

# Path Dependent Addressing for 1394 Bridges

David Wooten  
Compaq Fellow  
Compaq Computer Corporation

# Definition

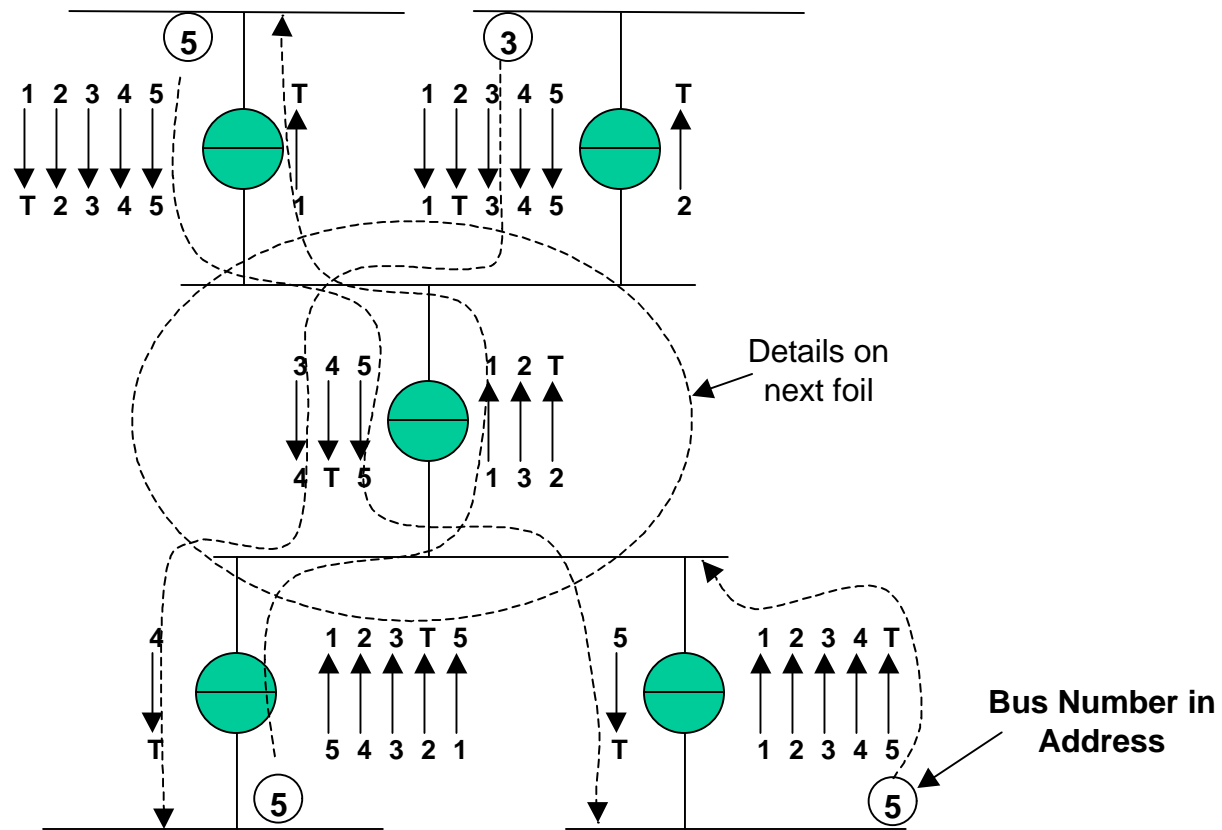
Path addressing is simply virtual bus numbers. When bus numbers are virtualized, it looks as if the address is indicative of the path that the packet is taking rather than a virtualization of the bus number.

Advantage of virtual bus numbers is that the numbers can be assigned based on information available on a bus segment rather than information available on nodes on other buses (e.g., the Prime portal).

