# IP Multicast and the MBone:
## Enabling Live, Multiparty, Multimedia Communication on the Internet

Steve Deering
Xerox Palo Alto Research Center
deering@parc.xerox.com
http://sandbox.xerox.com/deering

December 1995

# Multicasting: Transmitting a Packet to a Group

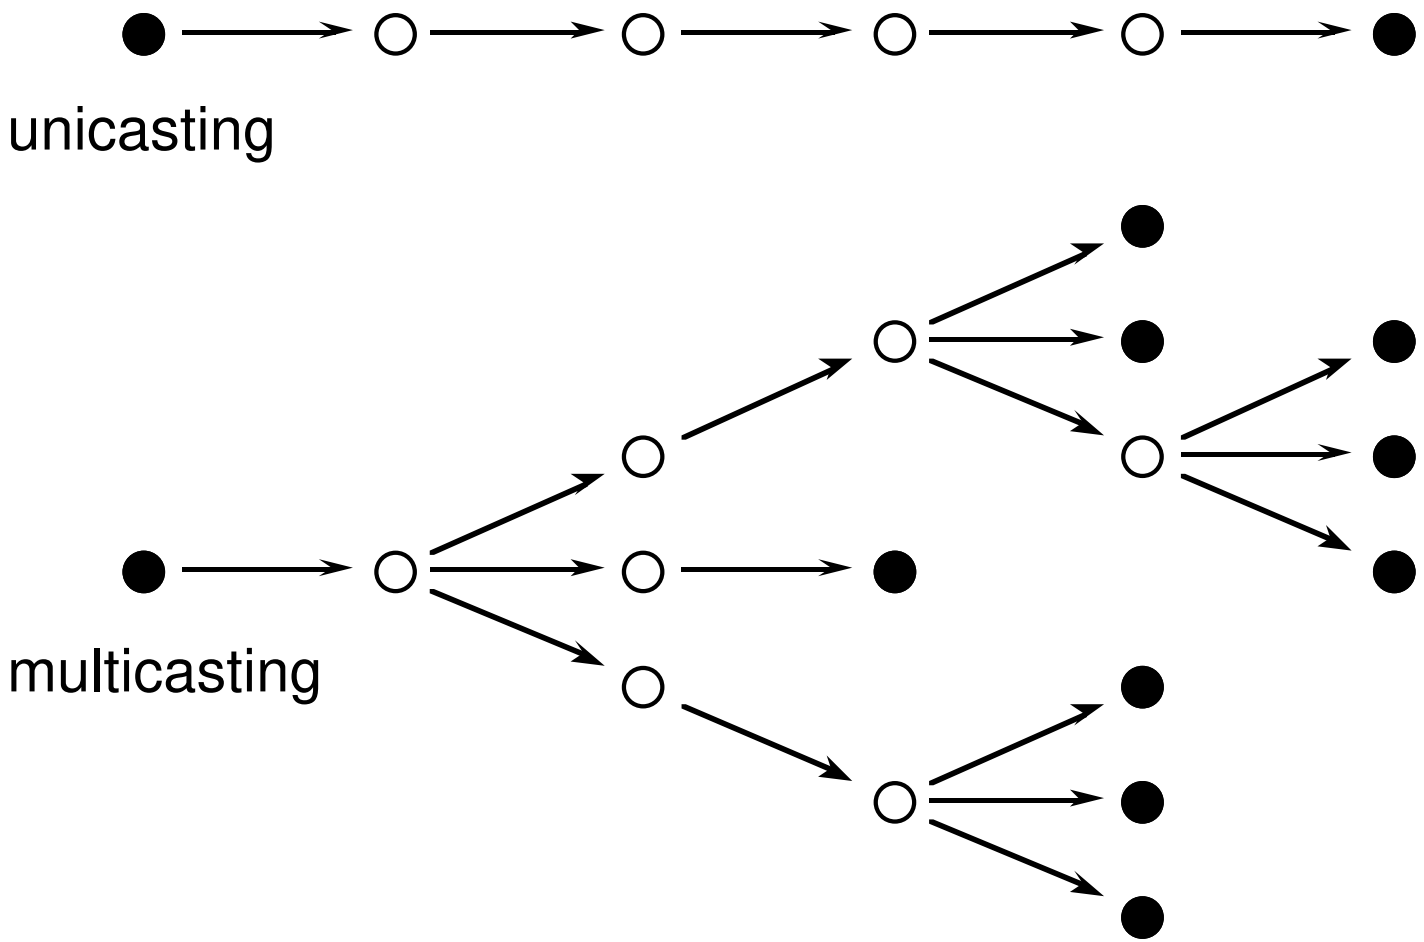

# IP Multicasting

- each group identified by a single IP multicast address
- groups may be of any size
- members of groups may be located anywhere in the Internet
- members of groups can join and leave at will
- senders need not be members

