

404 NOT FOUND? A QUEST FOR DTN APPLICATIONS

Jörg Ott
Aalto University Comnet

MobiOpp – 16 March 2012

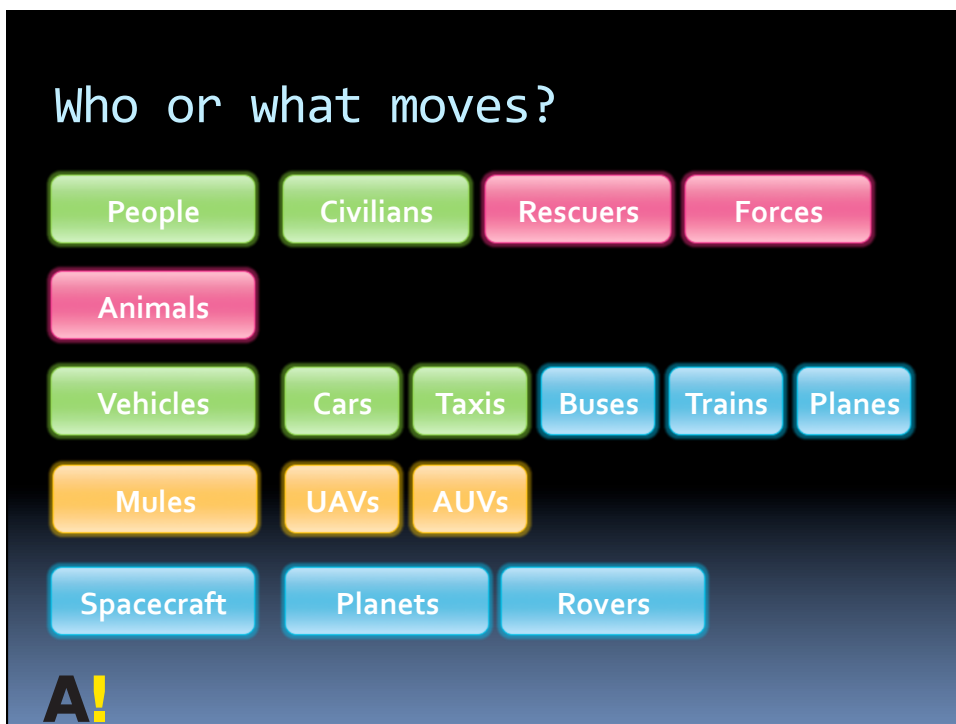



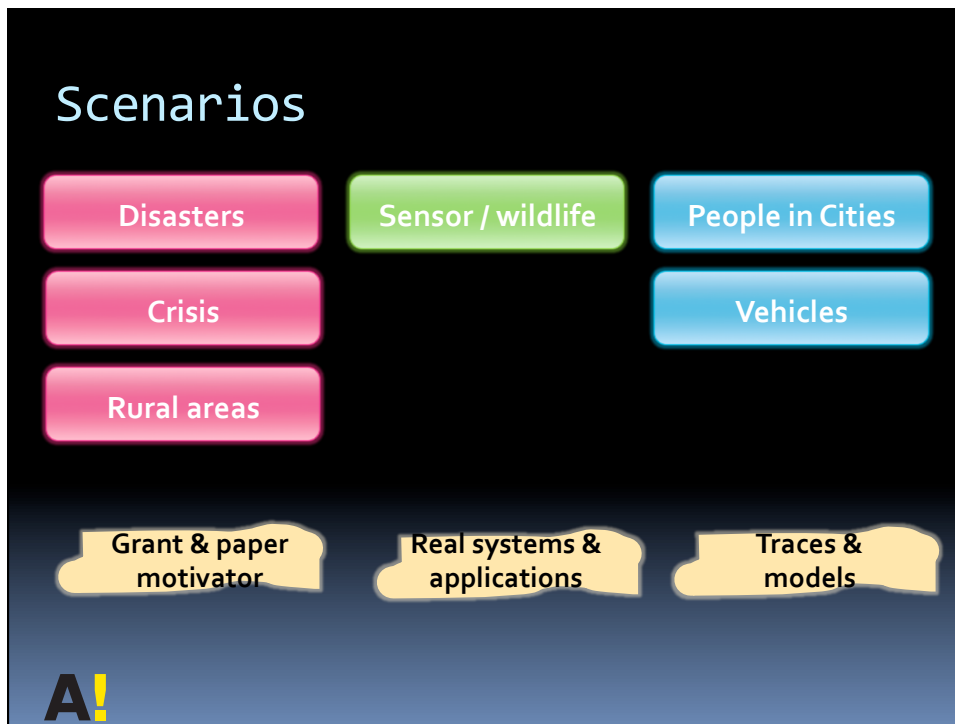
Roadmap

- Evolution of mobile opportunistic communication
- Applications and their characteristics
- Two case studies
- Some conclusions





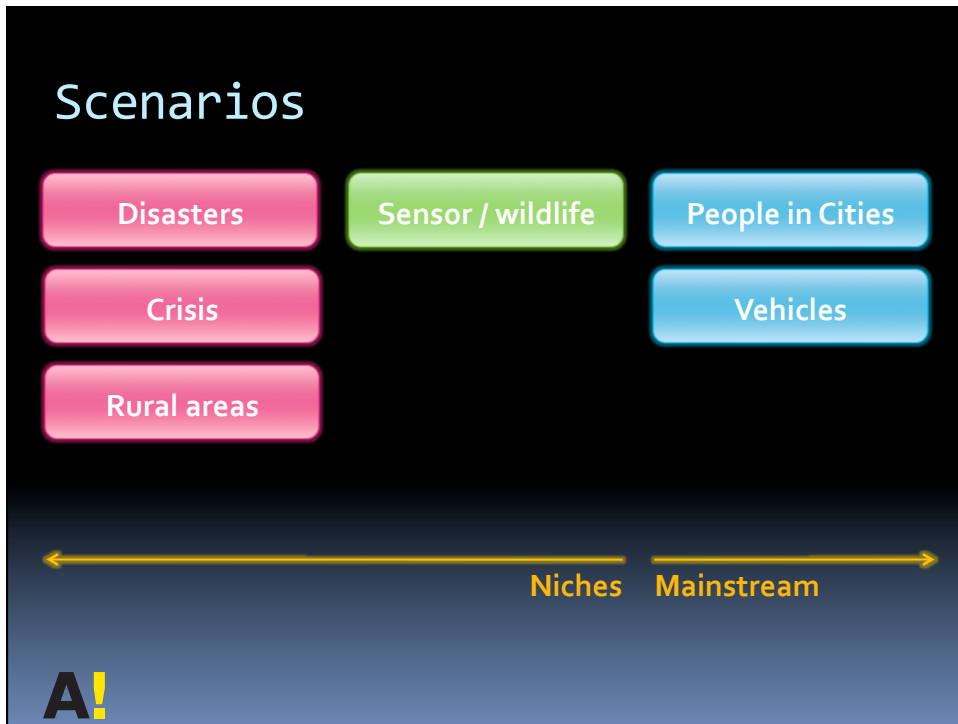




What do we observe?

- We have lots of motivation for some areas
- We have specific applications in another
- We focus our tracing and modeling on a third one
- Is this sensible from an evaluation perspective?
- Do we know what we are doing?
- Need to understand what we design for

A!



TURNING TO APPLICATIONS...

Traditional: We tried them all

- Email & messaging (Jabber)
- SSH
- Web access
- Blogging
- Web search
- Environment search
- Content caching
- Twitter
- Audio & video

Designed

Simulated

Implemented

Tried

A!

Applications

There are quite a few...

Learned a lot about protocol design

Performance isn't competitive

A!



A note on users

- Key: user expectations!
- Governed by the Internet
 - You can get (usually) to everything that exists
 - You can get there right away
- Instant predictable outcome: result or error
 - 404 Not Found
 - Server not found

A!