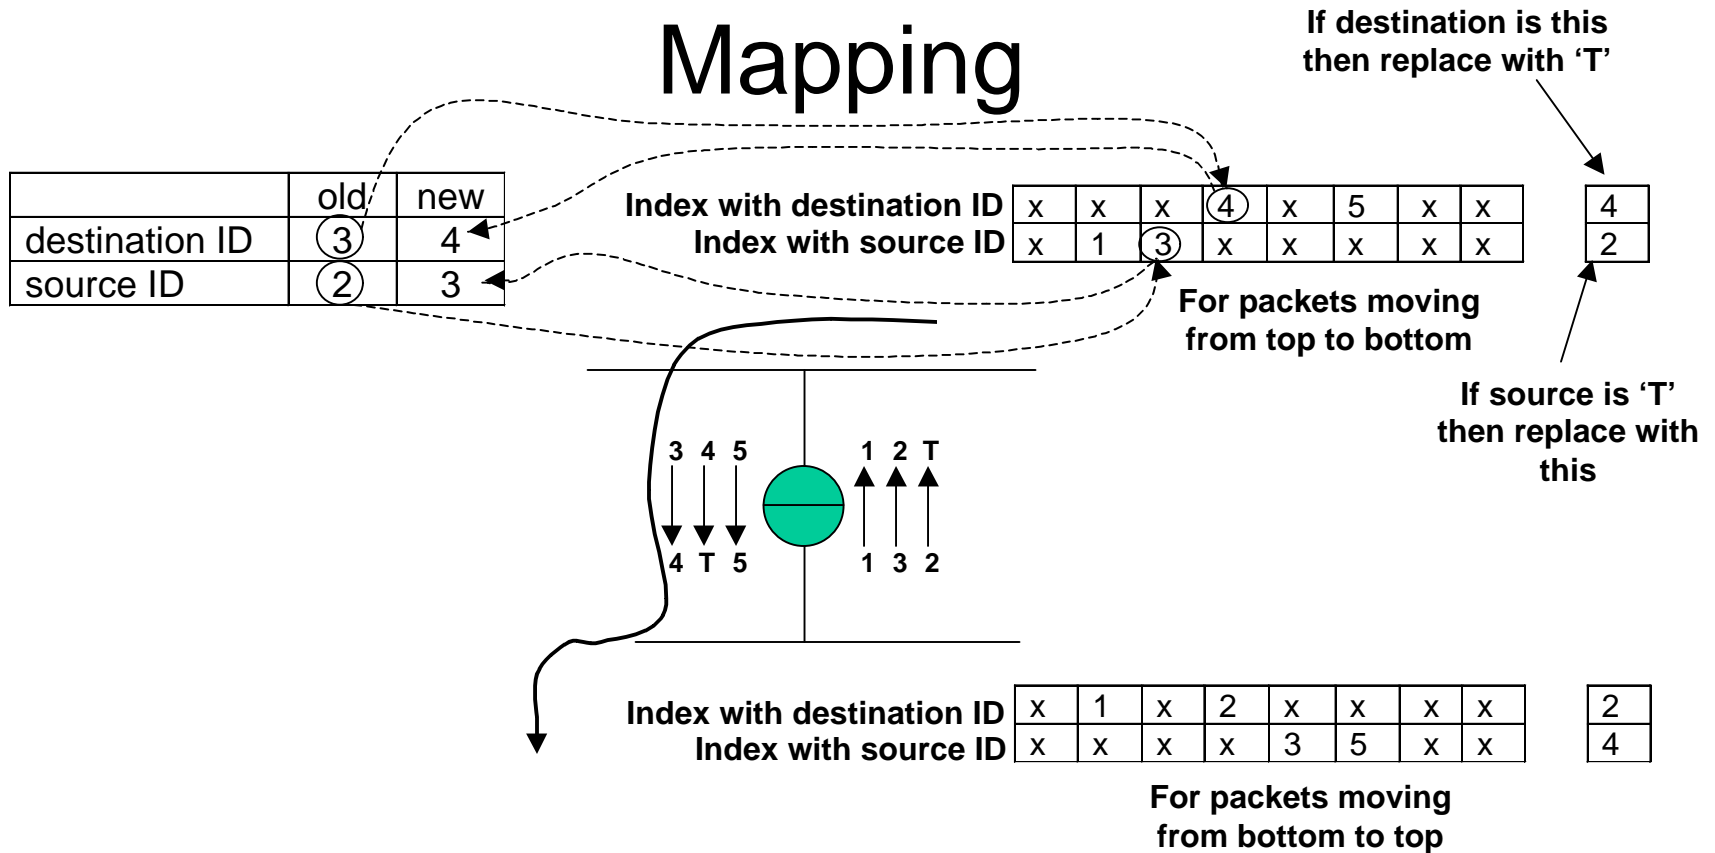
Allows us to take full advantage of persistent node ID's. Once the address of a node is known, it doesn't change even if buses are added or removed.

Another advantage of virtual bus numbers is that since each portal has a different notion of the addressing of a bus, they can have a different notion of the node ID's. This eliminates the need to coordinate persistent node numbers between portals.

# Path Addressing



# Mapping



As packet moves through bridge, both the source and destination ID's get modified.

