Distributing Content Updates Over a Mobile Social Network

Stratis Ioannidis

Augustin Chaintreau

Laurent Massoulié

Opportunistic Communication

- Bluetooth and/or WiFi enabled devices
- Communication opportunities whenever people meet

Application: Distributing Content Updates

- A wireless service provider distributes content that changes dynamically.

Application: Distributing Content Updates

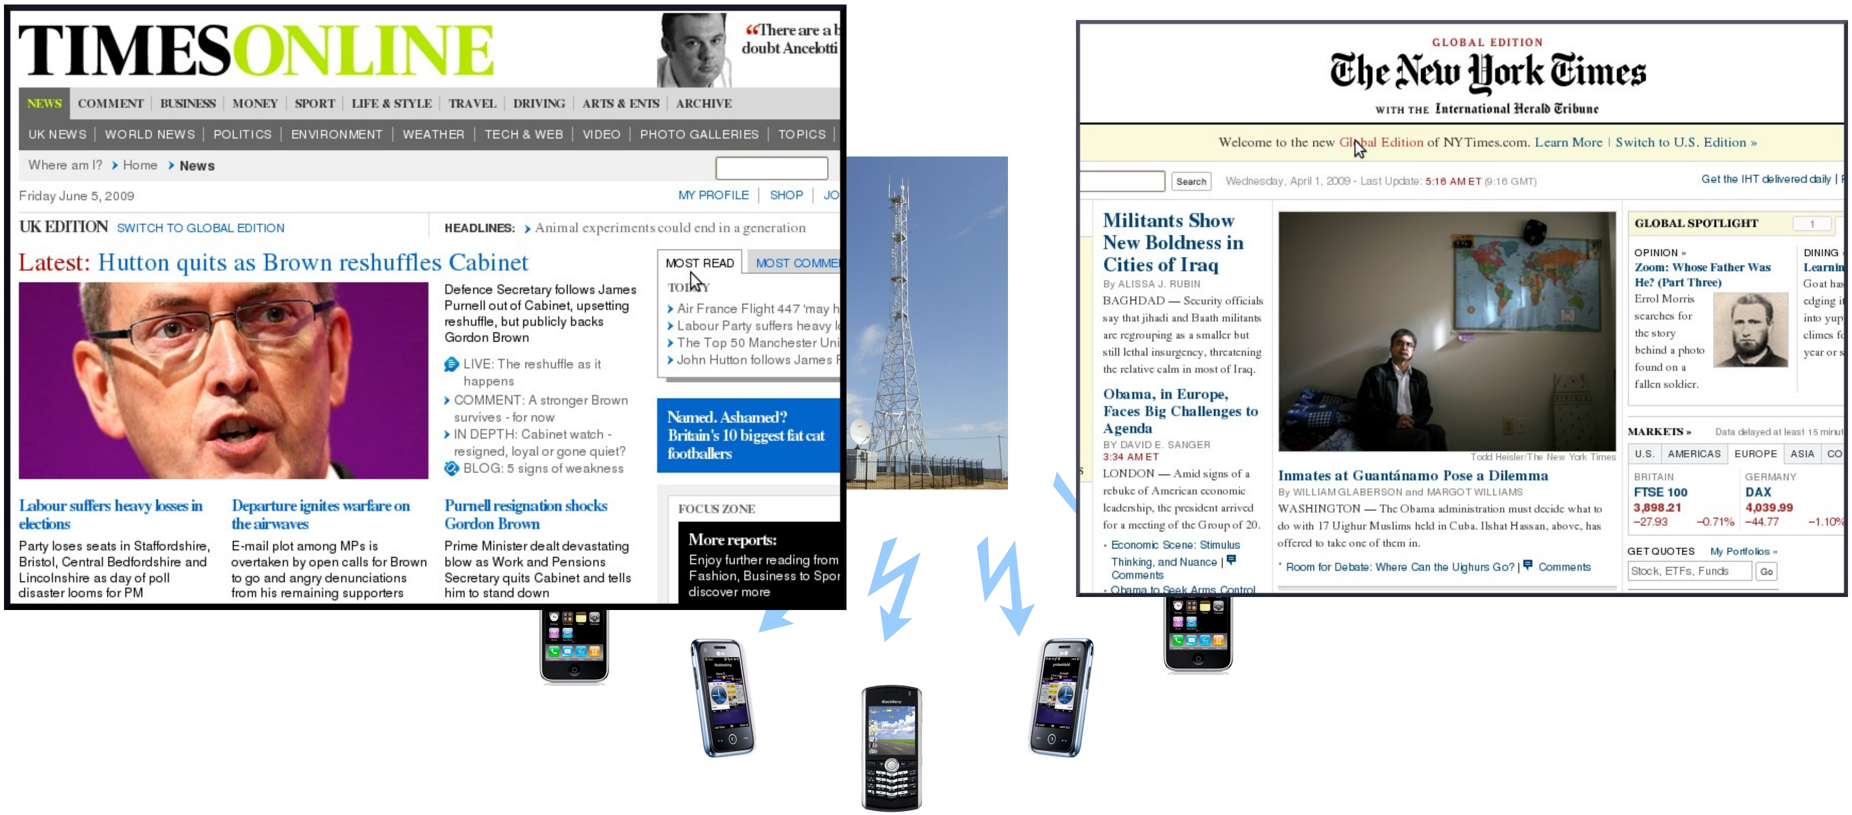

- News

Application: Distributing Content Updates

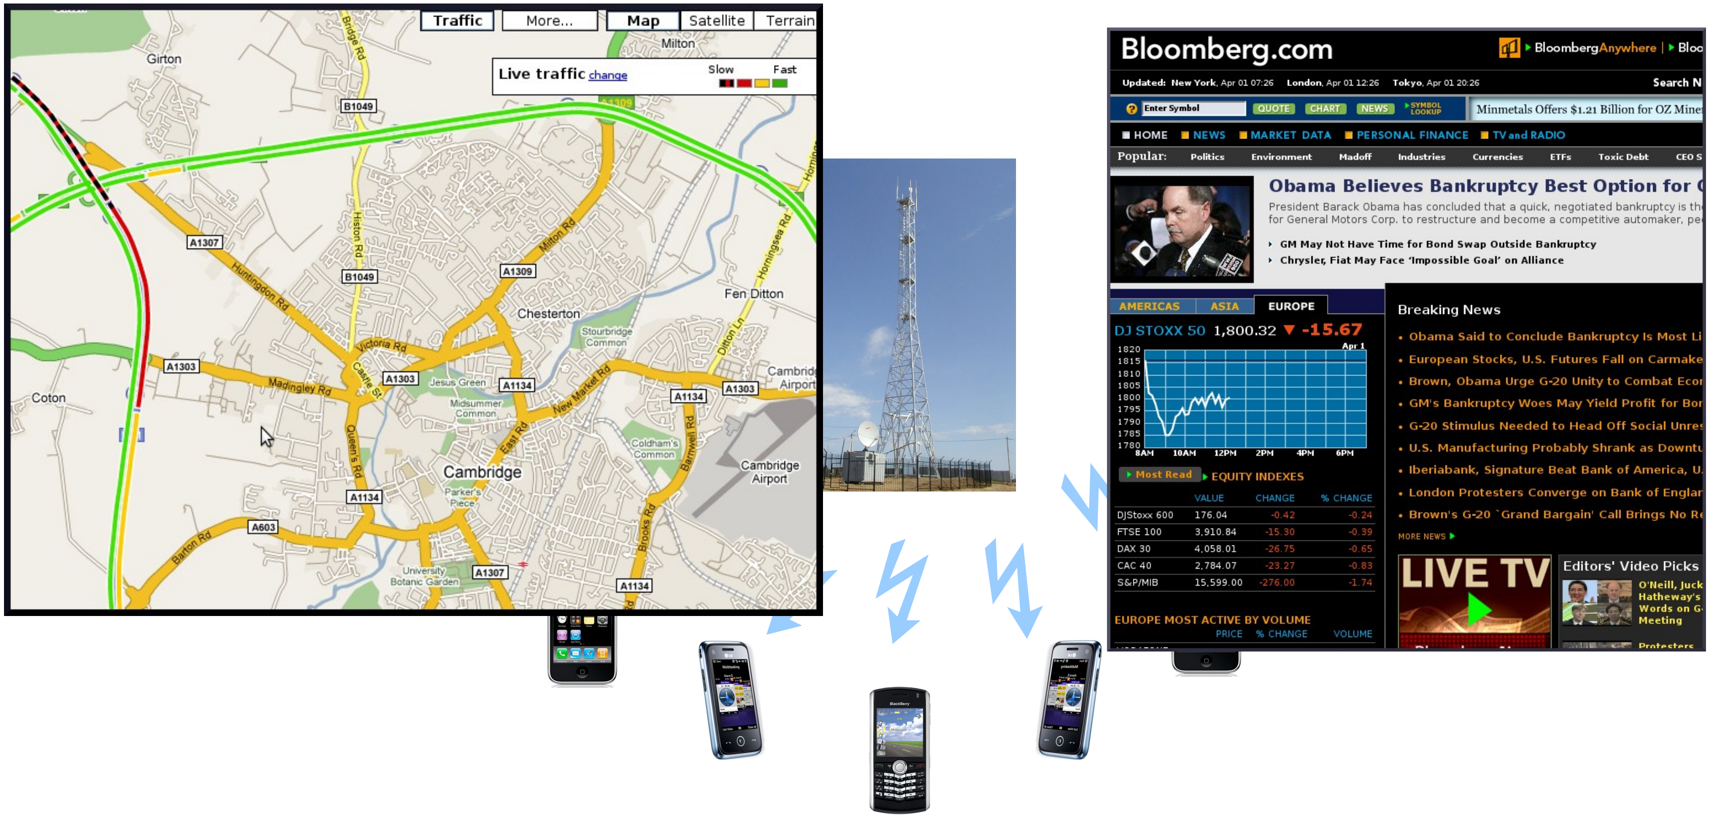

- News
- Traffic information, stock prices

Application: Distributing Content Updates

- News
- Traffic information, stock prices
- Blogs, *etc.*

Users Can Share Updates

- User with most recent content pushes it to the one whose content is outdated

Questions

Goal:

- Make content at users as “fresh” as possible

Questions

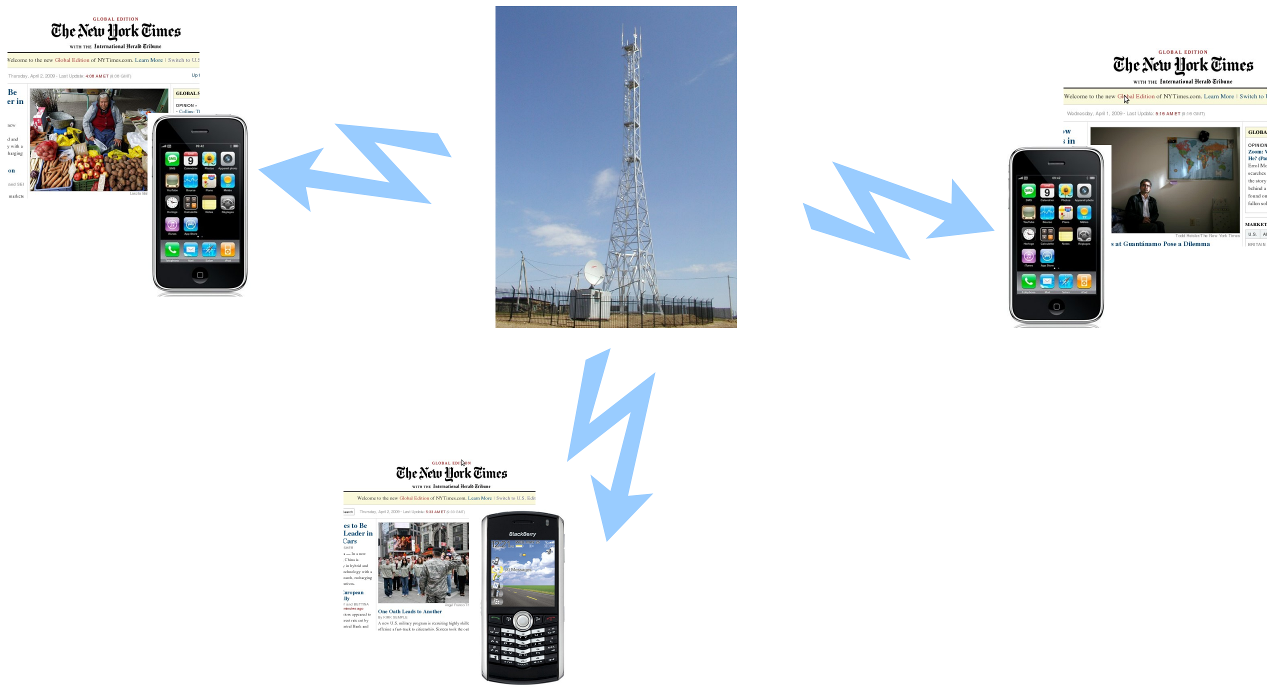

Optimality: Allocation of downlink rate

- *E.g.*, push updates to “social” users more frequently?

Questions

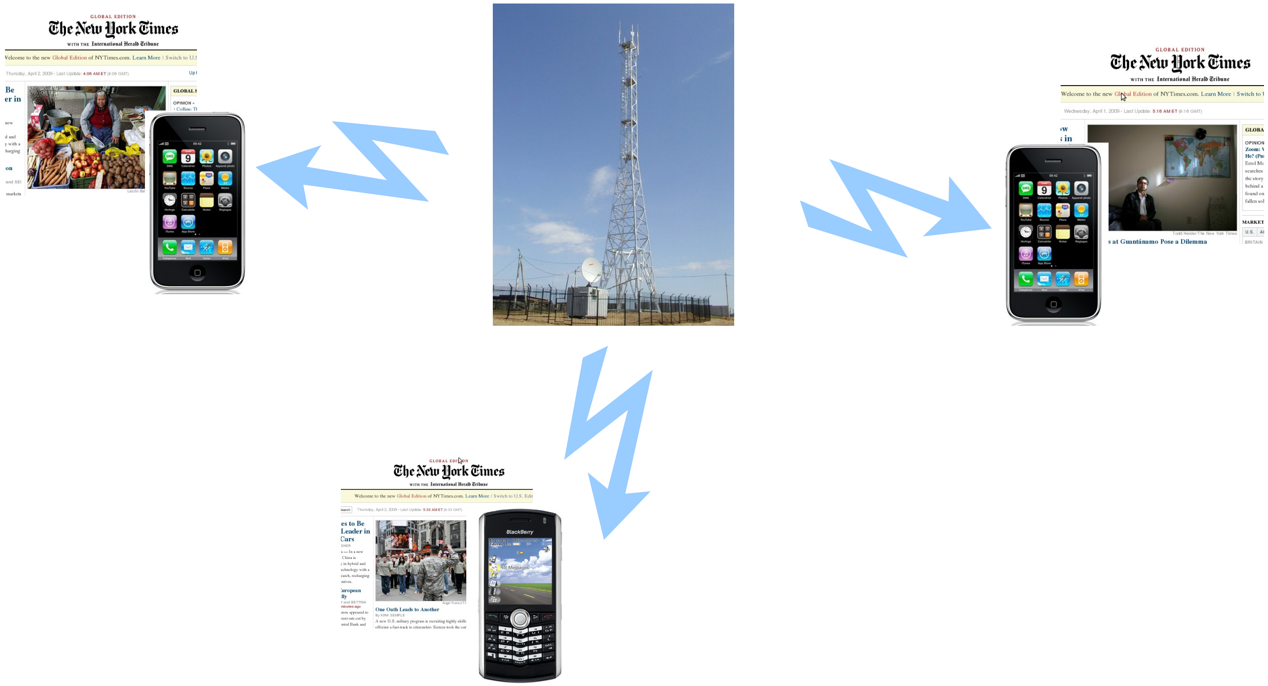

Scalability: How does age grow with user population n?

- *E.g.*, $O(n)$? $O(\log n)$?

Our Results

Optimality

- The service provider can compute an optimal allocation of its downlink rate.

Scalability:

- The age of content at users will grow slowly (as $\log n$) if the social network is an *expander*.

THOMSON TECHNICOLOR

Model: Update Distribution Process

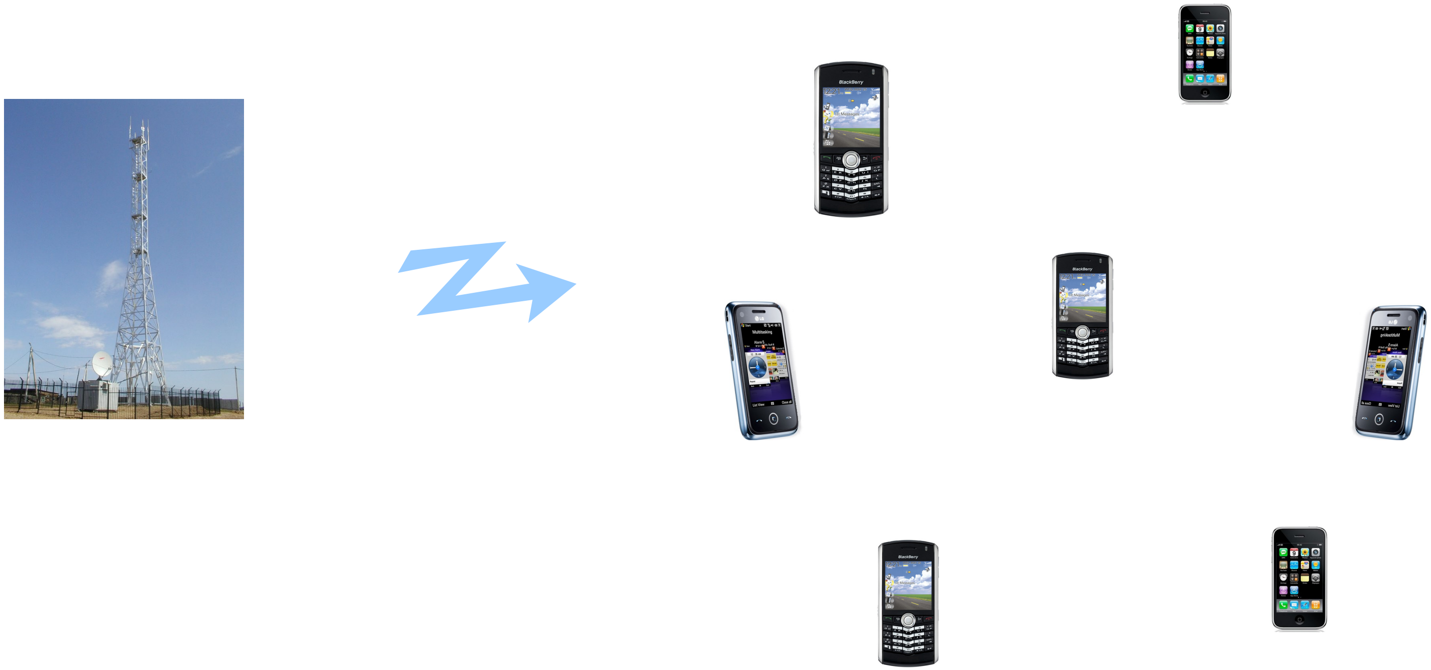

n users

Model: Update Distribution Process

Each user i receives updates according to Poisson process with rate x_i, where

$$\sum_{i=1}^{n} x_i \leq \mu$$

Model: Contact Process

time

E.g.:

- Independent inter-contact times (but not necessarily identically distributed)
- Random trip mobility [Le Boudec and Vojnovic '06]
- Group mobility [Hong et al.'99, Musolesi and Mascolo '07]
- Markov chain mobility [*e.g.*, Nicholson and Noble '08]

Model: Contact Process

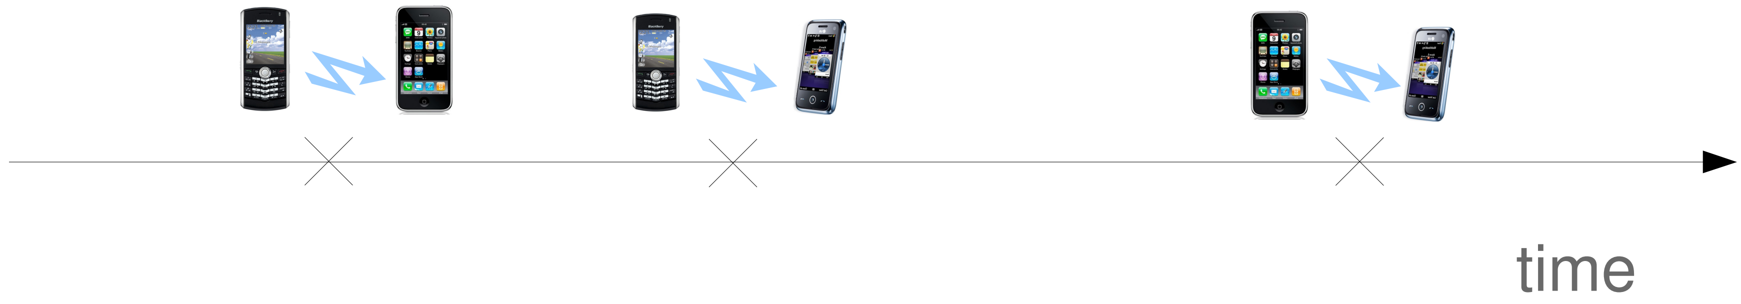

E.g.:

- Independent inter-contact times (but not necessarily identically distributed)
- Random trip mobility [Le Boudec and Vojnovic '06]
- Group mobility [Hong et al.'99, Musolesi and Mascolo '07]
- Markov chain mobility [*e.g.*, Nicholson and Noble '08]

General model: joint contact process is *stationary ergodic.*

Model: User Utilities

$u_i(Y)$: Utility of user i when the age of its content is Y.

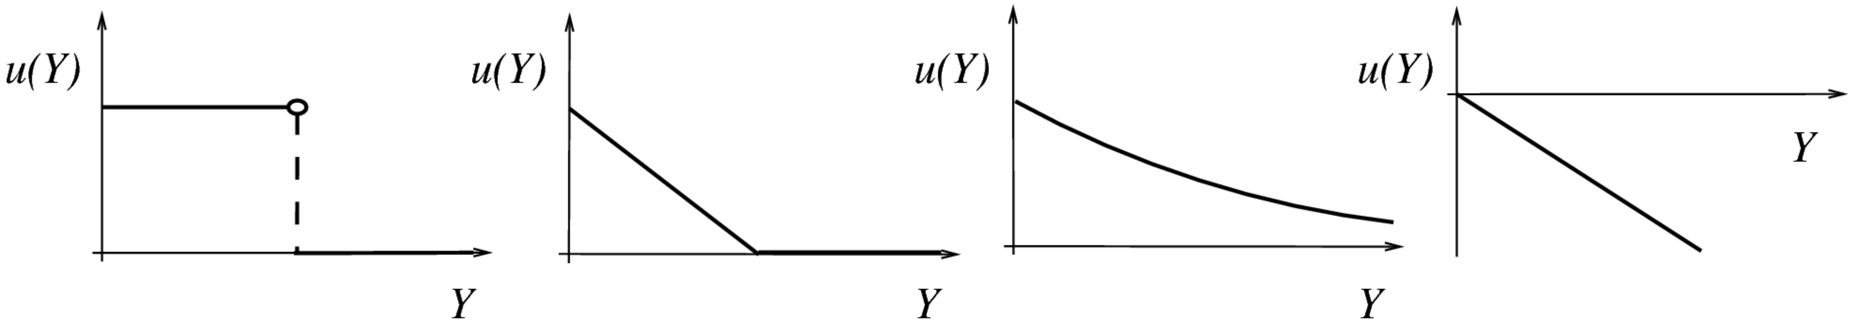

Fresher content is better.

SOCIAL WELFARE MAXIMIZATION

Maximize: $$f(\vec{x}) = \sum_{i=1}^{n} \mathrm{E}_{\vec{x}}[u_i(Y_i)]$$

subject to: $$\sum_{i=1}^{n} x_i \leq \mu,$$

$$x_i \geq 0, 1 \leq i \leq n,$$

where

Y_i : the age of content at i (in steady state),

$u_i(Y_i)$: the utility at i,

x_i : the injection rate at i,

μ : the downlink capacity.

THOMSON TECHNICOLOR

Main Result

Theorem: SOCIAL WELFARE MAXIMIZATION is a convex optimization problem.

Main Result

Theorem: SOCIAL WELFARE MAXIMIZATION is a convex optimization problem.

Hence, it can be solved by gradient descent.

Main Result

Theorem: SOCIAL WELFARE MAXIMIZATION is a convex optimization problem.

Hence, it can be solved by gradient descent.

Algorithms for estimating the gradient

- **Centralized**: The service provider knows all user utilities and collects user contact logs
- **Distributed**: The gradient is estimated by the users in a decentralized manner

Decentralized Estimation of the Gradient

- Objective function: $$f(\vec{x}) = \sum_{i=1}^{n} \mathrm{E}_{\vec{x}}[u_i(Y_i)]$$

Decentralized Estimation of the Gradient

- Objective function: $$f(\vec{x}) = \sum_{i=1}^{n} \mathrm{E}_{\vec{x}}[u_i(Y_i)]$$

- Each user i estimates $\dfrac{\partial \mathrm{E}_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$, for all j, in a distributed manner.

THOMSON TECHNICOLOR

Decentralized Estimation of the Gradient

- Objective function: $$f(\vec{x}) = \sum_{i=1}^{n} \mathrm{E}_{\vec{x}}[u_i(Y_i)]$$

- Each user i estimates $\dfrac{\partial \mathrm{E}_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$, for all j, in a distributed manner.

- These estimates are reported to the service provider, who then computes

$$\frac{\partial f(\vec{x})}{\partial x_j} = \sum_{i=1}^{n} \frac{\partial \mathrm{E}_{\vec{x}}[u_i(Y_i)]}{\partial x_j}, \qquad \text{for all } j.$$

An Unbiased Estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

User j:

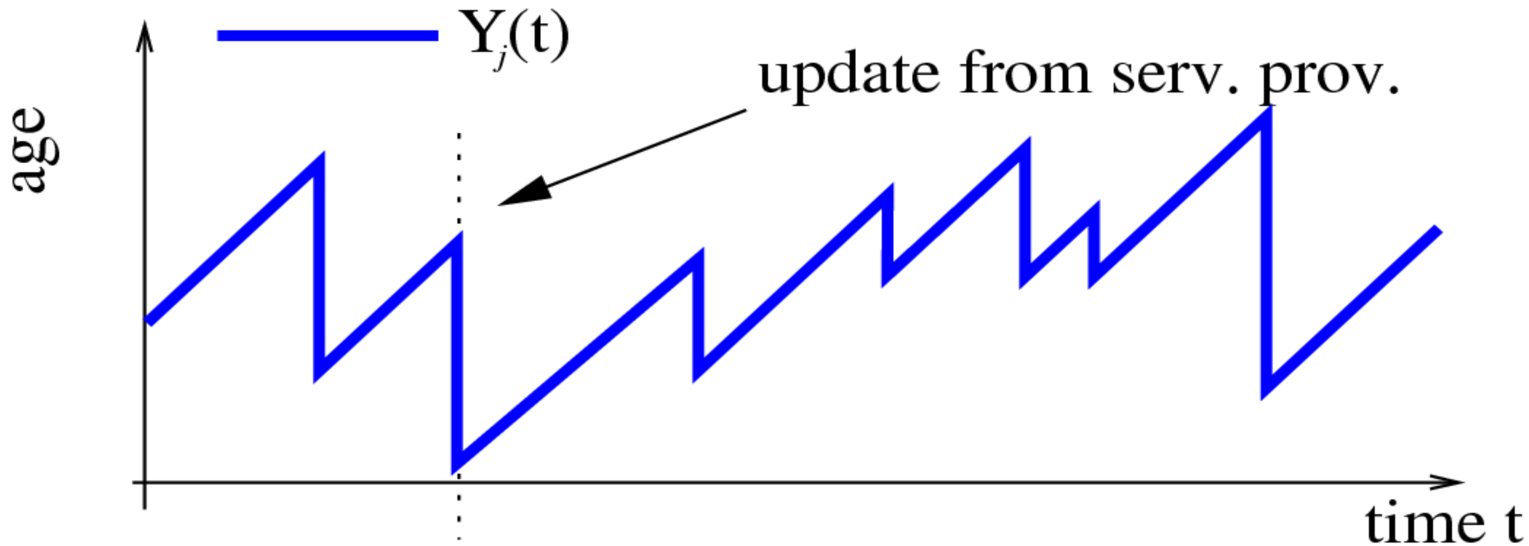

The age $Y_j(t)$ increases between encounters, and becomes zero when a new update is received from the service provider.

An Unbiased Estimator of $\frac{\partial \mathrm{E}_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

User j:

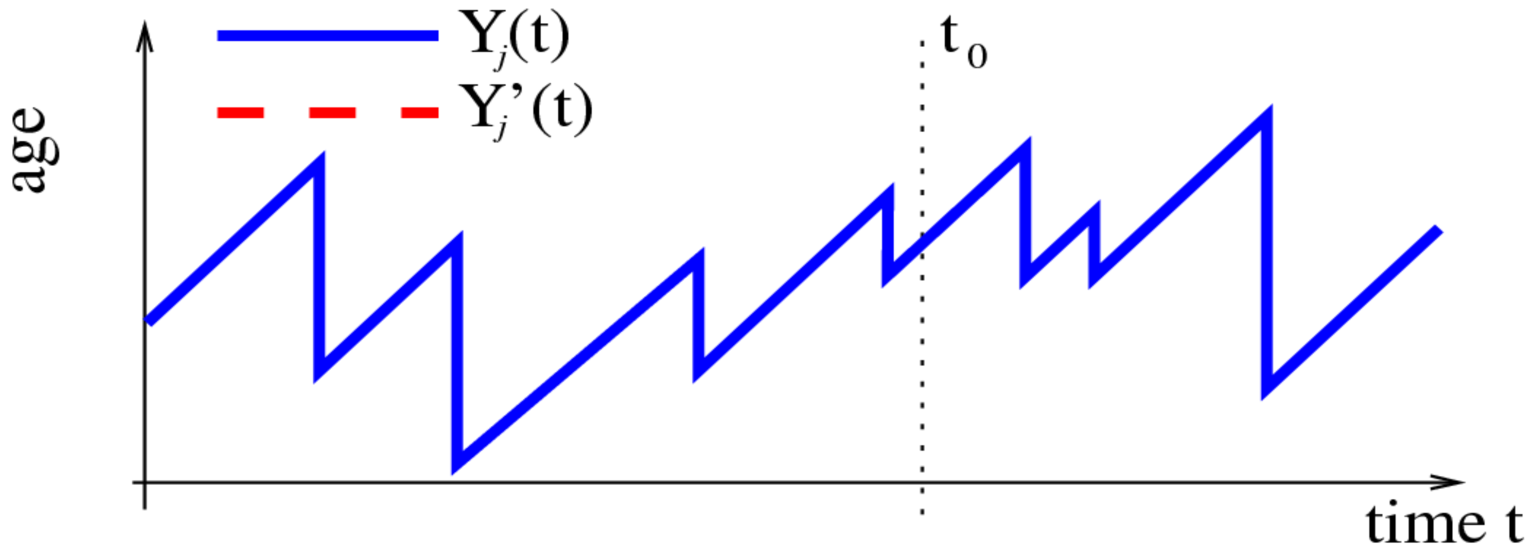

At some time t_0, user j pretends it receives an update from the service provider.

An Unbiased Estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

User j: 

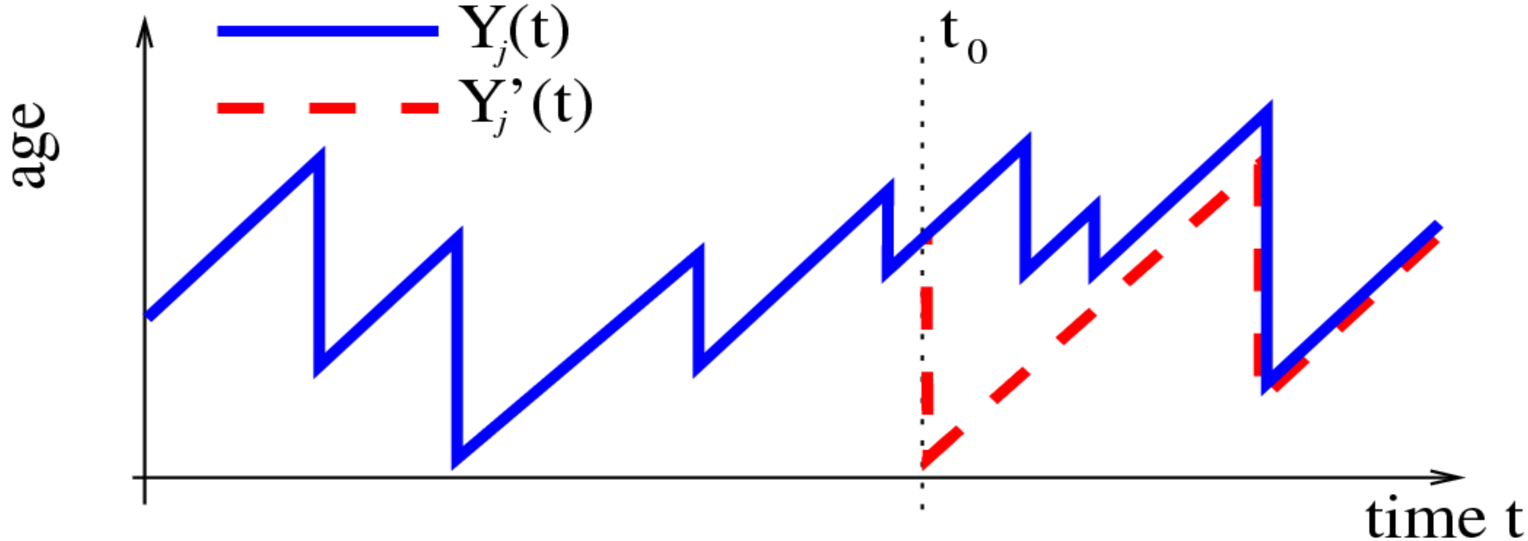

From that point on, j maintains a "**dummy**" age process $Y_j'(t)$, along with its actual age process $Y_j(t)$.

An Unbiased Estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

User j:

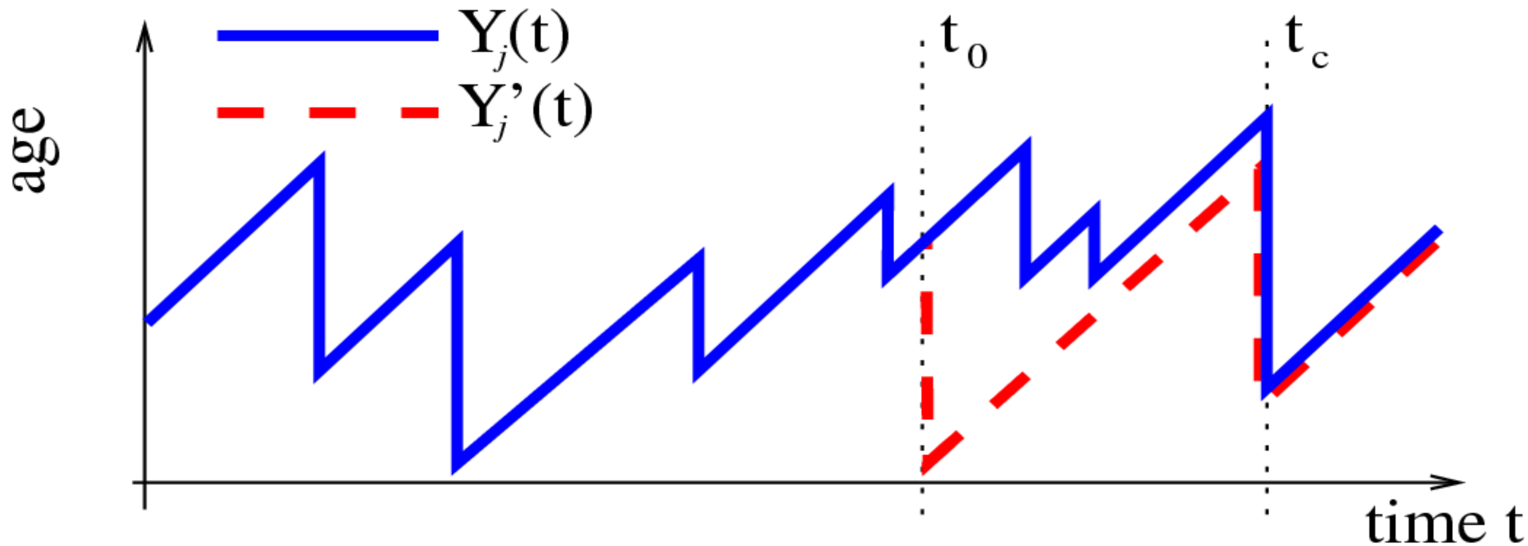

Eventually, $Y_j'(t)$ and $Y_j(t)$ coincide.

