TDK - Team Distributed Koders
Distributed Systems I

# Fairness in P2P Streaming Multicast

Team Members:

*Jason Winnebeck*
*John Kaeuper*
*Kumar Keswani*

Team Report IV
2/21/07

# Presentation Topics

- Goal
- Accomplishments
- What we could not implement
- Software Inputs (Test Cases)
- Software Outputs (Test Results)
- Future Work
- Lessons learned from project
- Demo

# Goal

*Explore the effectiveness of various mechanisms for enforcing fairness and incentivizing social welfare in a multi-tree peer-to-peer multicast system by using the ideas discussed in our research papers. The evaluation would be done with a discrete event simulation.*

# Accomplishments

- Individual packets forwarded between nodes
- Percentage of packet loss on per node basis
- Implemented Fairness Mechanisms
  - Debt Maintenance
  - Ancestor Rating
- Detection of Freeloaders in the network
- Identifying Freeloaders refusing to forward content to children
- Implemented Tree Reconstruction mechanism
- Evaluation using Discrete event simulation

# What we could not implement

- Taxation scheme presented in our third research paper
- Freeloaders refusing children
- Pastry Type of Tree Manager

# Software Inputs

- List of nodes and their configuration
  - Behavior algorithm
  - Percentage of packet loss
  - Inbound and outbound bandwidth capacity
- Global Parameters
  1. Number of stripes
  2. Stripe bits per second
  3. Node packet size (bytes)
  4. Simulation Duration (seconds)
  5. Random Seed
  6. Tree Manager implementation
  7. Tree Manager reconstruction rate (seconds)

# Software Outputs (Test Results)

- Percentage of packets received by node behavior group over time
- Cumulative distribution function (CDF) of nodes by behavior group and debt level at end of simulation
- CDF of nodes by behavior group and negative confidence at end of simulation

# Inputs used for Simulation Test Cases

- Constant Input Parameters:
  - 10 stripes
  - 100 kbps / stripe
  - 512 bytes / packet
  - 1560 second simulation
  - Node inbound bandwidth = 1000 kbps
  - Packet Loss = 0.5%
  - RandomTreeManager always used
- Parameters Varied:
  - Node outbound bandwidth
  - Tree Reconstruction Rate
  - Number Nodes (and Node type distribution)
  - Debt and Ancestor Rating Thresholds

# Results for Debt Maintenance

- Even with no packet loss and no freeloaders, normal nodes lose many packets:

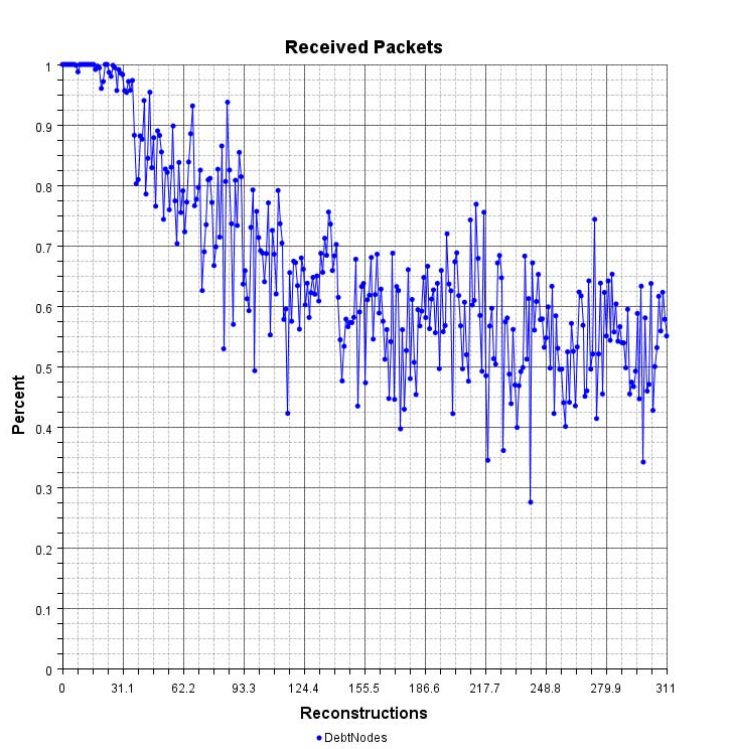


- Even with no packet loss, Nodes' reputations suffer – any time packet sent, a debt is incremented and Nodes may be falsely blamed
- Maybe a bug in our code?
- If Reciprocal Requests also used, could resolve issue