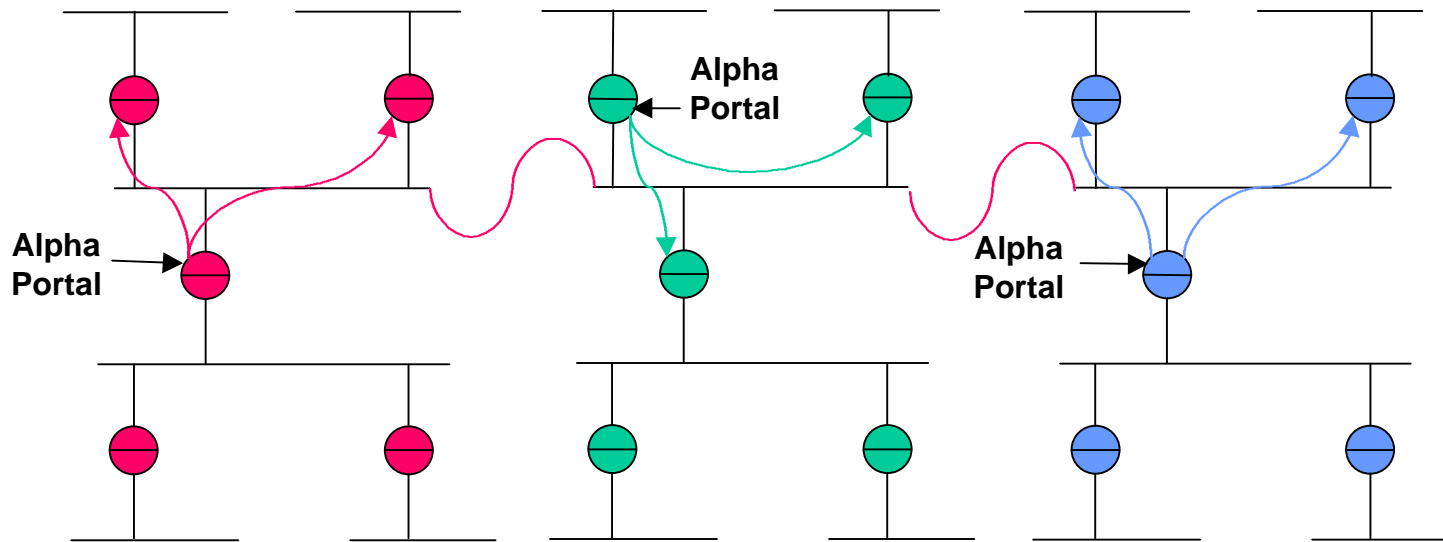
# Joining Buses - Intro

When buses are joined, some or all of the outbound address mappings may be the same. The portals must eliminate the duplicates. Also, for each outbound mapping in the 'other' group, a notification of a new bus needs to be sent to the other portals in the old group to indicate that a new bus exists. This function is similar to what needs to happen when an address is assigned to a bus (previous presentation).

# Eliminating Duplicate Mappings

The process of eliminating duplicate mappings is complicated by the fact that there may be multiple net additions.

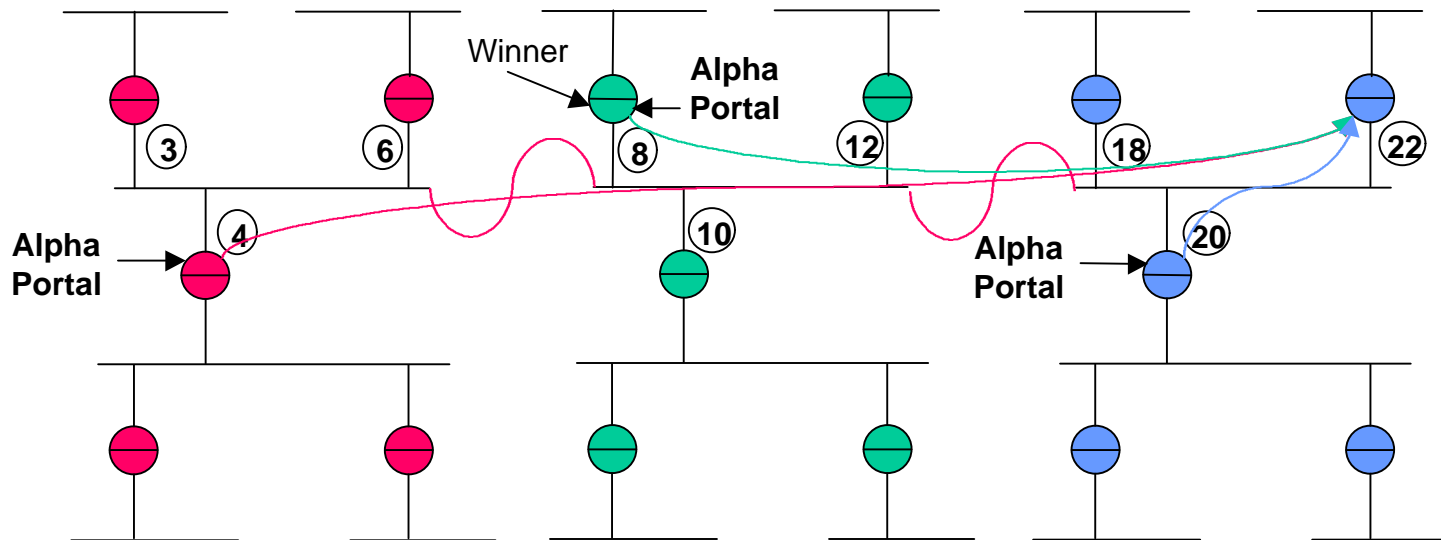
# Joined Buses (1)



Alpha Portal on each old segment collects outbound bitmap from all other portals in its old grouping.