An Unbiased Estimator of $\frac{\partial \mathrm{E}_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

User j:

User i:

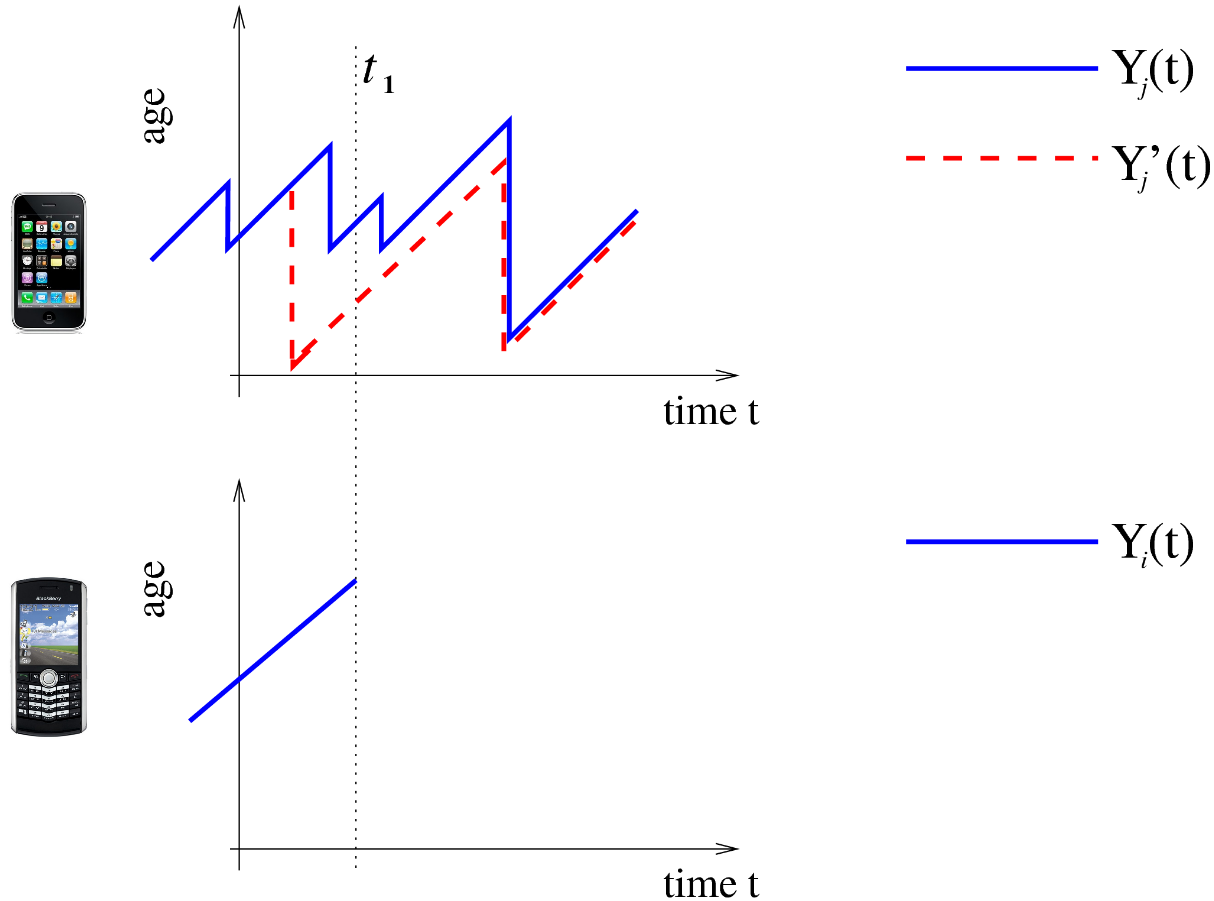

User i meets j at time t_1, while j is maintaining a “dummy” age process.

An Unbiased Estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

User j: 

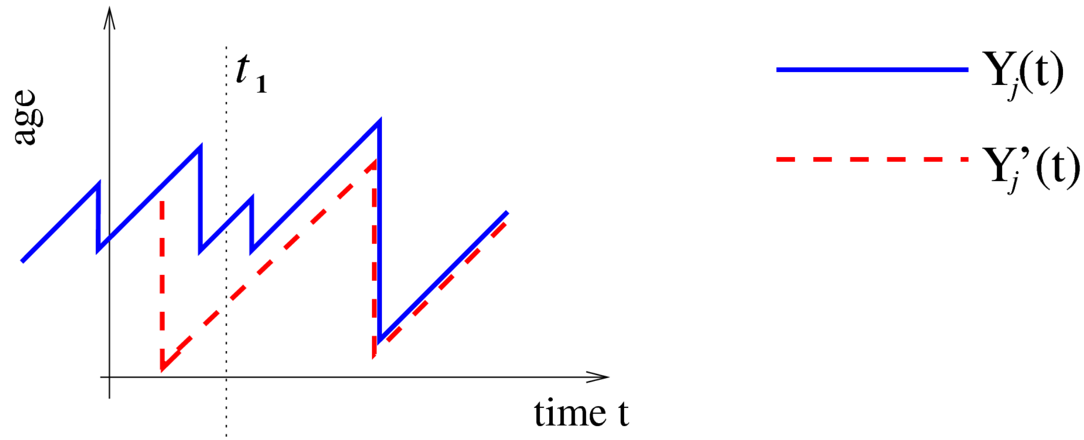

User i: 

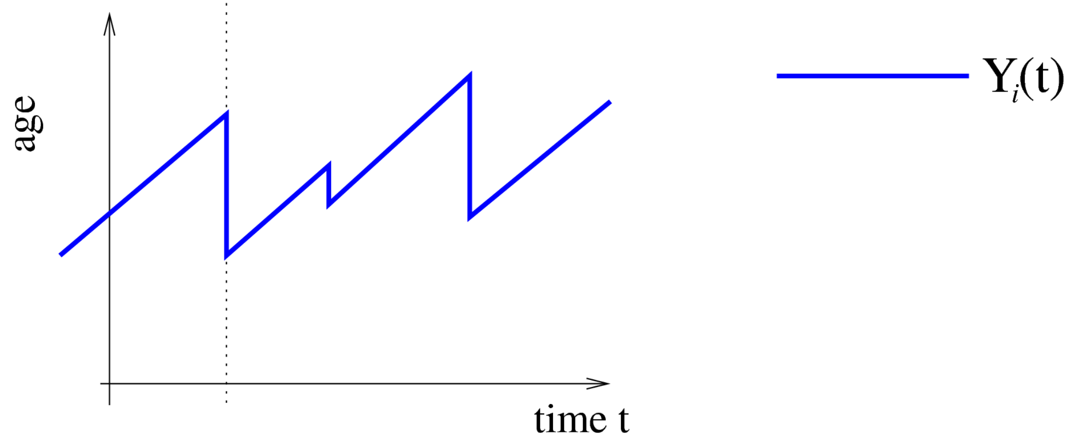

User i receives content from j, changing its actual age $Y_i(t)$.