# Results for Debt Maintenance (cont'd)

- Debt Threshold Tradeoff (Tree reconstruction every 5 sec)
- 95 Normal and 5 Freeloaders
- Outbound bandwidth - 4 children / node

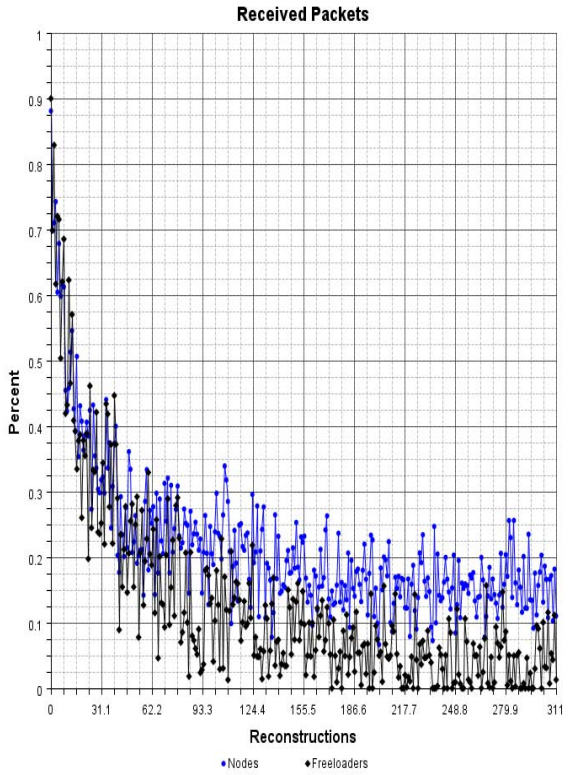


Debt threshold=200

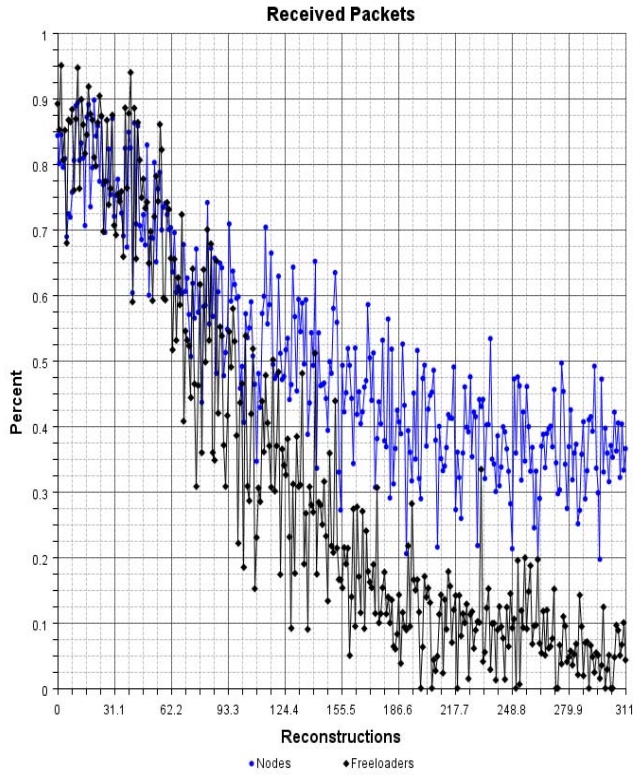


Debt threshold=550

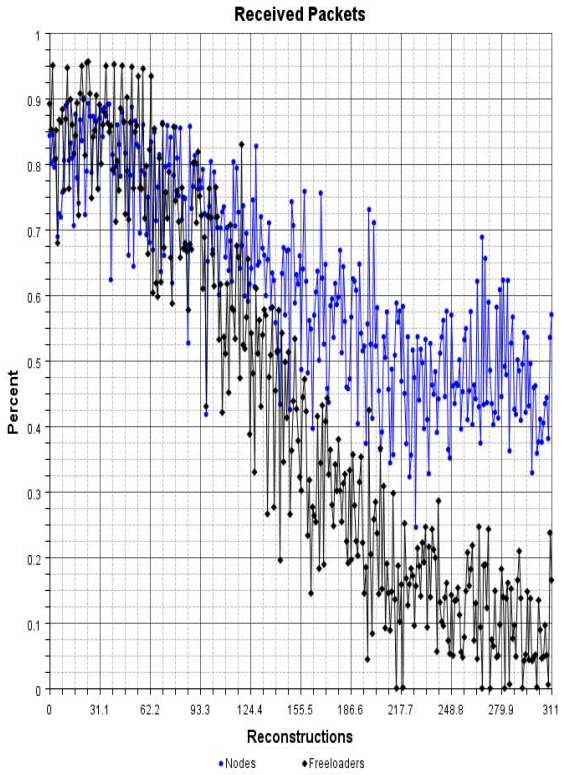


Debt threshold=750

# Results for Debt Maintenance (cont'd)

- Tree Reconstruction Rate more important than Debt Threshold:
  - This is result for Tree Reconstructions every 1 second
  - Other input parameters are the same