# MBone?

MBone = Multicast Backbone

An infrastructure for delivering IP Multicast packets

A virtual network, overlaid on the Internet

# Components of the MBone

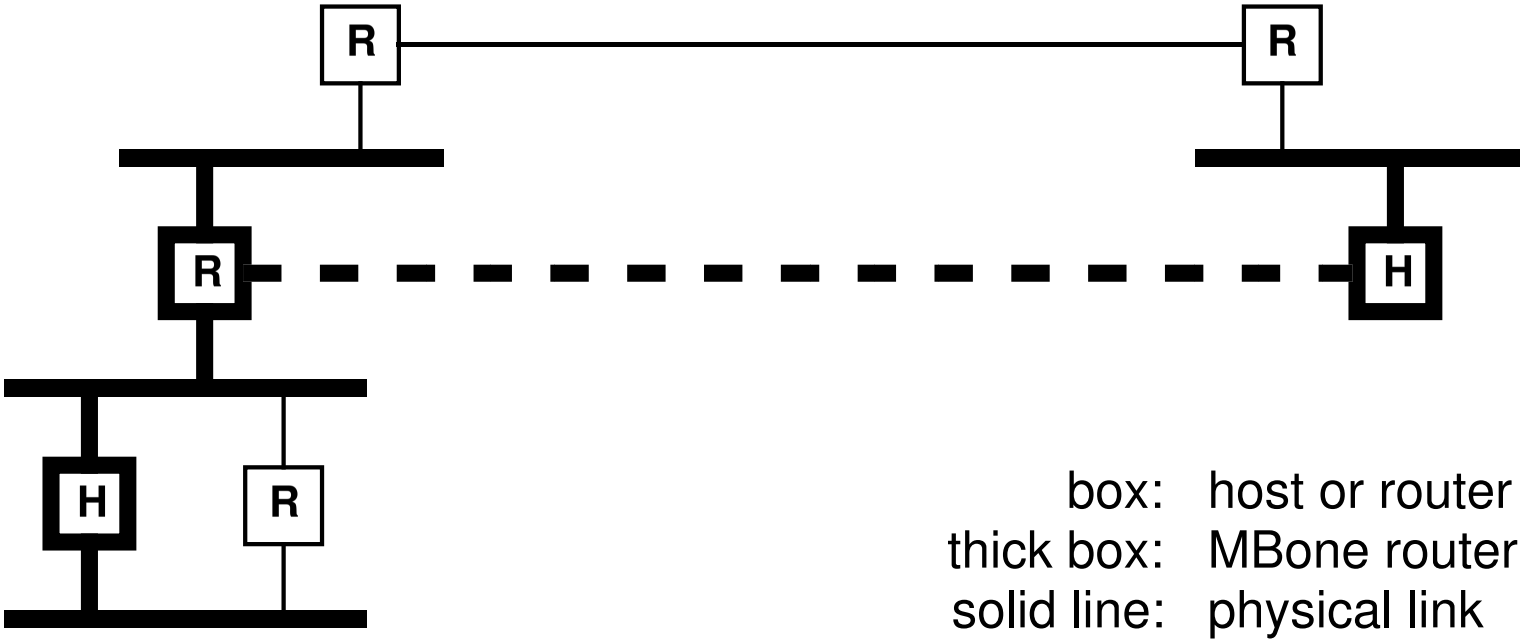


box: host or router
thick box: MBone router
solid line: physical link
dashed line: tunnel
thick line: part of MBone