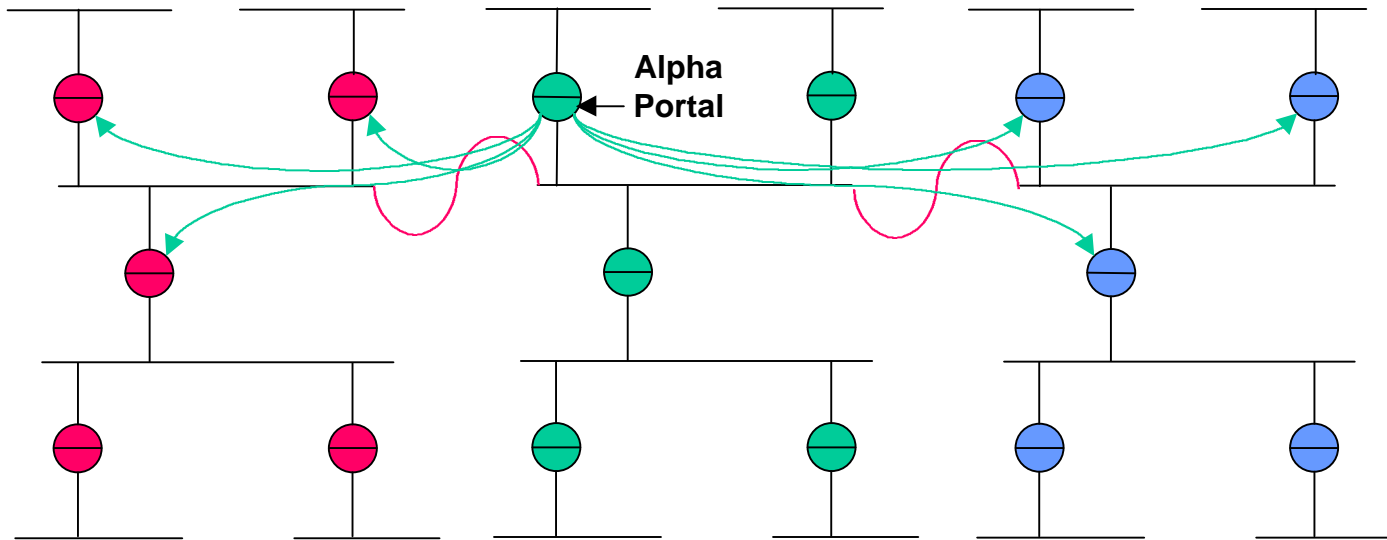
It is likely that this mapping will be current so some Alpha portals may be able to skip this step.

# Joined Buses (2)



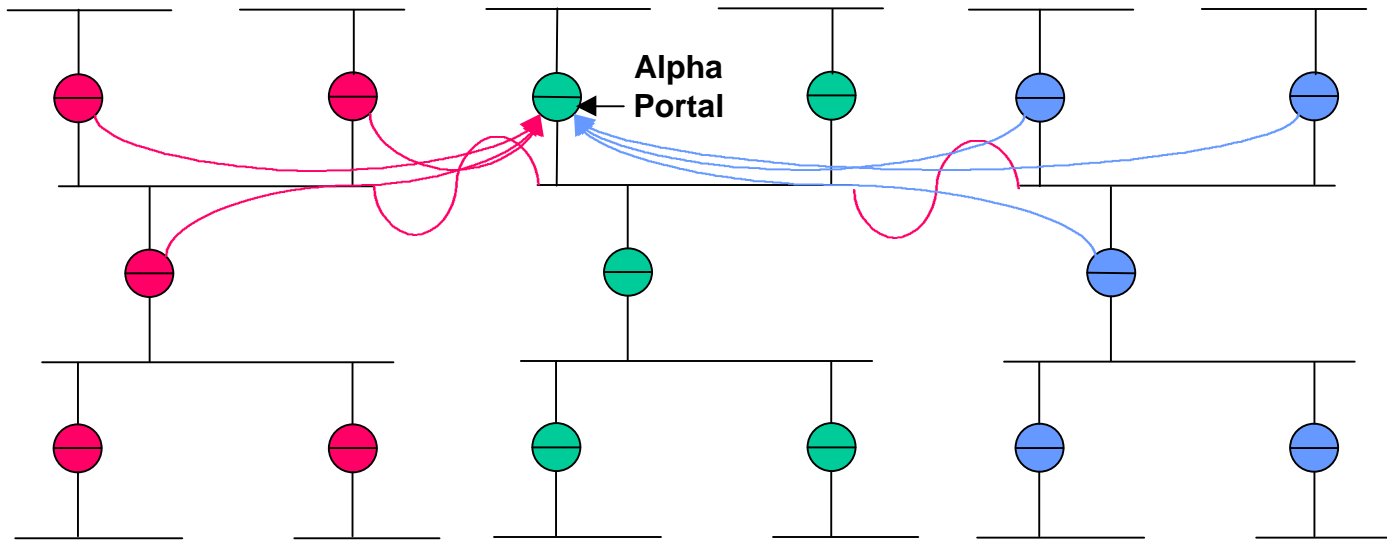
Portal with highest node number on new bus becomes controlling portal. Each Alpha tries to become 'controlling' portal (compare and swap node number -- value reset after bus reset in which change in portal configuration is detected).

