E6998-02: Internet Routing

Lecture 20
IP Multicast Routing Protocols

John Ioannidis

AT&T Labs – Research

ji+ir@cs.columbia.edu

Announcements

Lectures 1-20 are available.

Homework 5 will be out soon.

I'll be grading all your homeworks next week.

Joel Gottlieb from AT&T Research is giving the guest lecture on Tuesday the 19th.

Al Broscius from Morgan-Stanley may be giving the guest lecture on Thursday the 21st.

Sparse vs. Dense Topologies

- Protocol performance depends on “density”.
- Dense topology: large fraction of hosts belong to multicast groups.
  - DVMRP
  - MOSPF
  - PIM-DM
- Sparse topology: small fraction of hosts belong to multicast groups.
  - CBT
  - PIM-SM

Joins

- How do group members join a WAN Multicast group?
- Implicit joins:
  - Sender-initiated.
- Explicit joins:
  - Client-initiated.

Implicit Joins

- *Broadcast-and-prune / flood-and-prune*.
- Sender initiates session.
- Router uses reverse-path-broadcasting.
  - Sends packets to all interfaces but the upstream.
- Initially, all internet routers get the traffic.
- When a router with no group members in its attached LANs, and no downstream routers gets a packet, it sends back a *prune message*.
- Prune messages propagate back to the source.
- Taking entire branches off the multicast tree.

