

Distributed Optimization of Event Dissemination Exploiting Interest Clustering

Max Lehn, Robert Rehner, Alejandro Buchmann

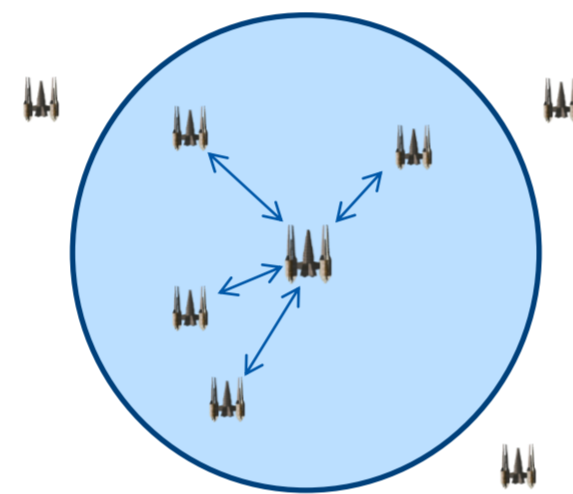

Goals

Interactive real-time online applications (e.g., games) need **timely many-to-many event dissemination**

Latency is critical → direct communication

↕ trade-off

Bandwidth is limited → aggregation




Further Challenges

- High dynamism in interest sets
- Heterogeneity in interest and capabilities
- No guarantees wrt. delivery, latency
- Streams of small update events

Existing Concepts


Application layer multicast

- + Efficient & scalable message dissemination
- No prioritization
- Group operations (join/leave) expensive



Publish/subscribe


- + Abstraction & decoupling, receiver-based selection
- Brokers are necessary infrastructure & bottleneck
- Subscription updates expensive



Context-aware [1] / parametric [2] pub/sub

- + No need for full re-subscription
- ~ Systems needs to be aware of context changes

```
subscribe(type = Action
  A myPos.x - 10 ≤ x ≤ myPos.x + 10
  A myPos.y - 10 ≤ y ≤ myPos.y + 10)
```



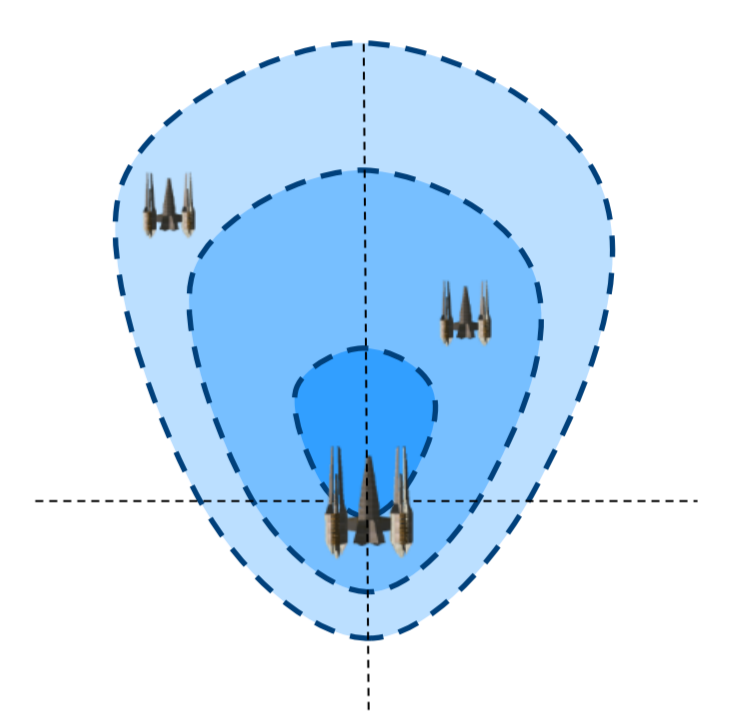

P2P gaming overlays (VON [3], pSense [4])

- ~ Interest management specific to virtual environments
- + Optimized for latency
- Event dissemination does not scale well

InterestCast

Interest-based Interface

Application-specific, continuous **interest level** is assigned to each neighbor



Interest Locality

Observation: interest in virtual worlds is local

Locality property: **transitivity**

$$B \in I_A \wedge C \in I_B \Rightarrow C \in I_A \text{ with high likelihood} \quad I_X: \text{interest set of node } X$$

