Karlsruhe Institute of Technology

Active Shared Access Network Congestion Avoidance in Heterogeneous Application-Layer Multicast

Capacity Sharing Workshop, Stuttgart, October 13, 2011
Christian Hübsch

Institute of Telematics, Karlsruhe Institute of Technology (KIT)

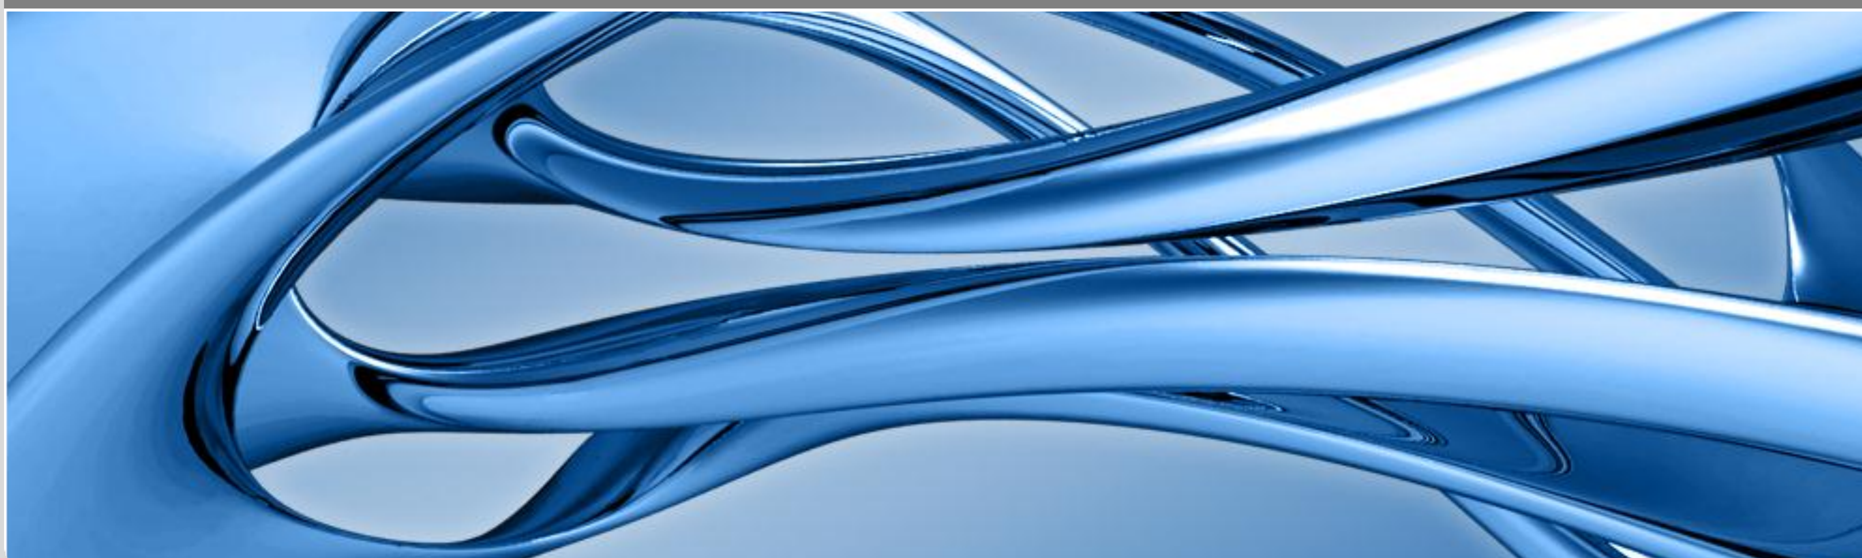

Mobile Communications: Blessing & Curse

- More and more mobile ES
- Higher Bandwidth Contents
- Leads to Problems

Server Side

Customers Angered as iPhones Overload AT&T

By JENNA WORTHAM
Published: September 2, 2009

COMMENTS
E-MAIL
PRINT
REPRINTS

Slim and sleek as it is, the iPhone is really the Hummer of cellphones.

Enlarge This Image

Michael Appleton for The New York Times
AT&T monitors its network from its operations center in Bedminster, N.J.
More Photos »

Multimedia

It's a data guzzler. Owners use them like minicomputers, which they are, and use them a lot. Not only do iPhone owners download applications, stream music and videos and browse the Web at higher rates than the average smartphone user, but the average iPhone owner can also use 10 times the network capacity used by the average smartphone user.

"They don't even realize how much data they're using," said Gene Munster, a senior securities analyst with Piper Jaffray.

Access Network Side

2010 — x26 — 66% Video → 2015

Solutions in the Scenario

- You could…
  - enhance infrastructure or add new mechanisms,
  - adapt the stream (e.g. Layered Coding)
  - Use P2P mechanisms to flexibly adapt multicast dissemination paths
  - Integrate and exploit alternative comm. channels