A note on performance

- Users aren't patient
 - If they want something, they want it now!
- Delivery times > 12h are pointless
 - You'll have WLAN before
- Delay = $f(\text{geo distance})$
 - Don't try to compete with infrastructure
- Low delivery rates?



A!

Two conclusions

Find applications that don't keep the user waiting!

Find those niche(s) where delay is acceptable
(e.g., don't involve a user in the first place)

A!

What does this mean?

- Applications operate invisibly in the background
- Respond immediately upon user attention
- Completeness, congruence & repeatability are not essential
- Provide best-effort value-add
- Reflect this in the UI

A!

Why bother?

- No connectivity (to infrastructure) needed
- Location privacy
- Content “privacy”
- No centralized censorship
- Geographic validity
- Temporal validity
- No user identification



Some suitable applications

- Content sharing: PodNet
- Volatile sharing: SCAMPImusic
- Geographic sharing: Floating Content
- Social networking: D-Book, GridNet, PeopleInHere
- Generic cellular offloading



APP 1: FLOATING CONTENT

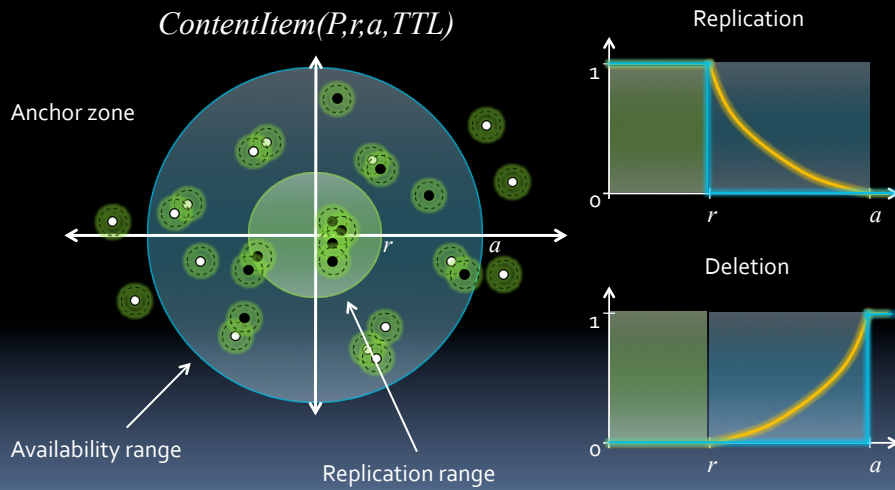
<http://www.floating-content.net/>

Floating Content

- Ad-hoc local social network-style information sharing: Digital graffiti w/o servers and infrastructure
- Leave notes, comments, photos, etc. in places
- Define reach (area of interest) and lifetime
- Prioritize contents for replication: inversely proportional to reach, size, and lifetime
- Inherently best effort


A!

Floating Model



A!

Simple Analytical Model: Black Box

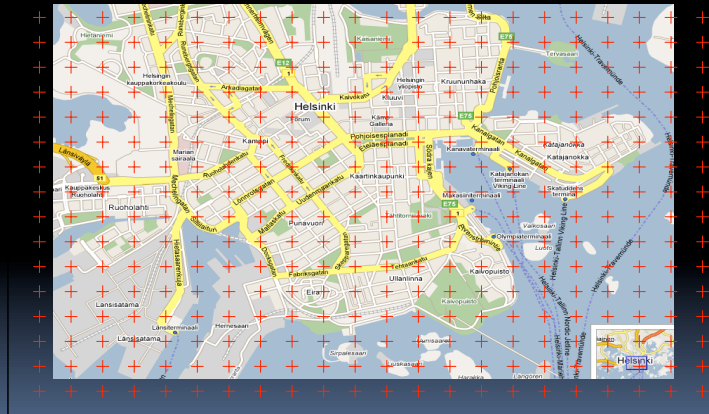


Average number of contacts of each node while in the anchor zone.

A!

Helsinki city evaluation


ONE – 126/252/504 nodes – 10/50m radio – 2 Mbit/s – 0.1-1 MB items – TTL=1h




A!