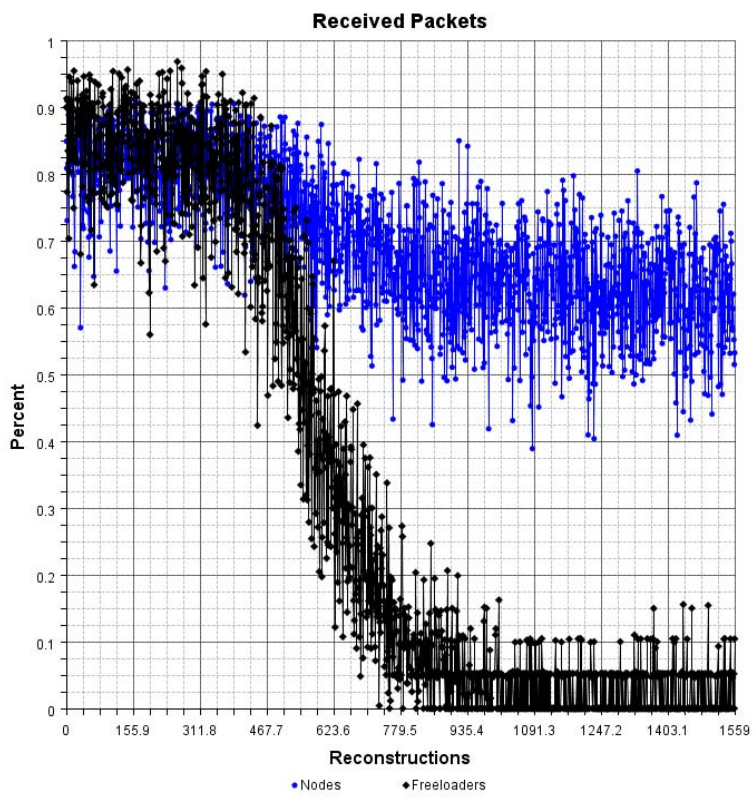

# Results for Ancestor Rating

- With no packet loss and no freeloaders, Nodes receive all packets

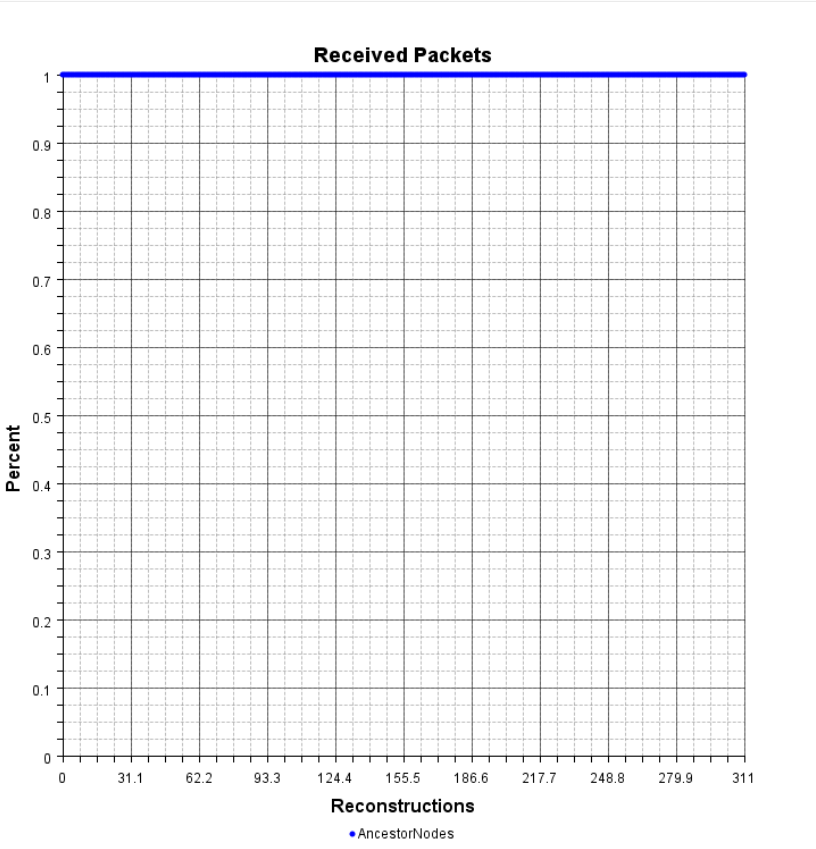


- If there is no packet loss, Nodes' reputations never suffer, unlike in Debt Maintenance
- With packet loss (no freeloaders), > 90% packets always received

# Results for Ancestor Rating (cont'd)

- 95 Normal and 5 Freeloaders
- Outbound bandwidth - 4 children / node
- Ancestor Rating Threshold = -200