An Unbiased Estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

User *j*:

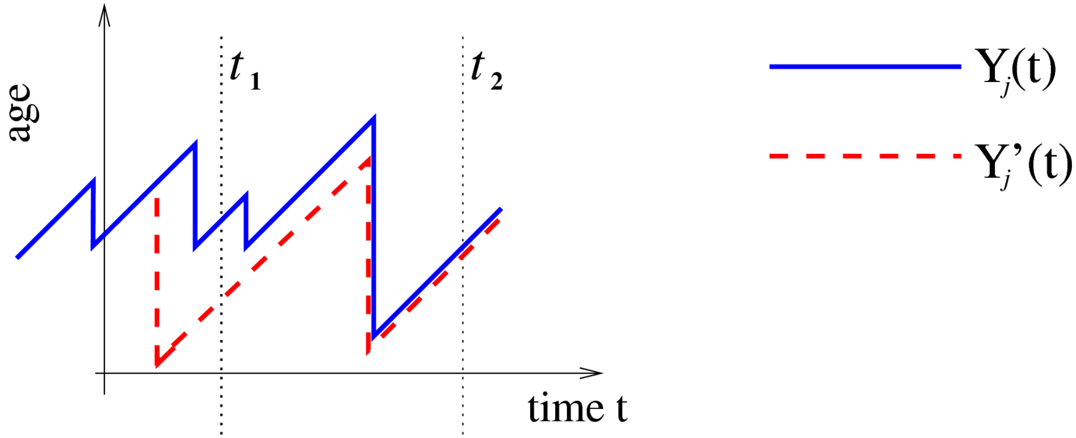

User *i*:

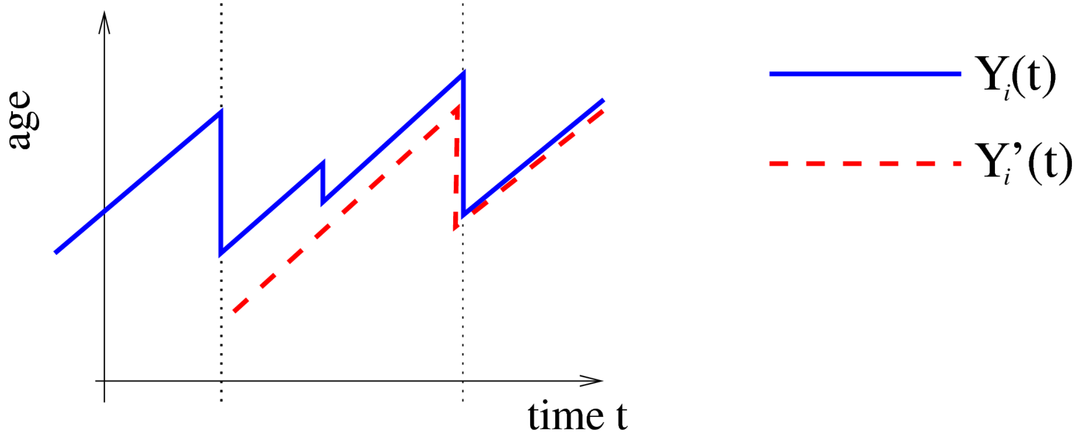

User i also starts a "dummy" age process $Y_i'(t)$. Again, $Y_i'(t)$ and $Y_i(t)$ will eventually coincide.

THOMSON TECHNICOLOR

An Unbiased Estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

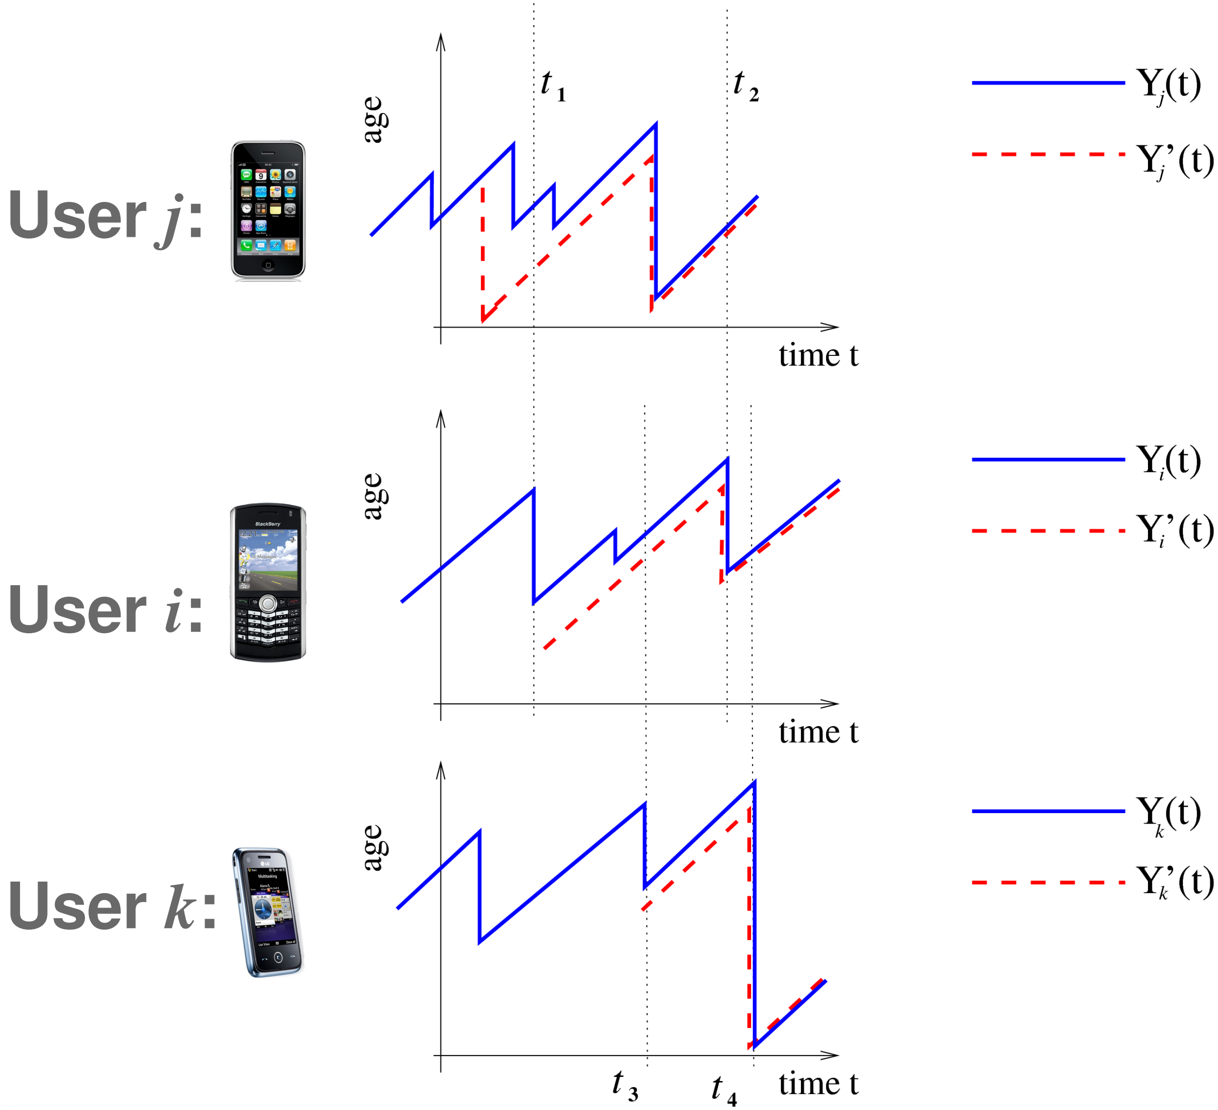

If i meets a third user k, while maintaining a "dummy" age process, k also creates a dummy age process $Y'_i(t)$.

THOMSON TECHNICOLOR

An Unbiased Estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

User j:

User i:

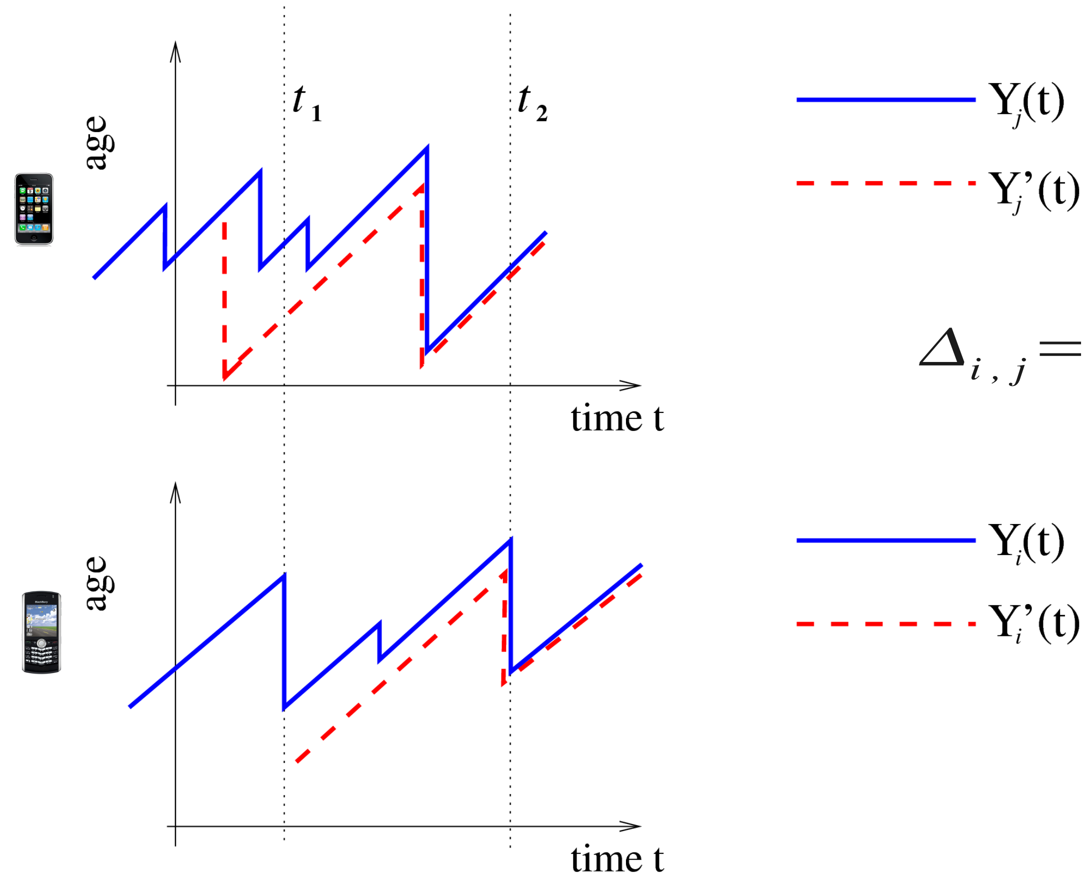

$$\Delta_{i,j} = \int_{t_1}^{t_2} u_i(Y_i(s)) - u_i(Y_i'(s))\, ds$$

[Bremaud, Vasquez-Abad 92]:

The quantity $\Delta_{i,j}$ is an unbiased estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$.