Implicit Joins: RPB

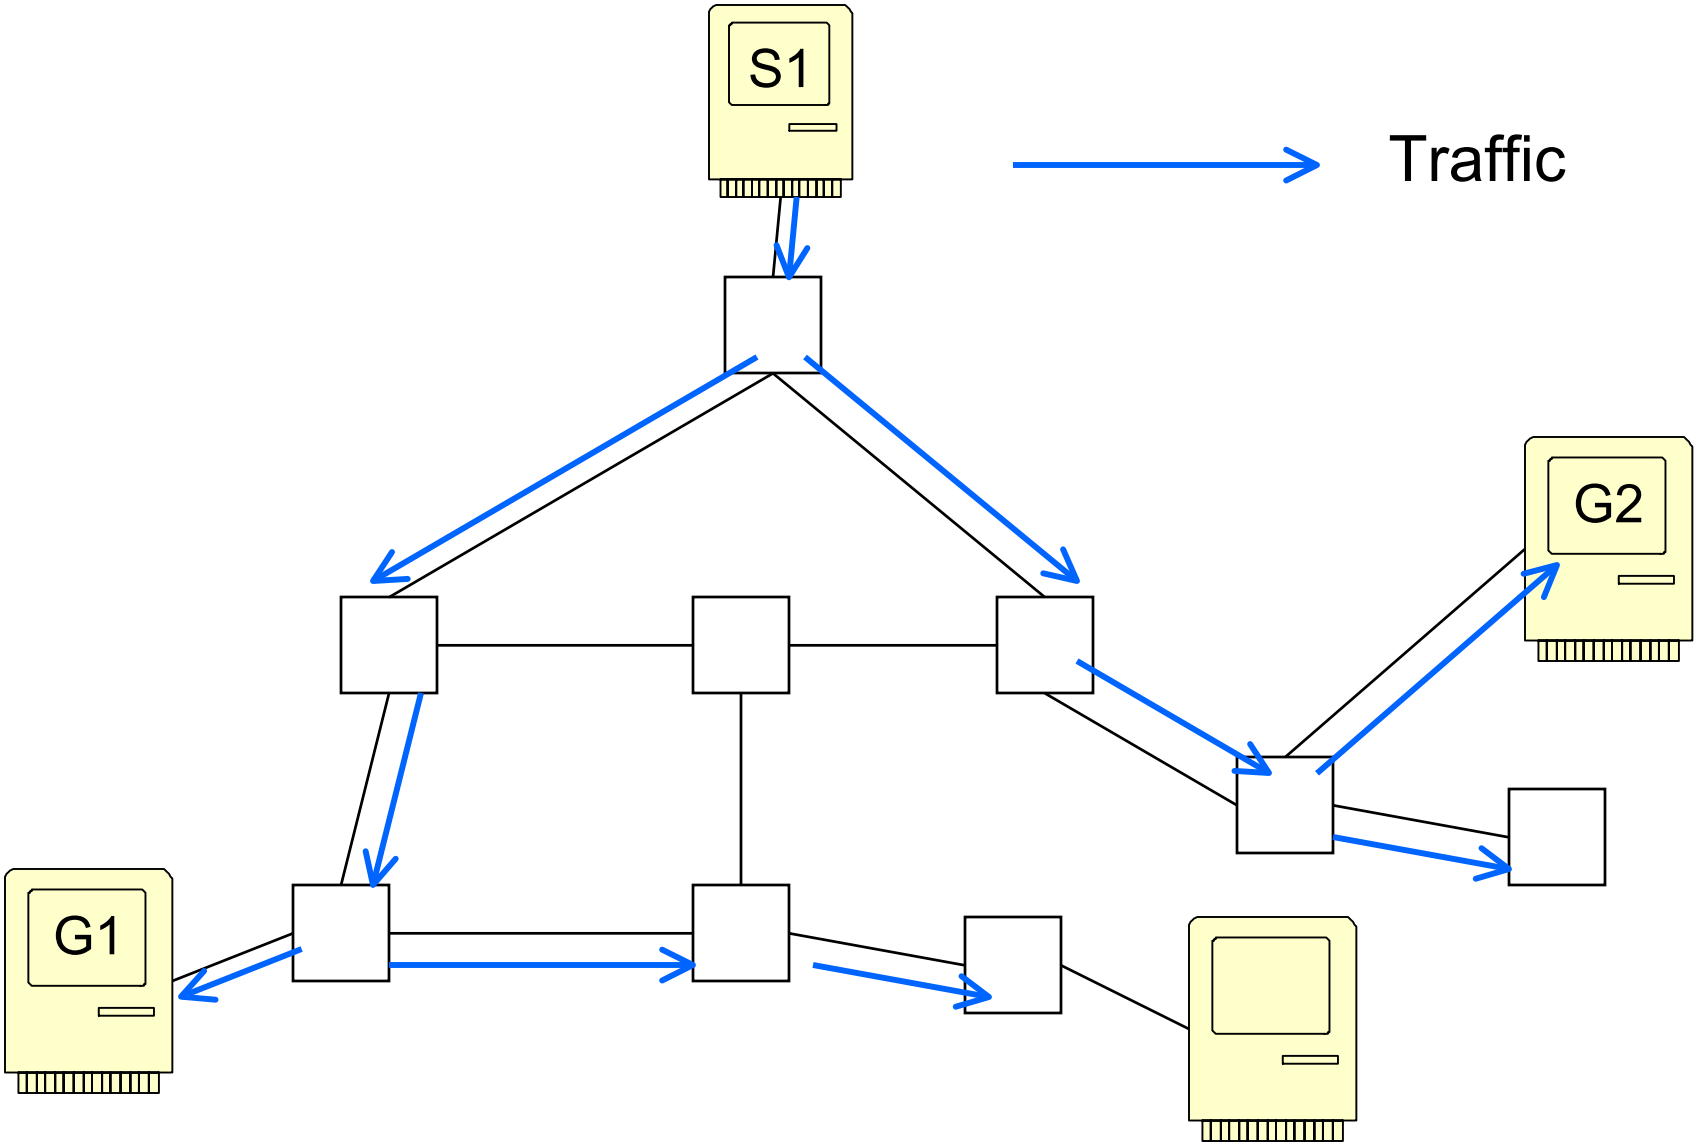

Implicit Joins: Prunes

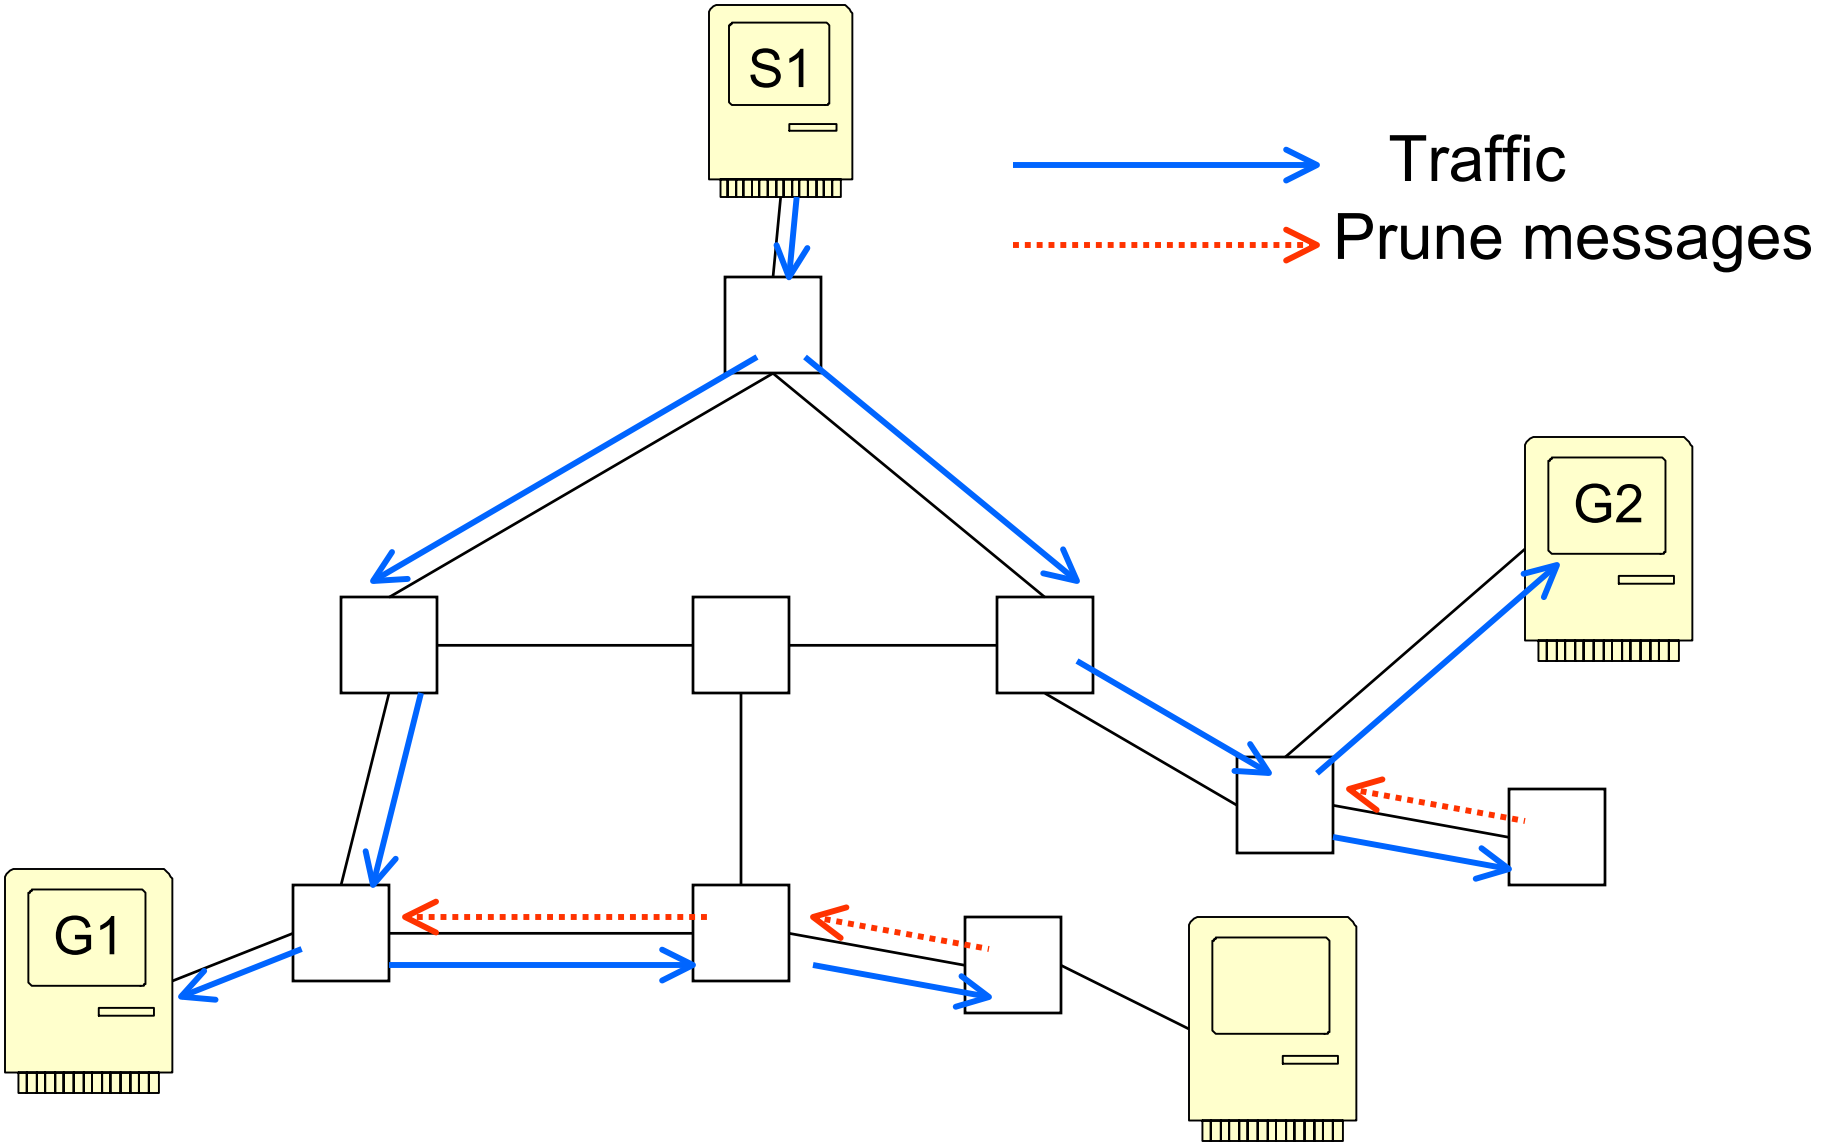

Implicit Joins: Resulting Tree

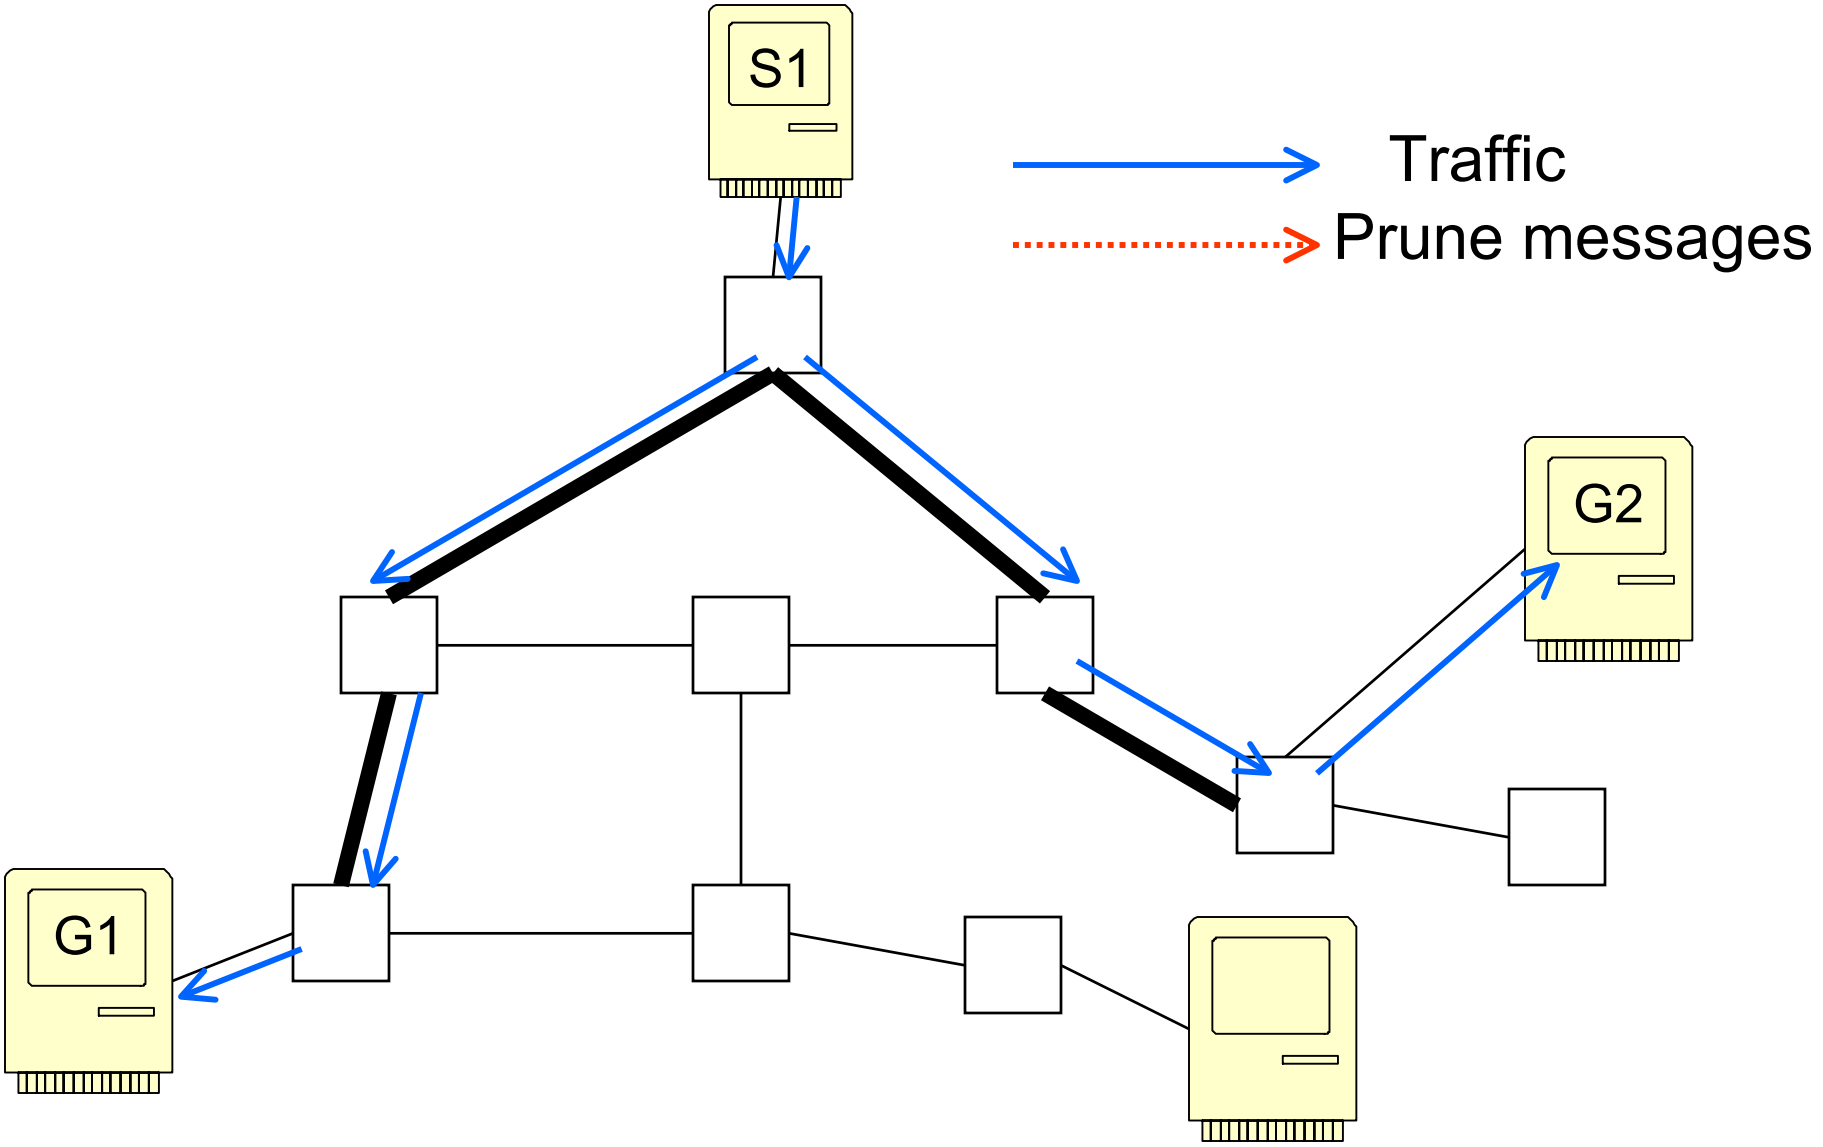

Implicit Joins, cont'd

- Each router is in either the *prune* or the *forward* state.
- Prune state has timer.
- When timer expires, router moves to forward state.
  - This way new group additions can be discovered.

- B&P is better suited to dense topologies.
