IEEE 802 Networks — MAC and Physical protocols

General

The IEEE 802 standards define a set of protocols for Local Area Networks and Metropolitan Area Networks. The protocols share several features —

1. They share a common philosophy of 48-bit addressing
2. They are designed for data transfer over essentially reliable broadcast networks and provide error detection but but not error recovery.
3. All protocols present a similar interface to the Logical Link Control (LLC) sub-layer of the OSI protocol stack.
4. They use similar methods for carrying data of other protocols.

The original standards cover three quite different LAN protocols 802.3 (Contention Bus), 802.4 (Token Bus) and 802.5 (Token Ring) and have now been extended to many other protocols, especially for larger MANs.

The IEEE 802 protocols provide very powerful error detection at the MAC layer (a 32-bit CRC), but no explicit error recovery—that must be handled by higher layers such as the Transport layer. There are two reasons for this –

1. Errors on a LAN are very rare, perhaps one in 10^{12} or 10^{13} bits. MAC or Data Link Layer error recovery requires an ACK/NAK response to each message, which is an unreasonable overhead if almost all messages are correct.
2. The LANs operate in broadcast environment, where all stations listen to recognise their address. If a message has an error, that error could have been in its Source Address, so a retry request could go to the wrong station. We just do not know who sent an erroneous message.

In both cases it best to just discard the message and allow higher layers to initiate any necessary recovery.

In each protocol it is necessary to distinguish between Data or LLC frames (which transfer user or LLC data) and MAC frames (in rings) which maintain the ring structure.

General Frame Format

preamble
Destination Address
Source Address
Data
… … … … …
Data
Checksum
Trailer

Frames contain a preamble specific to the implementation, and then Destination Address, Source Address, Data, Checksum and possible Trailer. Some protocols include a Frame Control octet.

The most-significant bit of the destination may be a 1 to force a Group address (a broadcast or multicast address, as opposed to an Individual address) and the second bit is 0 for a globally administered address, or 1 for a 48 bit locally administered address. Local administration is assumed if 16-bit addresses are used. An address of all ones is a broadcast address, recognised by all stations on the network.

The standards suggest the format for the 48 bit globally administered addresses

I/G	G L	region subadr	segment subadr	station sub address
x	*x*	6 bits	8 bits	32 bits
		← region →		
		← individual segment adr →		
← individual station address →				

IEEE 802.3 Contention Bus CSMA/CD

This is very close to Ethernet, except that the data is preceded by a 16 bit length field, which replaces the Ethernet type field. With suitable choice of type field Ethernet and IEEE 802.3 can coexist on the one net. The original physical specifications differed slightly, but Ethernet has now been brought into line with 802.3. Octets are transmitted low-order bit first (except for the FCS which is sent high-order first).

preamble	1010101 … … 0101010 56 bits)
Start Frame	10101011
Destination	LSB first (=1 for group adr)
Source	
Length	Length of data, excluding padding
Data	Data, up to 1500 bytes, may be
… … … …	padded up 46 byte min. length
FCS Frame	32 bits

The basic parameters for the standard 10Mb/s implementation are –

slotTime	512 bit times (51.2µs)	
interFrameGap	9.6µs	gap between frames
AttemptLimit	16	maximum number of retries
backoffLimit	10	no. of retries while increasing delay
jamSize	32 bits	bits forced on to medium after collision
maxFrameSize	1518 octets	
minFrameSize	512 bits (64 octets)	
addressSize	48 bits	

There is no explicit "trailer" or "end delimiter" – the "carrier" or transitions which define bits on the cable just stop.

