Growing the MBone: Scaling, Performance, and Security Issues with Internet Multicasting

Steve Deering
Xerox Palo Alto Research Center
deering@parc.xerox.com

INET '94
Prague, Czech Republic
June 17, 1994

Multicasting: Transmitting a Packet to a *Group*

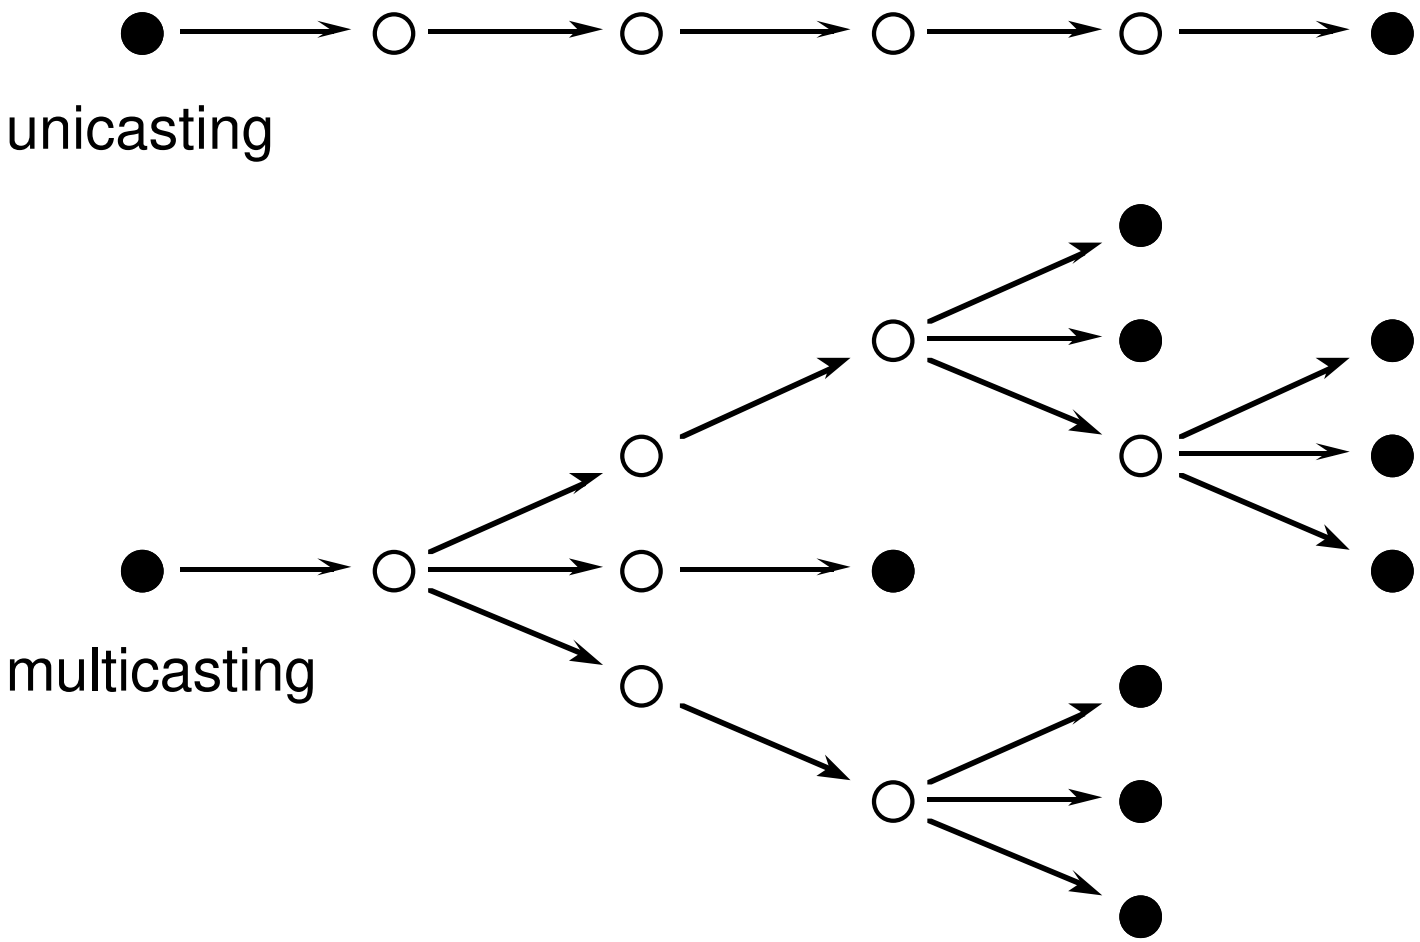

MBone—the Multicast Backbone

- an experimental IP multicast infrastructure
- subnets interconnected by multicast routers and "tunnels" (IP encapsulation)

