School of Computer and Communication Sciences

Distributed Information Systems Laboratory (LSIR)

Wojciech Galuba, Karl Aberer, Zoran Despotovic, Wolfgang Kellerer

Bandwidth-efficient, delay- and loss-tolerant overlay routing

Problem

- Messages routed in overlays can be dropped or delayed
  - Significant cumulative impact over longer paths
- Existing protocols rely on message redundancy to mask routing failures
  - High bandwidth consumption

Failure prediction

- Feedback is binary – routing success or failure
- All nodes keep a **success estimator per each (next-hop, destination) pair**
- Success estimators: exponential running averages of message delivery rates
  - Updated when new feedback is available
- **Routing: pick the next hop with the highest current estimate**

Solution

Forward feedback protocol (FFP)

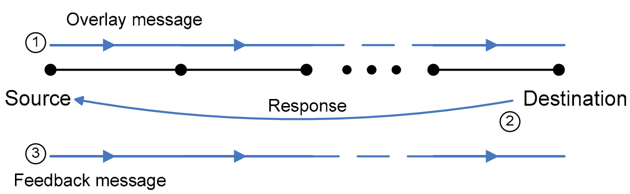

Idea: **Feedback follows the same route as the overlay message. Feedback tells forwarders whether they did a good job forwarding.**

Mechanisms

- Each node locally gathers feedback
- Based on feedback nodes learn to predict failures of their neighbors
- When one next hop becomes faulty alternative next hop is selected
- **Network-wide effect: good paths rapidly reinforced with positive feedback, traffic redirected on failures**

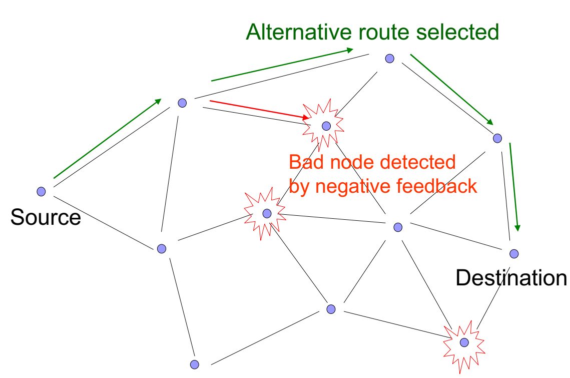

Results

- Chord deployed on 350 PlanetLab nodes
- Setup: every 5 min. further 10% of the nodes start dropping overlay messages
- Compared FFP to 4-way iterative, 4-way multipath and regular routing (NM on plots)
  - FFP uses up to 5 times less bandwidth
  - FFP reaches comparable success rate
- Similar results for message delaying nodes

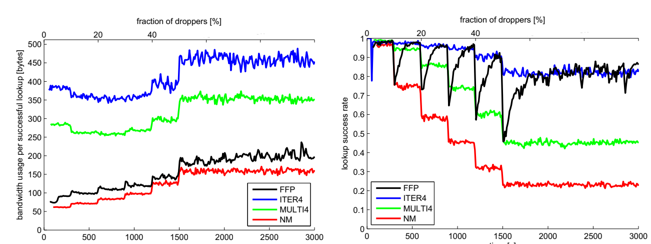

Properties

- **low overhead** – feedback messages are one-bit only (amortized), plenty of room for optimization
- **scalability** – local state is $O(\log^2 N)$ in terms of the network size
- **zero-knowledge routing** – FFP is capable of „cold-starting", no knowlede of neighbors' addresses or their routing tables necessary
- **universality** – works in any network, not only overlays
  - A secure variant of FFP developed for MANETS, outperforms other routing protocols

The work presented in this poster was partly carried out in the framework of the EPFL Center for Global Computing and was supported by the Swiss National Science Foundation. Parts of this work were done in collaboration with NTT DoCoMo Euro-Labs, Munich.

Contact: Wojciech Galuba
wojciech.galuba@epfl.ch