The basic operation is –

- A station which wishes to transmit senses the medium and defers or waits until 9.6µs after the medium becomes idle (the interFrameGap delay).
- When the medium has been idle for 9.6µs, the station may begin transmitting the 64 bits of the preamble and Start Frame Delimiter, followed by the remainder of the frame, up to and including the FCS. The "carrier" is turned off immediately after the FCS has been sent, with a single transition if necessary to return the signal to the zero or "resting" level. Short LLC data must be padded up to a data size of 46 octets; the length field defines the *valid* data.
- The FCS is a CRC code, generated by the polynomial below, and sent MSB first. The register is initially set to all 1s and each bit is complemented as it is transmitted. The receiver shift register for a good packet (including the packet CRC within the checksum) contains
  11000111 00000100 11011101 01111011.
  The CRC polynomial is $x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^{8}+x^{7}+x^{5}+x^{4}+x^{2}+x+1$.
- The medium is sensed during transmission and any interference noted as a "collision". When a collision is detected jamSize random bits are transmitted and the transmission stopped. The actual jam bits are not defined, but *must not* be the FCS corresponding to the partially transmitted frame. This ensures that the whole medium is filled with random data.
- When a collision is detected the station delays and then retries, deferring according to the normal access rules until the medium is idle. The delay is an integer number r slot times where, for the n-th retry, r is a randomly distributed integer such that $0 \le r \le 2^k$, where k = min(n,backoffLimit), up to a maximum of attemptLimit attempts. The possible delay thereby increases from 51.2µs up to 1024 slotTimes (≈52ms) at the 10th attempt and remains at that value for a further 6 tries, after which a failure will be reported to higher layers of the protocol.
- Frames which are not an integral number of octets are truncated to an octet boundary and reported as AlignmentError. Those shorter than minFrameSize are discarded without error (or sometimes reported as a "runt packet"). Frames longer than maxFrameSize *may* be reported as lengthError.

The "Truncated Binary Exponential Backoff" algorithm for retries minimises the retransmission delays for light or moderate loading, but reduces the offered traffic during heavy loading. It therefore combines good throughput under light loading with good stability under heavy loading.

A minimum packet 64 octets (51.2µs) is sufficient to fill the medium so that all collisions are detected. The maximum packet length ensures that no transmission can overwhelm the network.

The 802.3 protocol has no MAC frames (Medium Access Control). These are needed in the 802.4 and 802.5 protocols to coordinate stations on the ring; with 802.3 there is no such coordination.

IEEE 802.4 Token Bus

This is an implementation of a token bus; it is physically similar to IEEE 802.3 with a passive communication bus, but the stations are logically in a ring and circulate a token to delegate control of the network. The standard covers data rates of 1, 5, and 10Mb/s, in all cases using Manchester, or Phase, encoding. All stations physically on the network can receive transmissions, but only those in the logical ring can transmit. A "response window" is defined as some integral number of octets related to the network size and speed. It is basically the time during which an immediate reply is expected from another station.

The frame format is –

preamble	
N N 0 N N 0 0 0	Start delimiter
F F M M M P P P	Frame Type, MAC
Destination Address	0x… Global; 1x… local
Source Address	
Data	LLC data and / or
… … … … …	MAC control
Data	
Checksum	
N N 1 N N 1 I E	End Delimiter (E = Error)

The preamble is an integral number of octets and is at least 2µs long. The two delimiters contain pairs of bits (N N) which violate the normal rules for Manchester coding – they do not contain transitions in the middle of the bit cells, but do have a transition on the boundary of the two bits. A "J" bit has no preceding transition, while a "K" bit has a preceding transition. These non-data symbols, first a "J-bit" and then a "K-bit", uniquely identify the delimiters.

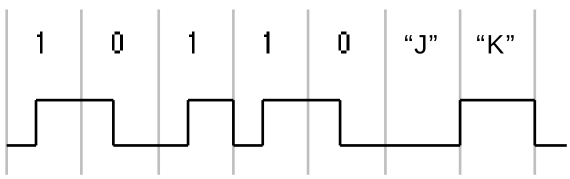

The Frame types are

00	Medium Access Control (MAC) control frame
01	Logical Link Control (LLC) data frame
10	Station management data frame
11	Special purpose data (reserved)

The LLC frames are the normal data, while the MAC frames provide the control signalling to coordinate the bus operation.

The Intermediate (I) bit of the End delimiter indicates that more frames follow in this sequence, while the Error bit (E) is set by a repeater which detects a bad Frame Check.

A station may terminate a frame prematurely by sending the “abort sequence” SD ED.

Abort	N N 0 N N 0 0 0	N N 1 N N 1 I E

802.4 Medium Access Control

The *descending* order of station addresses determines the logical ring around the network. Each station knows its own address (TS – this station), its successor (NS – next station), and its predecessor (PS –previous station). Stations become members of the logical ring only if they choose to respond to a Solicit_Successor message.

