Converging an Overlay Network to a Gradient Topology

Håkan Terelius†, **Guodong Shi**†, **Jim Dowling**‡, **Amir Payberah**‡, **Ather Gattami**† and **Karl Henrik Johansson**†

†ACCESS Linnaeus Centre, KTH Royal Institute of Technology

‡Swedish Institute of Computer Science

Introduction

This work investigates the topology convergence in the gossip-based gradient overlay network.

- An *overlay network* is a virtual network built on top of another network, which structure can be utilized to improve application performance.
- The goal of the *gradient overlay network* is to arrange the nodes in layers, based on the their local utility values.

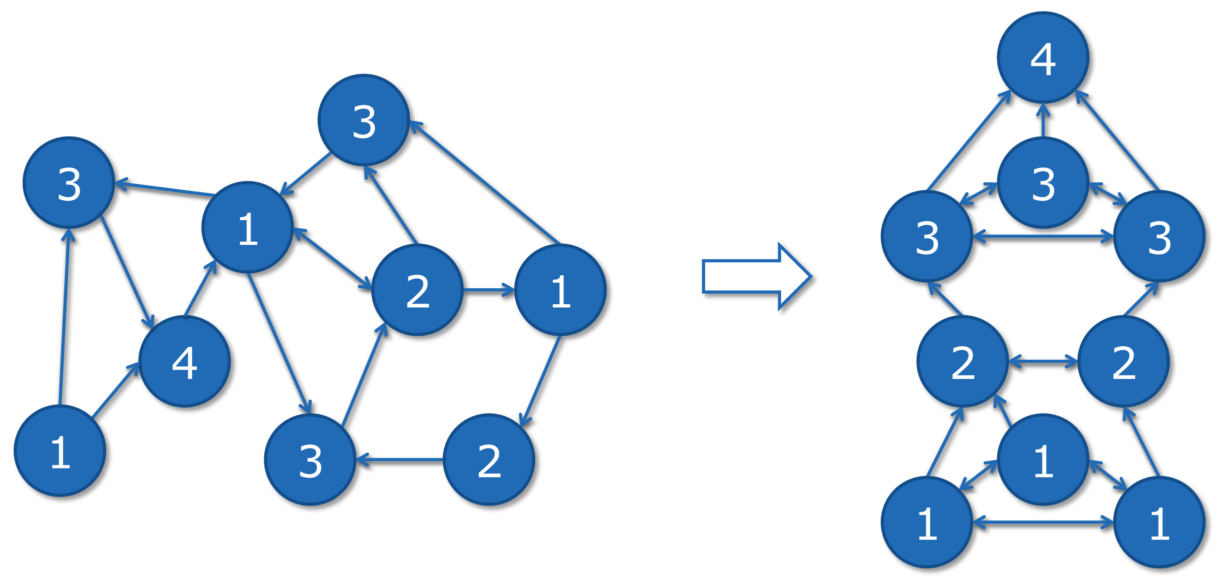

Network Model

- Consider a directed graph *G=(V,E)* of *N* nodes.
- Each node *v* is given a fixed utility value *U(v)* by the application.
- Node *v*'s out-neighborhood is denoted by N_v.
- The out-degree $m_v=|N_v|$ is bounded by *m.*
- Node *v*'s *preference function* $f_v(u)$ is defined by

$$f_v(u) = \begin{cases} -\infty & \text{if } U(v) > U(u) \\ -U(u) & \text{otherwise} \end{cases}$$

- Each node *v* is maximizing the sum of its preference function over N_v under a degree constraint:

$$\max_{|N_v|=m_v} \sum_{u \in N_v} f_v(u)$$

- Each node *v* update its out-neighborhood $N_v(t)$ at the discrete times *t=1,2,….*

Overlay Algorithm

The network topology is updated using a random peer sampling service according to the following algorithm:

foreach time $t = 1, 2, \ldots$ **do**
  foreach node v **do**
    Choose a node $u \in V$ with uniform probability p_t, $0 \le Np_t \le 1$
    if u is an improvement of v's neighborhood **then**
      Discard v's worst neighbor
      Add u to v's neighborhood
    end
  end
end

