Interconnection Technologies

Bridging II

Consider...

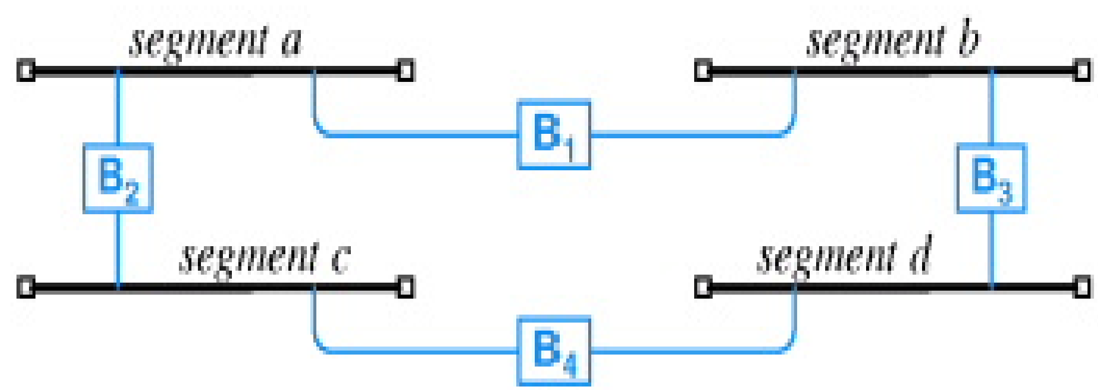

Transparent bridging loops

- Bridges will *flood* unlearned destinations
  - With bridge loops, frames proliferate
- Bridges will learn source addresses
  - With bridge loops, address table *thrahses*
- Bridges will forward unlearned destinations
  - With bridge loops, frames can loop indefinitely

Bridge loop solutions

- Absolve never to have loops
- Allow loops, but prevent looped frames
- Do not allow loops, but allow redundancy

Spanning tree algorithm

- Ensures just one single path between any two hosts
- Disables active redundancy/loops
- Bridges must cooperate to discover and form loop−free paths
- Bridges should (de)activitate paths when necessary
- Standardized in IEEE 802.1D

