Defining the IP Multicast Service Model

Mark Handley

USC/ISI

The Multicast Service Model

1. ***Senders send to a Class-D destination address.***
2. ***Receivers express an interest in receiving traffic destined for a Class-D group address using IGMP.***
3. ***Routers do what is required to ensure data travels from senders to receivers.***

[3] is not really properly defined anywhere, and this causes problems for designing complex protocols such as RMPs that need to be able to assume something about how data is distributed.

Potential Problems

Two that are cause for concern are:

- **TTL scoping**
- **Bursty sources**

TTL Scoping

There are two aspects to TTL scoping:

- senders restricting TTL to achieve locality of distribution within a larger group
- routers implementing thresholds to constrain groups

The latter should be phased out in favour of administrative scoping, but the former is very useful, especially for RM applications.

The Requirement

To be able to constrain selected traffic to a part of the network that is "topologically close" to the sender.

This is a requirement for almost all RM schemes that implement any form of hierarchy, or wish to scalably locate group members that are likely to be receiving similar network performance from a third party.

Not all multicast routing protocols implement TTL scoping usefully.

Sparse-mode PIM

When a new sender sends, initially data is encapsulated to the Rendezvous Point

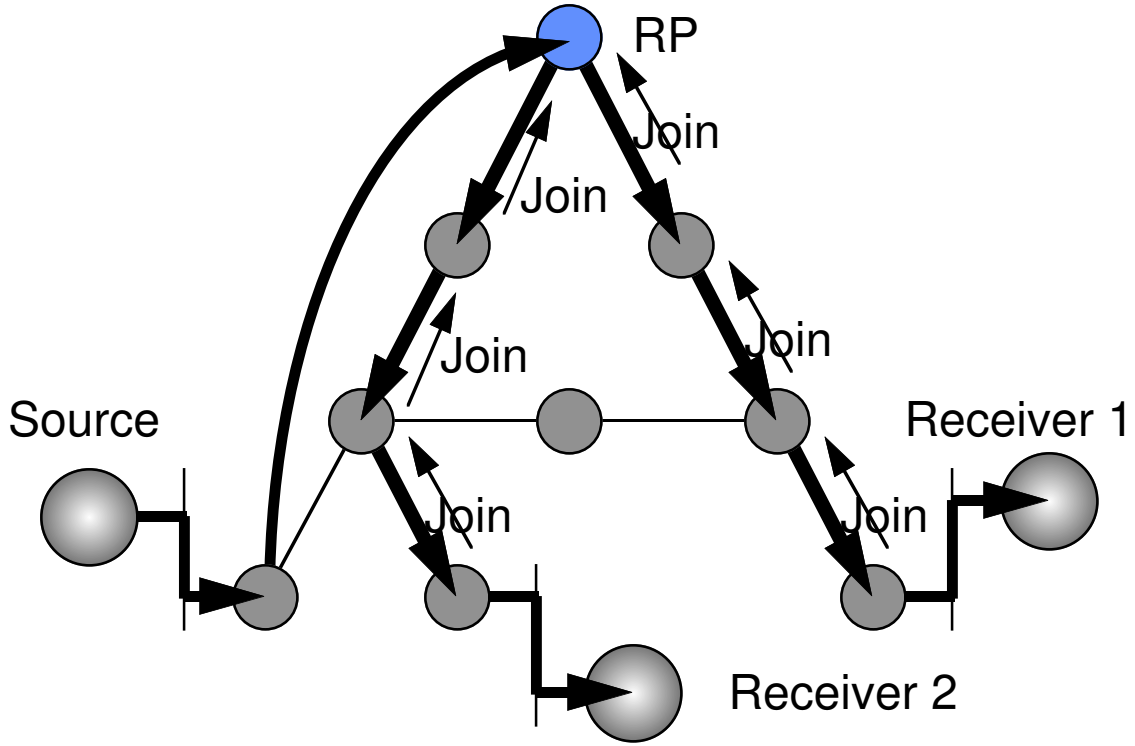

Sparse Mode PIM

There are two choices in SM-PIM for how to encapsulate:

- **Encapsulate with the original TTL**
  - in which case the packet doesn't reach anyone until the TTL is sufficient to reach the RP, an d then it starts to find sites close to the RP not the sender.
  - distributed expanding ring searches turn into a delayed implosion at the RP
- **Encapsulate with a fixed high TTL**
  - in which case all packets with TTL>=2 reach the RP
  - distributed expanding ring searches turn into an implosion at the RP

Sparse Mode PIM

- **There are ways to fix SM-PIM to support a manner of TTL scoping, but they involve changes to the PIM architecture and protocol.**
- **Also if the PIM RP chooses to inititate a source-specific join back to the source the problem goes away, but doesn't help RM schemes because we don't want the implosion that is required to initiate this behaviour, nor the pointless routing state it induces.**

Bursty Sources

- **Nowhere does it say what acceptable behaviour is for the first packet(s) sent from a sender.**
- **Is it acceptable to drop those packets whilst initiating a lookup for (for example) a Rendezvous Point in a distributed database?**
- **If this was acceptable behaviour, then RM applications that send a small number of packets occasionally (eg SRM NACKs or Session Messages) would deterministically have those packets dropped.**

Summary

The IP Multicast Service Model should be written down in enough detail that applications know what behaviour to expect irrespective of the multicast routing scheme, but with enough freedom to allow creative solutions to multicast routing.

So, what behaviour do we actually want to specify???