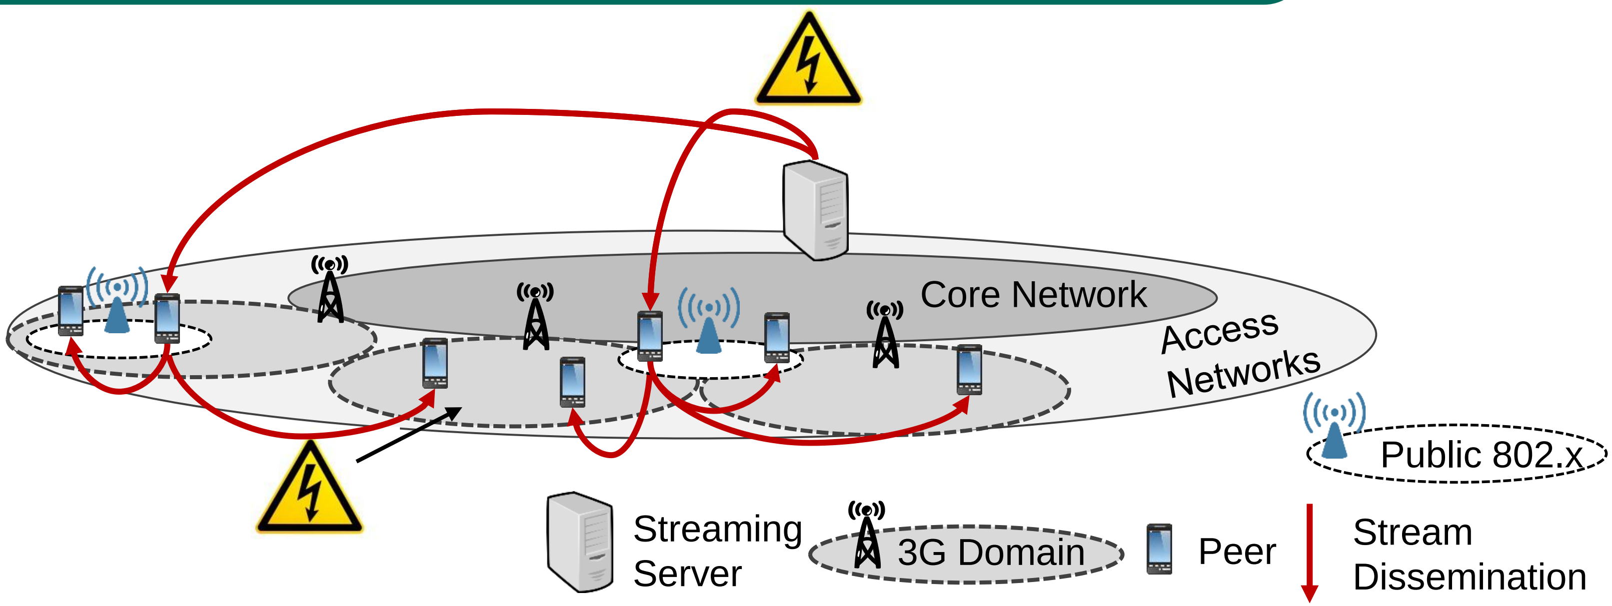

Use P2P mechanisms to flexibly adapt multicast dissemination paths

Integrate and exploit alternative communication channels

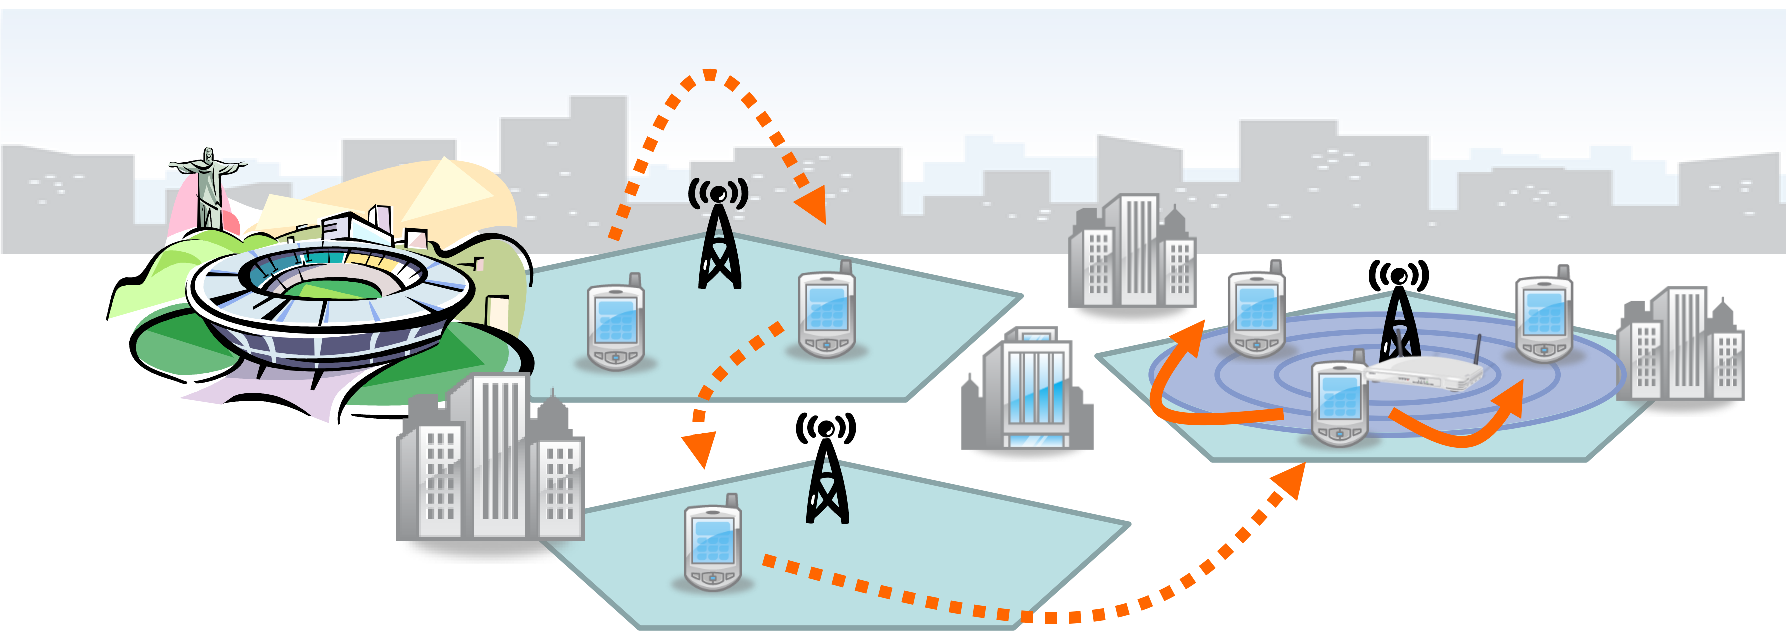

Capacity Matching

- Field with *m* cellular domains
  - „Shared Access Medium Domains (SAM)“
- *n* Peers
- Tree-based Multicast
- Topology Matching?
- **Capacity Matching!**

SAMs: Shared Communication Capacities!