A *Pruned tree* topology

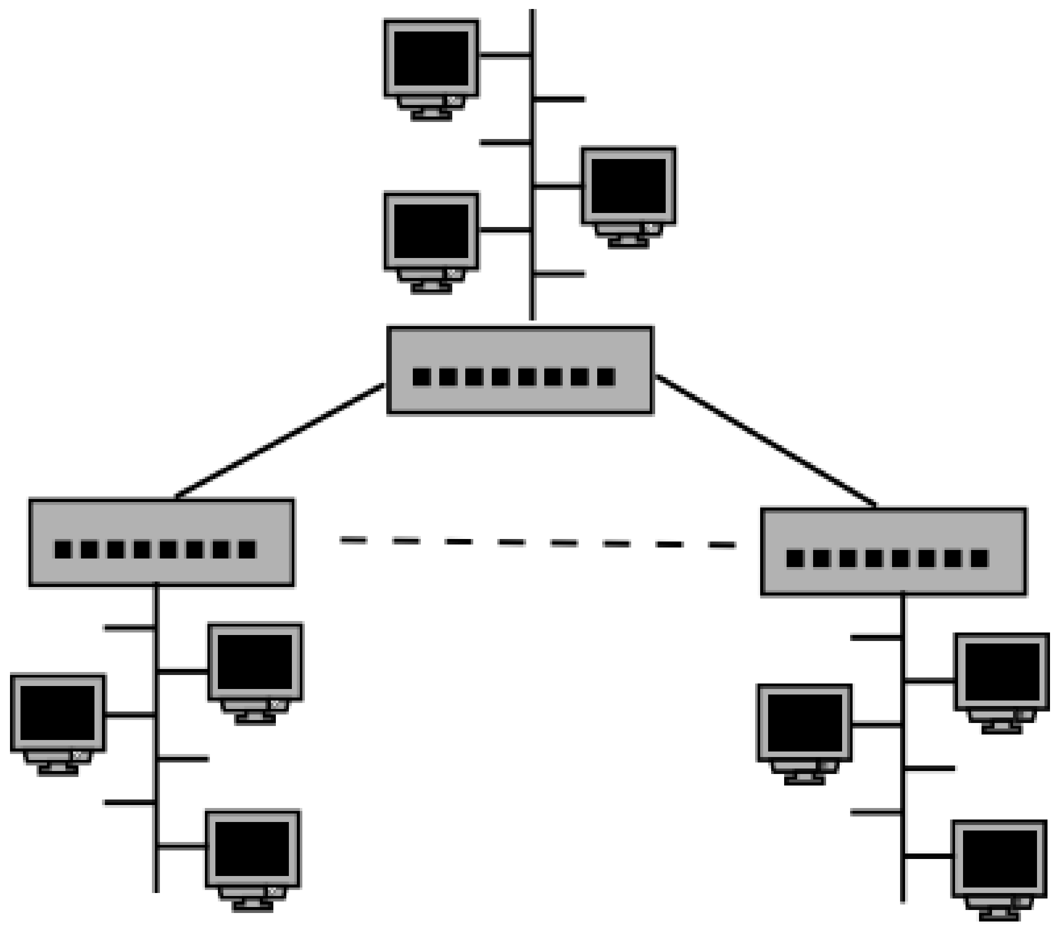

The root bridge

- Think of an upside down tree...
- Elect a single *root* bridge
- Root bridge is logical center of the catenet
  - Not necessarily the physical center
- Any bridge could potentially be root bridge
- Root bridge can change over time

Designated bridges

- One designated bridge per branch (LAN)
- Only bridge to forward traffic for the branch
- Root bridge is designated bridge for all its LAN attachments
- Non−root bridges may be designated bridge for zero or more LAN attachments

Root ports

- Designated bridges have only one root port
  - Except the root bridge which has no root ports
- A port that forwards traffic to the root bridge
- Traffic received on a bridge’s root port will be forwarded to one or more of its designated ports
  - As determined by normal forwarding rules

Designated ports

- A bridge port that forwards traffic away from the root bridge onto the LAN link (branch)
- All root bridge ports are designated ports
- Non−root bridges may have 0 or more designated ports

Bridge identifier (ID)

- Each bridge is assigned one bridge ID
- Uniquely identifies a bridge in the catenet
- 64 bit field = 48 bit MAC plus 16 bit priority
  - Priority field allows an administrator to control which bridge becomes the root bridge

Port identifier (ID)

- Each bridge port is given a 16−bit ID
  - 8−bit ID and an 8 bit priority field
- Locally significant (not globally unique)
- Ports are number 1 to N
  - Limits bridge to 255 ports or less

Link/path costs

- Each port has an associated link
- Each link has an associated speed in bps
- An inverse cost is associated with link speed
  - Example:
    - 10 Mb/s = 100, 100 Mb/s = 19, 1 Gb/s = 4
- Spanning tree is built using *shortest path cost* to root rather than strictly shortest path
- Can be used to alter the path to root

Calculating the spanning tree

- Elect a root bridge
- Elect designated bridges and ports
- Maintain the topology over time

Elect a root bridge

- Initially each bridge considers itself root
- Bridges send BPDU frames on LAN links
  - Bridge and port ID of sending bridge
  - Bridge and port ID of bridge considered root
  - Root path cost for the sending bridge
- Bridge with the lowest ID becomes root
- Bridges leaving/entering catenet may cause the spanning tree to reconfigure

Elect the designated bridges

- Initially each bridge considers itself the designated bridge
- Bridges send BPDU frames LAN links
  - Bridge and port ID of sending bridge
  - Bridge and port ID of bridge considered root
  - The root path cost for the sending bridge
- Lowest path cost to root wins

Elect the designated port

- Port from the designated bridge on the LAN
- In case a bridge has more than 1 link on the LAN, port ID breaks a tie

Bridges select root port

- Each non−root bridge selects root port
  - Based on lowest path cost to root bridge
- In case a bridge has more than one 1 link on the LAN, port ID breaks tie

