A Case for End System Multicast

Yang-hua Chu, Sanjay G. Rao, and Hui Zhang

Presented by: Pratik Fegade

IP Multicast

- Best effort multicast service provided by the IP layer
- Drawbacks:
  - Stateful
    - Vulnerable to flooding by malicious sources
    - Difficult to implement reliability and the like on top
    - Distributed consensus required to obtain group address
- End system multicast helps mitigate these

Performance Tradeoffs

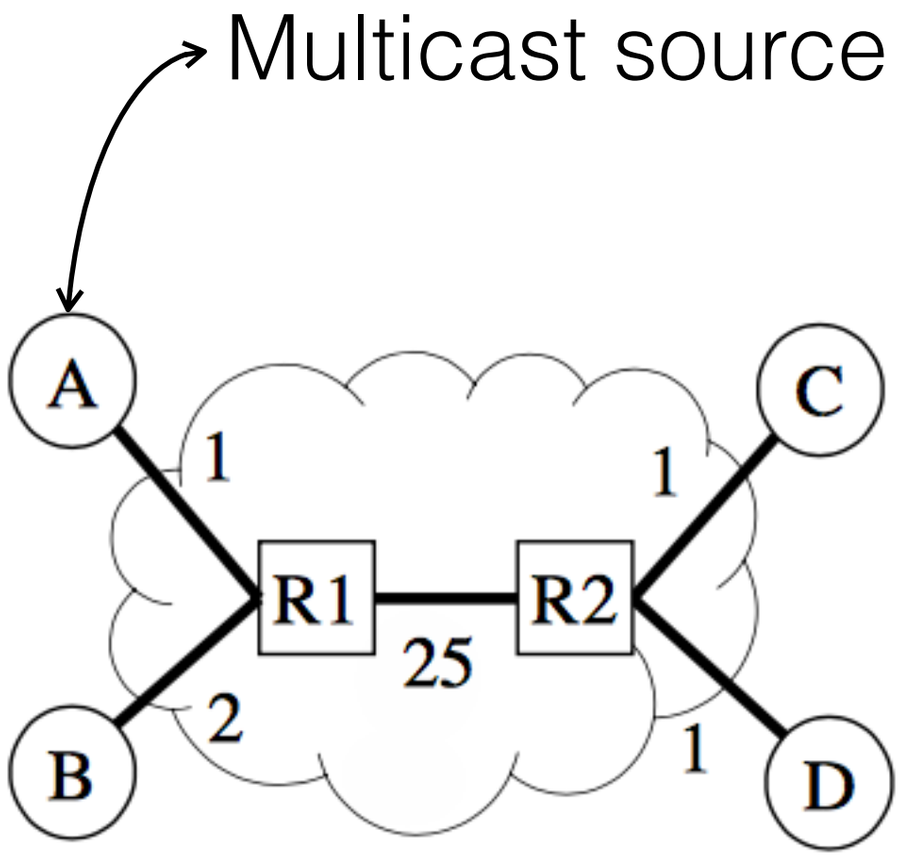

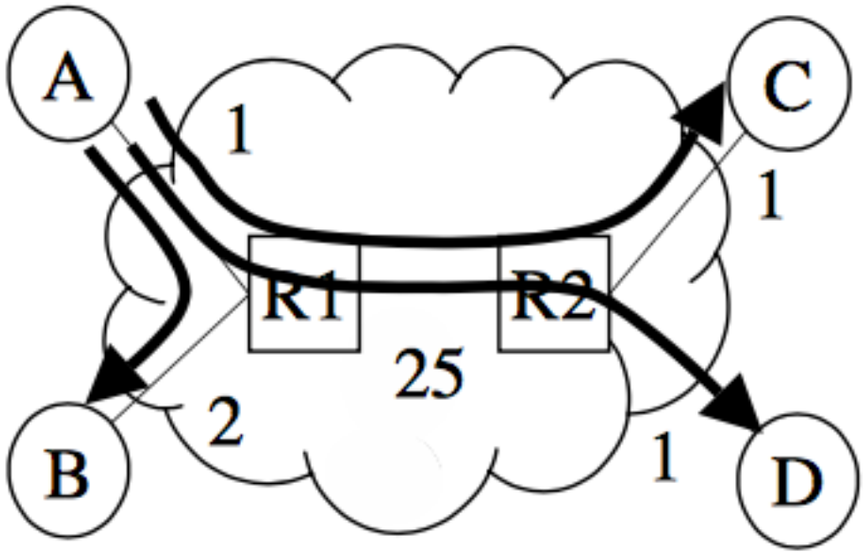

Naive Unicast

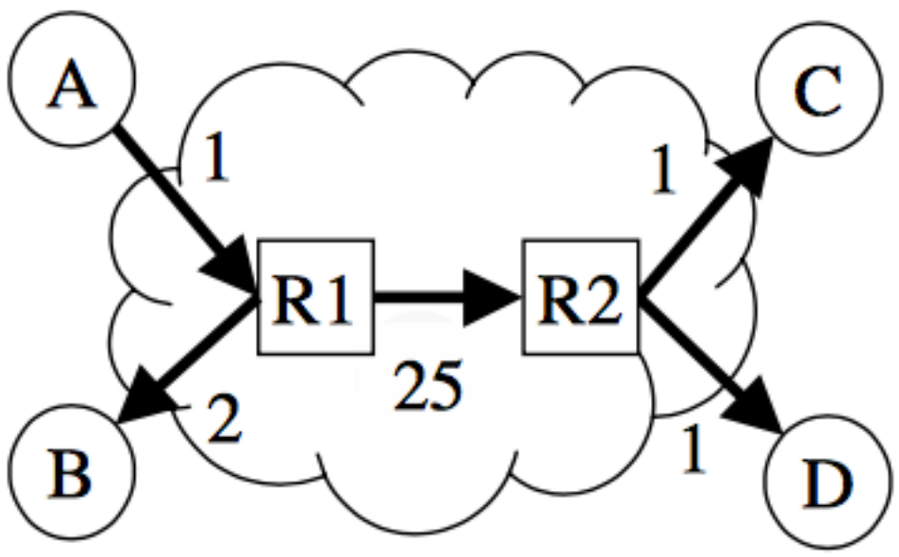

IP Multicast

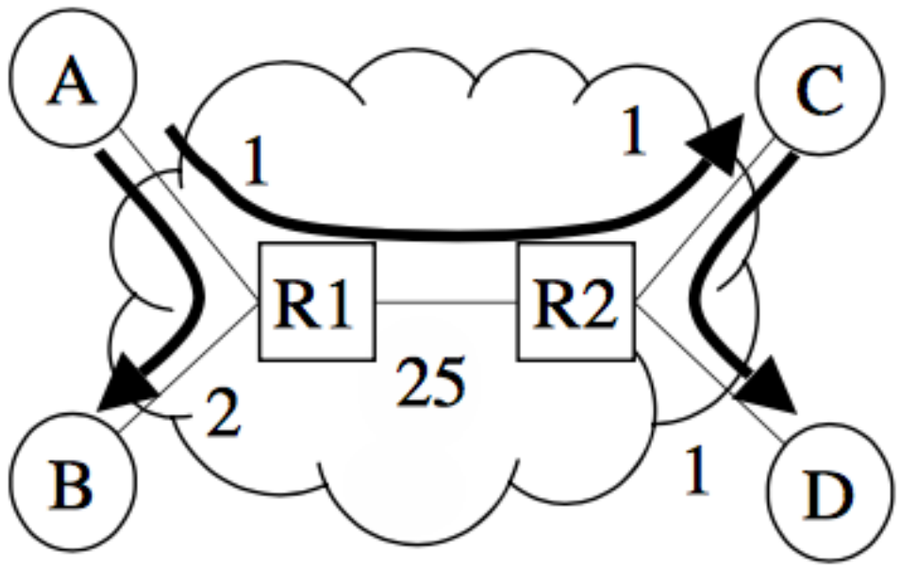

End System Multicast

Narada: An Overview

- Construct an overlay (a mesh) between the end systems
  - Used for distributed control and coordination