# Joined Buses (3)



Winning Alpha Portal writes to all portals on the victim buses giving the address of the new Prime Portal (this write discloses the node number of the new Alpha Portal.)

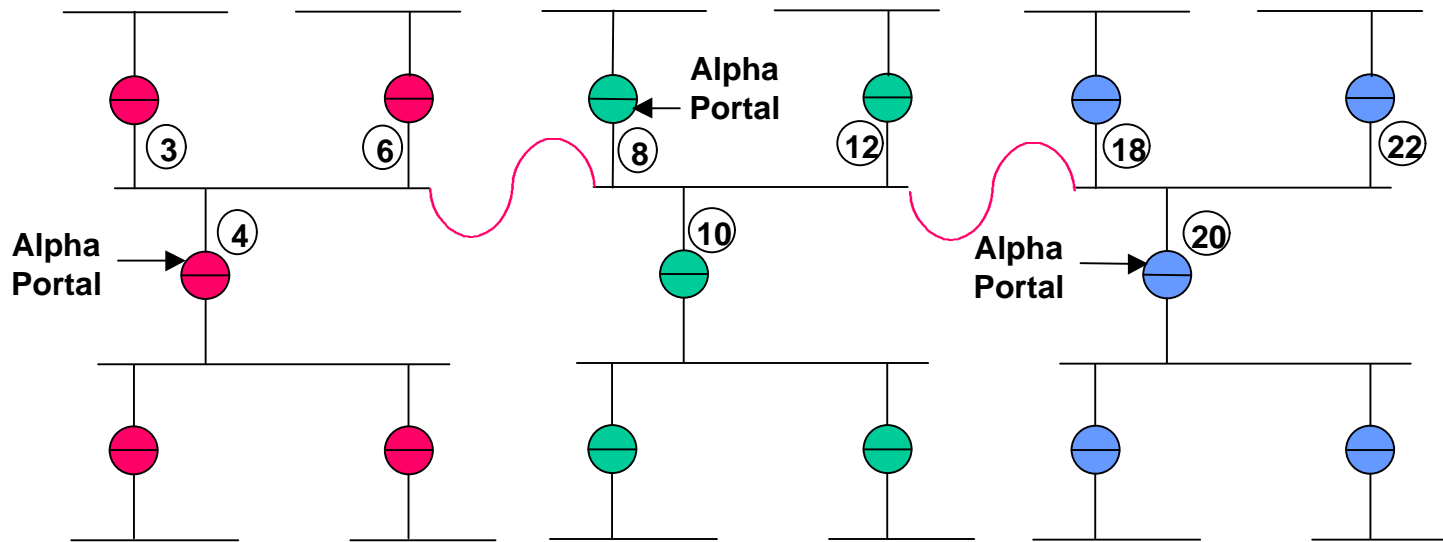
# Joined Buses (4)



Each portal on the victim net, do bit test-and-set against the outbound map in the new Alpha. The return value indicates which of the addresses are duplicates and which are not.