Bridge port state

- Disabled – broken or administratively down
- Blocking – not part of spanning tree
- Listening – participating in election process
- Learning – building address table
- Forwarding – active in spanning tree

Spanning tree illustrated

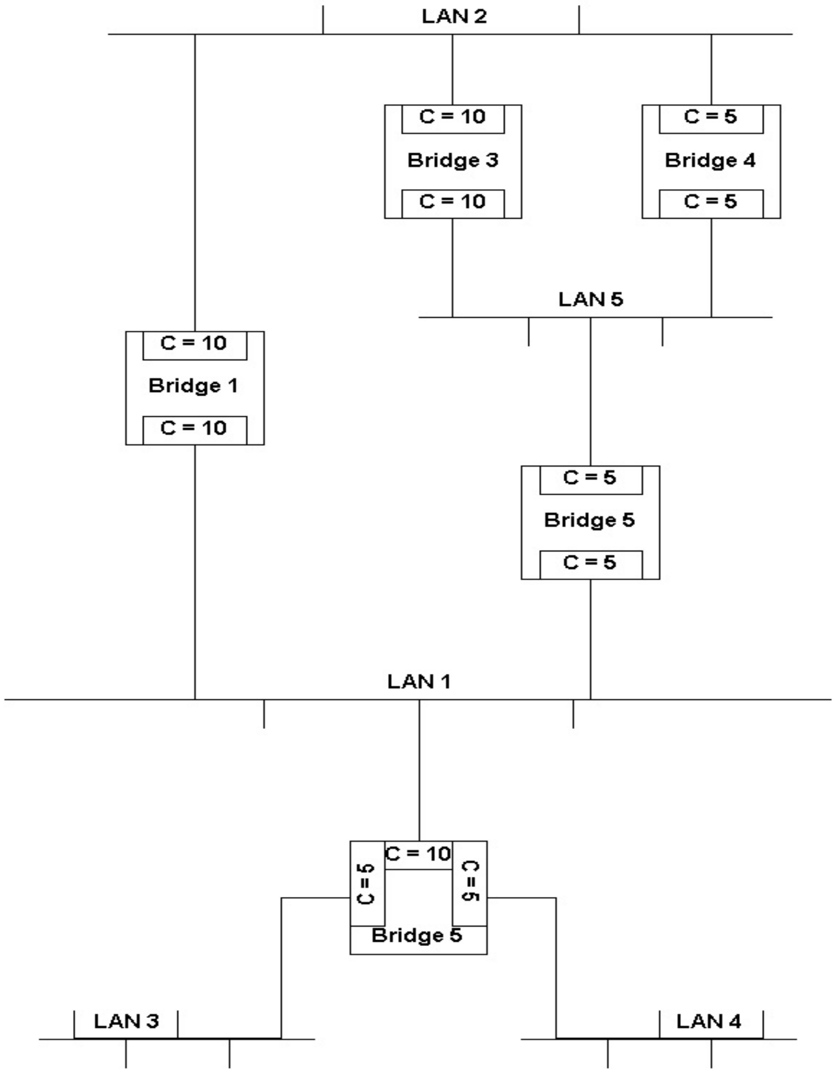

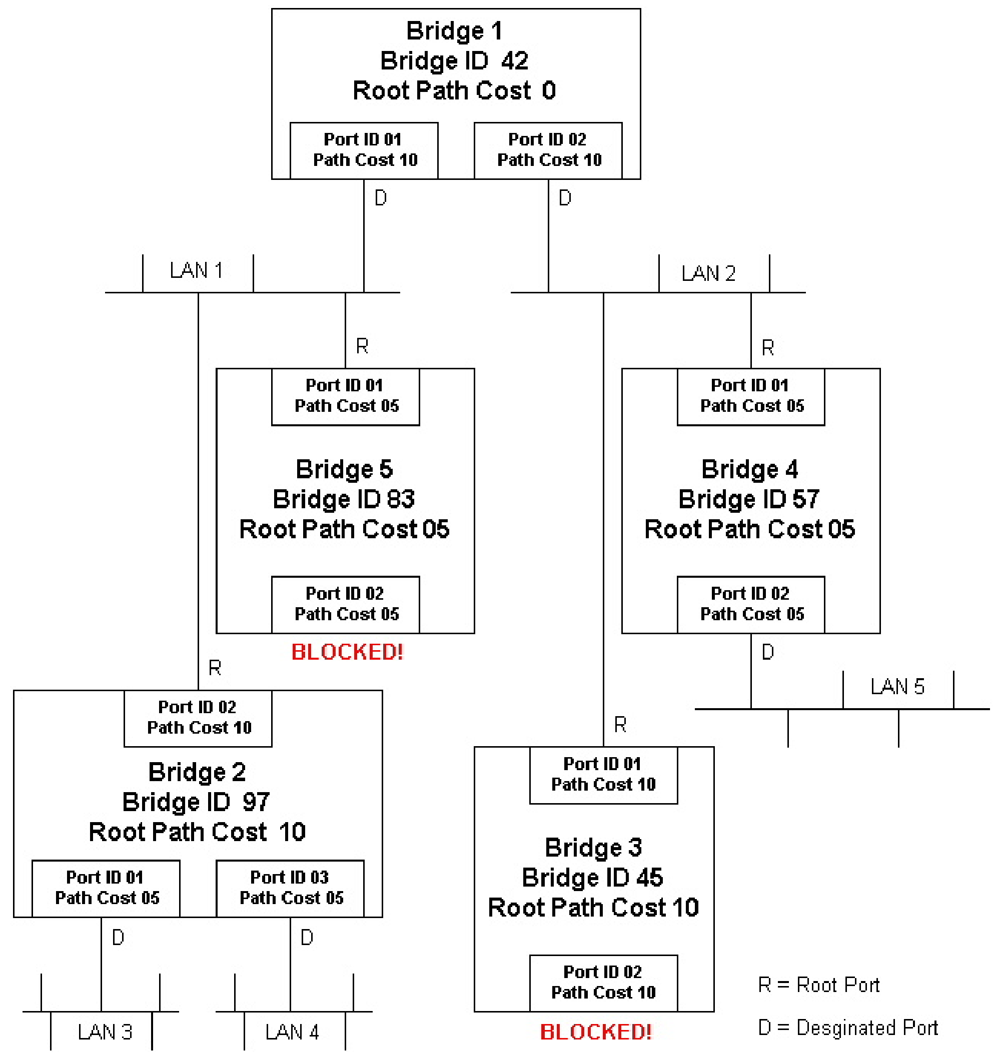

Bridge protocol data unit (BPDU)