An Unbiased Estimator of $\frac{\partial E_{\vec{x}}[u_i(Y_i)]}{\partial x_j}$

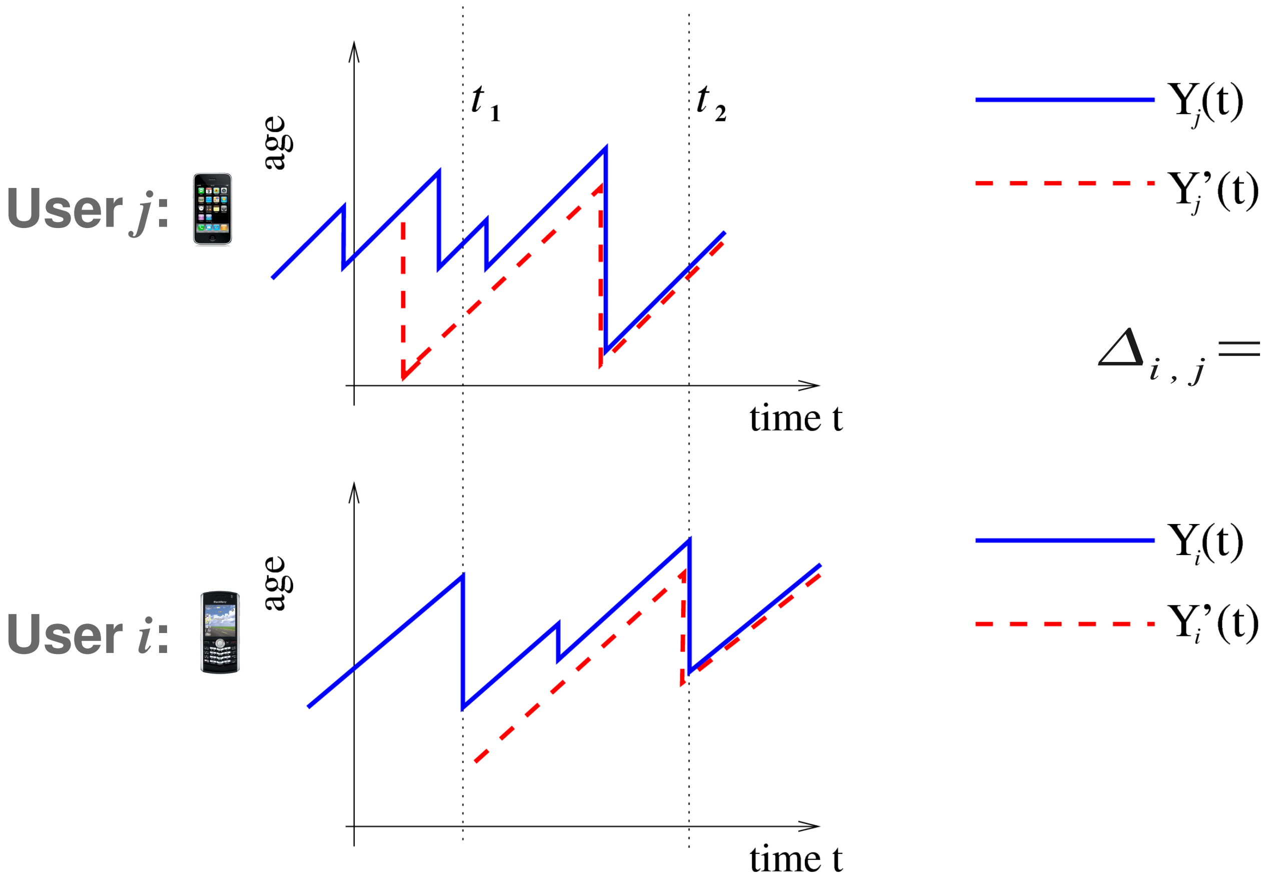

$$\Delta_{i,j} = \int_{t_1}^{t_2} u_i(Y_i(s)) - u_i(Y_i'(s))\, ds$$

$\Delta_{i,j}$ is similar to the *sample path shadow price* of a packet in a congested network [Gibbens and Kelly, 99].

Empirical Study

Infocom 06: 10 hours, 98 users

MIT: 80 days, 95 users

User utilities:

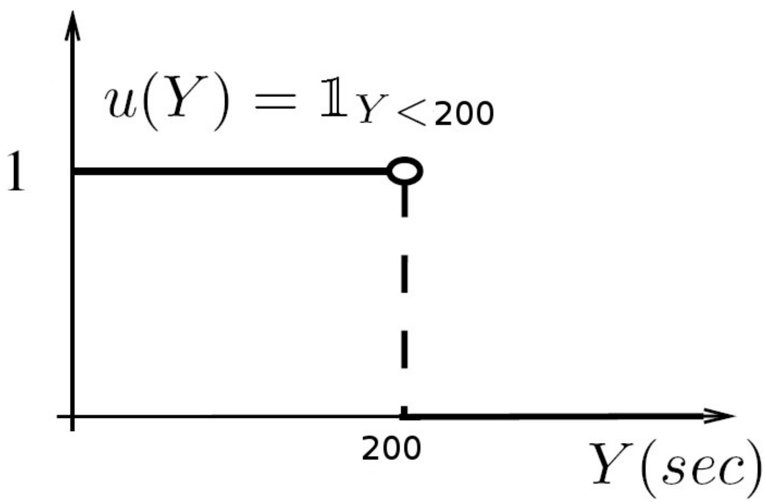

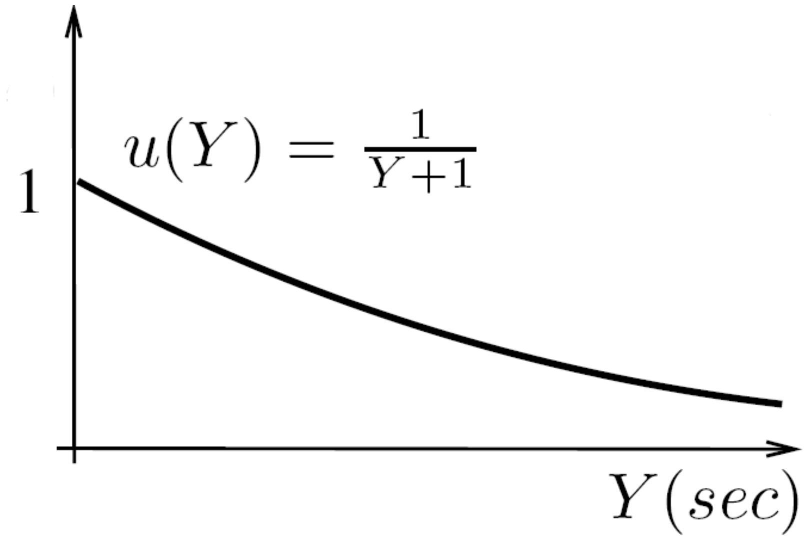

Optimal Allocations (Infocom '06)

