Source Address). When received at switch-1, switch-1 picks off the low order 12 bits (the low order because D=0) as the Logical Address. Even though these 12 bits contain a logical partition between the bits holding the "switch number" and the bits holding the "port number", all 12 bits can be used together as an index into the look-up table. Because this switch is not the switch specified in the Destination Logical Address, the look-up table will provide one or more possible routes to the Destination end-point

(we will assume that it picks output port-2 on switch-1). The I-Field is then passed, unchanged, to switch-2. None of the switches change the I-Field as the packet is routed. A similar action occurs at switch-2, passing the I-Field to switch-3. When switch-3 receives the I-Field, its look-up table yields the output port number specified in the Logical Address (the 9 of 3-9), and selects output port-9.

In summary:

Ctl = L = 0 (not a locally administered I-Field)
 VU = xx (Vendor Unique bits ignored)
 W = x (this example is not concerned with width)
 D = 0 (direction bit)
 PS = 01 or 11 (logical address)
 C = x (this example is not concerned with Camp-on)

At host-A and input to switch-1: I-Field = Ctl, 1-1, 3-9
 At input to switch-2: I-Field = Ctl, 1-1, 3-9
 At input to switch-3: I-Field = Ctl, 1-1, 3-9
 At host-B: I-Field = Ctl, 1-1, 3-9

Host-B can send a reply packet back to host-A by complementing the D bit (direction) in the Ctl field. Now the switches will use the left 12 bits of the Routeing Control field as the Logical Address and the right 12 bits as the Source Address. (Note that the actions taken by the receiving node are the same for both source routing and logical addresses.) In summary, it would look like:

Ctl = L = 0 (not a locally administered I-Field)
 VU = xx (Vendor Unique bits ignored)
 W = x (this example is not concerned with width)
 D = 1 (direction bit *Note that it changed*)
 PS = 01 or 11 (logical address)
 C = x (this example is not concerned with Camp-on)

At host-B and input to switch-3: I-Field = Ctl, 1-1, 3-9
 At input to switch-2: I-Field = Ctl, 1-1, 3-9
 At input to switch-1: I-Field = Ctl, 1-1, 3-9
 At host-A: I-Field = Ctl, 1-1, 3-9

Note that the fabric-specific address does not need to be explicitly broken into switch number and port number fields, it could be all one field. Each switch still determines the path to take to the final Destination, i.e., which output port to use.

A.3.2 Host-specific addresses

In the following example, the host-specific addresses specify unique hosts. For example, the 12-bit host-specific Logical Address for host-A in figure A.1 could be hexadecimal "AAA". Both host-B and host-C would use this address when establishing a connection with host-A. In the following example, the host-specific address for host-B is hexadecimal "BBB". Note that the host-specific addresses are independent of the fabric end-point addresses.

If hosts A and B have addresses "AAA" and "BBB", then to go from host-A to host-B, host-A would set the I-Field = Ctl,AAA,BBB (this assumes that the direction bit = 0).

When received at switch-1, switch-1 uses the low order 12 bits (the low order because D=0) as an index into a look-up table to obtain the possible output ports associated with the Logical Address. Switch-1 could determine that output port-2 can be used to reach Destination BBB. Switch-2 would then use the same process to determine that it should use output port-9. Alternatively, switch-1 could have chosen to use output port-7; passing the connection through switch-4 to get to switch-3.

In summary:

Ctl = L = 0 (not a locally administered I-Field)
 VU = xx (Vendor Unique bits ignored)
 W = x (this example is not concerned with width)
 D = 0 (direction bit)
 PS = 01 or 11 (logical address)
 C = x (this example is not concerned with Camp-on)

At host-A and input to switch-1: I-Field = Ctl, AAA, BBB
 At input to switch-2: I-Field = Ctl, AAA, BBB
 At input to switch-3: I-Field = Ctl, AAA, BBB
 At host-B: I-Field = Ctl, AAA, BBB

Host-B can send a reply packet back to host-A by complementing the D bit (direction) in the Ctl field. Now the switches will use the left 12 bits of the Routeing Control field as the Logical Address and the right 12 bits as the Source Address. (Note that the actions taken by the receiving node are the same for all of the routing schemes described.) In summary, it would look like:

Ctl = L = 0 (not a locally administered I-Field)
 VU = xx (Vendor Unique bits ignored)
 W = x (this example is not concerned with width)
 D = 1 (direction bit *Note that it changed*)
 PS = 01 or 11 (logical address)
 C = x (this example is not concerned with Camp-on)

At host-B and input to switch-3: I-Field = Ctl, AAA, BBB
 At input to switch-2: I-Field = Ctl, AAA, BBB
 At input to switch-1: I-Field = Ctl, AAA, BBB
 At host-A: I-Field = Ctl, AAA, BBB

Host-specific addresses have the advantage that when hosts are moved from one fabric end-point to another, the host-specific address does not change and hence the address tables in the other hosts do not need to be changed. Note that the look-up tables within the switches do need to be updated to reflect the new route to the host. Some disadvantages of host-specific addresses might be longer fabric reconfiguration times to update the switch tables, and problems with keeping the host-specific addresses unique (i.e., no duplicate addresses).

Fabric-specific addresses have the advantages of guaranteeing address uniqueness and quicker fabric reconfiguration.

By specifying the address sub-fields of the I-Field as unpartitioned, the sub-fields can be used for fabric-specific or host-specific addresses.