Token

The “token” is a special MAC control frame; a station with the token is entitled to transmit or otherwise control the bus. When it has finished transmitting it will send the token on to its successor and then monitor traffic to sense that the token has been accepted. If there is no traffic it will first retry passing on the token, then attempt to remove the successor, and finally to reinitialize the ring.

Token	preamble	
	N N 0 N N 0 0 0	
	0 0 0 0 1 0 0 0	
	Destination Address	DA is NS (Next Stn Adr)
	Source Address	
	Checksum	
	N N 1 N N 1 I E	

The protocol defines the **slot time** as, informally, the greatest time within which a station should expect a response. It is formally defined as–

$$\{ [2 \times (\text{Transmission\_path\_delay} + \text{Station\_delay}) + \text{Safety\_margin}] \,/\, \text{MAC\_symbol\_time} + 7 \} \ \mathbf{div}\ 8$$

where MAC_symbol_time = the time to send a single bit on the physical medium

and Safety-Margin = at least one MAC_symbol_time

A station which receives a valid token records the token Source Address as its Previous Station (PS, or predecessor).

Entry to the Ring

Stations within the ring will periodically issue one of the “Solicit_Successor” MAC control frames. A station which wishes to enter the ring will reply if its address is in the requested range for the message type, sending a Set_Successor message with its own address. A Solicit_Successor_1 is sent by most stations and requests a reply in the range NS < adr < TS, ie between the next station and current station.

Solicit_Successor messages always have the Next Station (NS) in the Destination Address, even though that station will never respond to the message.

Solicit_Successor_1	preamble
	N N 0 N N 0 0 0
	0 0 0 0 0 0 0 1
	Destination Address
	Source Address
	Checksum
	N N 1 N N 1 I E
	delay of one response window

The lowest numbered station on the network will issue a Solicit_Successor_2 message. Waiting stations with an address less than TS reply immediately, while those with an address greater than NS reply after one response window.

Solicit_Successor_2	preamble
	N N 0 N N 0 0 0
	0 0 0 0 0 0 1 0
	Destination Address
	Source Address
	Checksum
	N N 1 N N 1 I E
	delay of two response windows

If several stations respond to the Solicit_Successor (shown by a garbled reply), the requesting station issues a Resolve_Contention frame – the responding stations wait randomly for up to 3 response windows before replying with their Set_Successor messages.

Solicit_Successor_1 messages expect a response from *within* the range TS → NS, while Solicit_Successor_2 messages expect a response from *outside* the range TS → NS.

Resolve_Contention	preamble
	N N 0 N N 0 0 0
	0 0 0 0 0 1 0 0
	Destination Address
	Source Address
	Checksum
	N N 1 N N 1 I E
	delay of four response windows

A station eliminates itself if it hears an earlier response; the soliciting station repeats the operation until it receives a unique reply and then passes the token on to the new successor. Contending stations will wait until the next solicitation of their address range before trying again to enter the ring.

Initialization

The initialization sequence is entered when a station wishes to enter the ring and senses that the medium is idle for a set period of time. It will then send a Claim_Token MAC control frame with an information field 0, 2, 4 or 6 slot times long, depending on 2 bits from its address, wait one slot time and sense the medium. If the medium is busy, there must have been another station trying to initialize and the station defers. If the medium remained idle the attempt is repeated until all bit pairs of the station address have been used (ie 8 or 24 attempts in all). A final attempt will be made with a random response window, to allow for another station with the same address (which is a fault condition). It is expected that only one station will survive this process and will then assume that it has the token. Other stations are then added in the normal way with "solicit successor" MAC frames.

Claim_Token	
	preamble
	N N 0 N N 0 0 0
	0 0 0 0 0 0 0 0
	Destination Address
	Source Address
	Arbitrary data, length of 0, 2, 4, 0 or 6 slot times
	Checksum
	N N 1 N N 1 I E

Exiting

A station wishing to leave the logical ring will first obtain the token and then send its predecessor (PS) a Set_Successor MAC control frame specifying its own successor (NS). The predecessor will then update its tables accordingly and the exiting station is removed from the ring. Alternatively, it may just turn off and expect other stations to recover from its absence.