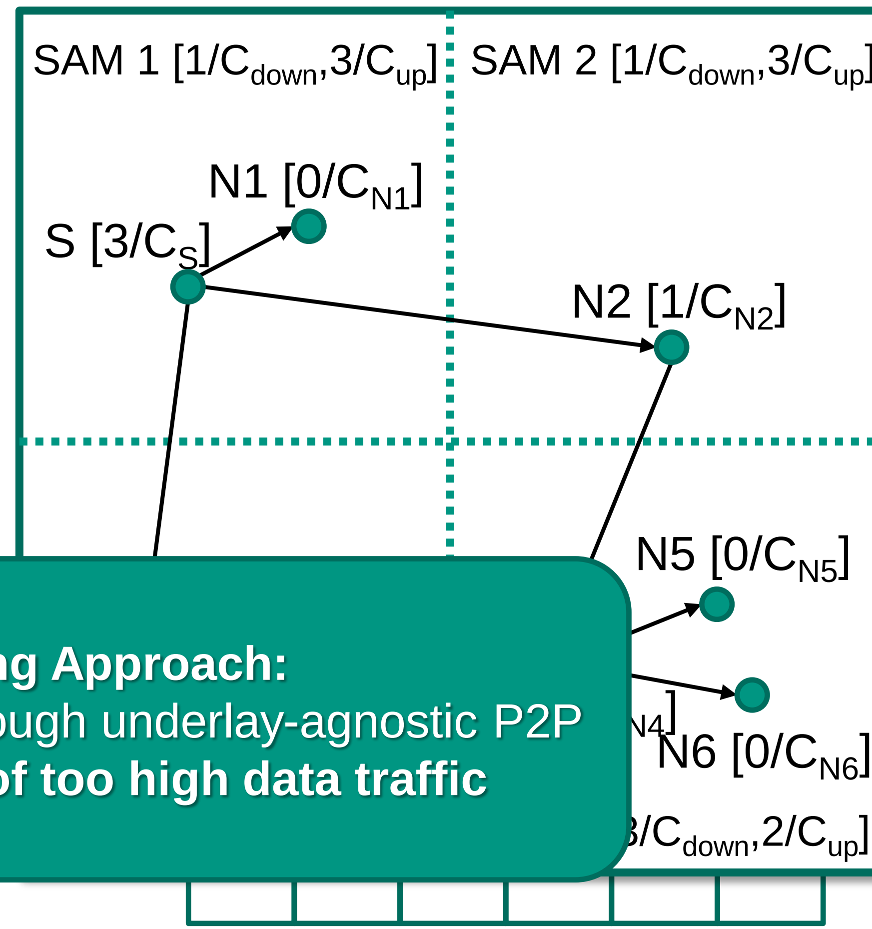

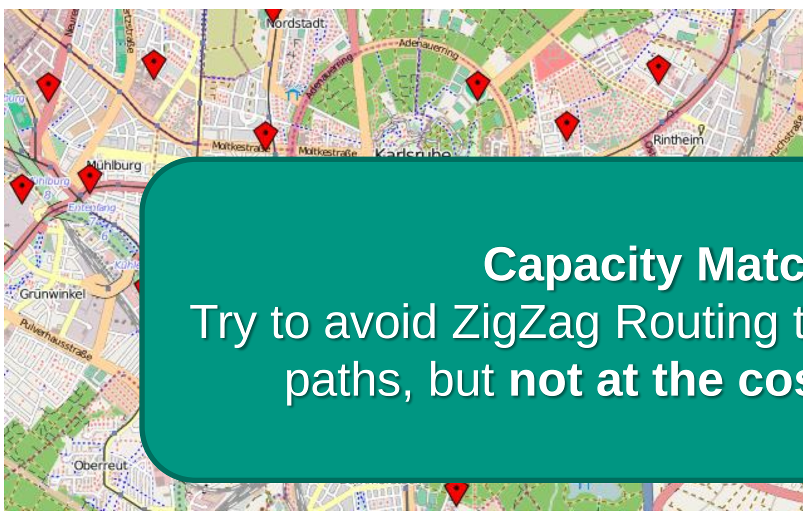

Capacity Matching Approach:
Try to avoid ZigZag Routing through underlay-agnostic P2P paths, but **not at the cost of too high data traffic**

Capacity Matching in Cellular Networks

- Goal: Tree-based dissemination (distributed and self-organizing)
  - Avoiding Access Network Congestions
- Metrices of Interest:
  - Maxmimum Dissemination Delay among all peers
  - Resulting Traffic Load in involved Access Networks (SAMs)
- Optimizational problem with > 1 Metric
  - NP-hard Problem
- Heuristic to find „near to optimal“ solution
- Weighted Sum
  - Transforms multi-dimension to single-dimension

$$\Psi(v) = (1-\alpha) * \text{Delay} + \alpha * \text{Load}$$

Delay → measureable

Load → e.g. IETF Alto (Application-Layer Traffic Optimization)

Distributed ALM Protocol

- Peers $V=\{v_1,\ldots v_i,\ldots, v_n\}$
- Peer Sampling:
  - No global knowledge
  - Random Subsets (RanSub, 2002)
  - Periodic, randomized Sampling
    → Random Subset $R = \{r_1,\ldots,r_j\}$
- Peer v_i calculates $\Psi(v_i \cup R)$
  - Parent node change
- Problems:
  - P2P induces even more traffic through forwarding
  - Only these outgoing edges adaptable