Helsinki city evaluation

Criticality factor

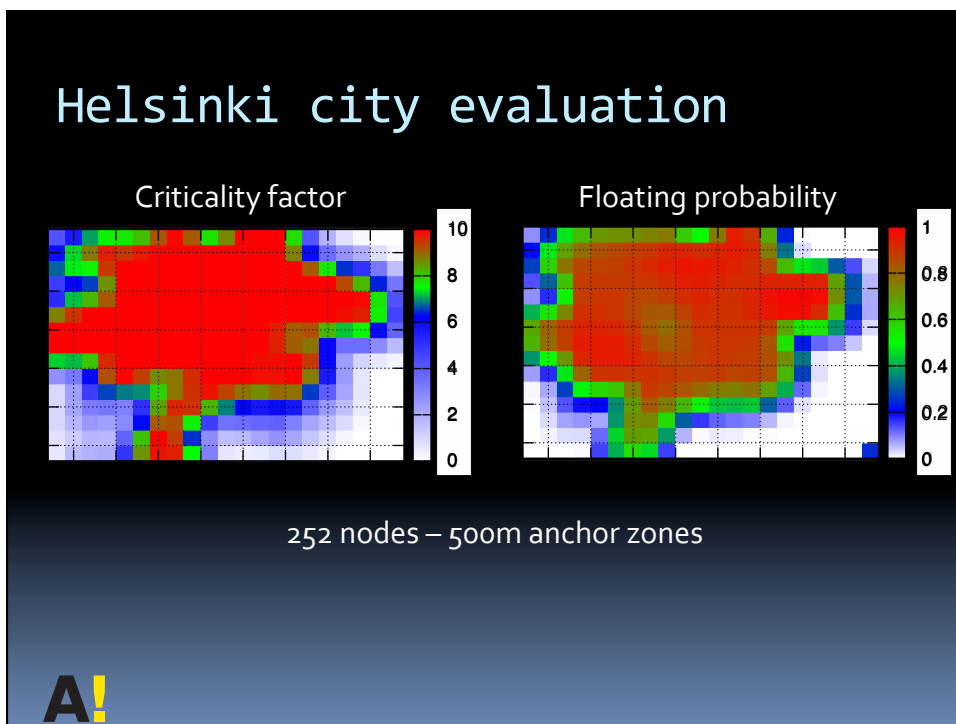
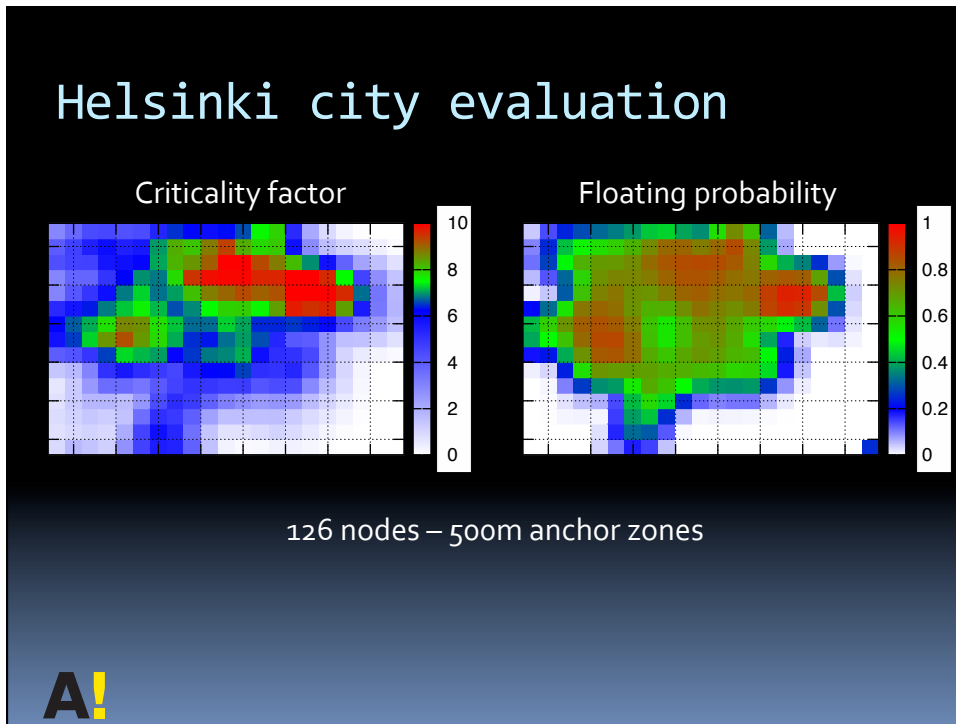


Floating probability




126 nodes – 200m anchor zones

A!




