IP Multicast over ATM

Greg Minshall
minshall@ipsilon.com

March 25, 1996

Abstract

This presents 4 viewgraphs I showed at the Los Angeles IETF, in the IDMR session, which show how we run IP multicast over ATM. In the interest of brevity, this is somewhat simplified; full details are available to anyone interested.

1 Introduction

There is a default VCI which is used for transmitting SNAP-encapsulated IP packets over an ATM link. In the diagram that follows, the *circles* represent ATM switches; the *rectangles* represent computers (which we term *IP switch controllers*) which run standard IP unicast routing (OSPF and RIP, say) and multicast routing (DVMRP, say) software, as well as other standard IP software (such as SNMP, etc.); and the *square* is an IP host.

The solid lines are (ATM) links; the dotted lines represent the flow of data (*downstream*; left-to-right); the dashed lines represent the flow of control messages (*upstream*; right-to-left).

2 Hop-by-hop forwarding

In this figure, IP datagrams are being multicast to some multicast group from a system to the left of the figure; the host is one receiver of the multicast traffic; there is another receiver of the multicast traffic in the direction of the lower right-hand corner of the diagram.

As can be seen in the figure, the multicast traffic is received via the default VCI at each system. This means that the traffic is being hop-by-hop forwarded by each of the IP switch controllers (the rectangles).

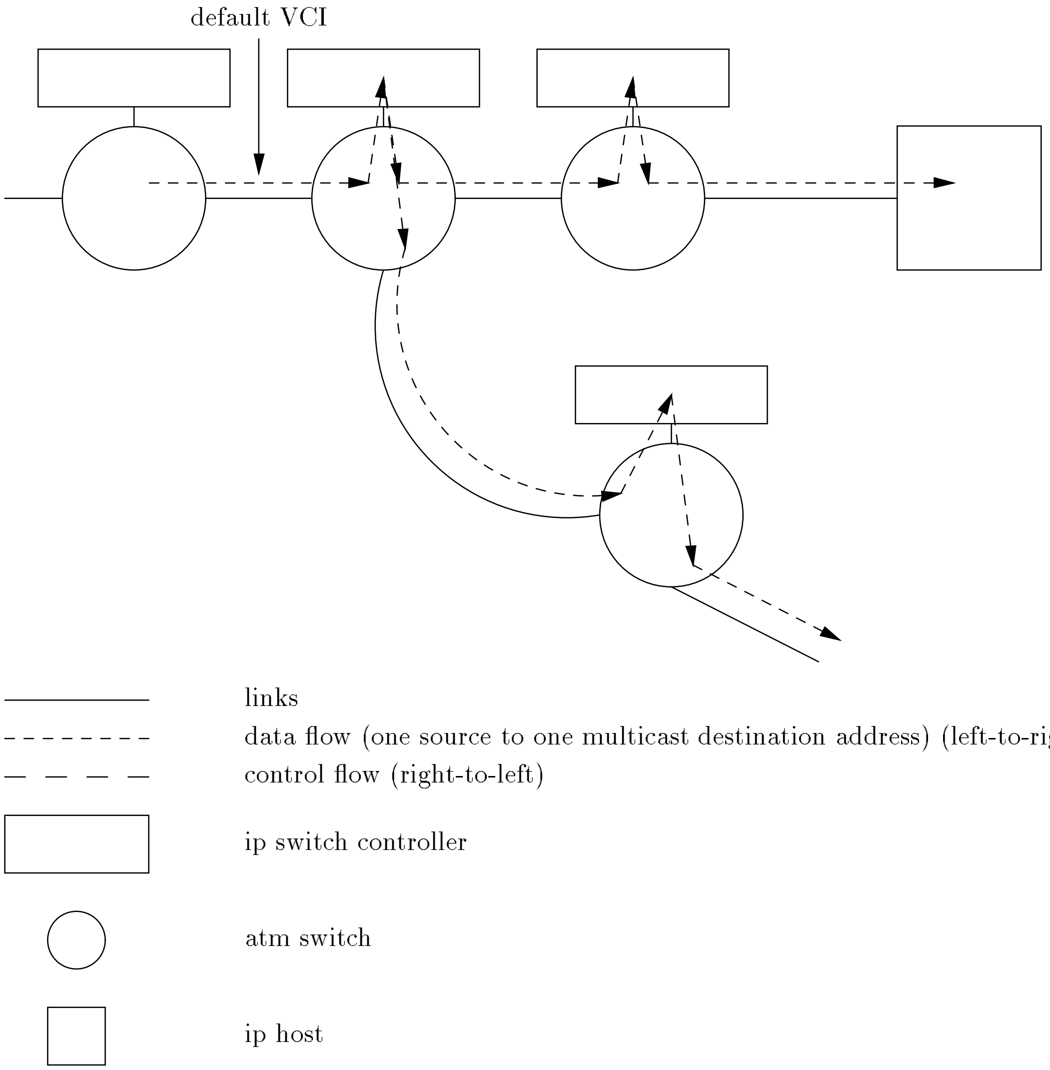

3 Send redirect upstream

When an IP switch controller or IP host decides that there are likely to be enough multicast packets sent from the source host to the multicast group to make it worthwhile to optimize these packets, it generates a *redirect* message to its upstream neighbor. This redirect message (similar to an *ICMP host redirect*) requests that the upstream neighbor transmit future packets from that source host to that multicast group on a specialized (non-default) VCI; the specific VCI is chosen by the system sending the redirect and is included as a parameter in the redirect message.