# Growth of the MBone (number of subnets)

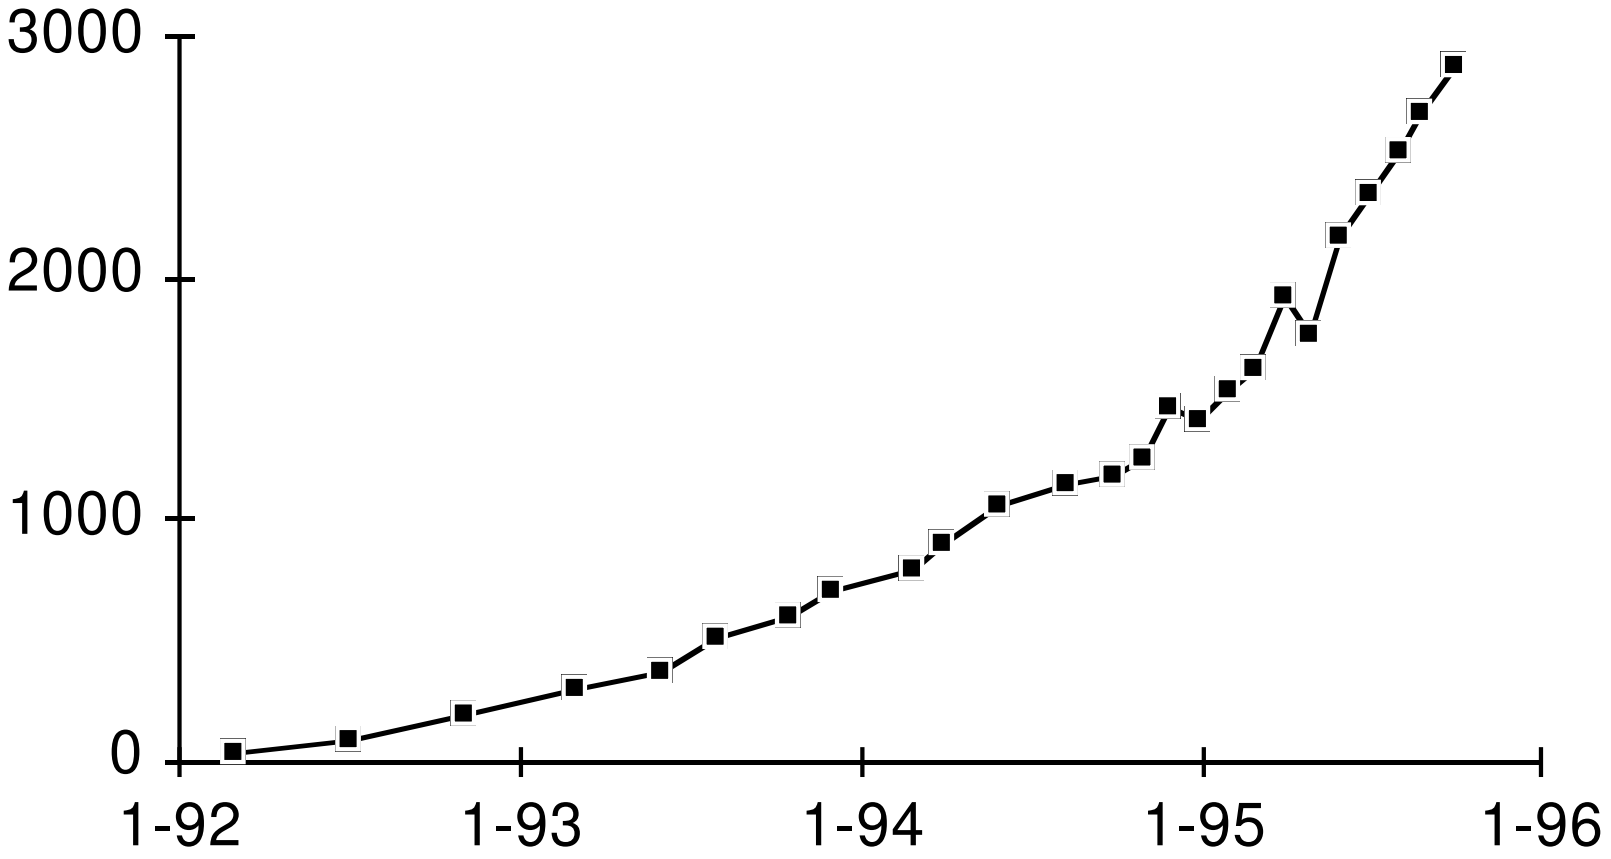

# Current IP Multicast Implementations

**shipping product**

Sun — Solaris 2.x

SGI — Irix

x86 — Microsoft, Windows 95

— FTP Software (host only)

— BSDI, NetBSD, FreeBSD

— Linux

Apple — Open Transport

NeXT — NeXTSTEP (?)

