ATM Routing and Switching

Adam M. Costello <amc@cs.berkeley.edu> 1999-Oct-08-Fri

host

switch

host

switch

switch

virtual

circuit

nputs

virtual

switch

path

switch

host

switch

host

outputs

A connection is a virtual path or virtual circuit. A virtual path acts like a single link carrying many virtual circuits, to reduce connection state in switches. A switch maps input (port,VPI,VCI) to output (port,VPI,VCI). Within a virtual path, the VCI is simply copied from input to output.

To set up a connection, the first switch chooses a course-grain path and sends a connection request to the next switch, which sets up tentative connection state, fills in finer-grain details of the path if necessary, and sends a connection request to the next switch... If a request is rejected, the requestor can try another path or send a rejection to the previous switch. If a path is found, a message is sent back along the path to confirm the tentative connection state.

Switches are grouped hierarchically according to their addresses. The figure below shows the view of switch 2.3.6. To connect to switch 3.7.6 it might choose the path (2.3.3, 2.9, 5, 3).



Each group elects a leader to exchange metric information (via flooding) with its peers in the higher-level group.

A single switch is a network of switching elements in a box. For example, we can connect twelve 2x2 switching elements (2 inputs and 2 outputs) to create an 8-port switch:

We also need a translation table at each input, to map input (VPI,VCI) to output (port,VPI,VCI). The port number can be prepended to the cell and used to route the cell through the switch. In the example, number the ports from 0 to 7 (000 to111 in binary). To get to output ABC from any input, choose output A of the first switching element, output B of the second, output C of the third.

Cells arriving at multiple inputs may all need to go to the same output, so we need buffers at the inputs, or outputs, or both. If the inputs are buffered, we have head-of-line blocking (cell destined for idle output stuck behind cell destined for busy output).

Two connections with different inputs and outputs should not interfere, but they might need to share an internal link. This can be solved by increasing the speed of the internal links and switch elements, or by increasing the number of stages to provide more paths. Dynamic routing spreads each connection over all paths to avoid concentrating traffic anywhere, but then cells must be put back in order at the output.

The number of switching elements needed for an n-port switch grows as n log n. Each switching element can be a crossbar, a bus, or a ring. For any of these, the complexity of a d-port switching element grows as d squared, which is why a single switching element is not used for a 65536-port switch.

4x4 crossbar (Each crosspoint can be connected or disconnected.)



4x4 bus Bandwidth must be O(d). Clock rate is O(1/d) (capacitance). So bus width must be O(d^2).



4x4 ring Bandwidth must be O(d) So d links each of width O(d).