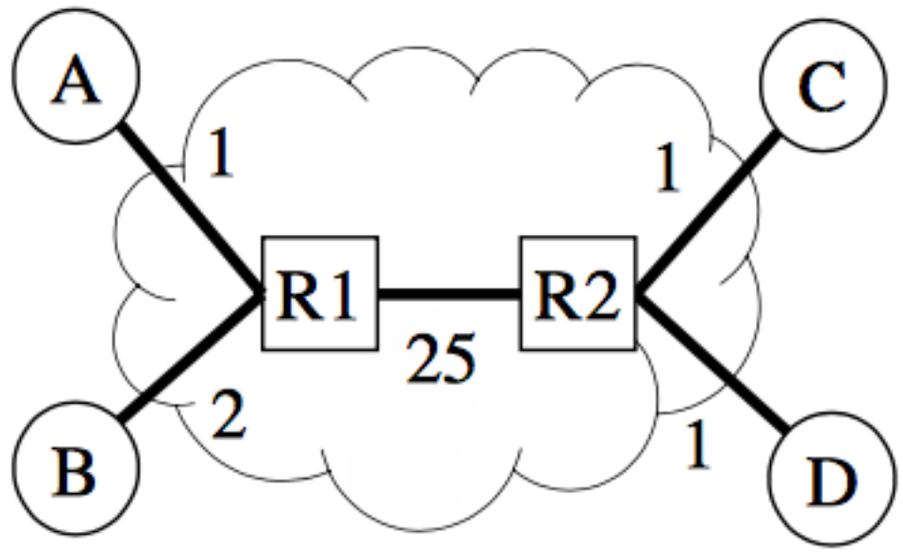

- Construct reverse shortest path spanning trees
  - Used for actual data transfer

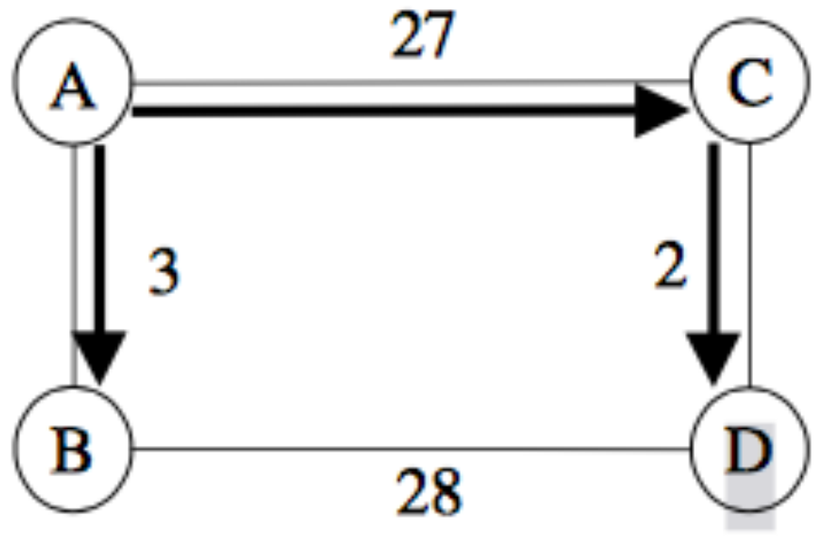

Narada: Design Goals

- Self-organizing
  - Fully distributed and robust to changes in membership
- Overlay efficiency
  - Efficient mesh construction
- Self-improving in an incremental fashion
  - Quality of mesh (and hence spanning trees) improves over time

Narada: Group Management

- Each member
  - maintains a list of all other members
  - periodically sends out this list to neighbours
    - Heart beat
    - Disseminating membership changes
- Repair mesh by probing and adding links
- New member requests current member(s) for membership

Incremental Mesh Improvement

- Probe members randomly to discover potentially better links
- Drop links perceived as not useful

Narada: Routing

- Run distance vector on the mesh
- Avoid count-to-infinity by maintaining shortest path
- Require leaving members to forward packets for some time after leaving to avoid transient losses

Narada: Evaluation

- 1024 end systems with a group size of 128
- All group join within the first 100 seconds with no changes later