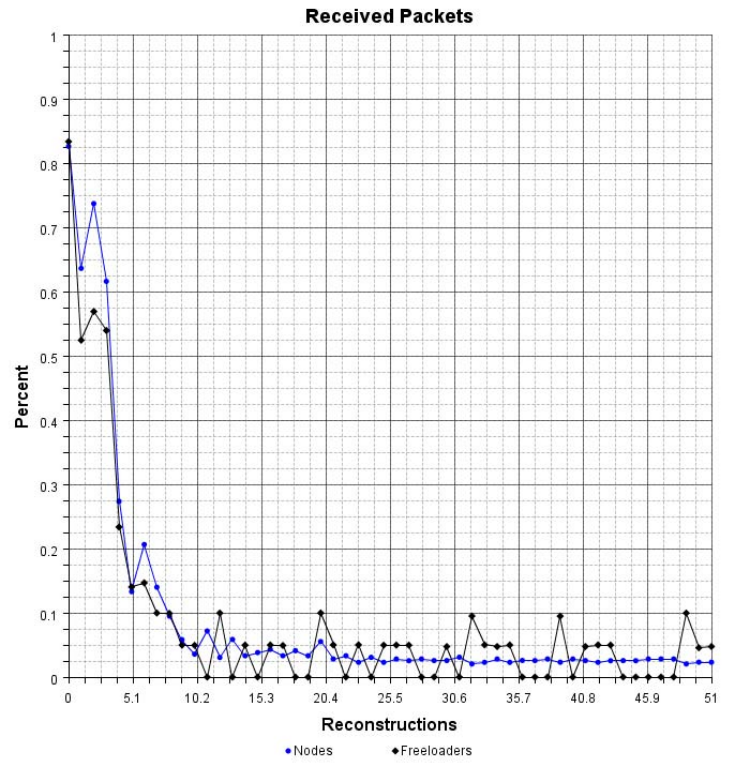


Tree Reconstruction Rate =every 30 sec

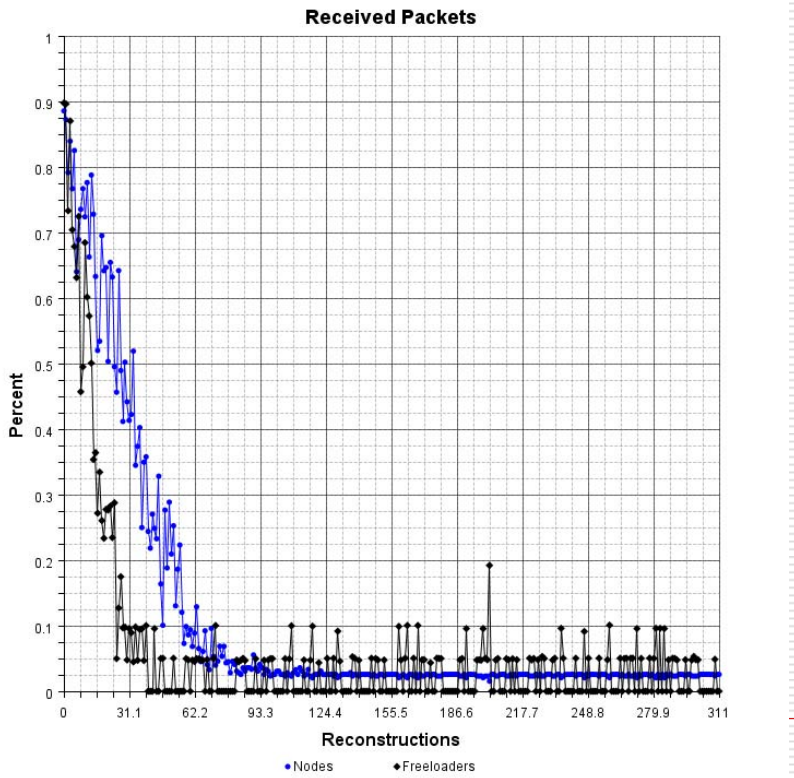


Tree Reconstruction Rate = every 5 seconds

# Results for Ancestor Rating (cont'd)

- 95 Normal and 5 Freeloaders
- Outbound bandwidth - 4 children / node
- Ancestor Rating Threshold = -550