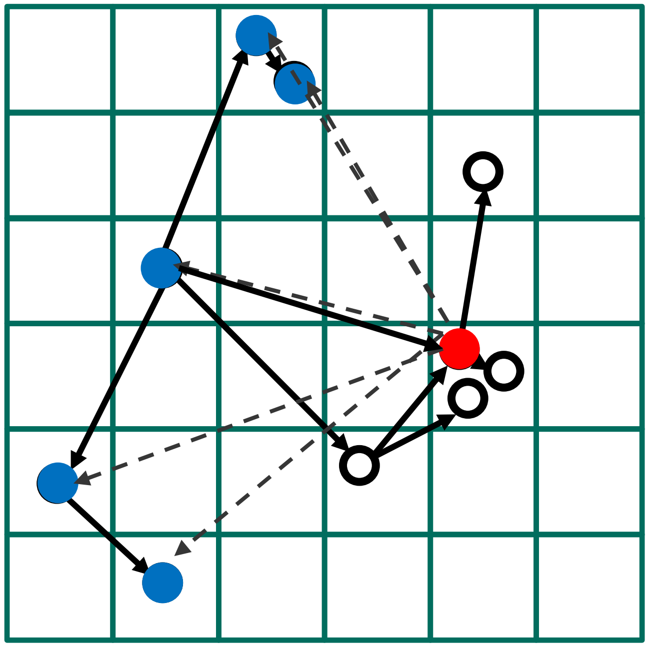

So: What can we win here, and at what price?