Metric: transitivity ratio or **clustering coefficient** (C)


$$C = \frac{\text{number of closed triplets}}{\text{number of connected triplets}}$$

Interest locality introduces routing optimization potential:

- Shift load to more powerful nodes
- Aggregate messages, save connection overhead

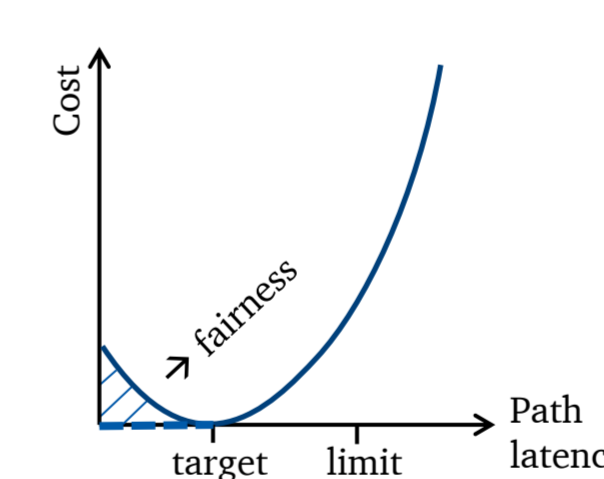

Local Optimization

- Each node **locally** evaluates utility
- Operations: **redirect** or **shortcut**