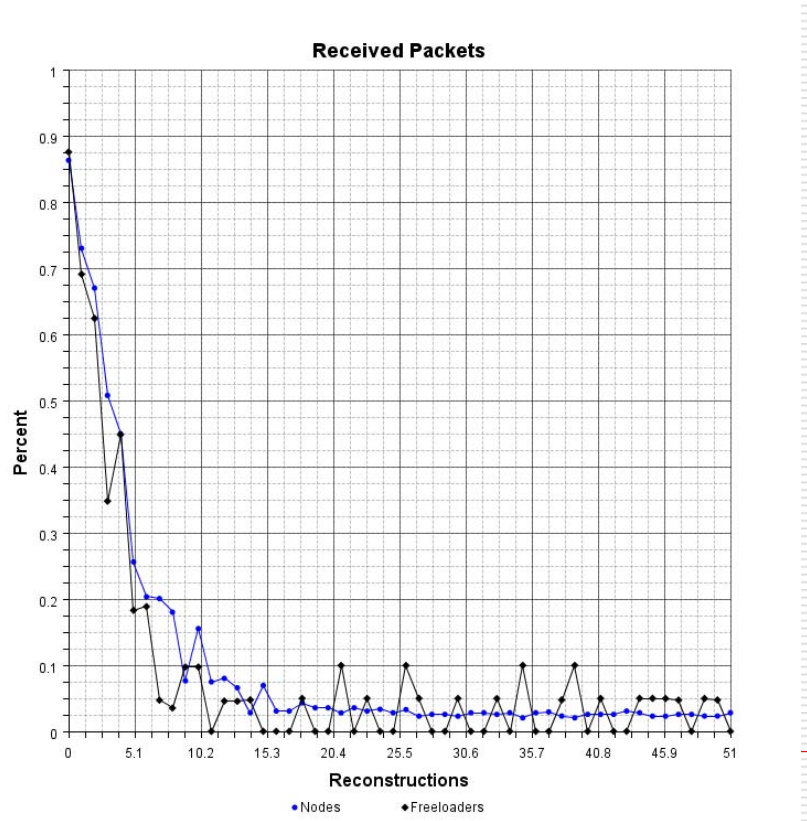


Tree Reconstruction Rate =every 30 sec

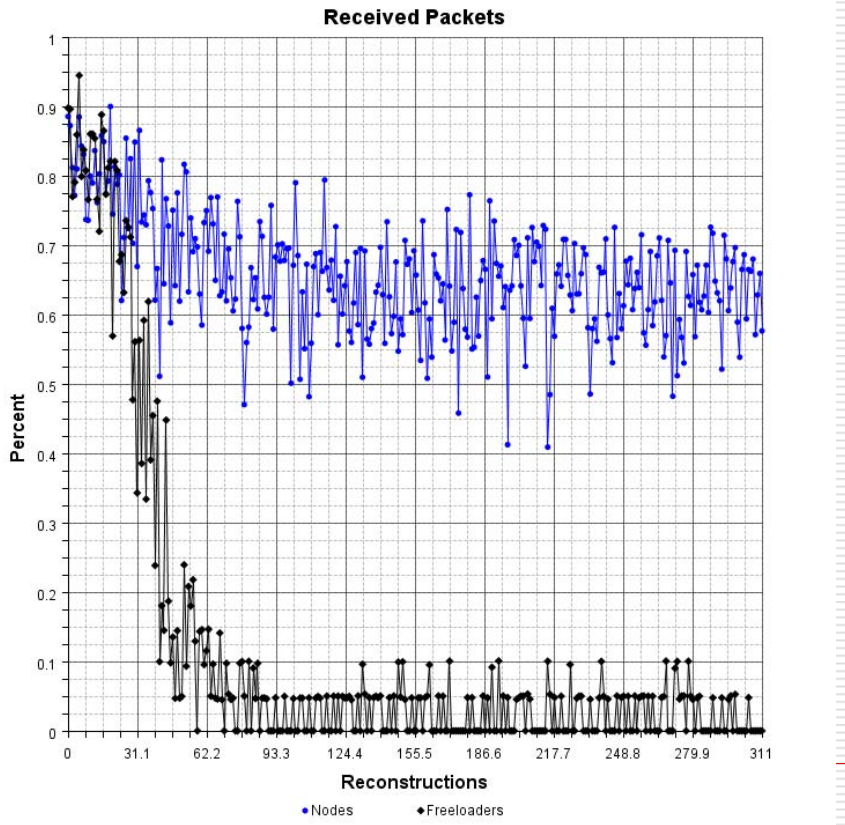


Tree Reconstruction Rate =every 5 sec

# Future Work

- Figure out problems with accuracy of detecting a freeloader and improving response
- Implementation of Publisher Taxation algorithm
- Other methods of fairness detection
  - Detection of Freeloaders refusing to accept children
  - Fairness detection under taxation
  - Sybil attack
- Pastry Type of Tree Manager
- Try further experiments with input parameters changing packet size, number of stripes, simulation time

# Conclusion from our Results

- Success of detecting freeloaders most sensitive to tree reconstruction rate and threshold values
- If tree reconstruction rate and threshold are optimal, the ancestor rating is much more effective than debt maintenance
- Node outbound bandwidths (number of children per node) are also influential
  - With more children, freeloaders are denied service faster, but normal nodes receive less packets
  - With fewer children, requires more time to deny service to freeloaders, but normal nodes suffer less packet loss
- Node numbers also influential