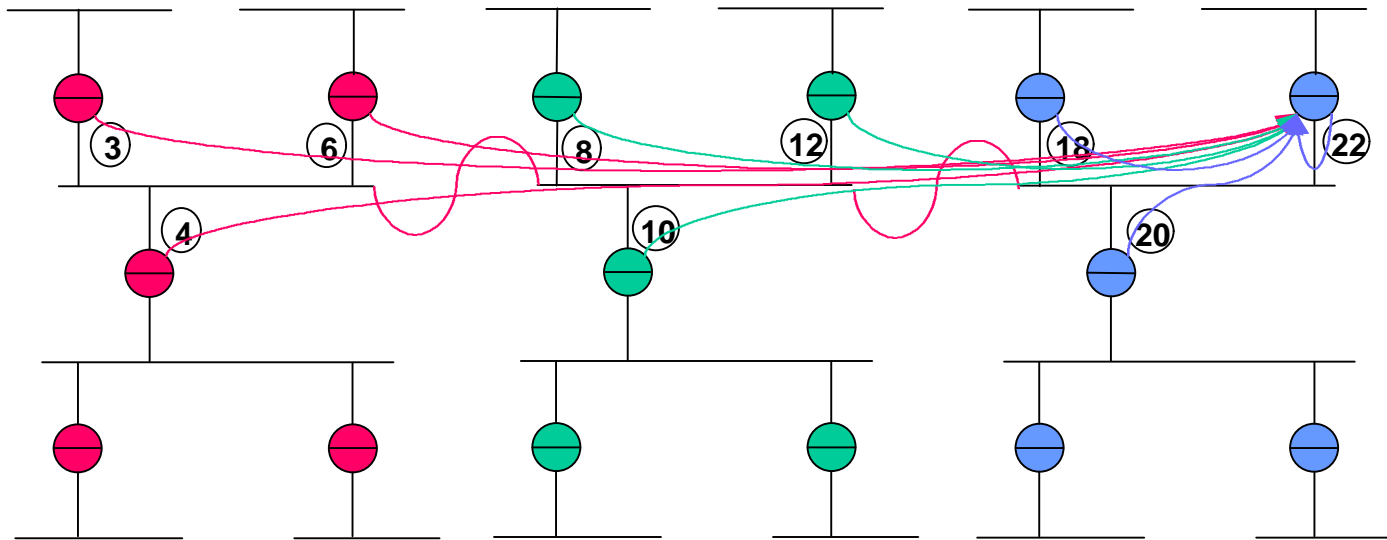


# Joined Buses (1)



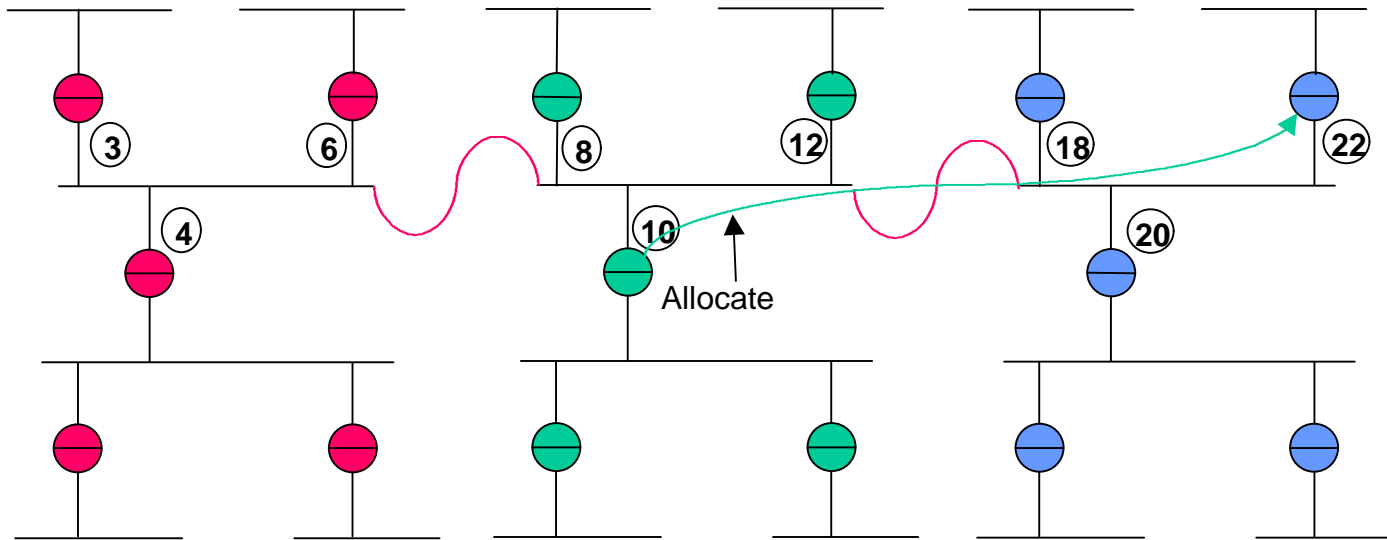
Need to pick a surviving bus. Send notification one the victim buses so that the acyclic tree is reestablished. This may not be necessary depending on how we propagate notification of new buses (might be able to have an acyclic tree centered on each bus so that the hop count to any node is minimized).

## Joined Buses (2)




Highest node number on new bus becomes arbitration portal. All portals attempt to reserve their outbound mappings by doing a test-and-set. If bit already set, then there is a duplicate outbound mapping.

# Joined Buses (3)



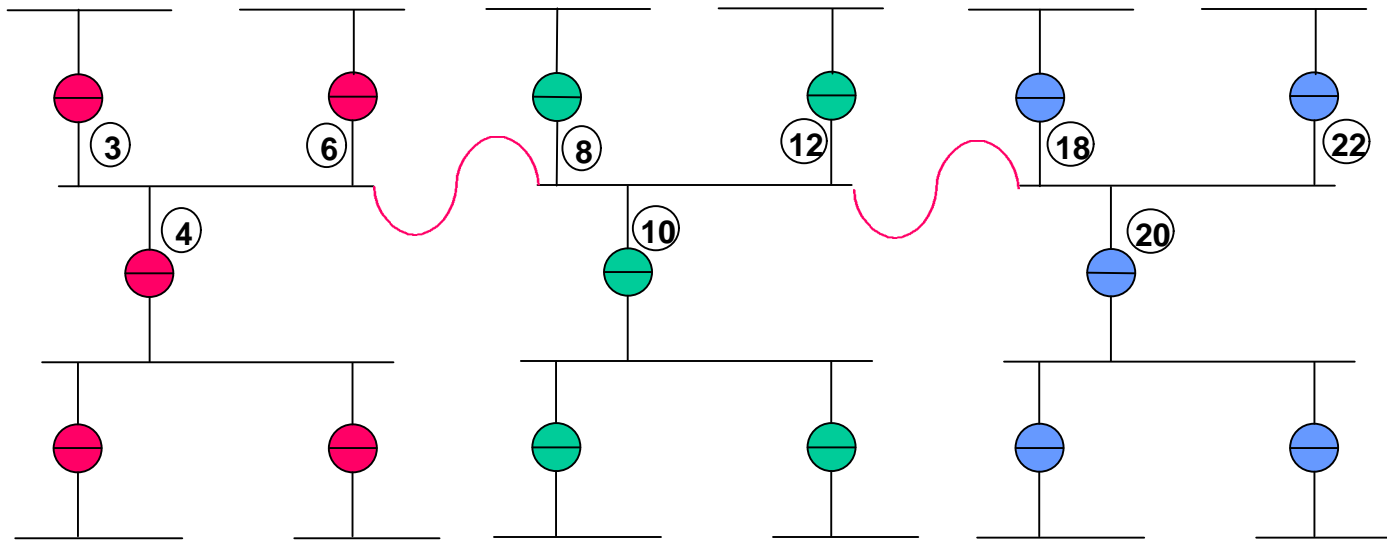
When a portal finds a duplicate mapping, it must arbitrate for an alternate mapping, again by doing a test and set at the arbitration portal

