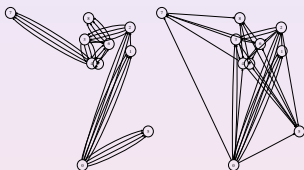


The performance of locality-aware topologies for peer-to-peer live streaming



R. G. Clegg (richard@richardclegg.org), D. Griffin, R. Landa, E. Mykoniati, M. Rio,
Dept. of Electronic and Electrical Engineering, UCL

Multiservice Networks, Cosener's House 2008

(Prepared using \LaTeX and beamer.)

Problem area

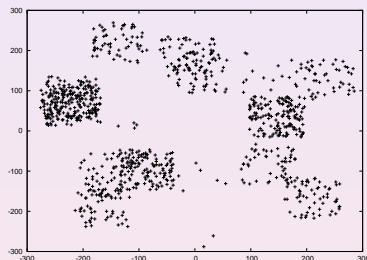
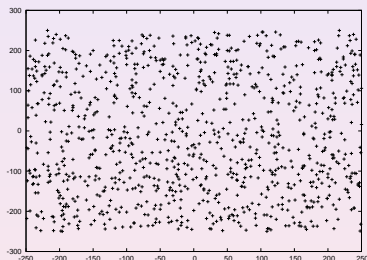
Motivation

- Current research interest in peer-to-peer live streaming.
 - Peer actions must be largely distributed.
 - Want low start-up and end-to-end delay.
 - Network co-ordinates give a distributed delay estimation tool.
 - Given delay info, how should peers choose partners?
-
- Want good end-to-end (peercaster to peer) delay, not throughput.
 - Want good reliability even in high churn.
 - Investigate this with simple low-parameter simulation.

Delay space

Delay estimate is distance in 2D Euclidean space (simplification of Vivaldi).

- 1 Flat peer distribution \mathcal{N}_F .
- 2 Loosely clustered peer distribution \mathcal{N}_L .
- 3 Tightly clustered peer distribution \mathcal{N}_T .



Experiment details

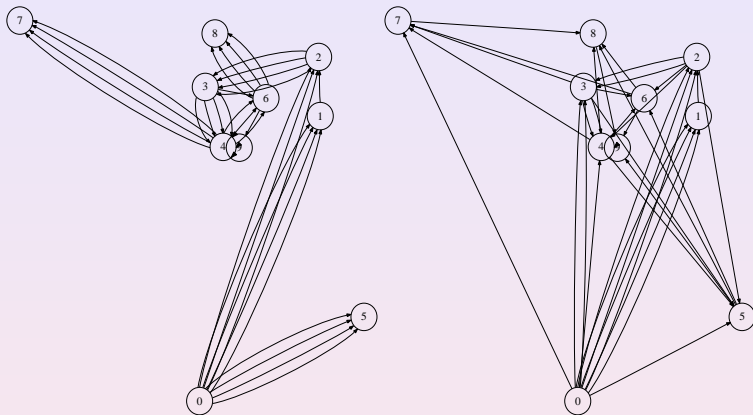
- Distribute $N + 1$ peers $(0, \dots, N)$ in the delay space and pick subset $n \leq N + 1$ for experiment.
- The stream has fixed bandwidth B . Peer 0 (peercaster) has some fixed upload capacity.
- Peers $i > 0$ randomly allocated some upload capacity from a distribution.
- Peers join in order and choose M (here 4) peers with spare upload (according to the **topology strategy**).
- Vary n , the peer distribution and the topology creation strategy.
- Repeat each experiment ten times to create a mean and a 95% confidence interval.

Topologies investigated

These strategies were investigated.

- **Local random** \mathcal{T}_R – M random peers selected.
- **Local closest first** \mathcal{T}_{C1} – M peer(s) with least delay to this peer.
- **Local closest with diversity** \mathcal{T}_{C2} – as above but M distinct peers if possible.
- **Local minimum delay first** \mathcal{T}_{D1} – M peer(s) with least delay to peercaster.
- **Local minimum delay with diversity** \mathcal{T}_{D2} – as above but M distinct peers impossible
- **Local small world** \mathcal{T}_S – This topology has $M - 1$ connections using \mathcal{T}_{C2} and one peer using \mathcal{T}_R .

Ten nodes connected with \mathcal{T}_{C1} and \mathcal{T}_{C2}



Metrics used – delay and vulnerability

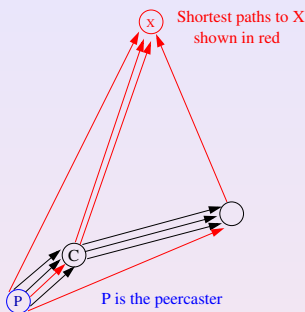
Let $D_i(j)$ be the delay from peer i using first hop on connection j and then shortest delay path. Let V_i (node vulnerability) be the maximum number of paths along $D_i(j)$ from i cut by the removal of one other node.