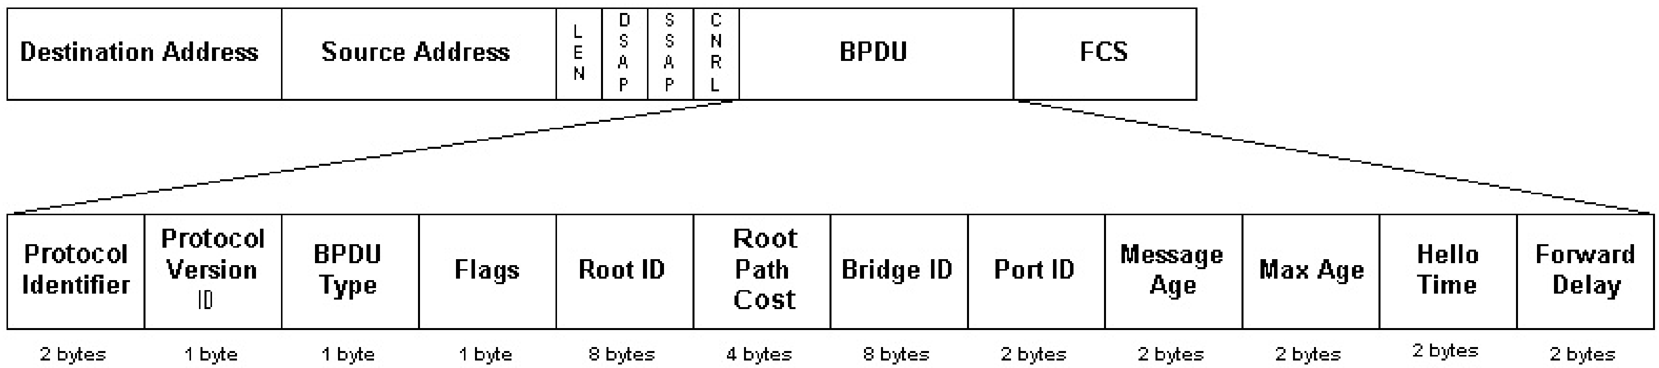

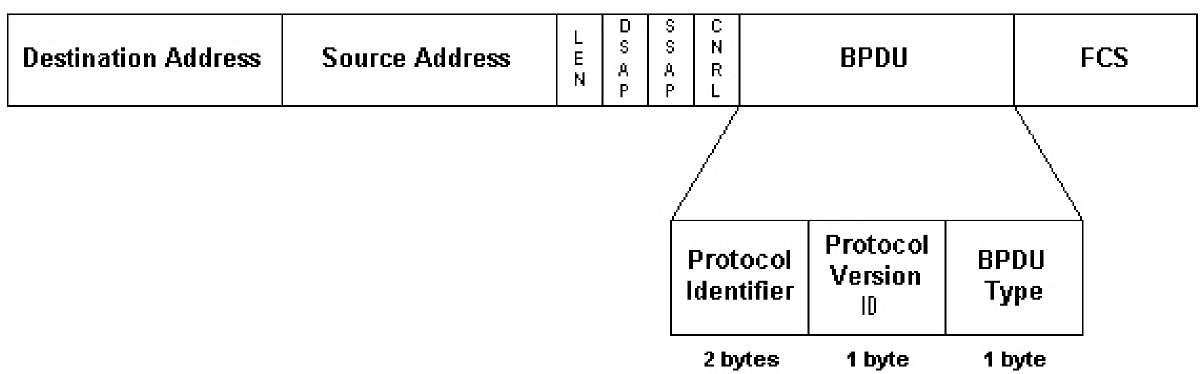

Source route bridging

- Intelligence moves to end stations
- Only used in token ring environments
  - Some FDDI, but non−existant for Ethernet
- Called bridging, but probably shouldn’t be
- Originally developed by IBM
- Mostly of academic interest now
- Uses a spanning tree for multicast/broadcast and route discovery traffic

Source route bridges

- Use routing information field (RIF) in frames
- No bridge address tables
- Hosts maintain address or rather route tables
- Must number rings and bridges in order to build a path (RIF)
- Allows for active loops

