Computer Networks

Multicast

David Henriques

November 8, 2010

Multicast

Problem: Many applications want to share data with several users

This is hard!

Proposed solution: IP Multicast

- uses address indirection
- builds source based tree with reverse path forwarding and pruning

Multicast

Problem: Many applications want to share data with several users

This is hard!

Proposed solution: IP Multicast

- uses address indirection
- builds source based tree with reverse path forwarding and pruning

Multicast

Problem: Many applications want to share data with several users

This is hard!

Proposed solution: IP Multicast

- uses address indirection
- builds source based tree with reverse path forwarding and pruning

Multicast

Problem: This doesn't really solve the problem!

- Requires "smart" routers.
- Huge scalability issues even for the simplest functional requirements
  - Packet acknowledgement
  - Error recovery
  - In order delivery
  - Congestion control
  - ...
- New models of service
  - Several senders
  - Replicated data
  - ...

Multicast

Problem: This doesn't really solve the problem!

- Requires "smart" routers.
- Huge scalability issues even for the simplest functional requirements
  - Packet acknowledgement
  - Error recovery
  - In order delivery
  - Congestion control
  - ...
- New models of service
  - Several senders
  - Replicated data
  - ...

Reliable Multicast

A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing

Reliable Multicast

Adds reliability to multicast (surprise!)

Very application focused (white board)

- Idempotent operations
- All users aware of all other users
- ...

Weak assumptions for reliability

- Eventual delivery
- No ordering
- Known participants
- ...

Reliable Multicast

State exchange between users

All packets have a source ID and a timestamp

Burden of error detection on the RECEIVER

When a packet loss is detected, a repair request is scheduled

- The node waits a random ammount of time
- If a request for the same data is heard in the meanwhile, cancels its own request
- Prevents a flood of requests in case of local catastrophe

Reliable Multicast

When a packet repair request is heard, an answer is scheduled

- The node waits a random ammount of time
- If an answer for the same data is heard in the meanwhile, cancels its own answer
- Prevents a flood of answers

Mechanism for adapting random waiting time based on timing of other requests/answers heard

Local recovery to reduce the number of messaes sent in case of local crisis.

Discussion

Popular demand: We want this studied in a "real" topology!

- Chain and star and tree topologies only cut it as proof of concept
- We know how to generate Internet-like graphs
- Power laws tell us there are a few nodes with very high fanout and yet there is a bounded degree assumption

This works well for whiteboard, but what about other applications?
Would a mechanism like this be able to accomodate other functionalities?
Disregards latency issues.

End System Multicast

A Case for End System Multicast

End System Multicast

Argues agains implementing multicast at IP level

- Complexity
- Scaling constraints
- Security

Narada abstracts physical topology and builds the spanning tree over the abstraction

- Still have scaling constraints
- Only considers a (random) subset of the links in the abstaction (mesh)
- Potential to chose a poor mesh

End System Multicast

"Nodes" keep a lot of state

- Knows about every other node
- Keeps last sequence number from each node
- Keeps last known activity of each node
- Need to send periodic refresh messages

We put up with this because

- The mesh repairs itself
- The mesh improves itself over time!
- Follows the end to end principle

Discussion

Main praise: Does not rely on IP multicast

Main concern: The approach does not scale! Will this affect its survivability?

- Usual examples of multicast applications involve many more than the supported number of users
- Examples of applications where it would work would be nice
- The project is dead (according to the site), which is not encouraging...
- There is a startup using this, which is...

Routing overhead is considerable, when compared to IP multicast, but then again, IP multicast is not available...