In this figure, we see a redirect message travelling from the second IP switch controller to the first IP switch controller.

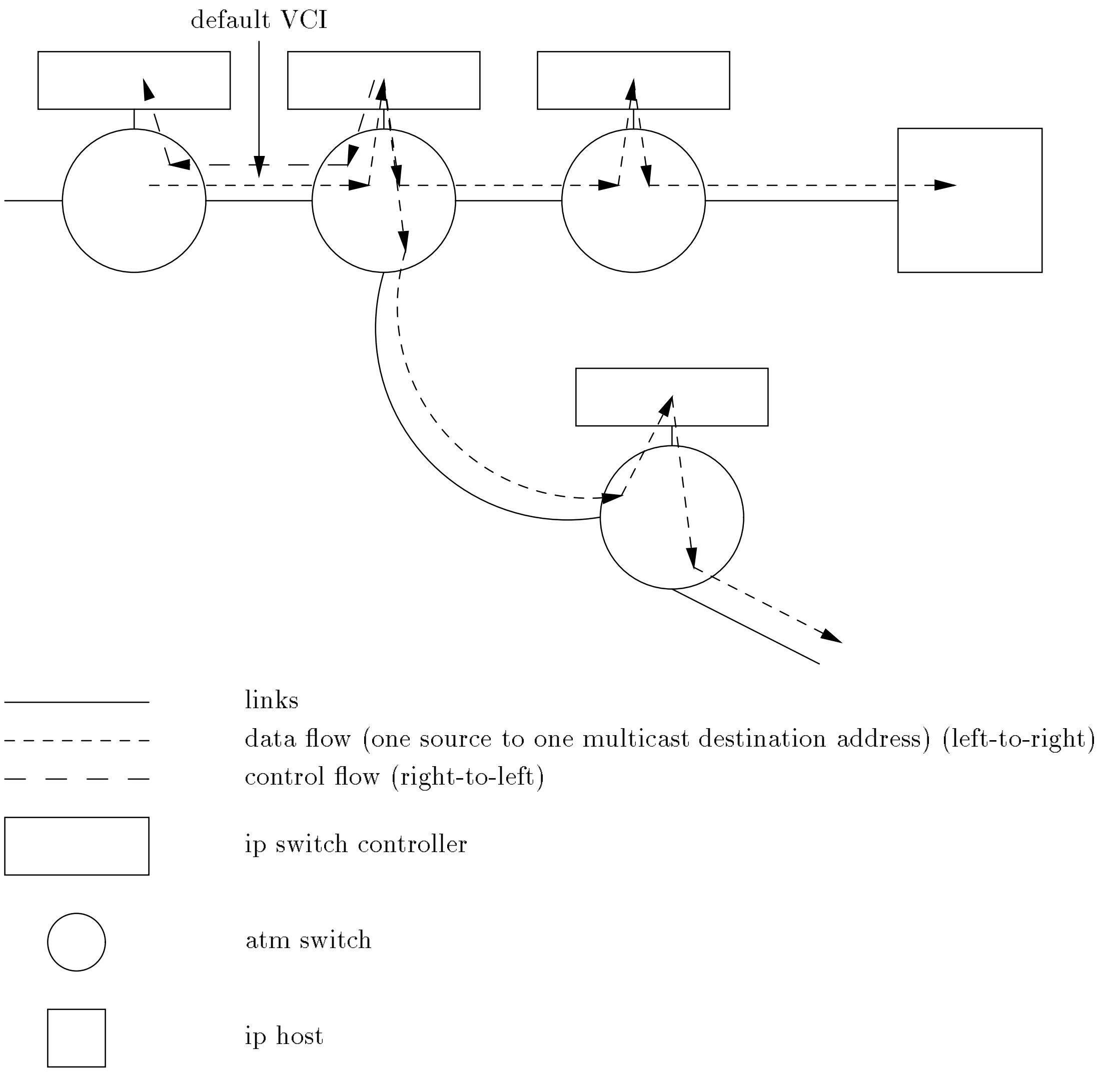

4 Downstream nodes send redirects

As the multicast datagrams propagate downstream, downstream neighbors send more redirect messages to *their* upstream neighbors; this is shown in this (messy) figure.