Set_Successor	
	preamble
	N N 0 N N 0 0 0
	0 0 0 0 1 1 0 0
	Destination Address (Src Adr of last frame rcvd
	Source Address
	New Value of NS (Next Stn)
	Checksum
	N N 1 N N 1 I E

There is a possible error here. With stations in A, B, C and D in logical order, if B decides to leave the ring it signals that A's successor should be C and then passes the token to C. If C also decides to leave at this time, it regards B as its predecessor (having just received the token from B)! C will attempt to set B's successor to D; the Set Successor will fail without notification. On the next token rotation A will fail to pass the token to C and the ring will reinitialise.

A possible solution to this unlikely event is a special "Set Successor with Token" frame. With such a frame, B would pass a SST frame to A, setting its successor and simultaneously passing the token but without setting A's predecessor. A would then pass the token to C.

Lost token

A token may be lost because of data corruption or noise on the network, or because a station has failed. The station which has just issued the token (presumably just before the corruption) senses the absence of activity, retries the token passing once, and then sends a Who_Follows MAC control frame containing its successor address (NS). All stations sense the frame and the one which recognises the address as its own predecessor (PS) will send a Set_Successor as an acknowledgement.

Who_Follows	
	preamble
	N N 0 N N 0 0 0
	0 0 0 0 0 0 1 1
	Destination Address
	Source Address
	test station address
	Checksum
	N N 1 N N 1 I E
	Three response windows

If there is still no response, the station can retry with a Who_Follows with the address of its successor. The next station on the logical ring will then respond with its own Set_Successor, nominating itself to the requesting station and thereby bypassing the station which has presumably failed.

If that attempt fails there is presumably a major failure and the logical ring will have to be reinitialized.

Priority

There is provision for a priority mechanism, according to the Priority field of the header, which defines a possible 8 service classes. The MAC sublayer uses only 4 access classes when sending information. Network capacity is allocated to the higher priority access classes, with lower priorities taking what is left. This is controlled by the station's Token Rotation Timer (TRT) and a fixed Target Token Rotation Time (TTRT). Stations not using priorities send all data at the highest priority.

Each access class (priorities 0, 2, 4 or 6) is allocated a "target" token rotation timer and each station measures the time for a token of that priority to circulate around the ring. A station which has just received the token first sends class_6 (or the highest priority) messages for up to **hi_pri_token_hold_time.** Messages queued for lower priorities may be sent to use the remaining difference between the current TRT and the TTRT; nothing is sent if TRT > TTRT. In effect the token is passed to up to 4 sub-stations, one for each priority class and each sends messages until its queue is empty or the time is exhausted. If time still remains when all queues are exhausted the station may elect to send a Solicit_Successor message.

Data Frame Formats

These are frames which transfer user or LLC data. The MMM field (MAC action) has the three possible values –

0 0 0 = request_with_no_response
0 0 1 = request_with_response
0 1 0 = response

A *request_with_response* packet to a station allows that station to reply immediately with a *response* packet; this constitutes a temporary delegation of the right to transmit quite apart from the normal token action. It is intended to allow an acknowledged connectionless service from the LLC layer. This feature does not have to be implemented.

LLC Data Frame	
	preamble
	N N 0 N N 0 0 0
	0 1 M M M P P P
	Destination Address
	Source Address
	Data — LLC data unit
	Checksum
	N N 1 N N 1 I E
M M M =	MAC action
P P P =	Priority

IEEE 802.5 Token Ring

The physical network is a standard token ring, where the medium can be anything suitable. As in 802.4, the delimiters use "non-data" bits – J has the same polarity as the preceding data bit, and K has the opposite polarity; neither has the mid-cell transition which Manchester encoding expects. Note too that the "ring" is often a star with go and return paths to facilitate bypassing and reconfiguration. Octets are transmitted most-significant bit first.

The format of the empty token is –

Token		
	J K 0 J K 0 0 0	SD Start delimiter
	P P P T M R R R	AC Access Control
	J K 1 J K 1 I E	ED End Delimiter

J & K	non-Data bits
PPP	Frame priority
T	Token bit (0 = token, 1 = other)
M	Monitor bit
RRR	Reservation
I	Intermediate frame (more to come)
E	Error detected

There are two basic frame types — MAC frames for maintaining the ring operation and LLC frames for transferring data (user or LLC control).

