

Scalable and Reliable Multicast

Contradiction in Terms ?!


Jörg Liebeherr

Polytechnic University (“Brooklyn Poly”)

Disclaimer

- When I talk about scalable and reliable multicasting
 - ... I don't mean the Scalable Reliable Multicast (SRM) Protocol
 - ... I do mean protocols that support large multicast groups

What is the problem ?




- **ACK Implosion:** Sender is overwhelmed with state

ICNP

3

How much scalability do we need?



ICNP

4

Solution 1: No Structure

- **“NACK Suppression”** (Ramakrishnan/Jain, XTP, SRM)
 - Back off for random time before sending a NACK
 - Think “p-persistence”



Time to recover data can be long if multicast group is large (i.e., does not scale)

- **Forward Error Correction**
 - Include redundancy in data



Loose less data by sending more data ?

ICNP

5

Solution 2: Ring Structure

- **Token Ring Scheme** (Chang/Maxemchuk, RMP)
 - Pass a token around in multicast group
 - Only the token holder recovers data



Scalable? No way!

ICNP

6

Solution 3: Tree Structure

- **Tree Hierarchy** (TMTP, log-based, RMTP, LORAX, etc.)
 - Request retransmissions only from parent node

Great if there is one sender. But how about multiple senders?

- **Subsolution 1:** Built one tree for each sender



1 million trees?

- **Subsolution 2:** “Rehang” a single tree



Unfair.

Unfairness grows exponential with number of nodes

ICNP

7

Feasible Solutions

- **Solution by Declaration**
 - “Problem is already solved”
 - “Problem never existed”
- **Solution by Redefinition :**
 - Redefine “Scalable:” “Groups of 10 are a crowd”
 - Redefine “Reliable” : “Maybe-once semantics”
- **Solution by Exhaustion:**
 - “Standardize!”
 - “It’s a Ph.D. topic!”

ICNP

8