Helsinki city evaluation

Criticality factor



Floating probability




504 nodes – 500m anchor zones




Helsinki city evaluation

Time to acquisition



Node coverage



Content sinks early or stays around long
 A passive buffer zone is sufficient (and simple!)
 Resource consumption-aware prioritization works



Floating API

- ↓ label = **post** ((B, meta), (L, a, r, TTL), <label>)
- ↓ **delete** (label)
- ↓ <(label, meta)> = **select** (meta, interval)
- ↓ ((B, meta), (L, a, r, TTL), <label>) = **get** (label)
- ↓ **watch** (meta, duration, wake-interval)
- ↑ **notify** (<(label, meta)>)

A!

Flea market application

- Simple auction variant using the Floating API
- A stationary seller posts offers for goods
 - Floating 500m around his position
- Pedestrians may place bids
 - Upon receiving the offer or another bid
- Effectively extends the reach of the seller
 - 50-75% of the bids make it to the seller
 - 85-95% of the highest bid value (assuming unit incr.)

A!

Nice properties

- Simple best effort geo cooperation model
- Workable already for modestly dense scenarios
- Independent of larger-scale mobility models
- “Routing” is trivial and scalable
- Built-in DoS protection and garbage collection
- Simple API to support a range of applications

A!




Mining


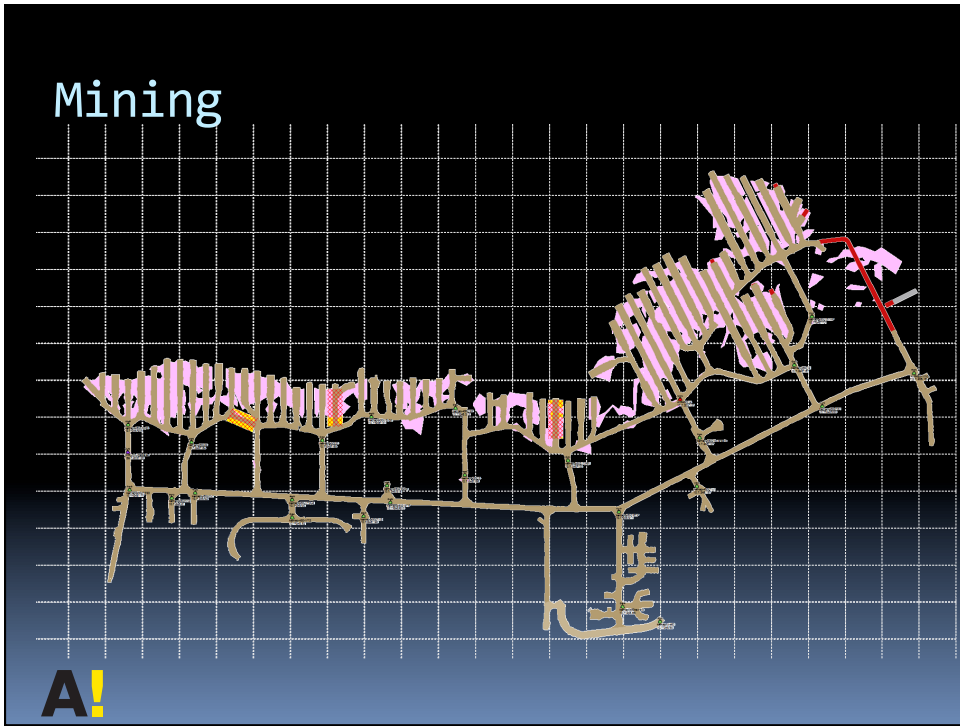
- 3000-4000 operating mines worldwide
 - Lifetime of a few years to tens of years
 - Ten to thousands of pieces of equipment per mine
 - Personnel operating in two to three shifts
- Development and production phases
 - Work cycles with specialized equipment
 - Tens of locations in various phases and stages
- Coordinating the fleet of equipment and personnel requires a robust communication in a challenging environment.