Main Results

Theorem (Convergence Condition)

The overlay algorithm converges to a gradient overlay network almost surely if and only if

$$\lim_{T\to\infty} \sum_{t=0}^{T} p_t = \infty.$$

Theorem (Convergence Rate)

The worst case expected convergence time is

$$\mathbb{E}[T] = \frac{1}{p}\sum_{n=1}^{m-1}\frac{1}{n} \le \frac{1+\ln(m-1)}{p}.$$

Proof Sketch

Consider each node's neighborhood, and rewrite the algorithm as a Markov chain. Notice that the transition probabilities are proportional to p_t. The results follow from spectral analysis of the Markov chain.

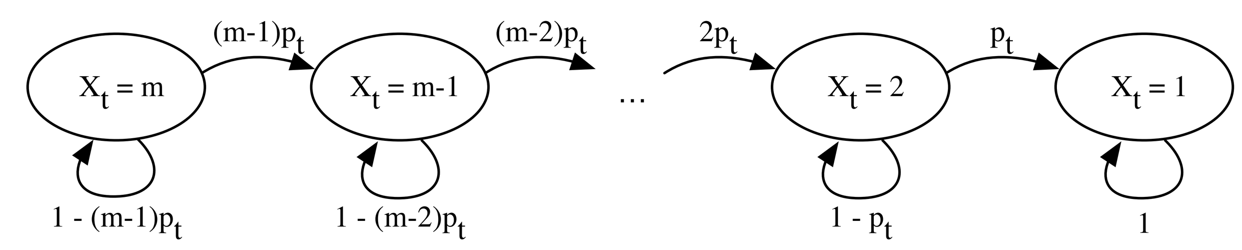

Live-streaming Application

The gradient overlay network is evaluated in a peer-to-peer live-streaming application called Glive.

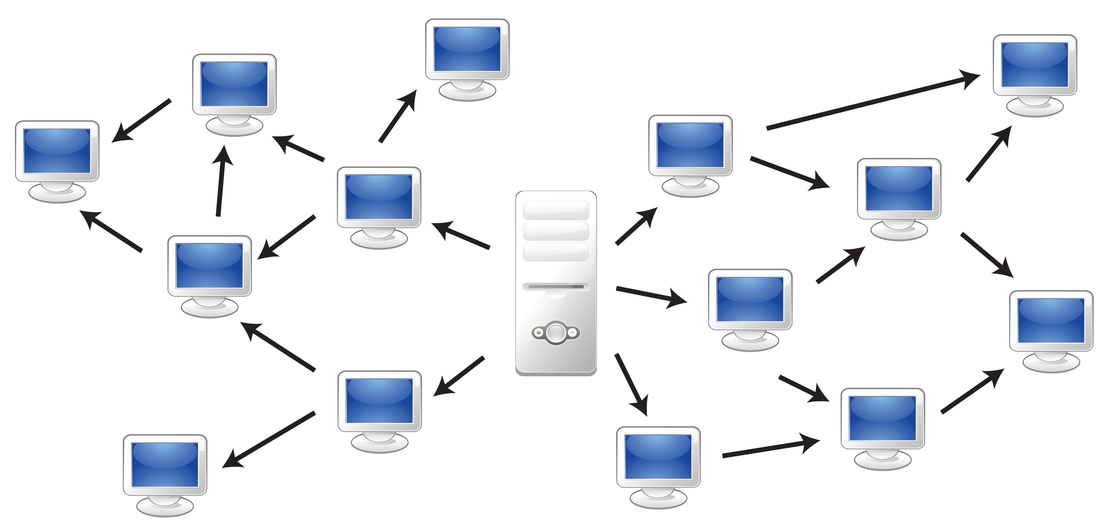

The gradient overlay network is built using the nodes' available upload capacity as their utility value, thus arranging the nodes with high upload capacity closer to the distribution source.