Load Balancing with Heuristic
Example: 500 Peers, 100 SAMs

$\Psi(v) = (1-\alpha) *$ **Delay** $+ \alpha *$ **Load**

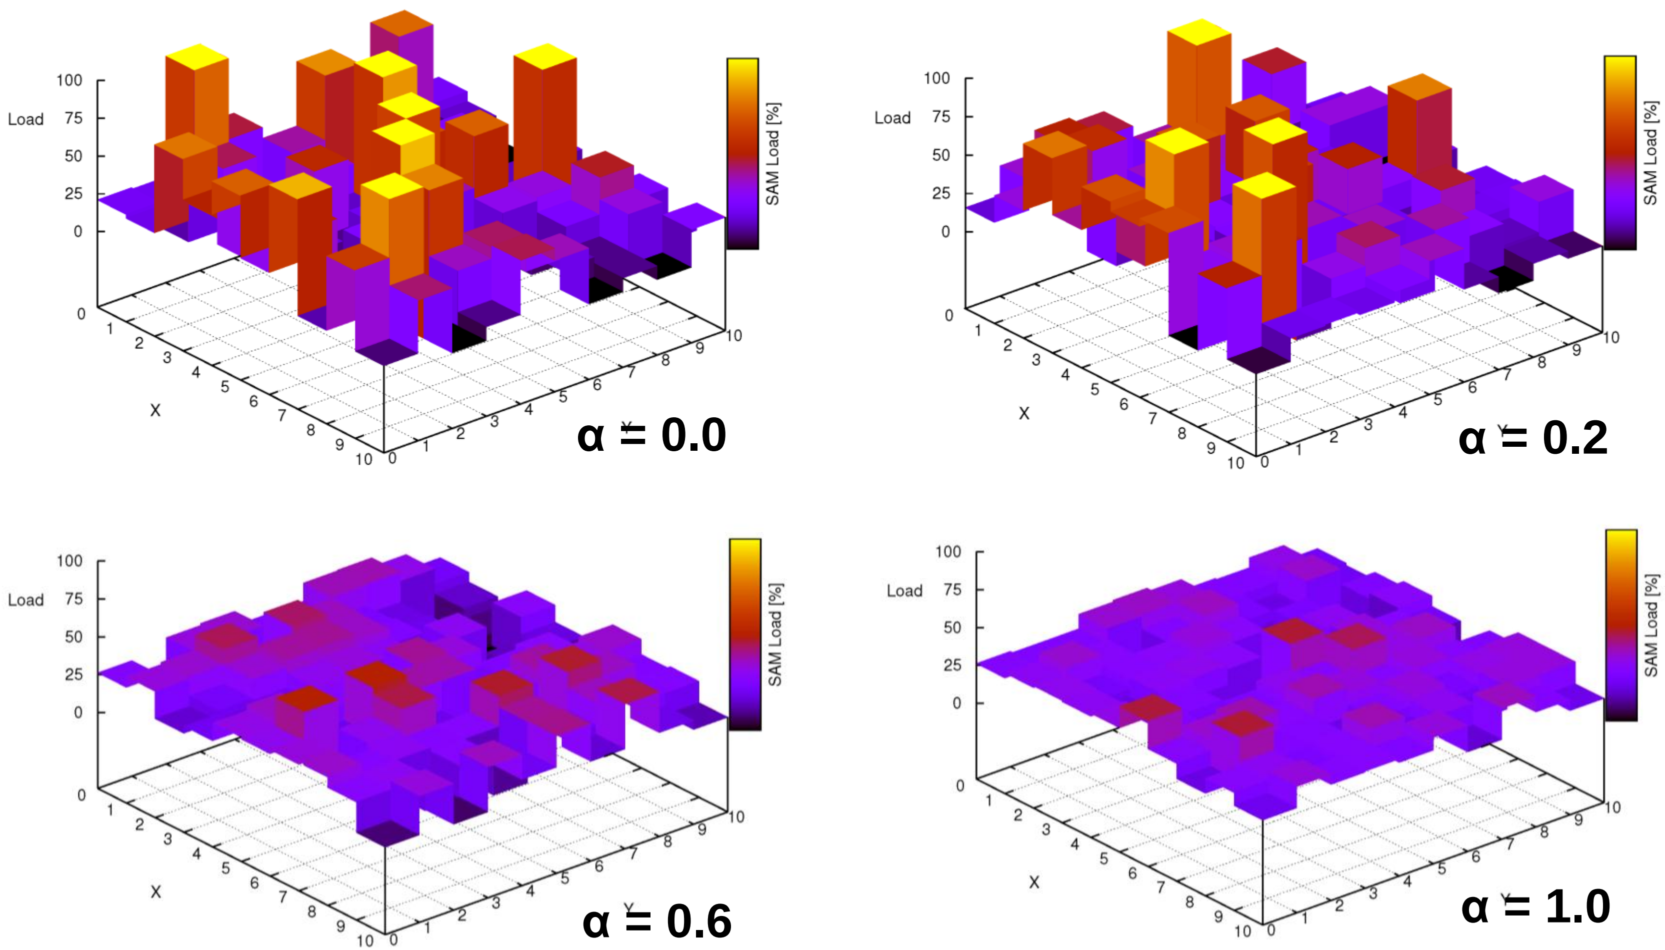

Used Metrices: Lorenz & GINI

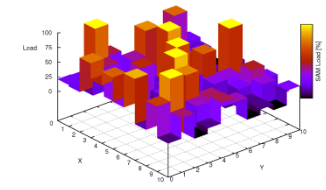

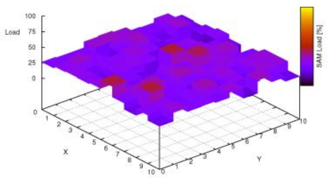

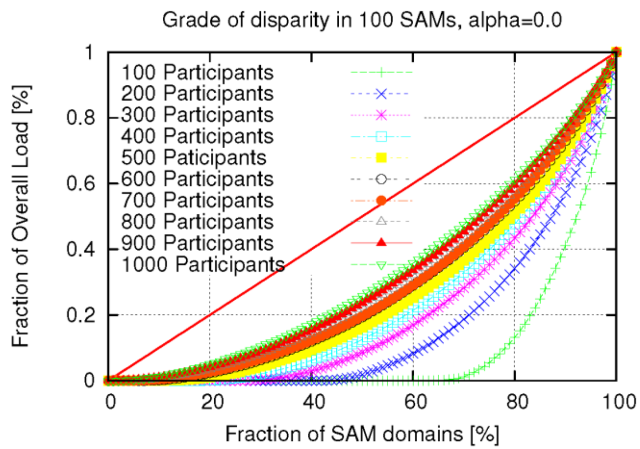

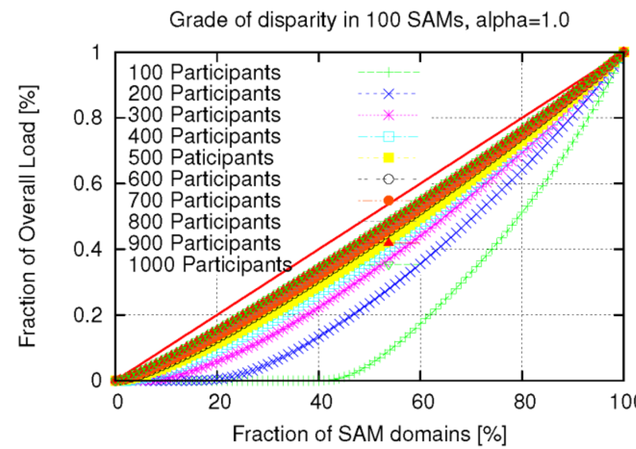

- Lorenz Curve: CDF of SAM Load Disparity

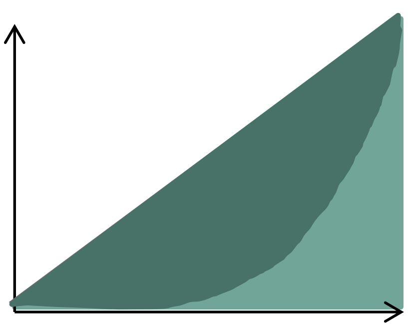

$$G^* = \left(\frac{2\sum_{i=1}^{n} i x_i}{n \sum_{i=1}^{n} x_i} - \frac{n+1}{n} \right) \cdot \frac{n}{n-1}$$

