

Observations on Filtering & Forwarding SR Classes

John Nels Fuller

Interconnect Solutions Architect

jfuller@computer.org

Filtering in Practice Today

- When a frame arrives at the port of a bridge the destination address, VID, and PCP fields are used with the filtering database to determine to which port(s) and queue the frame is to be sent (if any).
- A subset filtering database is normally kept in fast hardware registers for use by the data plane.
 - This normally includes DA, VID, and a port map (one bit per port) of where to forward frame.

Effects of SR Classes 1

- If we try to live with existing HW port map:
 - MMRP (or IGMP snooping) set bits as today.
 - SRP causes bits to be cleared for all ports until a successful reservation on a port sets that port's bit.
 - Best effort traffic classes stop flowing to DA+VID when SRP starts for DA+VID (even unicast DAs); resumes at a port when there is a successful reservation.
 - Only one stream per DA+VID can be managed (applies to both unicast and multicast DA's). See example later.

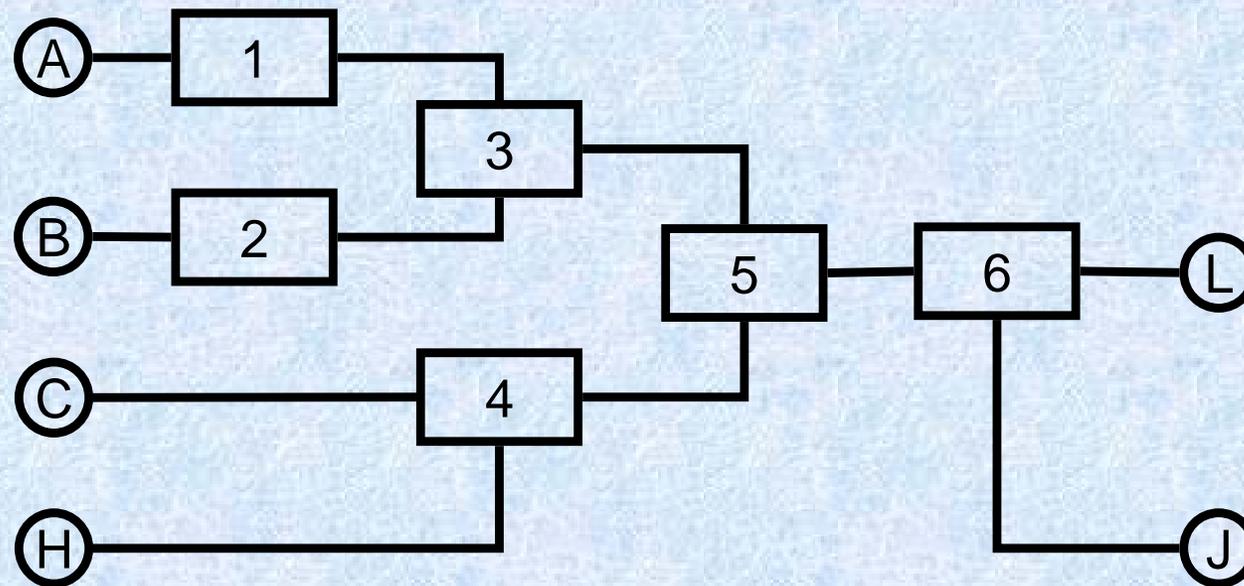
Effects of SR Classes 2

- If we expand HW port map to 2 bits/port:
 - MMRP (or IGMP snooping) set first bit as today.
 - SRP sets second bit for port upon successful reservation.
 - Best effort traffic classes are unaffected by SRP actions.
 - Only one stream per DA+VID can be managed (applies to both unicast and multicast DA's). See example later.