Proteon, Alantec

Cisco, Bay, 3Com

**patches available**

Sun — SunOS 4.x

DEC — Ultrix

— OSF-1

HP — HP-UX

IBM — AIX (host only)

# New Applications Enabled by the MBone

- **A/V teleconferencing:** vat, vic (LBL), nevot (UMass), ivs (INRIA), nv (PARC)
- **shared "whiteboard":** wb (LBL)
- **conference management:** sd (LBL), mmcc (ISI)
- **scientific data dissemination:** telemetry (Woods Hole), imm (UHawaii)
- **distributed interactive simulation:** NPSnet (NPS), games

# MBone Events

- **meetings:** IETF working groups, research collaborations
- **seminars & conferences:** UCB & MICE seminar series, National Press Club lectures, networking & other confs.
- **public information:** NASA Select, VIP visits, satellite weather images, live House & Senate broadcasts
- **entertainment:** Radio Free Vat, experimental films, established & aspiring rock bands
- **"in-house" use:** who knows?

# MBone in the News

- San Francisco Examiner, Dec 4/94
- Newsweek, Dec 5/94
- San Jose Mercury News, Jan 8/95
- New York Times, Feb 8/95
- Newsweek, Feb 27/95

# Experimental —> Production Service

- most large Internet service providers support MBone tunnels to some of their customers
- commercial routers starting slowly to supplant Unix boxes as MBone routers
- availability of MBone taken for granted by users, e.g, for IETF meetings, seminar series, Internet Multicasting Service
- much still to be done to make it a real production service!…