$$G^* \in [0, 1]$$

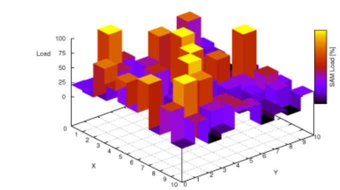

G* = 0.533

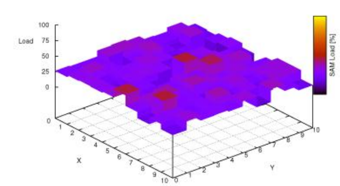

G* = 0.172

- GINI Coefficient: Scalar Value from Lorenz Curve

Trading Off Dissemination Delay Against Load

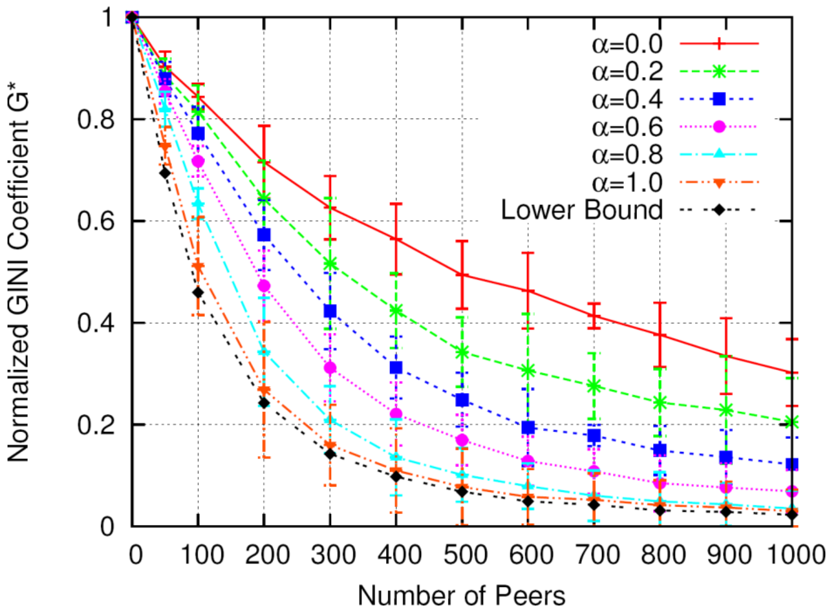

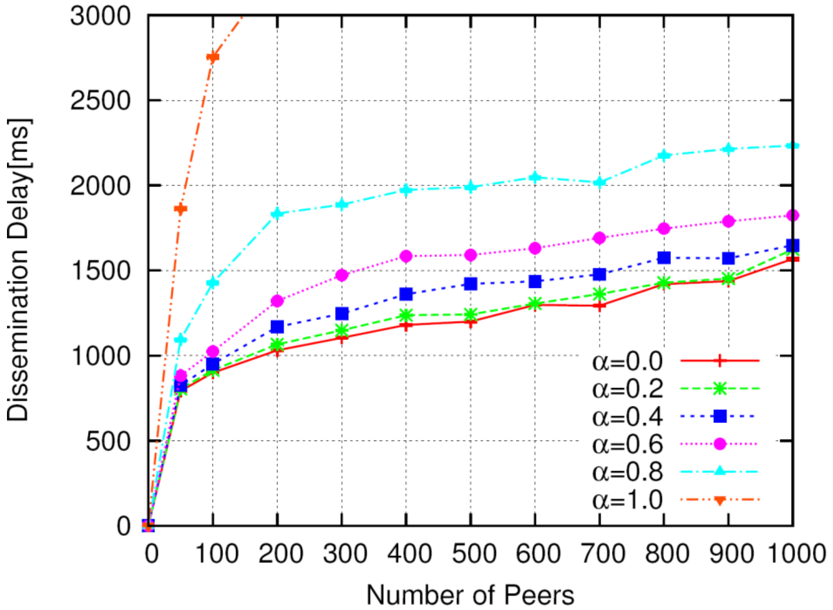

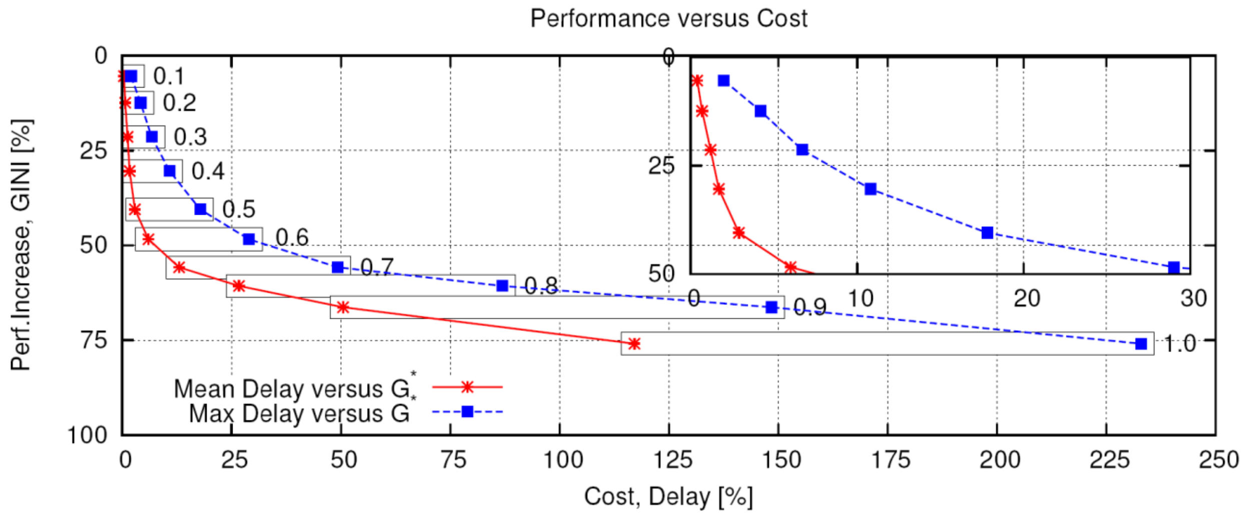

Insight here:
Trading Off between Delay and Load Balancing seems promising

Take Away Message So Far

Employing P2P Protocols in Tree-based Data Dissemination comes with benefits and drawbacks

- Benefits
  - Server congestion can be avoided by unburdening it from forwarding load
  - High flexibility can be achieved by end-system based decisions
    - Builds the basis for integration of further mechanisms

- Drawbacks
  - Users have to contribute to the system
  - P2P typically less robust than centralized approaches
  - P2P induces even more (outgoing) data traffic in access networks

- But…
  - If ISPs provide traffic information, the increased traffic can be handled
  - With further mechanisms, P2P benefits may outweigh the drawbacks…