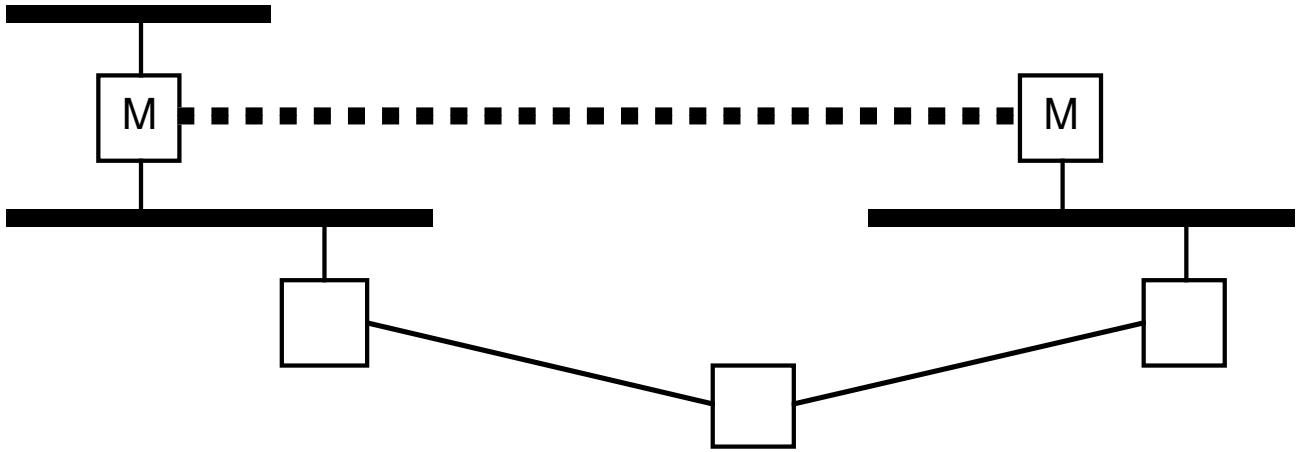

- forms a virtual network, overlaid on the Internet

Growth of the MBone (number of subnets)