# Multicast Routing — DVMRP Example

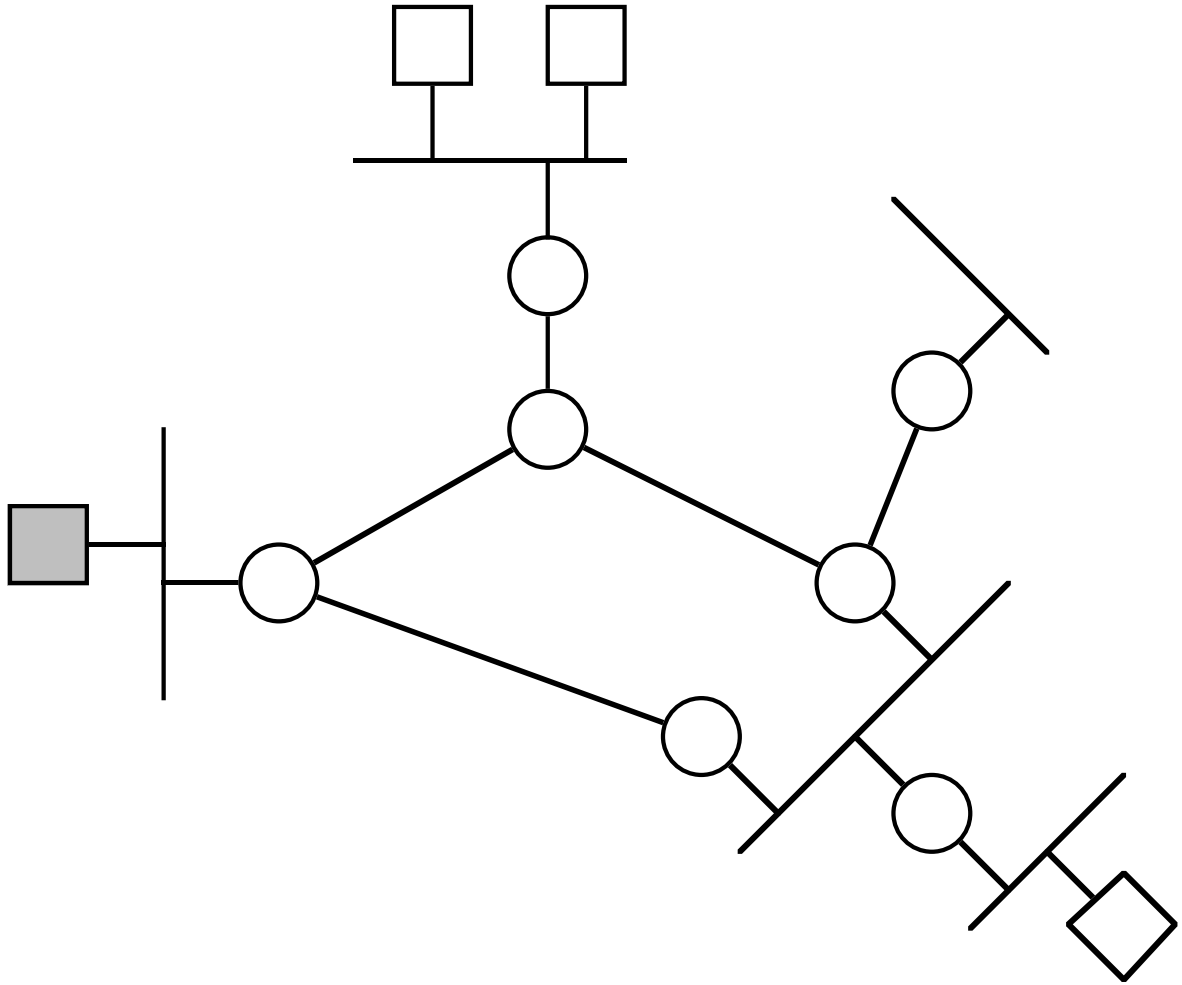

# Multicast Routing — DVMRP Example

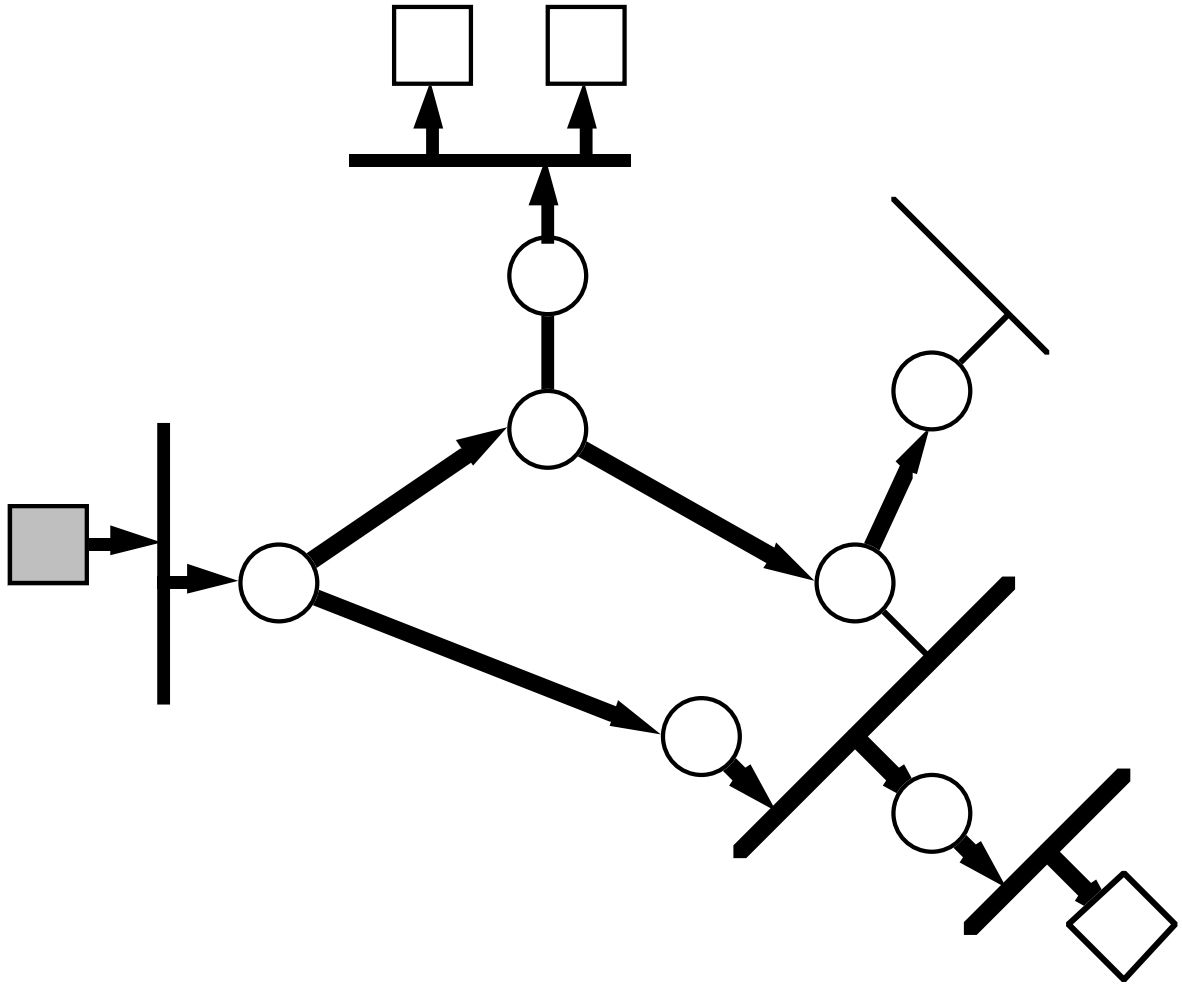

# Multicast Routing — DVMRP Example

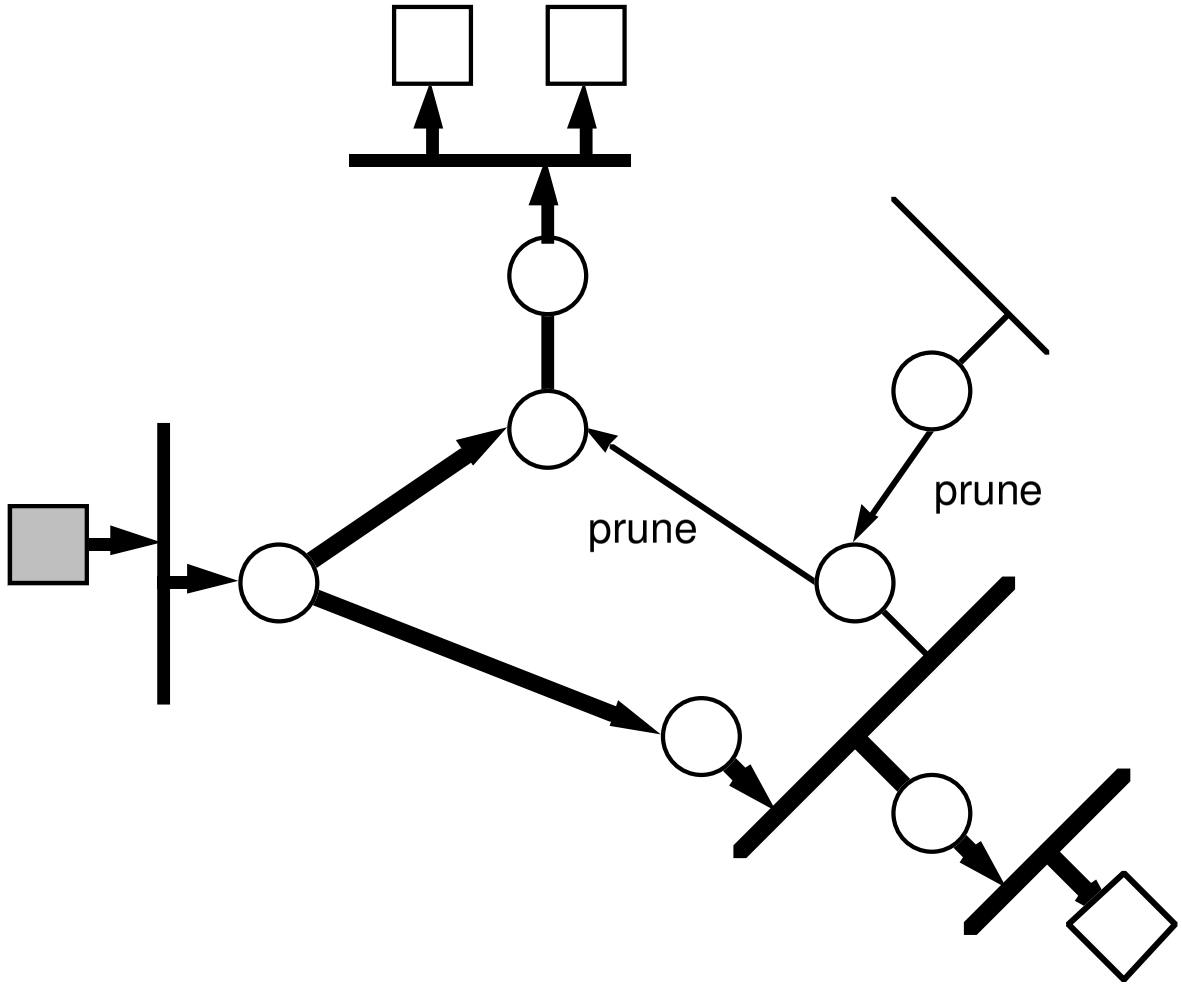