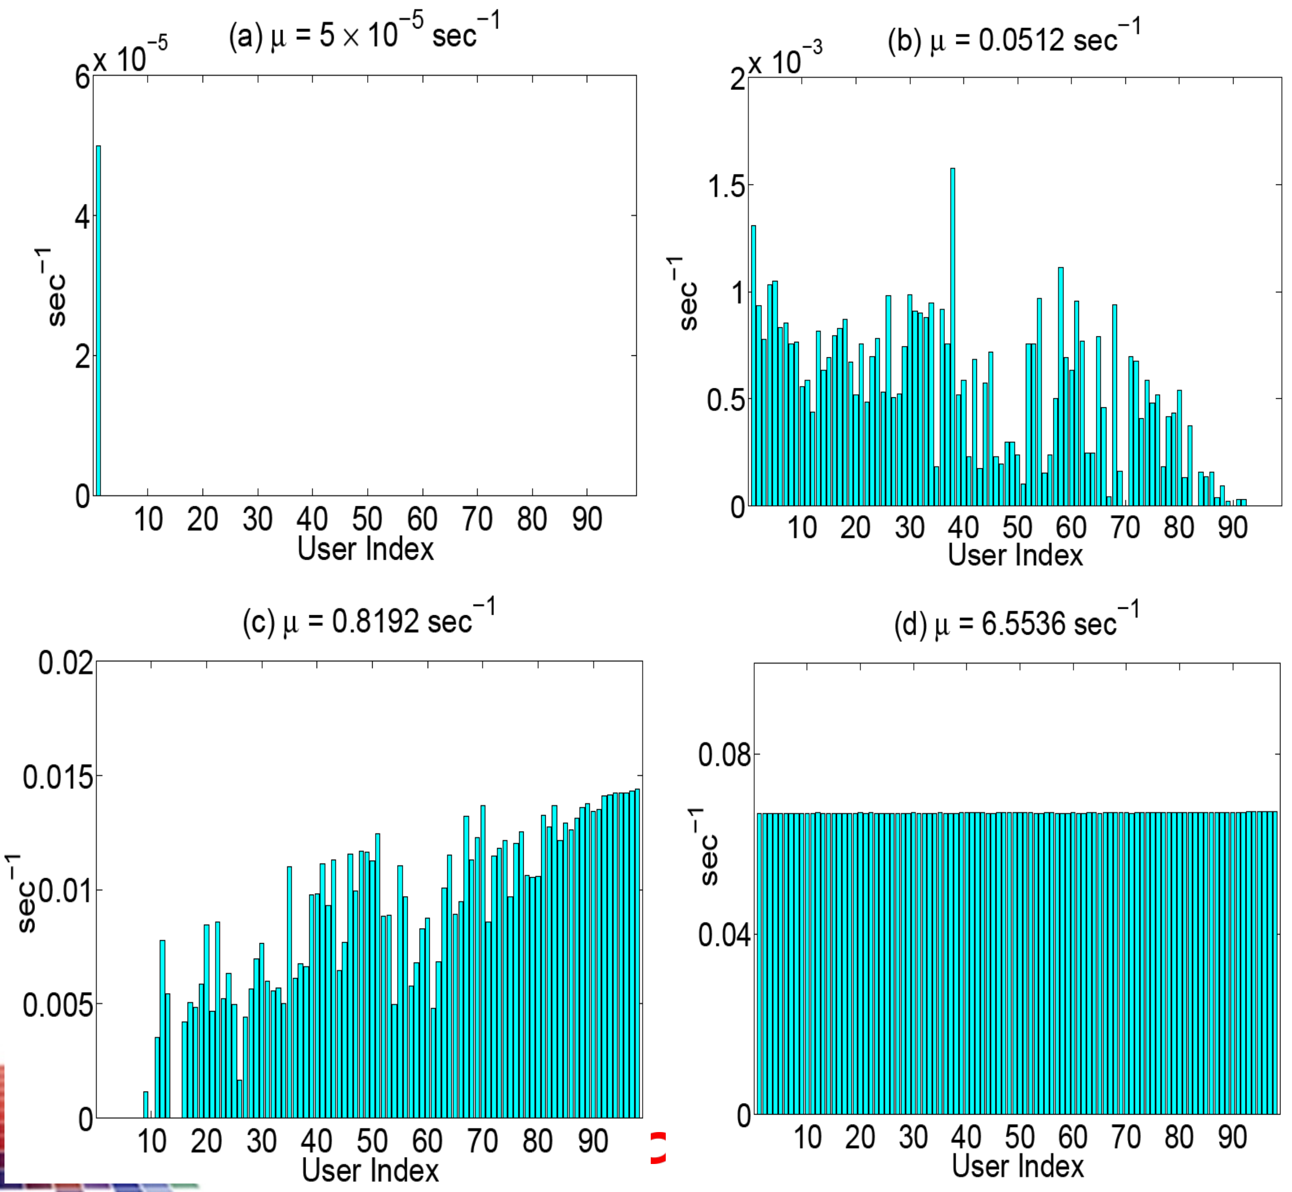

Optimal rate allocations for step-function utilities.

Users are indexed from most to least social.

Comparison to Heuristic Allocations

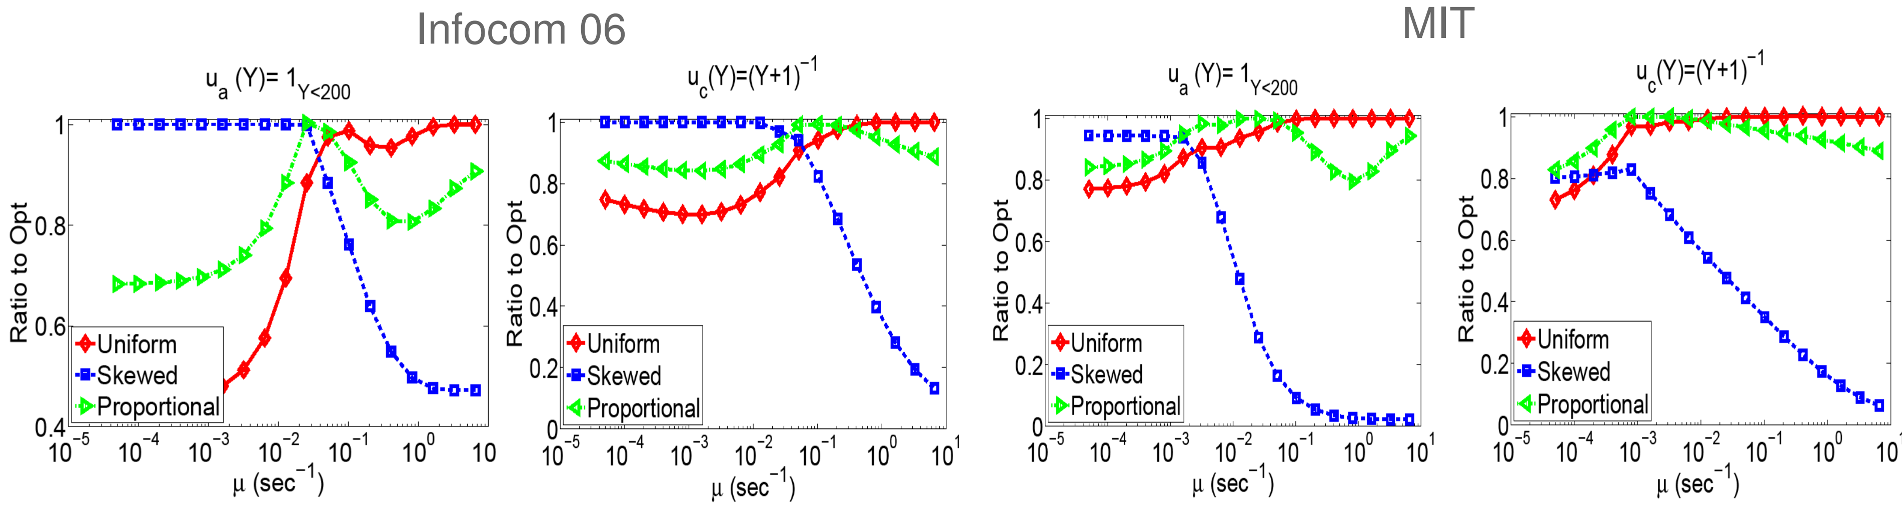

Ratio of social welfare of heuristic rate allocations to the optimal.

μ (sec^{-1})	No-sharing	Skewed	Proportional	Uniform	Optimal
0.0128	2.5%	34%	24%	30%	34%
0.0256	5.1%	42%	42%	37%	42%
0.1024	19%	46%	56%	60%	60%

Fraction of users with age below 200sec (Infocom 06).

THOMSON TECHNICOLOR

Conclusions & Future Work

Conclusions

- Service provider can distribute updates in optimal way.
- No need to know utilities and/or contact statistics!
- Optimal allocation can be non-trivial.

Future work

- Incentives to share
- Pricing
- Full decentralization-no service provider

Thank you!

S. Ioannidis, A. Chaintreau, and L. Massoulié: *Optimal and Scalable Distribution of Content Updates in a Mobile Social Network.* In INFOCOM 2009.

S. Ioannidis, and A. Chaintreau: *On the Strength of Weak Ties in Mobile Social Networks*. In SNS 2009.

A. Chaintreau, J.Y-Le Boudec, and N. Ristanovic: *The age of gossip: Spatial mean-field regime*. In SIGMETRICS 2009.

Scalability

Restricted Model:

- Users meet according to independent Poisson processes with rates $q_{ij} \geq 0$.
- Weighted graph $G(V,E)$ where
  - V = the set of users,
  - $E = \{(i,j),\ i,j \in V\}$,
  - The weight of edge (i,j) is $q_{ij} \geq 0$.

Weighted Edge Expansion:

$$h_G = \min_{A \subseteq V,\ |A| \leq |V|/2} \frac{\sum_{i \in A,\ j \in A^c} q_{ij}}{|A|}$$

Main Result

Theorem: If $x_i = \mu/n$, for all $i \in V$, (*i.e.*, the rate allocation is uniform) the expected age seen by any user $i \in V$ in steady state satisfies

$$E_{\vec{x}}[Y_i] \leqslant \frac{2}{\mu}\left(2e^{-1/2} + \log(n)\right) + h_G^{-1}\log(n)$$

where h_G the edge expansion G.

- If G is an expander, the expected age is $O(\log n)$.

- Considerable improvement compared to no-sharing ($O(n)$).