LLC frames have the format –

LLC Data Frame		
	J K 0 J K 0 0 0	SD Start delimiter
	P P P 1 M R R R	AC Access Control
	0 1 0 0 0 Y Y Y	YYY P D U priority
	Dest. Address	
	Source Address	
	Data	LLC data Unit
	Checksum	
	J K 1 J K 1 I E	ED End Delimiter
	A C r r A C r r	FS Frame Status

A	Address recognised
C	Frame Copied
r r	reserved for future use

LLC frames are used to transfer User data, or any other information passed in by the LLC layer. The MAC and LLC Frames have the Token bit set to 1; this bit is preceded by the priority field so that the receiving station is first sensitive to the priority. The originating station sets the A & C bits of the FS octet to 0. A station which recognises the DA (either individual or broadcast) will set A to 1, and will also set C to 1 if it copies the frame.

The A and C bits are set only if the frame is recognised as "good", or in the correct format with no errors. When the frame returns to its sender, the four bit combinations have the meanings —

A bit	C bit	Meaning
0	0	Destination absent
0	1	Error – impossible condition
1	0	Destination busy, no free buffer
1	1	Destination copied frame

MAC frames maintain the ring operation and have the format –

MAC	J K 0 J K 0 0 0	SD Start delimiter
Data	P P P 1 M R R R	AC Access Control
Frame	0 0 Z Z Z Z Z Z	Type & control
	Dest. Address	
	Source Address	
	Data	MAC data Unit
	Checksum	
	J K 1 J K 1 I E	ED End Delimiter
	A C r r A C r r	FS Frame Status

Data within the MAC frame are regarded as a "vector" preceded by 2 octets of length and 2 octets of type and containing a series of sub-vectors. The Z…Z bit values control handling of the frame. If Z…Z = 000000, the frame will be copied only a buffer is available. If Z…Z=000001, the frame must be copied if at all possible, including overwriting earlier information. Other values of Z…Z imply a broadcast function to all stations and are copied if a buffer is available.

Medium Access Control

A station must wait for a token before it can transmit. Having obtained the token it can then send as much data as it wishes, up to some limit set by a timer; if the timer expires the current frame must be aborted (by sending the pair SD ED) and the token passed on. Otherwise the station waits until it has received the start of the last frame which it sent, including the source address, before it generates a token to pass on. It then generates nulls or Idles (any combination of 0 or 1, not including an SD symbol)) until the end of the frame is received. More recent variants of 802.5 allow "immediate token release" with the token released as soon as the last message is finished, in contrast to the "delayed token release" of standard 802.5.

The Active Monitor times the ring by emitting bits timed by its own internal clock. All other stations recover a clock from their incoming data and use that to time their outgoing bits. In the absence of deliberate traffic, all stations emit Idles to maintain the ring timing.

Stacking Stations & Priority Transfer

A station can claim a token if the token priority is less than or equal to that of the message which it wishes to transmit. If a token is in use, the station may set the reservation bits to its desired message priority. The station which initiated the transfer will then generate a new token with priority equal to the maximum of the previous token level and the reservation level. It must remember the old priority (it is a "stacking station" which stacks the old priority) and drop the priority back when there are no messages in the network with a higher priority (it receives a token with a priority equal to that which it generated).

To illustrate, consider a situation where station B wishes to transmit at priority P_m = 4 and sees a message (busy token) from station A at ring priority P_r = 2. Then

1. B sets the busy token reservation bits R_r = 4
2. A receives the frame which it generated at P_r = 2, but now with R_r = 4. A interrupts its transmission, issues a free token with P_r = 4 and saves (stacks) its previous priority of 2.
3. B receives and acquires the token with P_r = 4 and transmits its high priority messages.
4. B releases a free token with P_r = 4.
5. A receives the free token with P_r = 4, the priority at which it issued a token, unstacks its previous priority of 2 and issues a new token at P_r = 2.
6. ——————————
7. Station B with a priority PDU to transmit will raise the reservation level R_r of a busy token to its PDU priority P_m (if the <u>incoming</u> $R_r < P_m$).
8. After a station has claimed the token, it may transmit PDUs which have $P_m \geq P_r$ until none remain or THT is exhausted. The PDUs are sent at priority P_r. The station then issues a new token at priority P_r.