# Multicast Routing — DVMRP Example

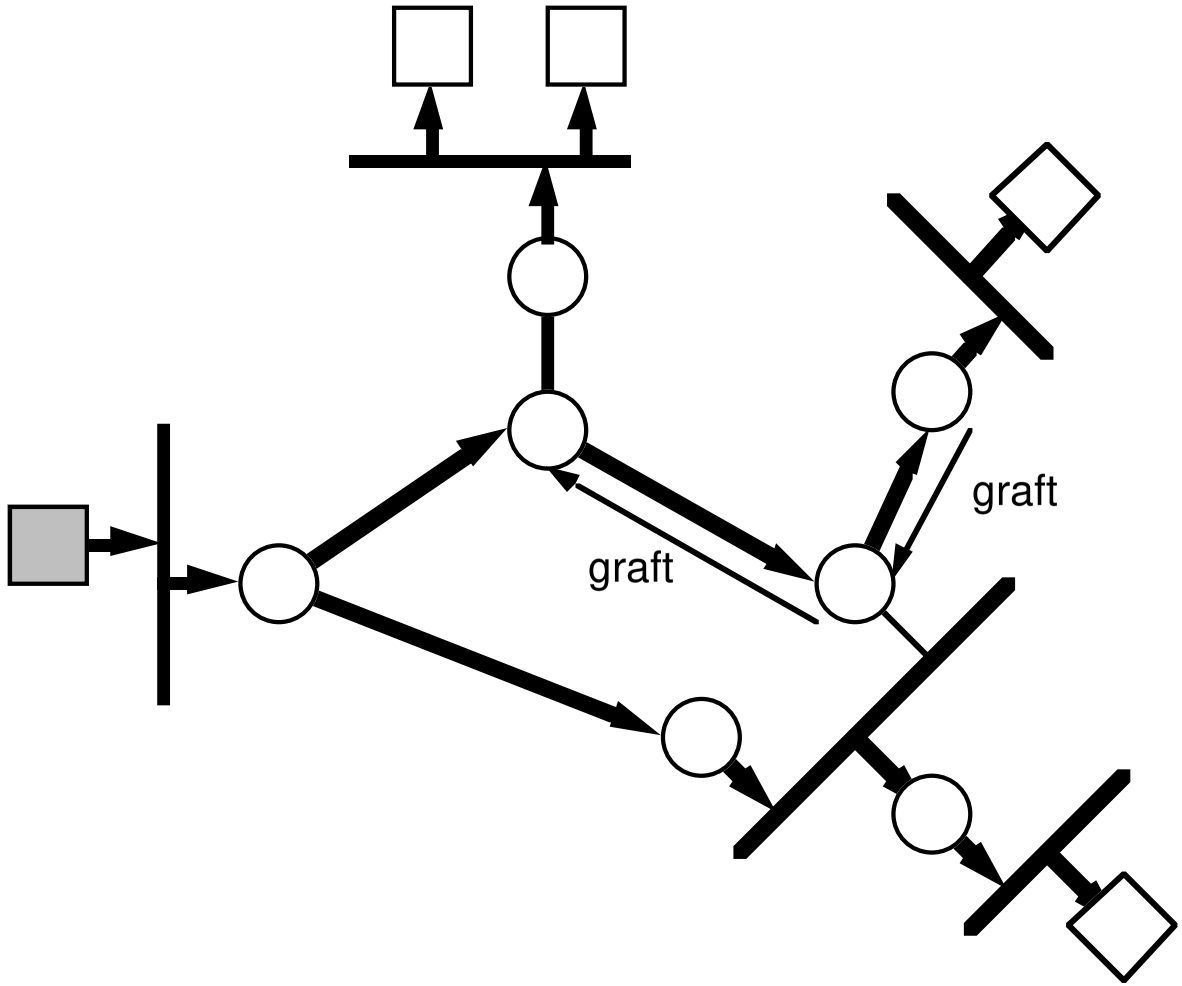

# Scaling Issues — Urgent

- most of MBone currently treated as a single, flat DVMRP domain — has routing table overhead O(#subnets)

  obvious fix: hierarchical multicast routing (like unicast)

  even better: eliminate separate topology-maintenance protocol for multicast, e.g., PIM or CBT — **but**, works only in domains without tunnels

