Content Dissemination in Mobile Wireless Networks

Chadi Barakat
Email: Chadi.Barakat@inria.fr
http://planete.inria.fr/chadi/

Context / Objectives

- **An activity running for couple of years, involving several people**
  - Amir Krifa, Francisco Santos, Benjamin Ertl, Thierry Turletti
  - Thrasyvoulos Spyropoulos (Eurecom)

- **To exchange data between mobile devices**
  - Without infrastructure – disaster, censorship, etc
  - Multi-hop mode: I give you, you give the others, and so on
  - Often called Delay Tolerant Networks, Pocket switched Networks, etc

- **Communication can be of different types**
  - Point-to-point: I send a content/message to someone
  - Broadcast
  - Publish-subscribe: One publishes, others search and request

Difficulties

- **Current Internet architecture does not work in this context**
  - No end-to-end path, no TCP/IP
  - Network topology changes frequently
  - Opportunistic encounters

- **Forwarding, replication and content-aware routing**
  - Either give the content to encountered devices and get rid of it
  - Or simply give them a copy
  - How to know?
  - Several solutions in the literature: spray-and-wait, age-based routing, community-based routing, routing by social networks, etc

- **Routing reduces the load, but does not provide explicit solutions in case of resource starvation**
  - TCP and buffer management vs. IP routing

Our framework

- **Transform the problem into a resource management problem**
  - Set a global objective for the network
    - Maximize throughput, Minimize delay, being fair, etc
  - Devices take local decisions that push the network to its desirable global behavior (and keep it there)
  - Two main decisions:

> **When there is a need to drop some data, drop the least useful first**
> **When there is a need to forward, forward the most useful first**

- **Control variable: The number of replicas …**
- **Control function: A utility per content**
  - The marginal gain/loss upon drop/replication