Also shown in this figure is the fact that the first IP switch controller is sending the packets for this multicast flow to the second IP switch controller on the specialized (non-default) VCI; the second IP switch controller, however, is forwarding the datagrams in this multicast flow to its downstream neighbors using the default VCI (since it has not yet received redirect messages from these downstream neighbors).

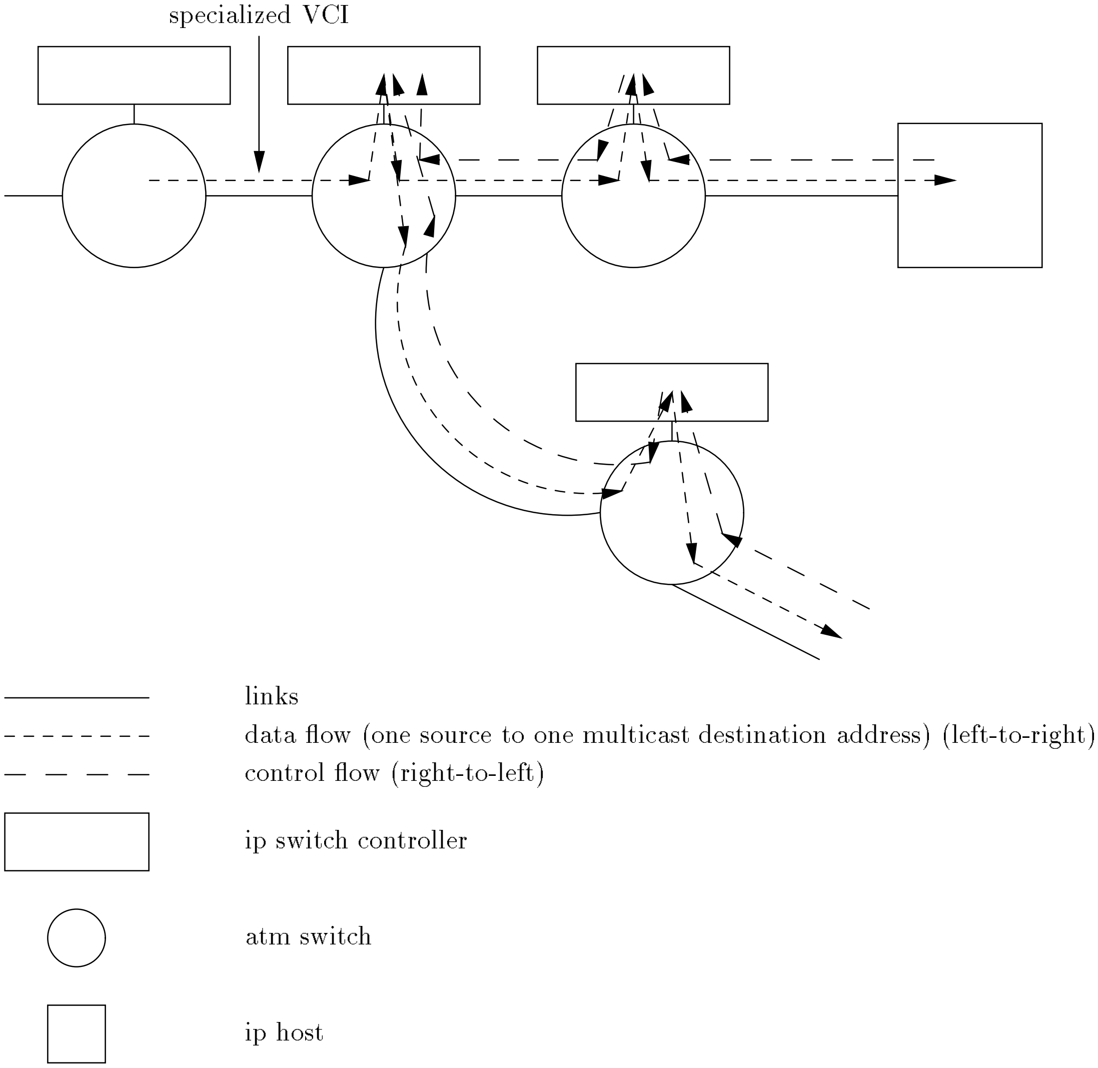

5 Fully switched

When an IP switch controller is receiving traffic on a given flow on a specialized VCI and transmitting the traffic on that flow on a specialized VCI, the IP switch controller can change the mapping tables in its attached ATM switch and cause the datagrams sent on that flow to be switched directly from the input port and VCI to the output port(s) and VCI(s).

This behaviour is shown in this figure, where both outgoing streams of data from the second IP switch are being switched from the input stream. (It is also possible to have some downstream neighbors receive their datagrams from the IP switch controller and others receive their datagrams directly from the ATM switch.)