  - With larger node numbers, harder to detect freeloaders
    - at about 500 nodes using ancestor rating, takes longer to deny freeloaders
    - at about 500 nodes using debt maintenance, simulation time was not long enough to detect any freeloaders

# Demo

- Sample Input File
- Sample Output Files
- How to use

# Research Papers


1. Castro, M., Druschel, P., Kermarrec, A., Nandi, A., Rowstron, A., and Singh, A. 2003. SplitStream: high-bandwidth multicast in cooperative environments. In Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles (Bolton Landing, NY, USA, October 19 - 22, 2003). SOSP '03. ACM Press, New York, NY, 298-313. DOI= http://doi.acm.org/10.1145/945445.945474
2. T. W. J. Ngan, D. S. Wallach, and P. Druschel. Incentives-Compatible Peer-to-Peer Multicast. In The Second Workshop on the Economics of Peer-to-Peer Systems, July 2004. http://citeseer.ist.psu.edu/ngan04incentivescompatible.html
3. Chu, Y. 2004. A case for taxation in peer-to-peer streaming broadcast. In Proceedings of the ACM SIGCOMM Workshop on Practice and theory of incentives in Networked Systems (September 2004). ACM Press, New York, NY, 205-212. DOI= http://doi.acm.org/10.1145/1016527.1016535