Algorithm in a nutshell

Beforehand

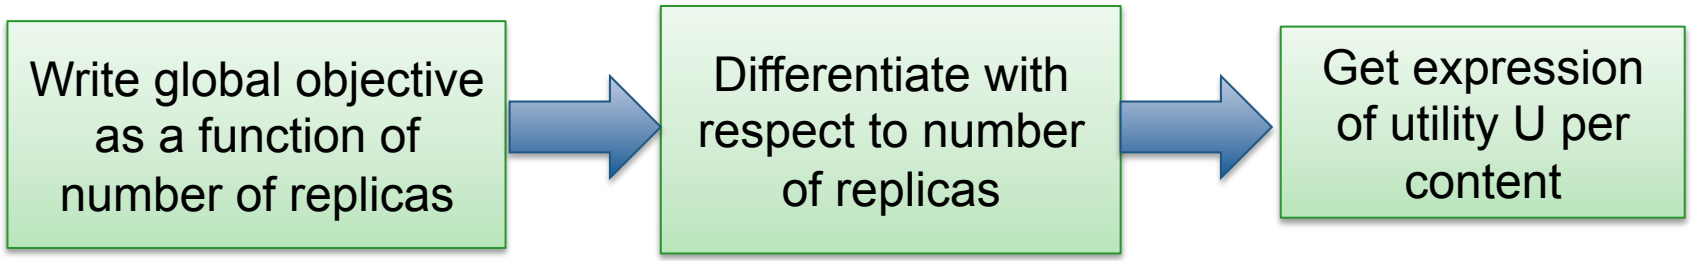

Locally, on the fly

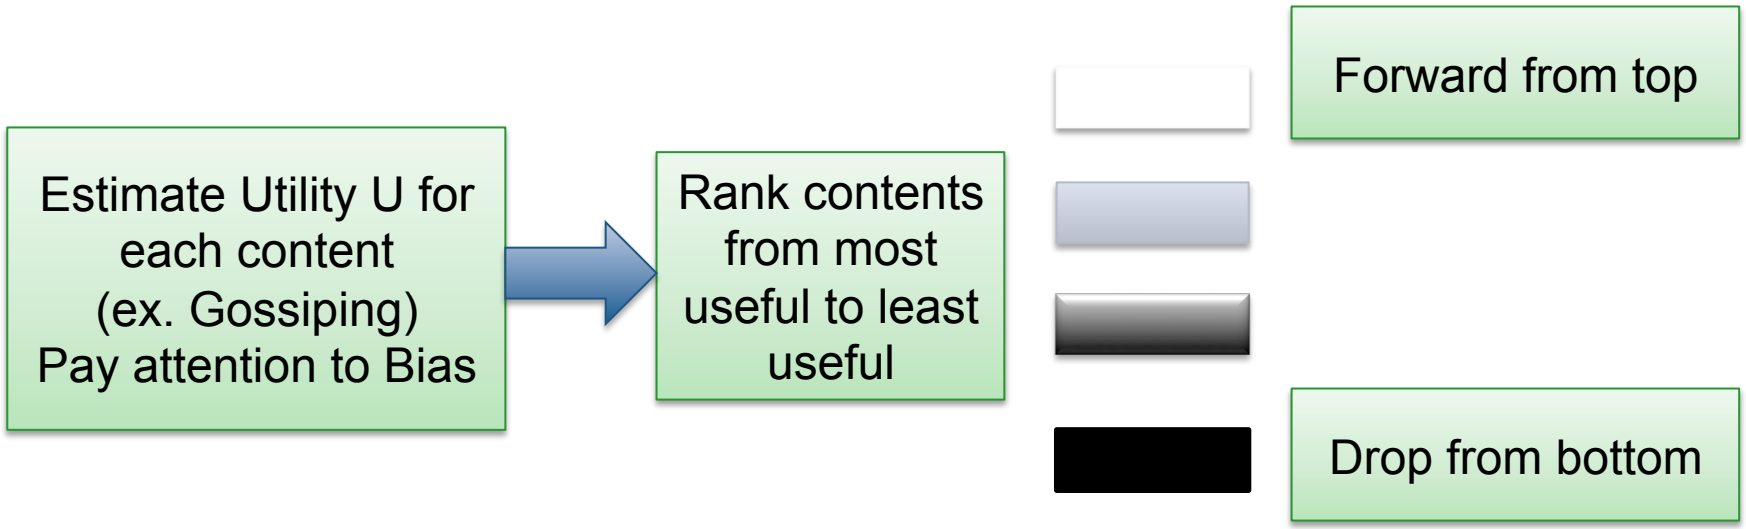

HBSD: The point-to-point case

- **History-Based Scheduling and Drop**
  - http://planete.inria.fr/HBSD_DTN2/
- **Contents/Messages appear at a device, try to find their way to some other destination device, then disappear**
- **Utility: A function of content age**
  - By gossiping, HBSD finds this function by itself
  - Same function for all contents
  - Age of content indicates its utility
- **Two variants: Maximum delivery and Minimum delay**

Further details: Amir Krifa, Chadi Barakat, Thrasyvoulos Spyropoulos, "Message Drop and Scheduling in DTNs: Theory and Practice", in IEEE Transactions on Mobile Computing, vol. 11, no. 9, September 2012.

Samples of utility functions

- For a lightly loaded network, utility decreases with age
  - Best is to schedule youngest first and to drop oldest first