# Joined Buses (4)



After getting an alternate mapping, the portal must inform each of the other portals in its group that the mapping has changed by giving the old and new mapping numbers.

# Joined Buses (5)



When a portal has resolved all its duplicates, it becomes a member of the surviving group.

# Effect of Reset or Loss of Port

If a bus reset occurs while the bus is grouped, each of the portals goes to the highest portal in the victim group to see if the numbering process is complete. If not, the portals remain grouped.

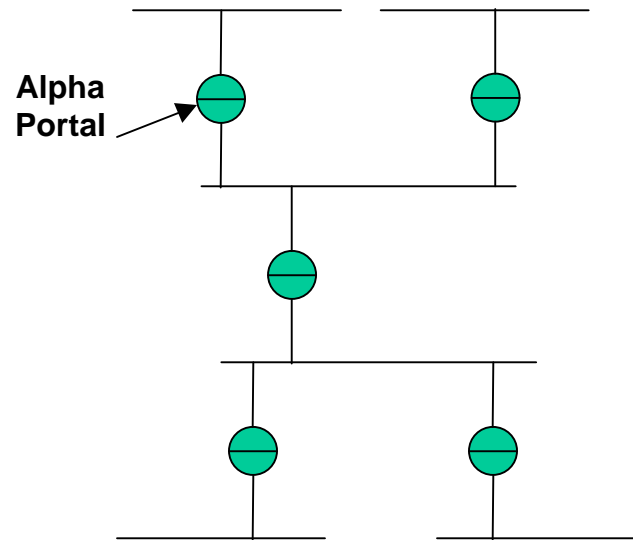
If the high portal in the victim group is removed during the renumbering stage, the new high portal takes over. It gets back to the known state by reading the outbound values from the remaining victim portals and reading the combined bitmap values from the high portal in the survivor group.

# Path Addressing, Address Assignment

The map that each node received at the end of the previous step, is the combined outbound map of all portals on the bus. For every outbound address, there must be a corresponding inbound address on every other portal. When a portal receives the map and the bus is reset, it checks the map to see where it needs to add inbound connections.


Following foils indicate how address assignment takes place in path addressed environment.

# Path Address Assignment (1)



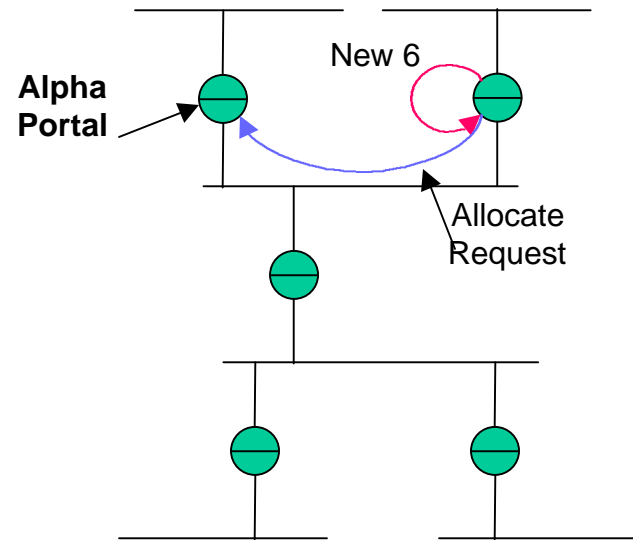
Each bus segment has an Alpha portal. That portal contains an up to date bit map of all outbound mappings that are used by the portals on that bus.