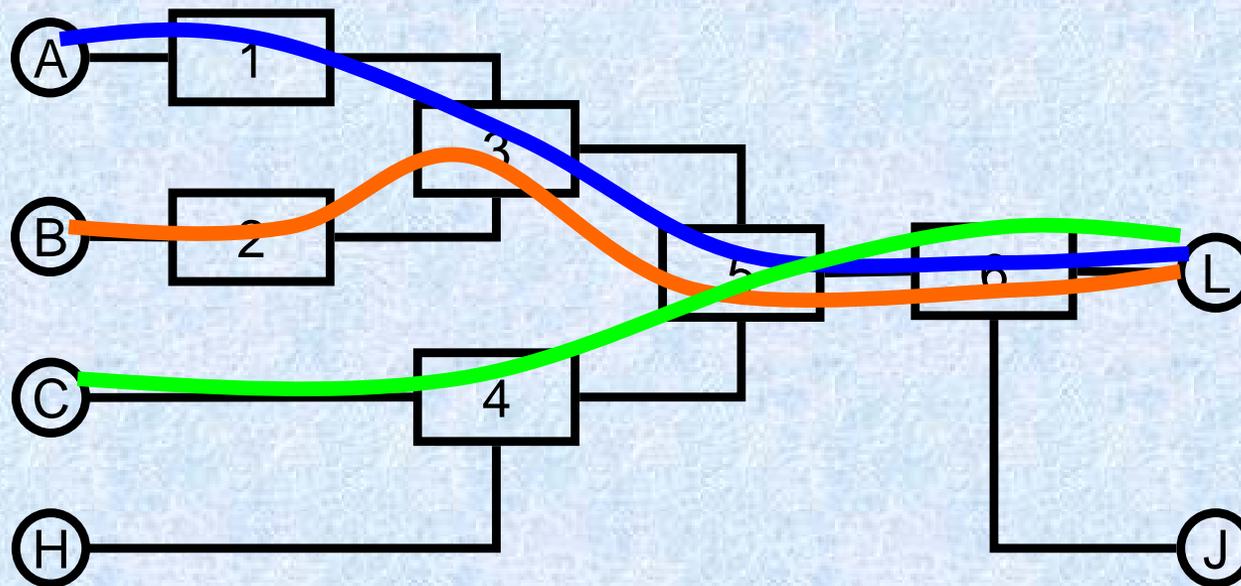
Effects of SR Classes 3

- If it is desired to have one stream of each SR class per DA+VID then HW filtering database must be expanded to include $1+n$ bits per port (where n is the number of SR classes supported on the bridge).
 - This seems to be a large price for small benefit.

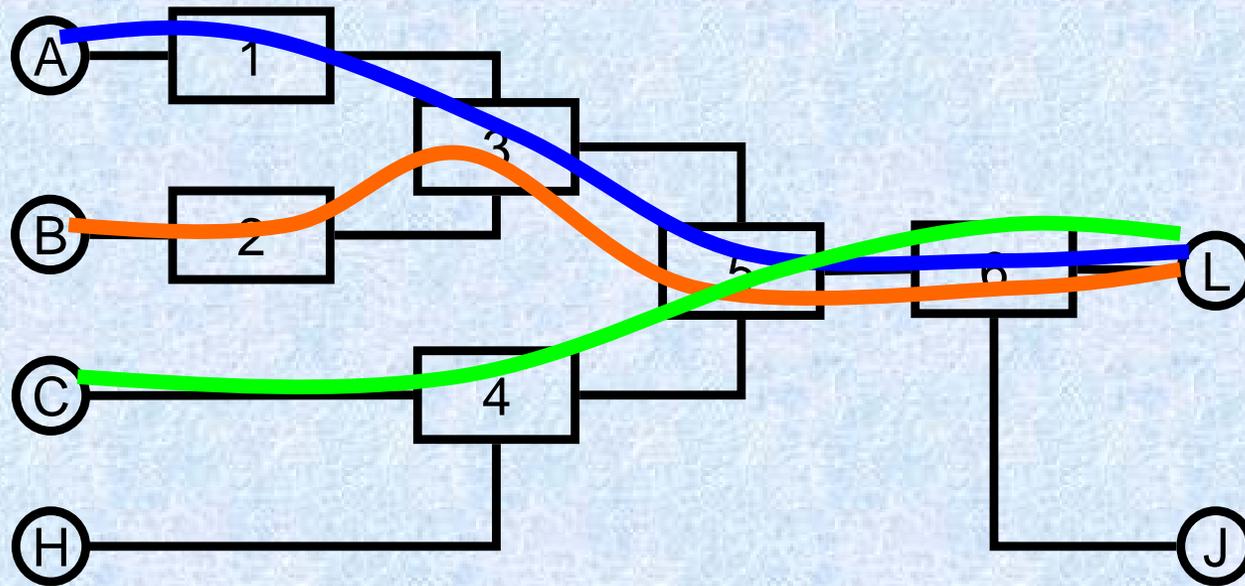
Example of Multiple Steams to Unicast MAC address