- Maintenance of prune states results in high resource utilization.
  - Why should a router that would never have any multicast members have to know about all this?

- DVMRP and PIM-DM are implicit-join protocols.

Explicit Joins

- Router gets IGMP message from one of its nodes.
- Sends a *graft* message upstream.
  - Propagated further upstream all the way to the source.
  - Routers in the path join the tree.
- Source must be known.
- When no more nodes remain, router prunes itself from the tree.

- Better suited to sparse topologies.

- MOSPF, PIM-SM, and CBT are explicit-join protocols.

Source-Based vs. Shared Trees

- *Source-based* trees: rooted at the source.
  - Separate tree for each multicast source.

- But: multicast group membership changes.
- Sources can also change.
- Or there can be multiple sources.

- However: there will usually be a shared subtree.

Shared Trees

- Many multicast trees share some routers.
- Tree is rooted at a shared router: *Rendezvous Point (RP)*.
  - Or *core*.
- Source registers with RP.
  - Source's router may have to find best path to RP.
  - RP may have to find path to each source.

- Shared trees are more scalable:
  - Preferred for sparse topologies (PIM-SM and CBT).

Multicast Scoping

- Some traffic may not be of interest to the entire network.
- TTL (ab)used to specify the *scope* of traffic.
  - 0: host
  - 1: subnet
  - 15: site
  - 63: region
  - 127: worldwide
  - 191: worldwide limited bandwidth
  - 255 Unrestricted

- Administrative scoping: RFC2365 (similar to IPv6 scoping).

Scoping in IPv6

- IPv6 multicast addresses are in the FF00::/8 range.

- FF00::/12 are well-known addresses
- FF10::/12 are transient addresses
- Scope:
  - 0, 3, F: reserved
  - 1: interface-local
  - 2: link-local
  - 4: admin-local
  - 5: site-local
  - 8: organization-local
  - E: global
  - 6, 7, 9, A, B, C, D: unassigned.