- **Mean minimum delay**

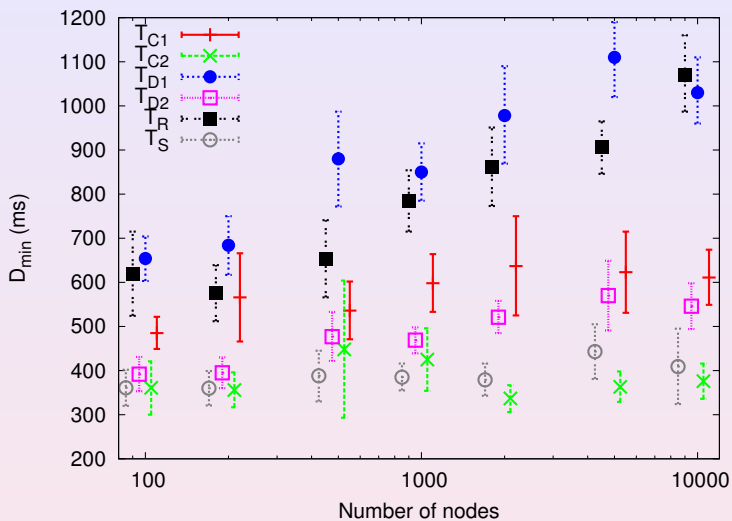
$D_{\min} = \sum_{i=1}^N \max_j D_i(j) / N$,
this is the mean of the minimum delay to the peercaster.

- **Mean node**

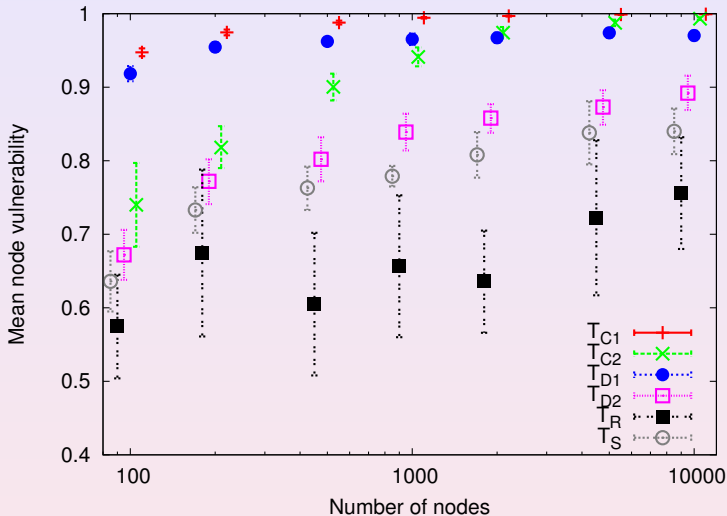
vulnerability $V = \sum_{i=1}^N V_i / NM$
– this is the mean proportion of its connections which each node could potentially lose by the removal of a single node.



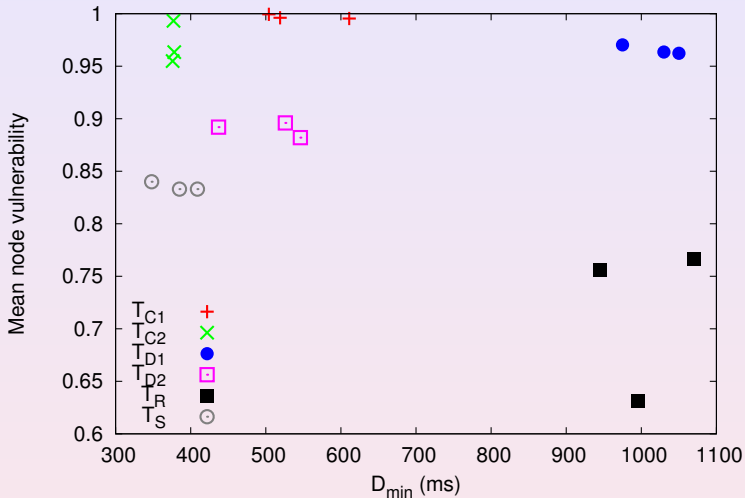
X has a node vulnerability of 2 when the node C is cut, two of the 4 red paths are cut as a result.

Results for D_{\min} on \mathcal{N}_F 

Results for \mathbf{V} (node vulnerability) on \mathcal{N}_L (loosely clustered)



Results for \mathbf{V} (node vulnerability) versus \mathbf{D}_{\min} all topologies $n = 10,000$



Conclusions and further work.

- The particular distribution of nodes seemed of lesser importance than the topologies.
- Topology strategies emphasising diversity performed better in most tests.
- Delay measure seem to scale well with size for the best policies.
- Much of the parameter space remains to be explored (reevaluating topologies).
- Need mathematical rigour but also to compare with a detailed simulation.
- See UK PEW paper for further details
www.richardclegg.org/pubs.