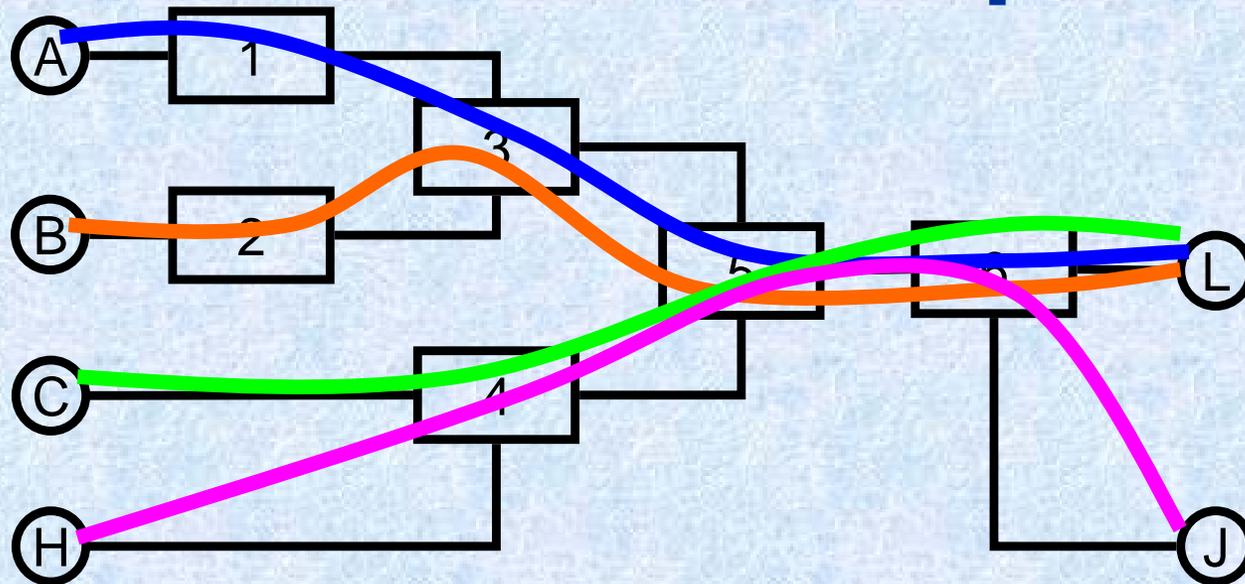
L is listening to unicast streams from A, B, & C



H & J want to set up a high rank steam, bridge 5 must make room by kicking stream B out



But B continues talking until it hears a Fail; even then a lump in bridge 3 continues to send the stream until it empties



The Result

- Even though bridge 5 selected stream B to delete, since streams A, B, & C are using L's address bridge 5 cannot determine to which of these streams a frame belongs.
- This is because they have the same destination address.
- The result is that for a period of time after stream B is deleted its bandwidth is still being used on the bus between bridges 5 & 6. This time depends on how many lumps there are in the path between B and bridge 5.