The gradient overlay network is compared against a random overlay network:

- *Playback continuity*: The nodes ability to play the video stream without interruptions when 1000 nodes join the network.
- *Playback latency*: The difference in playback point between a node and the source when 1000 nodes join the network, and then when 500 of the nodes leave the network after t=150 s.

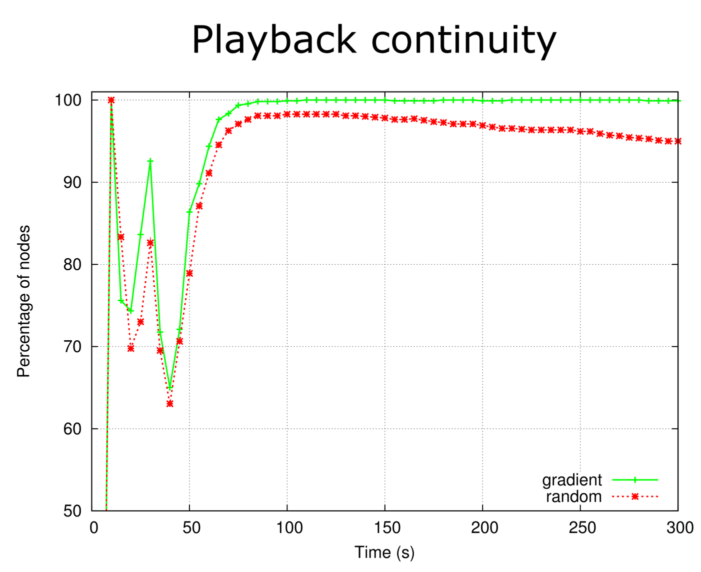

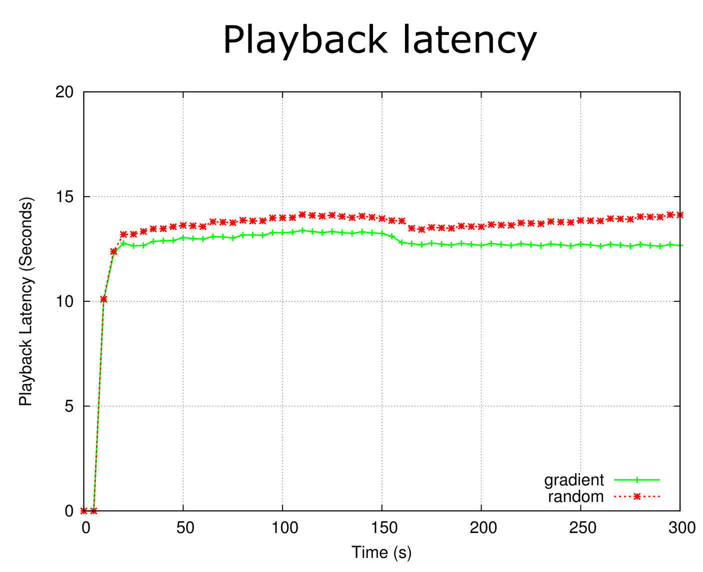

Our experiments on live-streaming applications demonstrates the importance of utilizing the structure in complex networks.

References

H. Terelius, G. Shi, J. Dowling, A. H. Payberah, A. Gattami, and K. H. Johansson, "*Converging an Overlay Network to a Gradient Topology*," http://arxiv.org/abs/1103.5678

A. H. Payberah, J. Dowling, F. Rahimian, and S. Haridi, "*Sepidar: Incentivized market-based p2p live-streaming on the gradient overlay network*," International Symposium on Multimedia, 2010.

M. Jelasity, S. Voulgaris, R. Guerraoui, A.M. Kermarrec, and M. van Steen, "*Gossip-based peer sampling*," ACM Trans. Comput. Syst., vol. 25, no. 3, 2007.

hakante@kth.se