Utility Function

- Link utilization
- Path latency



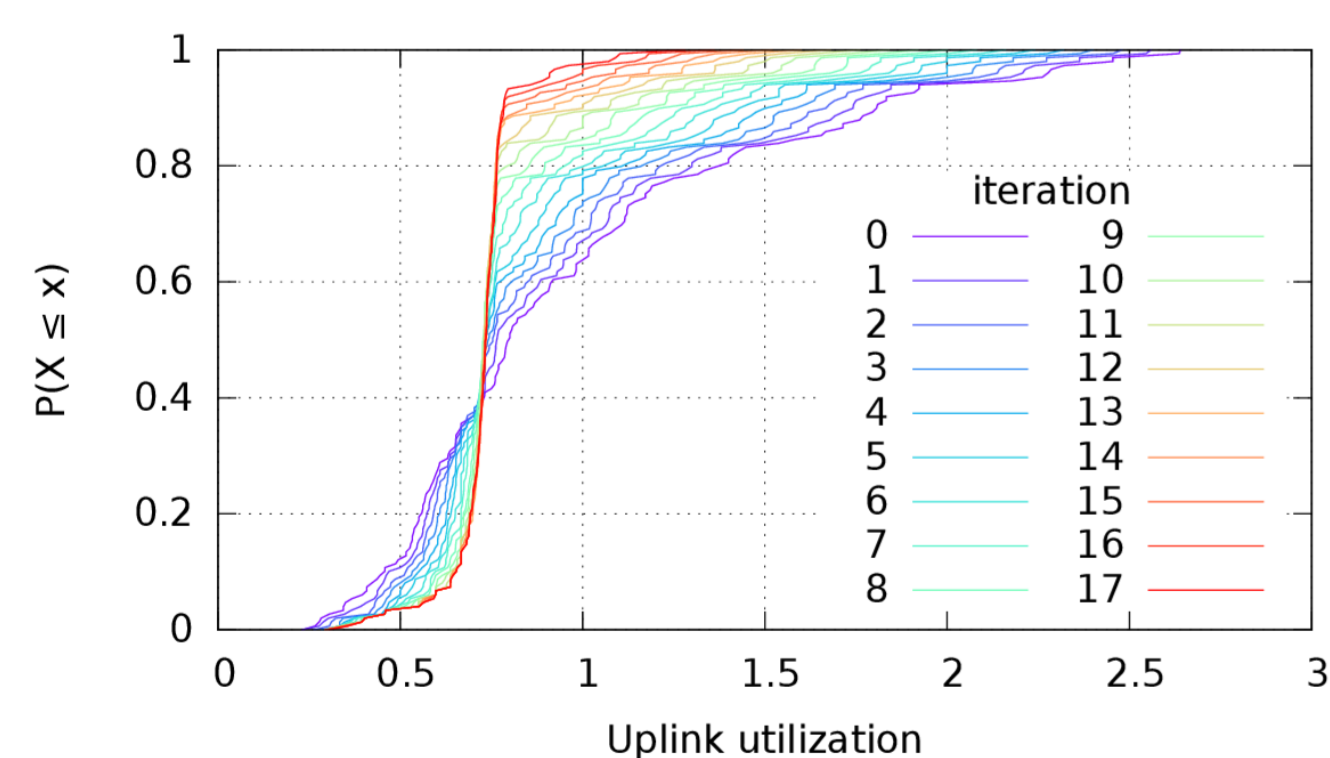
Finally: $utility = -c_{base}$

$$+ W_{bw,S} * (c_{bw,S,before} - c_{bw,S,after})$$

$$+ W_{bw,P} * (c_{bw,P,before} - c_{bw,P,after})$$

$$+ W_l * (c_{l,before} - c_{l,after})$$