Source route bridging illustrated

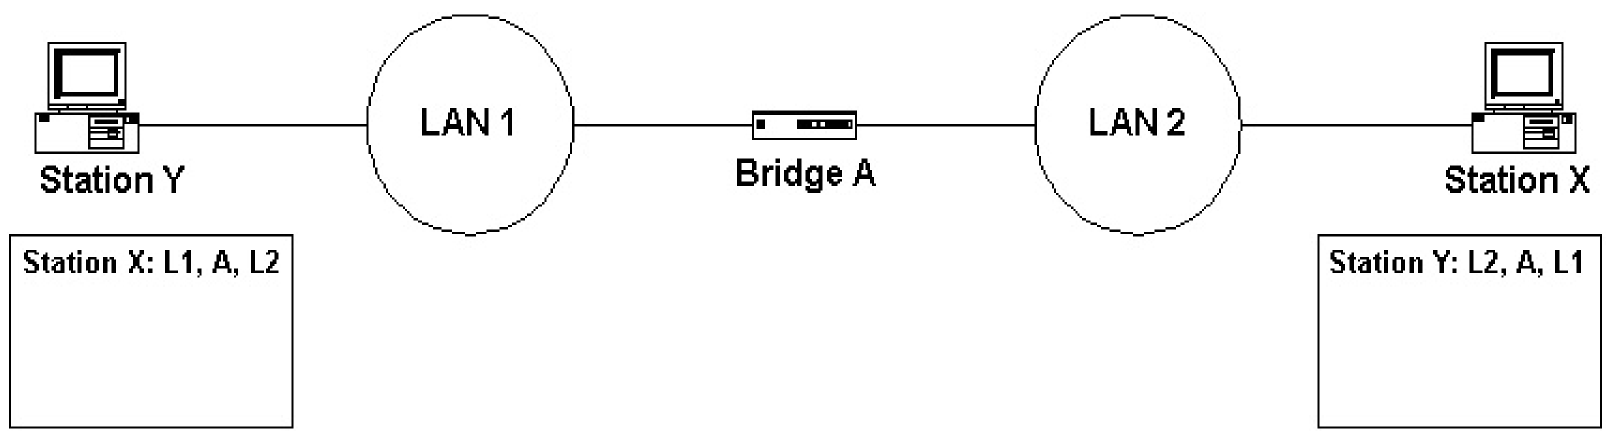

Source routed destinations

- Null
  - Destination on same LAN segment
- Non−broadcast
  - Frame includes a RIF
- All routes broadcast
  - Flooded to each LAN, bridges add to the RIF
- Single route broadcast
  - Use spanning tree, bridges add to RIF

Route discovery

- Source sends all routes broadcast
  - Destination responds with a single route return using the route used by the first frame received
- Source transmits a source route broadcast
  - Destination responds with a all routes broadcast
  - Source uses route from the first returned frame
- Routes can also be statically configured in end stations