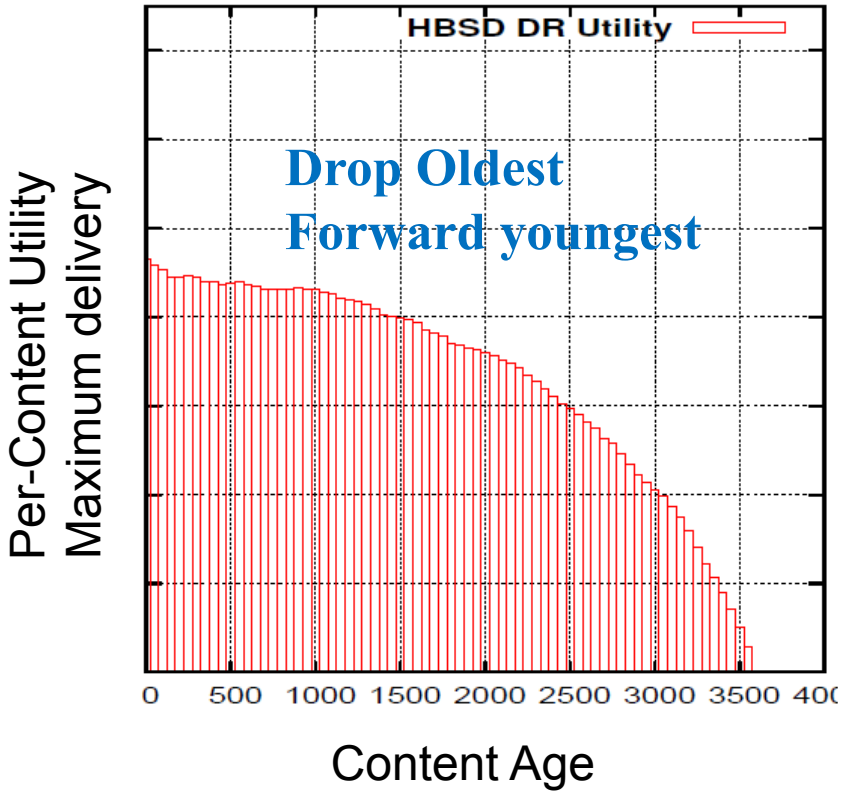

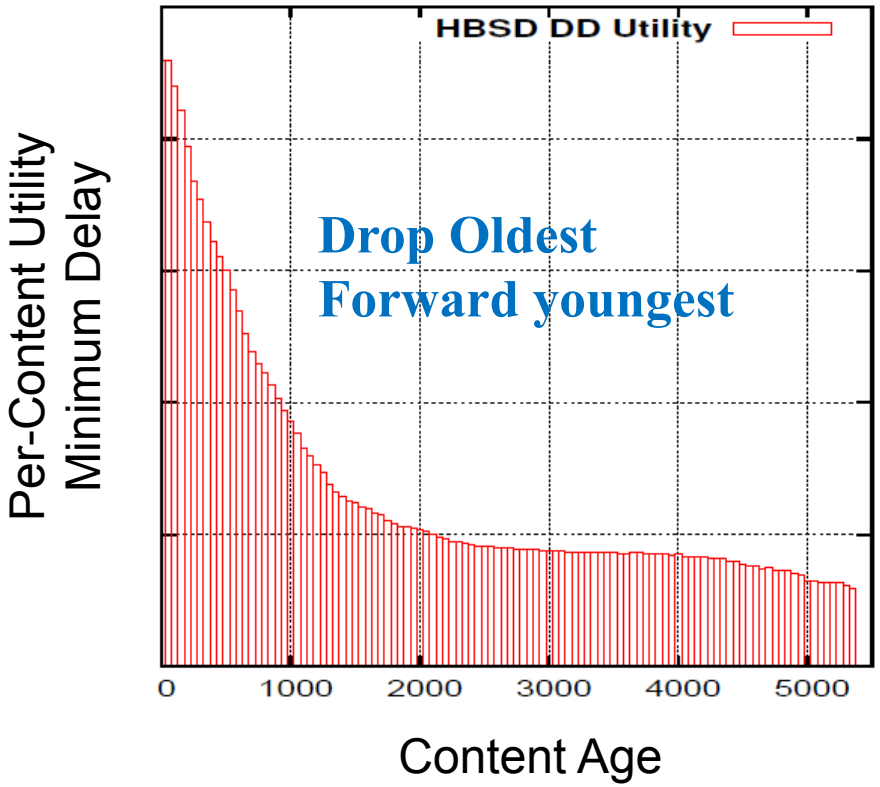

Samples of utility functions

- **When the load increases, the shape of utilities changes**
  - Simple policies are no longer optimal
  - HBSD can capture the optimal behavior whatever is the load