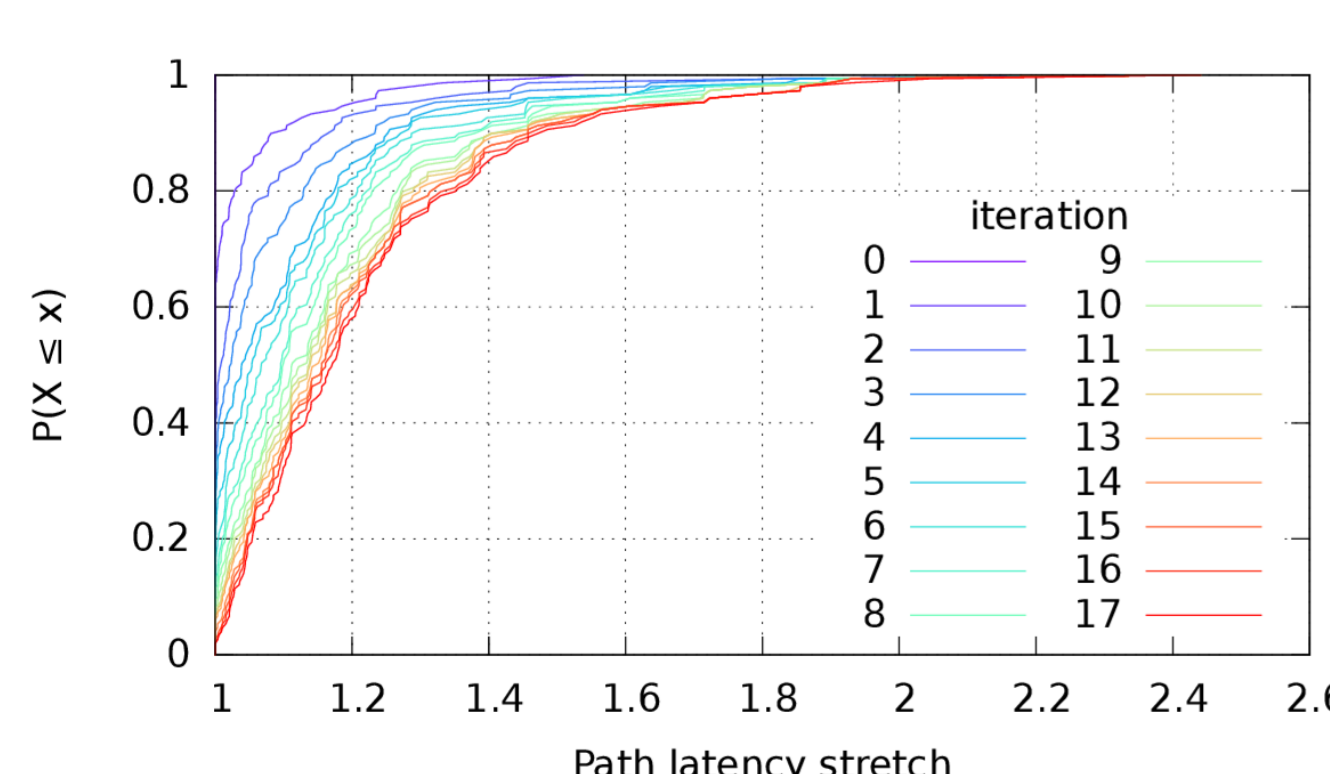
Evaluation



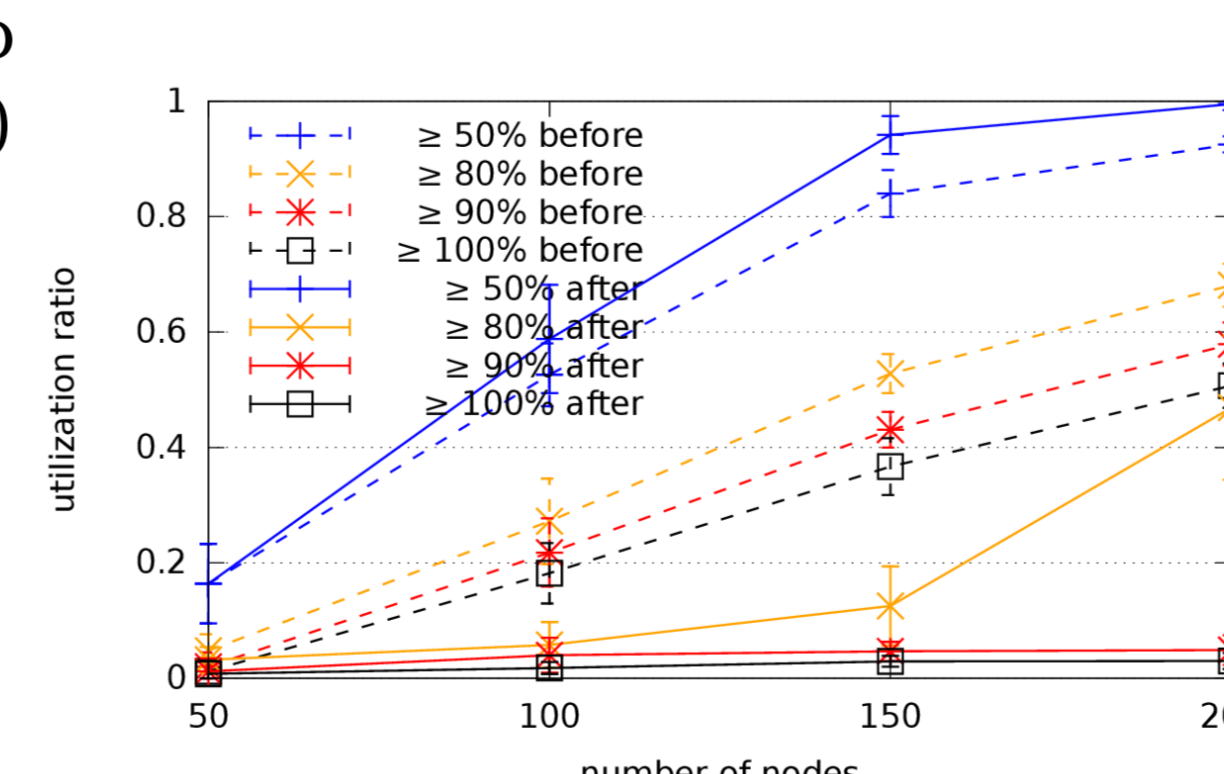

Virtual world interest graph

Link utilization distributions during a static optimization process


Optimization potential depends on the node fan-out (proportional to number of nodes)