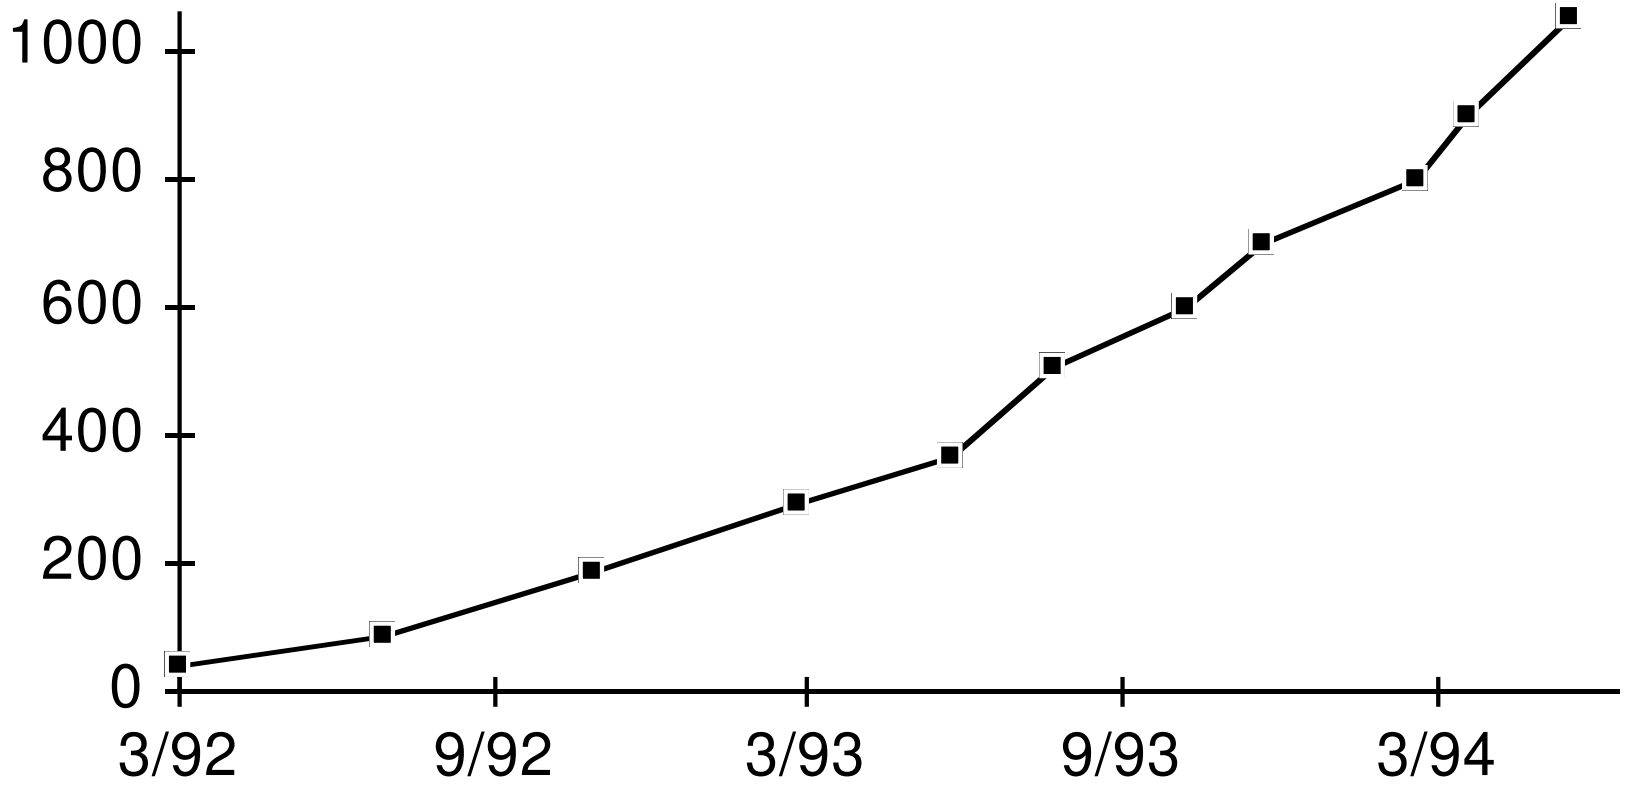

Growth of the MBone (cont.)

- started with 40 subnets in 4 countries;
  after 2 years, have >1000 subnets in >20 countries
- MBone routers are mostly Unix workstations (Sun, SGI, DEC, HP, IBM); some are “real” routers (Proteon, Alantec, cisco)
- MBone subnets are mostly Ethernets; some are other types of networks (FDDI, point-to-point links)

New Applications Enabled by the MBone

- **A/V teleconferencing:** vat (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:** PARC & MICE seminar series, National Press Club lectures, networking & other conferences
- **public information:** NASA Select, satellite weather images, VIP visits
- **entertainment:** Radio Free Vat, experimental films, aspiring rock bands
- **“in-house” use:** who knows?

The Story So Far…

- new applications and events <=> rapid growth of the MBone
- a remarkably successful experiment and technology demonstration
- **but**, what will it take to sustain the growth of the MBone and evolve it into a “production” service?…

Scaling Issues

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

  obvious fix: hierarchical multicast routing (like unicast), e.g., multi-level DVMRP, MOSPF, MBGP

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

Scaling Issues (cont.)

- 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 require occasional broadcasts to all routers

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

Performance Issues

- large delay (~5 minutes) from time last member on a subnet leaves a group until traffic stops being sent on that subnet

  fix: new version of IGMP, currently under test

- tunnels are performance bottlenecks and often multiply traffic on particular subnets

  fix: get multicast routing into the unicast routers

Performance Issues (cont.)

- 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!

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 prunes”

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

  fixes: ???

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, but work is just beginning on design, implementation, and testing
- still some hard research problems to be solved