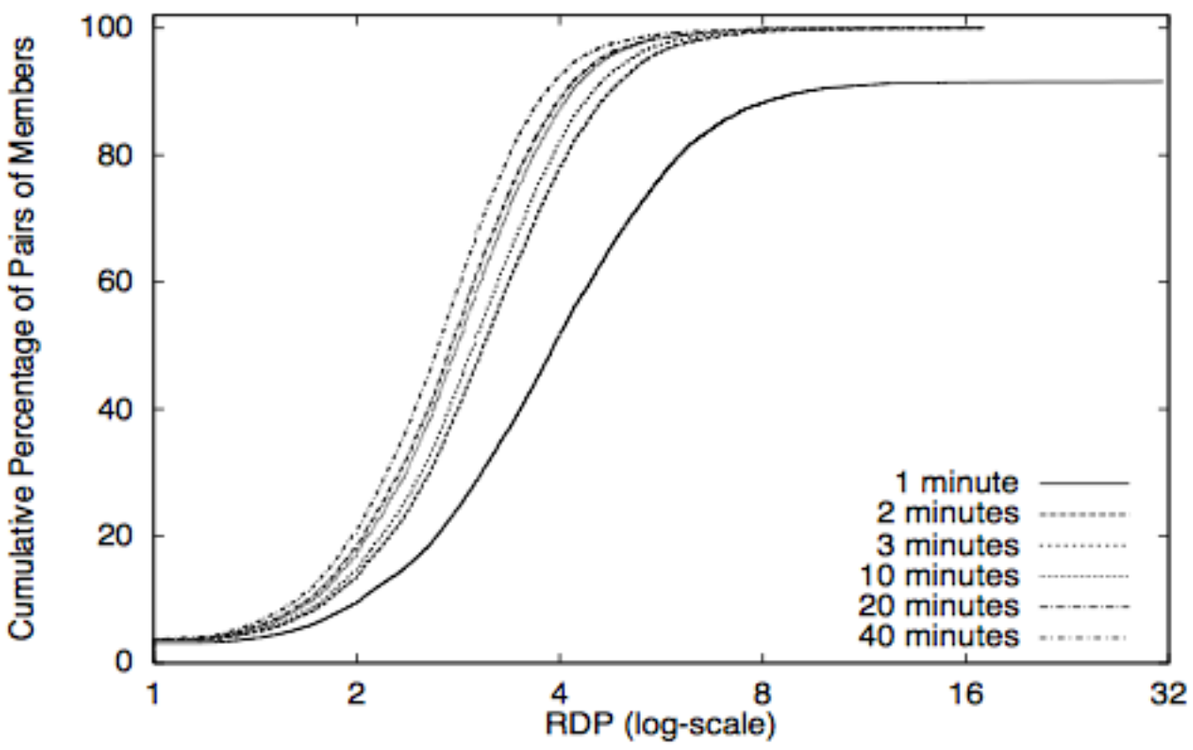
Cumulative Percentage of Pairs of Members
RDP (log-scale)
1 minute
2 minutes
3 minutes
10 minutes
20 minutes
40 minutes