Use P2P mechanisms to flexibly adapt multicast dissemination paths

Integrate and exploit alternative communication channels

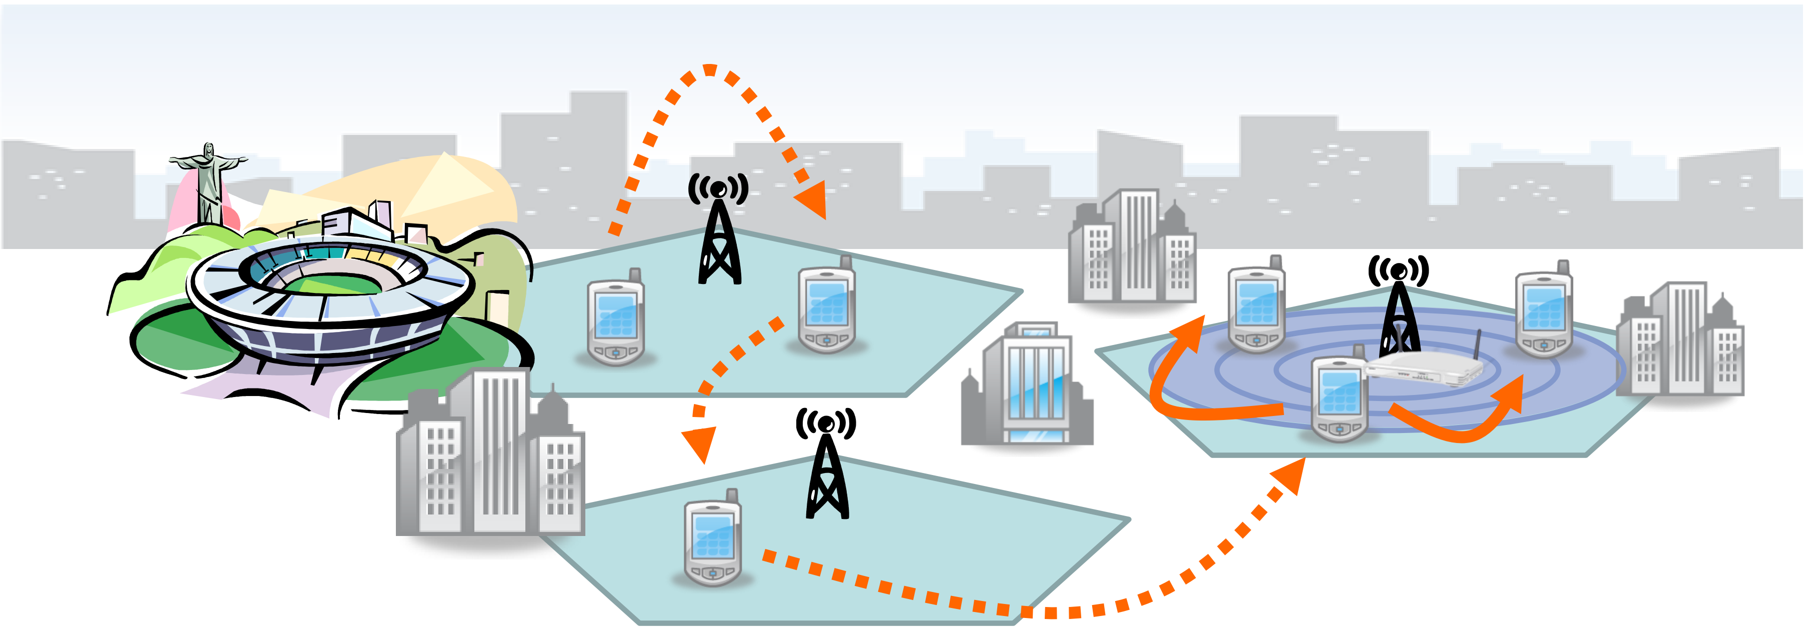

Motivation

- Sacrificing delay to copy with higher traffic load can alleviate the P2P drawback of newly induced traffic load
- But: To really lower the traffic load, alternative communication techniques have to be exploited
- Most modern mobile devices provide diverse possibilities here
  - IEEE 802.11x, Bluetooth, 3G/4G
- Goal:
  - Multiplex P2P edges to different communication domains, where possible

Not to see the wood for the trees…

- Goal: Find potentially reachable peers in alternative domains (e.g. IEEE 802.11x)

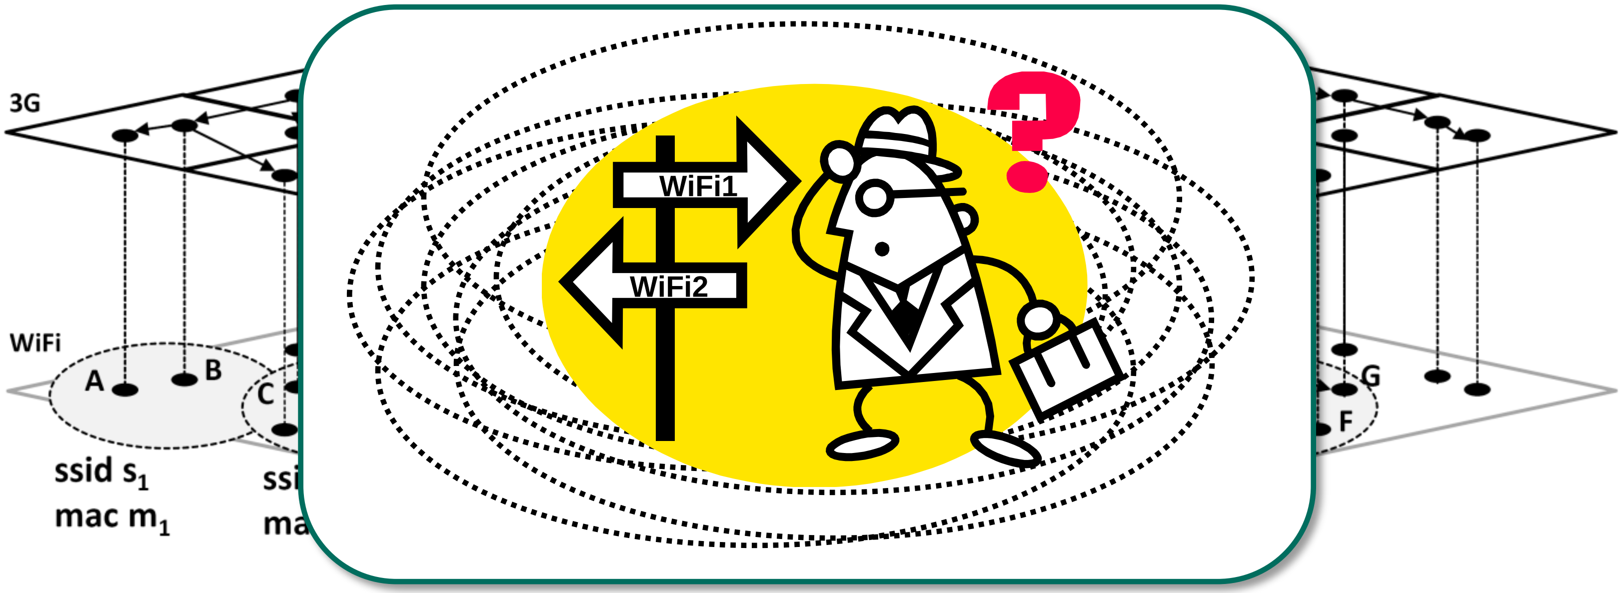

- Keep an eye on energy-efficiency and robustness
  - Ad-Hoc expensive→ Prefer infrastructure (Public WiFi)
- Promising: SSIDs per km^2 in Manhattan 2010: **~ 2K**
- But: How to synchronize?