# Short-Term Plan: Two-Level DVMRP

- Unix-based DVMRP (i.e., mrouted) at Level 2 (inter-domain)
- Variety of platforms and routing protocols at Level 1, e.g.:
  - — Unix, Alantec, Bay Networks, 3Com DVMRP
  - — Proteonm 3Com MOSPF
  - — cisco PIM
- Level 2 routing based on multicast domain IDs to retain topological independence from, and smaller tables than, unicast routing — requires encapsulation for L2 forwarding

# Why Two-Level DVMRP?

- easiest to do quickly — urgency of problem requires a quick solution, even if not the right long-term solution
- benefits:
  - — reduction of routing table & update message size
  - — relaxed limit on MBone diameter
  - — architected interoperation between diff. protocols, including better damage containment
- will serve as a testbed to debug hierarchical architecture before asking vendors to support L2 routing

# Scaling Issues — Longer Term

- multicast forwarding state is O(# active source-group pairs)

  — hierarchical routing achieves source aggregation
  — “scoped groups” limit range over which state must be mainatained
  — research problem: how to aggregate groups?

- current algorthms (DVMRP, MOSPF) require occasional broadcasts to all routers within a domain

  — fixed in new algorithms, currently under test (CBT — Ballardie et al; PIM — Estrin at al)

# PIM — Protocol-Independent Multicast

- eliminates broadcast behavior by using “rendevous points” (at a cost of vulnerability to partition failures)
- eliminates separate multicast routing table; uses unicast tables instead — simplifies policy, complicates deployment
- for low-traffic groups, can reduce forwarding state by keeping traffic on a single per-group tree (subsumes CBT)
- have preliminary ideas for per-group state reduction, still to be worked-out in detail

# Recent / Planned Changes to IGMP

lower leave latency:

augments timeout-based leave-detection with an explicit Leave Group message — reduces latency to < 1 second

source-specific Join / Leave

rendezvous-point identification, for PIM

# Congestion Issues

- most current multicast apps are not flow controlled —> congestion problems, and “hogging” of bandwidth

  short-term fix: multicast rate limits on tunnels and subnets (manually config’d)
  medium-term fix: “fair queueing” in unicast routers
  long-term fix, for some apps: resource reservation and admission control (e.g., RSVP)

- some current multicast apps are bandwidth hungry, esp. A/V

  fix: more bandwidth, of course!

# A Possible Alternative to Resource Reservation?

- (statistical) fair queueing in routers (Q: what granularity?)
- layered-encoding of data, with each layer sent to distinct multicast address
- ~~priority dropping of layers by congested routers~~
- adaptive joining / leaving of groups by receivers
- provision of adequate bandwidth

# Security Issues

- IP multicast service allows any host to join any group —> privacy concerns

  fixes: end-to-end encryption — need key distribution infrastructure; legal and licensing obstacles

  use scope-limited groups, where possible

  enhance the IP multicast service to support greater control over group membership (research problem)

# Security Issues (cont.)

- denial of service by “rogue” senders

  fixes: fair queueing mechanisms

  source-specific Leave

- problem of authenticating multicast senders — public-key methods work, but are too expensive at high rates

  fixes: ???

# Network Management Issues

finally have a decent set of diagnostic tools:

- mapping tools
- multicast traceroute
- SNMP support

problem: getting people to upgrade to latest versions

# Building on the Multicast Infrastructure

RTP — Real-Time Transport Protocol

- media & encoding identification
- multiplexing
- timing reconstruction & synchronization
- quality-of-service monitoring

SDP — Session Directory Protocol

- multicast address allocation & advertisement (informed, partitioned, random allocation)
- key distribution
- support for a range of conference “styles”

# Conclusions

- there are many issues to be addressed to enable continued growth of the multicast infrastructure
- solutions to most of the identified issues are known and are currently being designed, implemented, and tested
- still some hard research problems to be solved

a PostScript version of these slides may be fetched from:

ftp://parcftp.xerox.com/pub/net-research/mbone/mbone-talk-dec95.ps