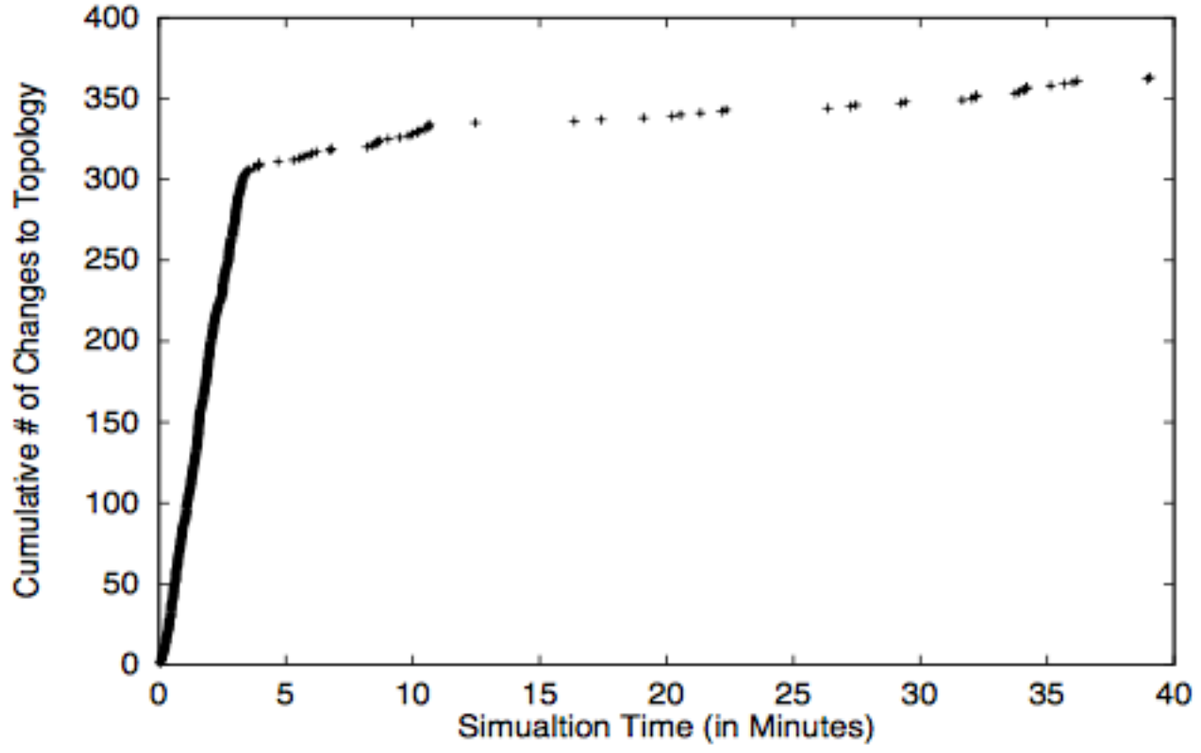
Cumulative # of Changes to Topology
Simualtion Time (in Minutes)

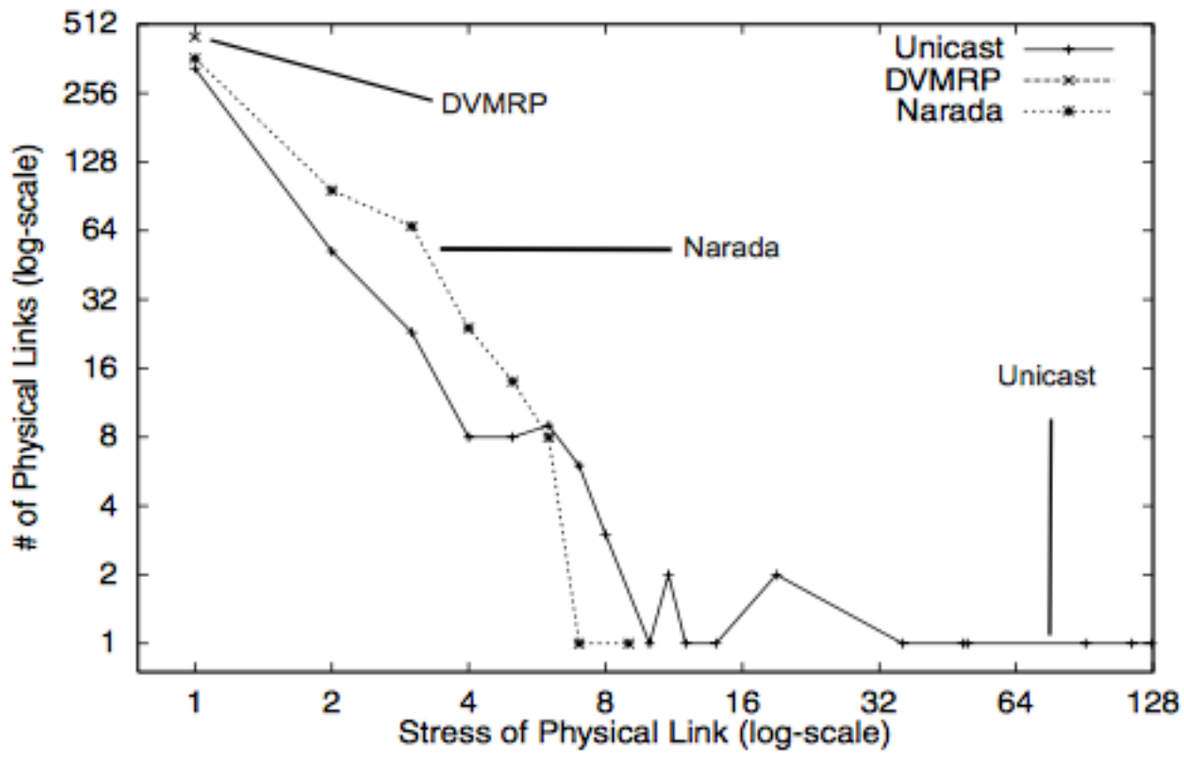
# of Physical Links (log-scale)
Stress of Physical Link (log-scale)
Unicast
DVMRP
Narada
DVMRP
Narada
Unicast