Standby Monitor Stations

All stations except the Active Monitor are "Standby Monitors". Each station periodically sends a "Monitor Present" MAC frame, either "Active Monitor Present" for the Active Monitor, or "Standby Monitor Present" for all other stations. Each Standby Monitor introduces a latency or delay of 1 bit into the ring — this delay must be as small as possible to avoid degrading the ring performance.

Active Monitor Station.

One station is designated as the Active Monitor Station, with several functions. The other stations are all Standby Monitors, ready to take over if necessary. Any station may become the Active Monitor by "winning" from a "Claim Token" on ring recovery or initialisation.

The Active Monitor has the special functions —

Latency. The Monitor (ie Active Monitor) station introduces a latency of at least 24 bits (the size of a token) so that a token does not destroy its own tail.

Lost Tokens. If no token or data is received for a time related to the maximum transmission time and ring latency, the Monitor will transmit idles for a time sufficient to clear the ring and then issue a token.

Circulating Busy Token. Whenever a busy token (a MAC or LLC data frame) passes the Monitor station, the Monitor will set its M bit. If it later detects a busy token with M=1 it clears the ring and issue a new token.

Duplicate Tokens. A transmitting station may receive a frame which does not contain its own address, or may receive a token with a corrupted header. In both cases it will refrain from generating a new token, thereby forcing a lost token situation, to be recovered by the Monitor station.

Neighbour Notification.

The Neighbour Notification procedure involves a station sending a broadcast message (essentially Here I Am, actually Active Monitor Present or Standby Monitor Present). Its downstream neighbour will recognize the address as its own and set the Addressed Recognized bit, and will note the sender as its upstream neighbour. Other stations, noting that the address has been recognised, will ignore the message. The neighbour notification process enables the ordering of stations on the ring to be established.

Medium Access Control Coordination.

Many of these activities are coordinated by MAC control messages with codes in the ZZZZZZ bits of the Frame Control octet. Other information is held in vectors and subvectors within the data field of the MAC frame (Vector Identifier, VI; SubVector Length SVL; SubVector Identifier SVI; and SubVector Value SVV). These values are given in IBM-format hexadecimal. Sub-Vector type 2 is always the Received Upstream Address (RUA).

At most one Vector may be present in a frame, but it may contain several sub-Vectors –

Octets of Information field

VL	2 octets	Length of whole vector
VI	2 octets	Type of whole vector
SVL	1 octet	Subvector 1
SVI	1 octet	
SVV	n octets	
SVL	1 octet	Subvector 2
SVI	1 octet	
SVV	n octets	
SVL	1 octet	Subvector 3
SVI	1 octet	
SVV	n octets	

Duplicate Address Test

During initialization a station may send a Duplicate Address Test message addressed to itself. If the frame returns with the Address Recognized bits set there must be a duplicate address, and the station becomes passive.

Duplicate Address Test	J K 0 J K 0 0 0	SD Start delimiter
	P P P 1 M 0 0 0	P P P is PDU priority
	0 0 0 0 0 0 0 0	FC DAT
	Dest. Address	MA curr stn address
	Source Address	
	X'0004'	VL = 4
	X'0007'	VI = 7 Dup Adr Test
	Checksum	
	J K 1 J K 1 I E	ED End Delimiter
	A C r r A C r r	FS Frame Status

Claim Token

A station which does not see an Active Monitor frame for a while will issue a Claim Token and wait for a Claim Token with its own source address – it then becomes the Active Monitor station. If it receives a Claim Token from another station it removes itself from contention and stops sending Claim Tokens.

Claim Token	J K 0 J K 0 0 0	SD Start delimiter
	P P P T M R R R	AC Access Control
	0 0 0 0 0 0 1 1	FC CL_TK
	Destination	all stns on this ring
	Source Address	
	X'0007' or X'000B'	VL = 7 or 11
	X'0003'	VI = 3 Claim Token
	SVL X'02'	SVI–1 Rcvd U A
	2 or 6 octet RUA	Rcvd Upstrm Adr
	Checksum	
	J K 1 J K 1 I E	ED End Delimiter (E
	A C r r A C r r	FS Frame Status

Purge

This is broadcast when the Active Monitor claims the token, or when it detects the Monitor Bit set and reinitializes the ring.