Not to see the wood for the trees…

- Idea: Efficient encoding of communication possibilities via Counting Bloom Filters

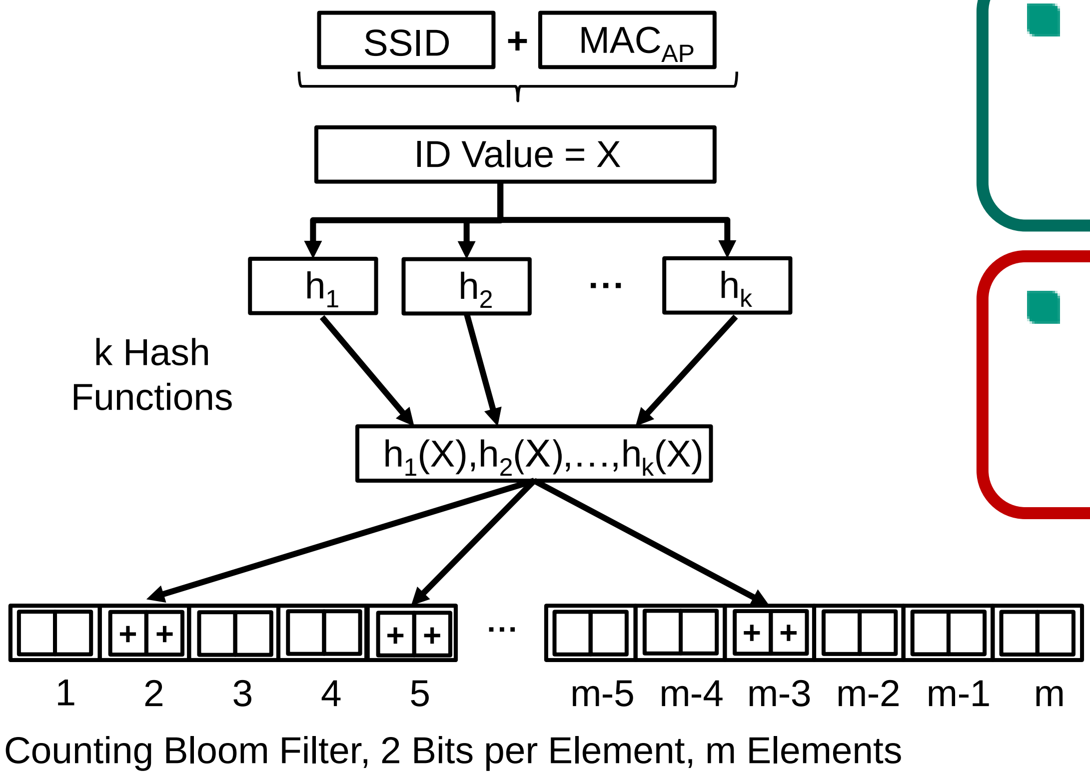

- Benefits
  - Aggregatable
  - Space-efficient
- Drawback
  - False Positives possible

Integrating WiFi with Bloom Filter Mechanism

- Integrate Bloom Filter in In-tree Control Traffic
- All Bloom Filters are aggregated by addition
- Aggregated Bloom Filter is sent down-tree to all peers
- If more than one entry in Bloom Filter:
  Check reachability and bargain roles

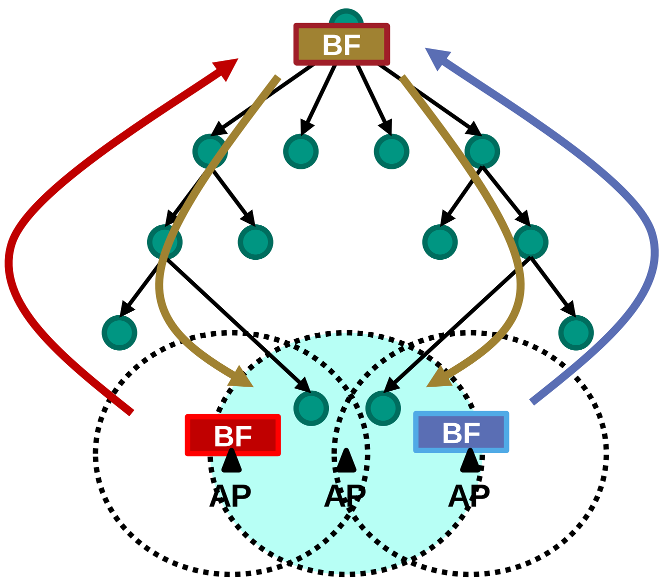

- Benefits
  - 64% Overhead savings compared to a list approach (with 1% False Positive Rate)
  - Try only when most likely successfull

Conclusion & Outlook

- Traffic load is current (and most likely persisting) problem in mobile communications
- P2P mechanisms to unburden server and provide high flexibility
- P2P also induces higher load
  - Can be handled if the network situation allows for it
  - Promising in typical cellular network environments
- Based on the P2P flexibility, further mechanisms can be integrated to really multiplex forwarding load via different technologies
  - Bloom Filter approach proposed

- Outlook
  - Integrate further technologies (Bluetooth, Femto-cells, …)
  - Build robust P2P protocol to cope with churn/mobility

Thanks for your attention.

Are there questions?