Internet deployment

- 13 hosts distributed throughout USA

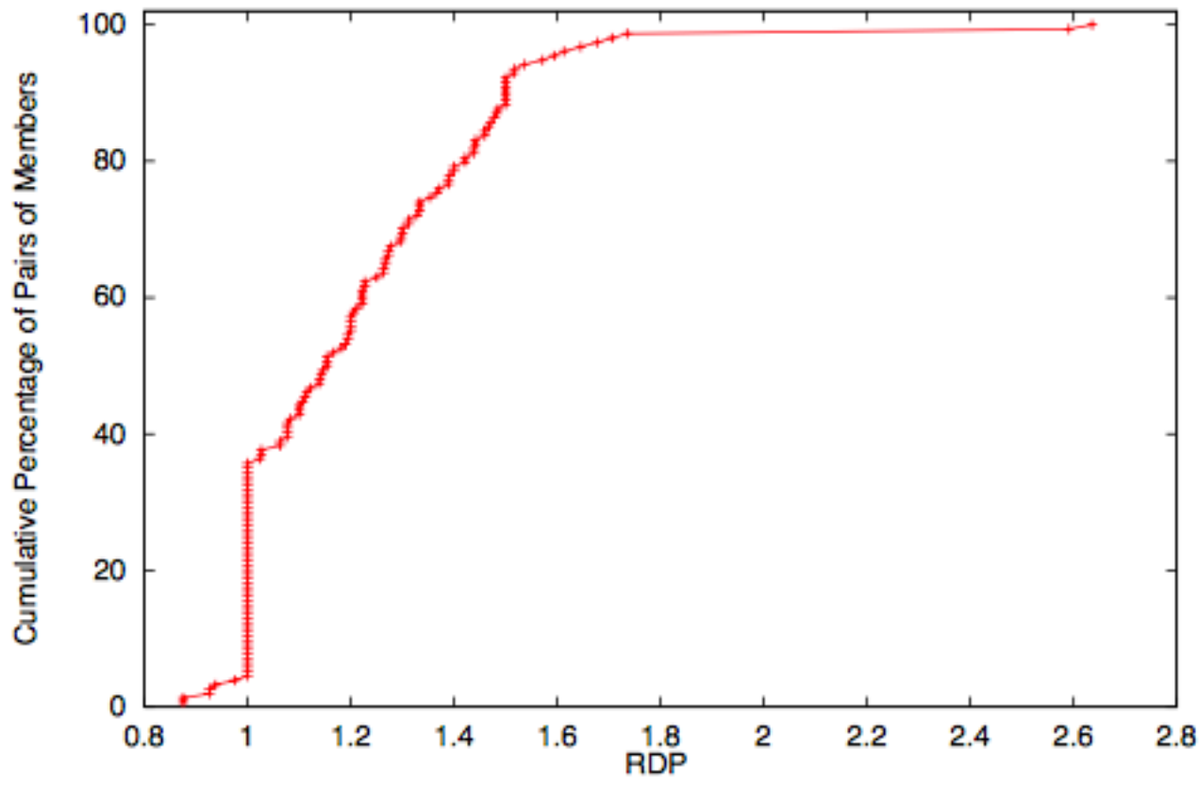

Thank you!