Purge (PRG)		
	J K 0 J K 0 0 0	SD Start delimiter
	P P P T M R R R	AC Access Control
	0 0 0 0 0 1 0 0	FC PRG
	Destination	all stns on this ring
	Source Address	
	X'0007' or X'000B'	VL = 7 or 11
	X'0004'	VI = 4 Purge
	SVL X'02'	SVI–1 Rcvd U Adr
	2 or 6 octet RUA	Rcvd Upstrm Adr
	Checksum	
	J K 1 J K 1 I E	ED End Delimiter
	A C r r A C r r	FS Frame Status

Active Monitor Present

The Active Monitor station periodically broadcasts an Active Monitor Present, and after the ring has been purged.

Active Monitor Present		
	J K 0 J K 0 0 0	SD Start delimiter
	P P P T M R R R	AC Access Control
	0 0 0 0 0 1 0 1	FC AMP
	Destination	All stations, this ring
	Source Address	
	X'0007' or X'000B'	VL = 7 or 11
	X'0005'	VI = 5 Active Monitor
	X'02'	SVI–1 Rcvd U Adr
	2 or 6 octet RUA	Rcvd Upstrm Adr
	Checksum	
	J K 1 J K 1 I E	ED End Delimiter
	A C r r A C r r	FS Frame Status

Standby Monitor Present.

Standby Monitor Present		
	J K 0 J K 0 0 0	SD Start delimiter
	P P P T M R R R	AC Access Control
	0 0 0 0 0 1 1 0	FC SMP
	Destination Adr	all stns on this ring
	Source Address	
	X'0007' or X'000B'	VL = 7 or 11
	X'0006'	VI = 6 S_M present
	X'02'	SVI–1 Rcvd U Ar
	2 or 6 octet RUA	Rcvd Upstrm Adr
	Checksum	
	J K 1 J K 1 I E	ED End Delimiter
	A C r r A C r r	FS Frame Status

This is periodically broadcast by each standby monitor station. As both "Monitor Present" messages are broadcast to all stations, the A and C bits will be set by the first station after the originator. Any station which receives a monitor present message with A & C = 0 is entitled to set its SUA (Stored Upstream Neighbours Address).

Beacon

The Beacon message is used in fault isolation. A station which senses a serious ring failure (no activity and no response to a Claim Token) may start Beaconing. It may then abort if no Beacon message is seen for a while. If a Beacon is seen then –

i. if its SA is MA (Source = My Address) the ring appears to be restored and the normal Claim Token may be entered, or

ii. if a Beacon is seen from another station, the station enters Standby State, expecting a token to be issued by the other Beaconing station.

Beacon		
	J K 0 J K 0 0 0	SD Start delimiter
	P P P 1 M 0 0 0	P P P is PDU
	0 0 0 0 0 0 1 0	FC BCN
	Destination	all stns on this ring
	Source Address	
	X'0009' or X'000D'	VL = 9 or 13
	X'0002'	VI = 2 Beacon
	SVL X'02'	SVI–1 Rcvd U Adr
	2 or 6 octet RUA	Rcvd Upstrm Adr
	SVL X'01'	SVI–2 Beacon type
	Beacon Type	See below
	Checksum	
	J K 1 J K 1 I E	ED End Delimiter
	A C r r A C r r	FS Frame Status

If the link has failed, the effect of the type (ii) Beacon is that all stations except that immediately following the fault enter Standby mode and all stations receive the address of that station and of its upstream neighbour just before the failure; the fault can then be located.

SVV-2	= X'0001'	(For future use)
	= X'0002'	Continuous J symbols received
	= X'0003'	Timer TNT expired during claiming token; no FR_CL_TK received.
	= X'0004'	Timer TNT expired during claiming token; FR_CL_TK received.

(A "J" symbol is defined to have no change from the previous bit; continuous J's then correspond to a broken link.)

Timers

The IEEE 802.5 standard defines various timers, with suggested default values

Timer, Return to Repeat (TRR) 2.5ms

A station which has completed transmission (empty PDU queue or THT expired) transmits fills and waits until it has received the last transmitted frame with its own address (MA), or until TRR has expired, in which case it reverts to repeating all incoming bits.

Timer, Holding Token (THT) 10ms.

A station, having acquired the token, may transmit for the time THT. It is expected to refrain from starting any packet which could not be completed before THT expires. If a station is transmitting when its THT expires it must *immediately* abort the current frame by sending the pair SD ED generate a new token and pass it on.