# Path Address Assignment (2)



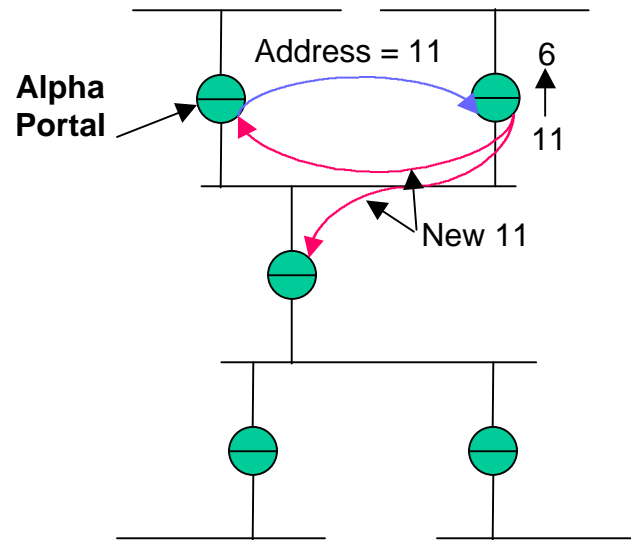
Notification of a new path (bus) will arrive from a portal connected to another bus. That notification will indicate the bus number that is used on the outbound path

# Path Address Assignment (3)



The new notification is sent to the co-portal. The co-portal requests a new, outbound path assignment from the Alpha portal.

# Path Address Assignment (4)

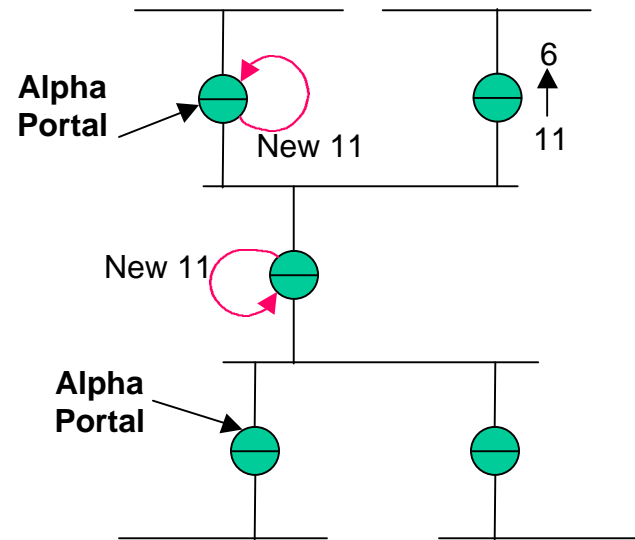


The Alpha portal returns an unused path number.

The requesting portal uses that number as an outbound mapping for the original number.


The requesting portal then sends a New indication to all other portals on the same bus.

# Path Address Assignment (5)



Process repeats.

# Path Address Assignment (6)



Eventually, there is a path from all buses, back to the original bus.

# Path Dependent Pros and Cons

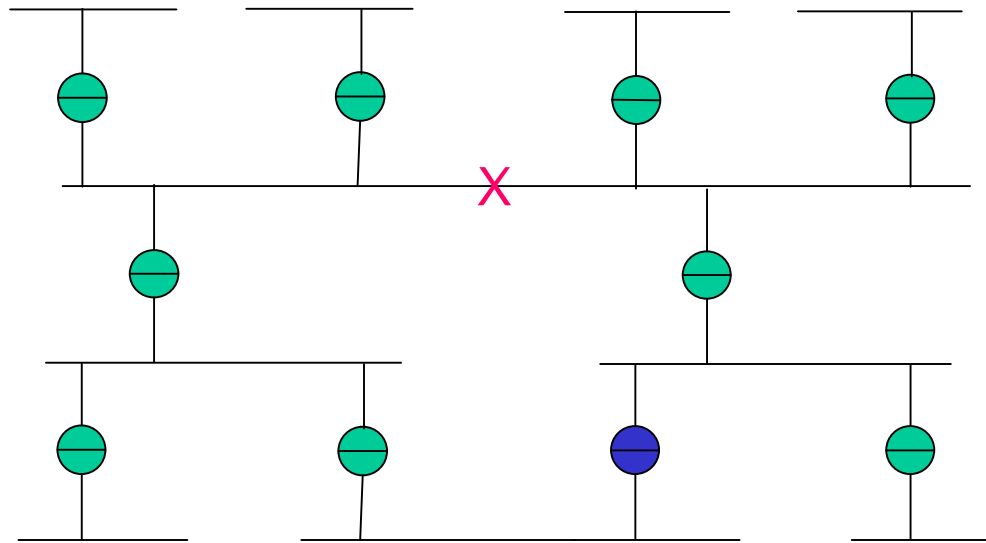
## Pro

- 1) address that is used for a node does not change as long as the path to that node exists.
- 2) because address to a node is dependent on path, there is no reason that the portals have to agree on the virtual node ID of a node.

## Con

- 1) can't pass the address of a node in a packet. Can only pass its GUID. Means an ARP-like facility is needed.
- 2) more mapping memory is needed for full implementation

# Reorganizing



When a loop existed in network and loop is physically broken, routing needs to be changed. Propagating the need to reroute virtual bus ID's is 'tricky'. Result is that most likely solution is that we would issue a Net Reset, which is the same thing that we would do with fixed numbers.

# My Conclusion

In many cases of topology changes, the action taken with virtual bus numbers and unified bus numbers is the same (Net Reset). In a significant number of cases the Net Reset generated in a unified numbering scheme is not all that disruptive. From this, I have concluded that the advantages of virtual bus numbers is not worth the effort.