Route discovery illustrated

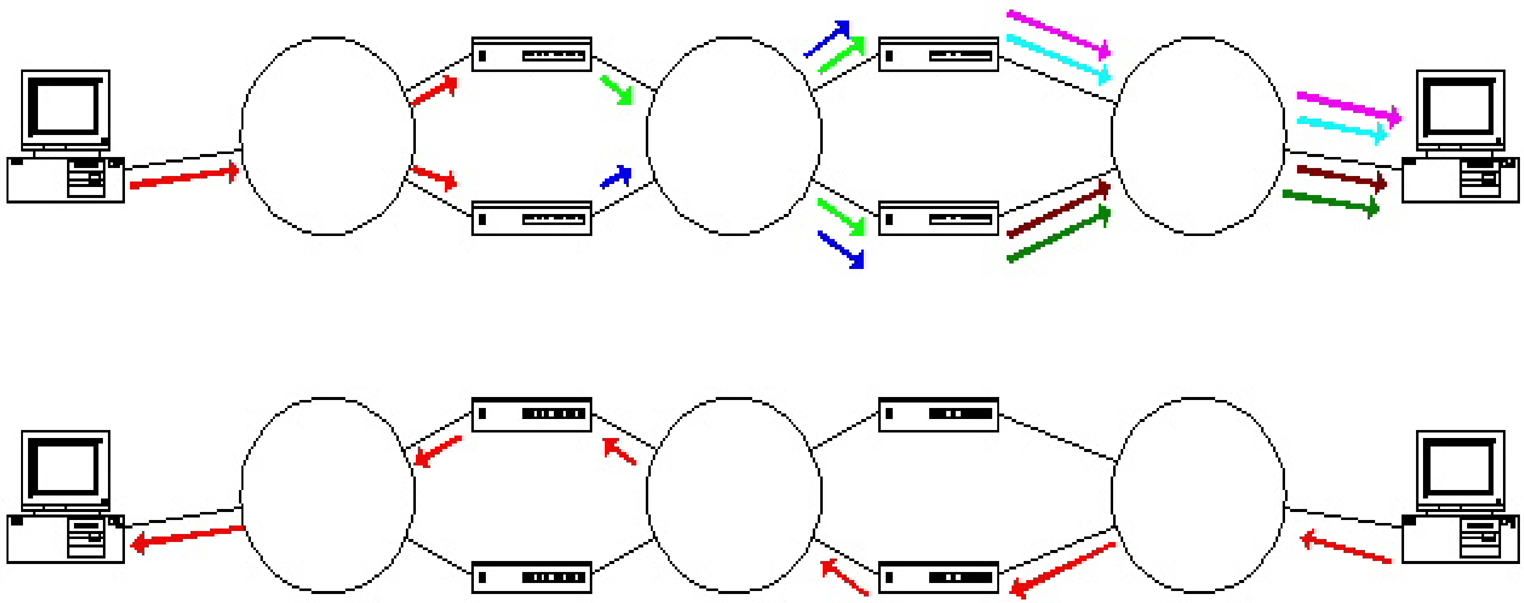

Routing information field (RIF)

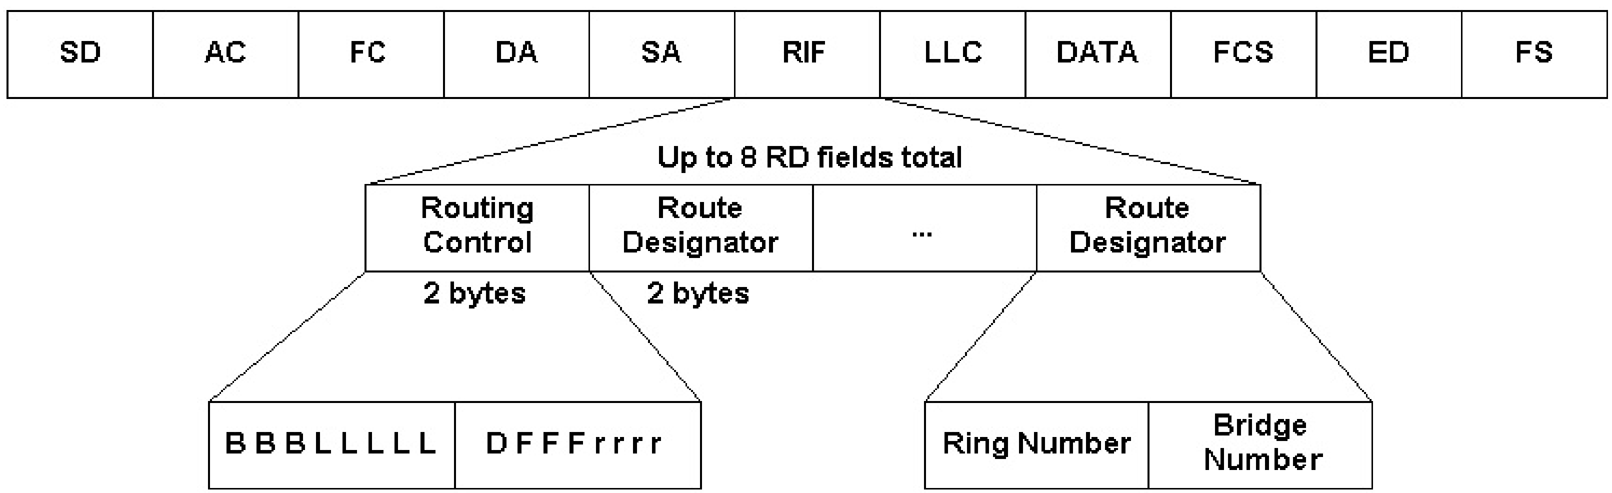

B = broadcast indicators
L = length indicators
D = direction bit
F = largest frame bits
r = reserved bits

Broadcast indicators

000 non-broadcast
100 ARB with SRR
110 SRB with ARR
111 SRB with SRR

Largest frame bits

000 up to 516 bytes
001 up to 1500 bytes
010 up to 2052 bytes
011 up to 4472 bytes
100 up to 8144 bytes
101 up to 11,407 bytes
110 up to 17,800 bytes
111 used in ARB frames

Final thoughts

- Token ring can use transparent bridging
- For all practical purposes, source route bridging and token ring are dead
- This is somewhat unfortunate, there are a lot of networking lessons to be learned – good and bad
- Yes, your instructor did a lot of this stuff