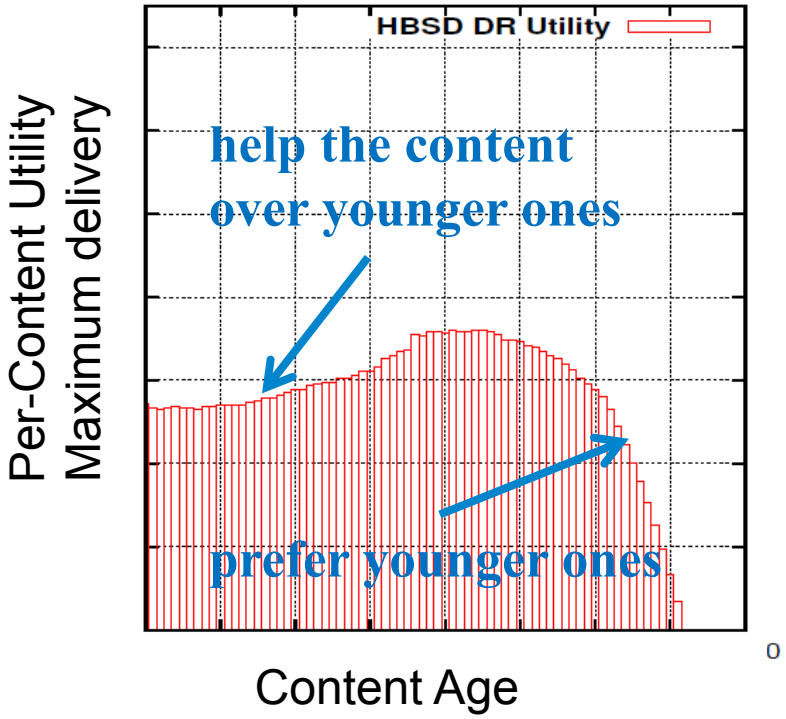

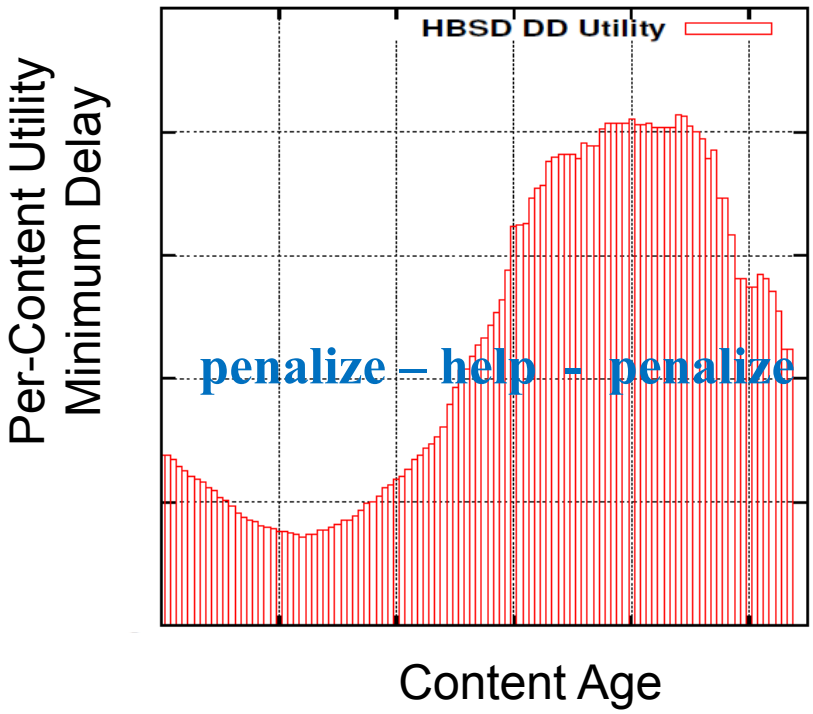

CODA: The publish-subscribe case (point-to-multipoint)

- **Content Optimal Delivery Algorithm**
  - http://hal.inria.fr/hal-00742734/
  - http://planete.inria.fr/Software/CODA/

- **Developed within the CCN/NDN framework**
  - Contents have universal names. Ex: \inria\team\diana\coda
  - Have different popularities (different request rates)
  - Users issue requests for contents, contents flow back, intermediate devices cache

- **Again, control the number of replicas per content**
  - Function of its popularity
  - Collaborative network-wide caching, taking mobility into account (vs. LRU, LFU)

- **First solution that maximizes the number of satisfied requests (throughput)**
  - Under some assumptions:

Utility ≈ miss rate = popularity – delivery

Maximize Throughput = Equalize miss rate

CODA: The publish-subscribe case (point-to-multipoint)

- Observed 50% more throughput than LRU
- And better protection of non popular contents
- Global network performance can be easily tuned

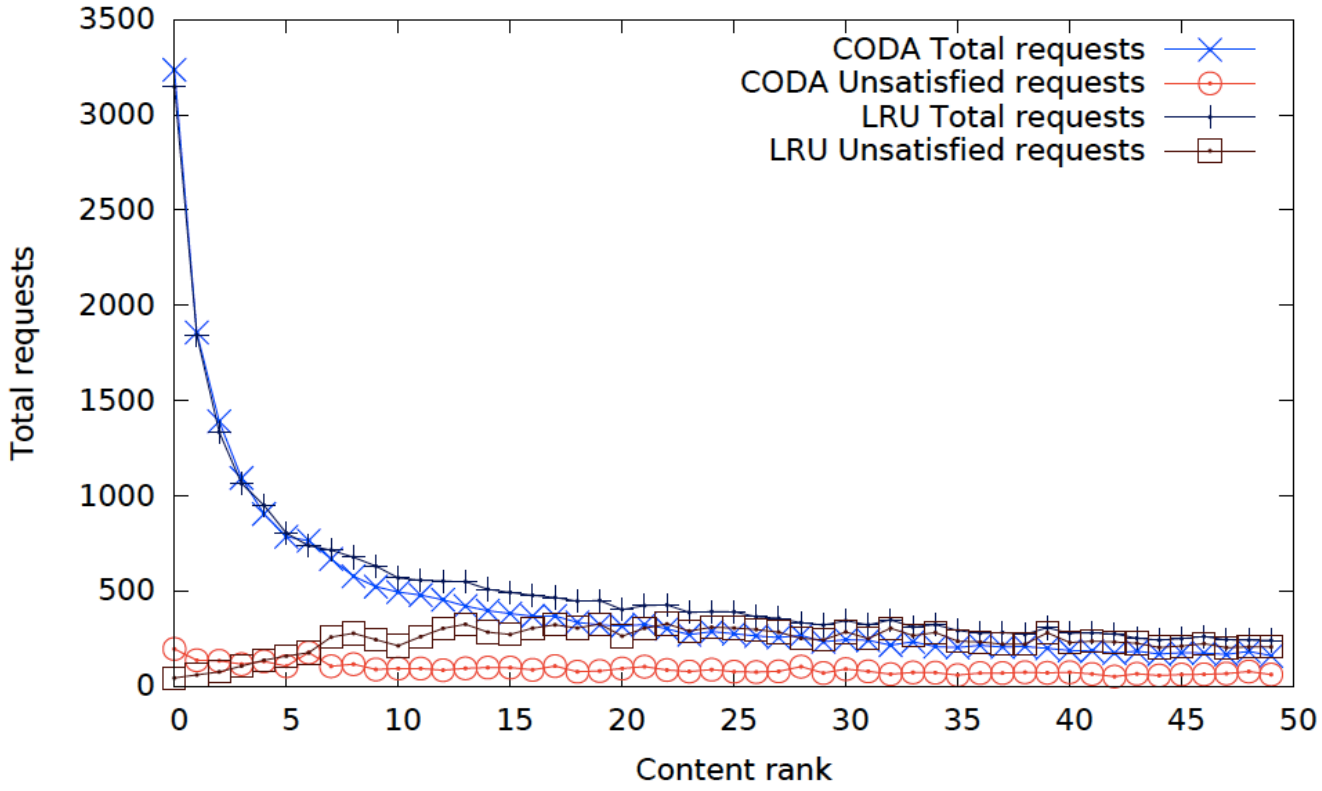

Open issues

Future activity will build upon CODA and transform it into a general solution for storage management and scheduling in mobile content-centric networks

- Scalability of signaling
- Fairness vs. Throughput
- Convergence vs. Reactivity
- Collaboration enforcement
- Heterogeneity of devices (battery, storage, etc)
- Large scale experiments

merci

www.inria.fr