Timer, Queue PDU (TQP) 10ms

A station which receives an AMP or SMP frame with the A and C bits equal to 0 should set TQP and transmit a SMP frame when TQP expires.

Timer, Valid Transmission (TVX) THT+TRR

If the monitor detects no transmission for a time TVX it immediately resets TNT and prepares to purge the ring when TNT expires.

Timer, No Token (TNT) 1s

If there are n stations on the ring, TNT = TRR + $n \times$ THT. TNT is reset when a token or frame is seen. When TNT expires the Monitor removes its latency buffer, clears TNT, and reverts to being a Standby Monitor for eventual ring recovery

Timer, Active Monitor (TAM) 3s

The expiry of TAM forces the Monitor to generate an Active Monitor Present frame.

Timer, Standby Monitor (TSM) 7s

Any Standby Monitor station may transmit a Claim_Token if it has not seen a token of frame for a time TSM.

Fibre Distributed Data Interface

FDDI is a development of token ring which runs at 100 Mbit/s, over rather larger distances than token rings (up to 200km total length) and is often used in campus backbones. Physically, FDDI uses either two counter-rotating rings, a primary ring and a secondary, or a single ring. In a typical campus network a double ring will be used for the backbone, connecting mostly bridges, while the local "spur" networks may use either single or double rings.

FDDI was originally designed for optical fibre, but for shorter distances an equivalent copper implementation is available (CDDI).

FDDI signal coding uses a 4B/5B code, with signalling at 125 MBaud to achieve 100 Mb/s data rate. Some of the code symbols are used as special "non-data" controls and delimiters. Data bytes (8 bytes) are divided into 4-bit "nibbles" for encoding into 5-bit symbols. The symbols shown are further encoded with NRZ-I for actual transmission.

4-bit data	5-bit symbol	4-bit data	5-bit symbol	Control signal	5-bit symbol
0000	11110	1000	10010	IDLE	11111
0001	01001	1001	10011	J	11000
0010	10100	1010	10110	K	10001
0011	10101	1011	10111	T	01101
0100	01110	1100	11010	R	00111
0101	01011	1101	11011	S	11001
0110	01110	1110	11100	QUIET	00000
0111	01111	1111	11101	HALT	00100

FDDI 4B/5B symbol encoding

The data frame and token formats are shown below. Note the difference between a *byte* (8 bits) and a *symbol* (4 bits).

Information Frame

PA	Preamble at least 16 IDLE symbols
SD	Start Delimiter (2 symbols - J, K)
FC	Frame Control (2 symbols)
DA	Destination address (2 or 6 bytes)
SA	Source Address (2 or 6 bytes)
INFO	Information, up to 4500 octets
FCS	Frame Check Sequence 4 octets
ED	End Delimiter (1 T symbol)
FS	Frame Status (3 symbols R & S)

Token

PA	Preamble at least 16 IDLE symbols
SD	Start Delimiter (2 symbols - J, K)
FC	Frame Control (2 symbols)
ED	End Delimiter (2 T symbols)

FDDI 4B/5B frame formats

The maximum size of an FDDI ring precludes the use of a token ring control protocol with several priorities as in 802.5 Token Ring. Instead FDDI takes over the Timed Token protocol from the 802.4 Token Bus.

Each station has two times, a Token Rotation Timer TRT which counts up and a Token Holding time THT which counts down. The network manager sets a Target Token Rotation Timer TTRT. The FDDI station –

1. Obtains the token
2. Sets THT = TTRT – TRT. (This sets THT to the remaining allowable token rotation time.)
3. Resets TRT = 0.
4. Transmits packets until THT = 0 or there is no packet left to send.
5. Releases the token to the next station.

An FDDI network also allows a synchronous data transmission mode for data with guaranteed throughput. In this mode a station negotiates with a network manager for a quota of fixed bandwidth; the result is that it receives a "Synchronous Allocation", held in its SA register as a transmission time. When a station receives the token it is entitled to transmit for the entire time held in the SA register, irrespective of the THT. (Overruns of THT are permitted.) Asynchronous traffic may be sent for any remaining time, up to THT.

FDDI uses immediate token release, forwarding a token as soon as the station finishes sending the last message of the transmission (in contrast to waiting for the message to be received back at the sender after one rotation around the ring – delayed token release).