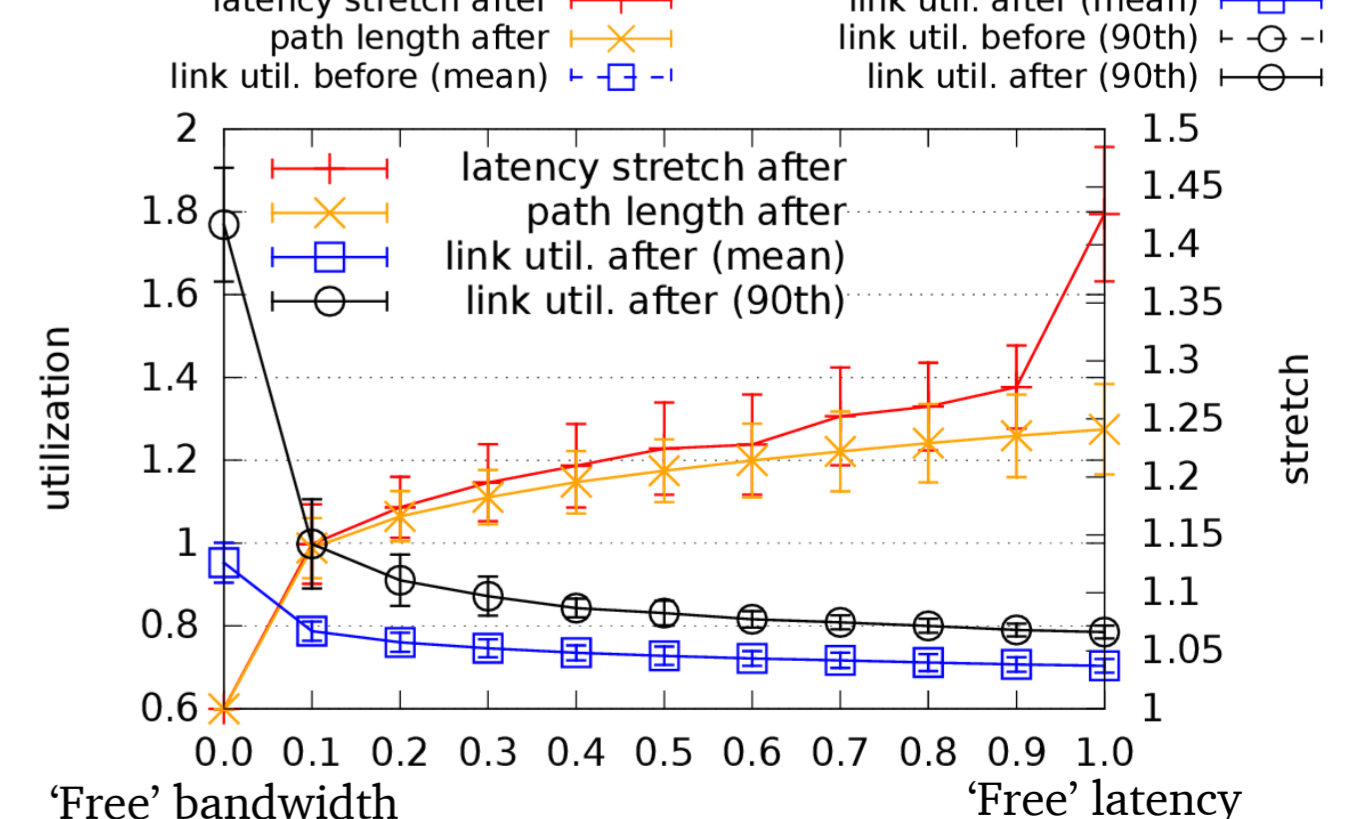
Path latency distributions during a static optimization process



The effect of clustering on the optimization potential (synthetic interest graph)



Tuning utility weights to trade off link utilization against latency



References:
 [1] Cugola G, Margara A, Miglavacca M.: *Context-Aware Publish-Subscribe: Model, Implementation, and Evaluation*. IEEE Symposium on Computers and Communications (ISCC'09). 2009:875–881.
 [2] Jayaram KR, Eugster P, Jayalath C.: *Parametric Content-Based Publish/Subscribe*. ACM Transactions on Computer Systems. 2013;31(2):1–52.
 [3] Hu S-Y, Chen J-F, Chen T-H.: *VON: A Scalable Peer-to-Peer Network for Virtual Environments*. IEEE Network. 2006;20(4):22–31.
 [4] Schmiege A, Stieler M, Jeckel S, et al.: *pSense - Maintaining a Dynamic Localized Peer-to-Peer Structure for Position Based Multicast in Games*. Eighth International Conference on Peer-to-Peer Computing (P2P'08). 2008:247–256.