A!

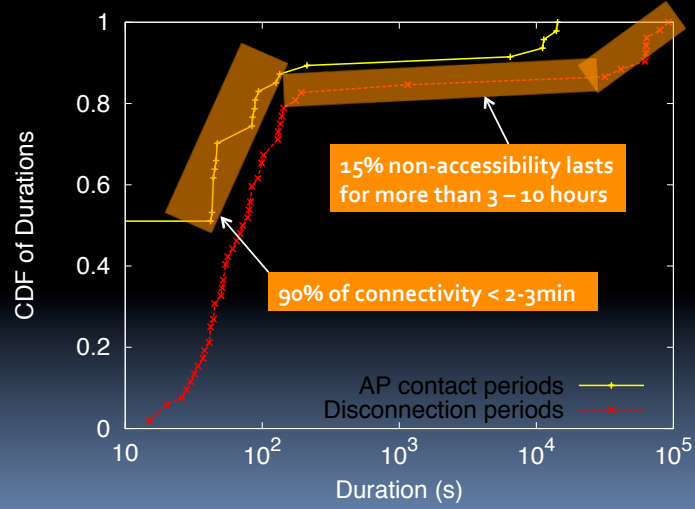
Mining



A!




Node mobility in mines



A!

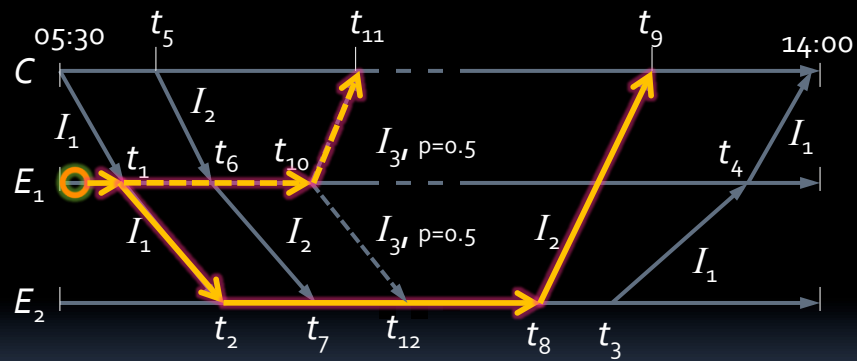
Node mobility example



I: Intermediary node, E: Equipment node, C: Control room, t: time

A!

Communication paths



A!

System properties

- Even a few hours delay is an improvement
- Limited degrees of freedom in mobility
- Routing is simple: Epidemic
- Cumulative data transmission + e2e ACKs
 - ACKs also serve as anti packets

A!

WRAPPING UP

Mainstream vs. Niche

- Generic devices vs. specific ones
- More heterogeneity
- Implementation obstacles by design
- More security and privacy concerns
- Potential legal caveats
- Incentives or force needed?
- User studies and large scale experiments needed

A!

Evaluation challenge: traffic

- There is lots of data on mobile user behavior...
 - ...when accessing the fixed Internet
 - ...when interacting via infrastructure
- Mobile opportunistic communication is clean slate
 - We have no real idea how they are going to be used
 - Nor which applications will emerge
 - How to make sensible assumptions about traffic load?
- Interaction of mobility and traffic generation

A!

Concluding thoughts

- New environments → new ways of thinking
 - Doing ok on some niches
 - Gradually becoming clueful on the mainstream
- Stuff doesn't need to be complex
- Understand your applications
 - Scenarios, traffic matrices, dependencies

A!

Perspectives

- We have enough generic routing protocols
- We don't have enough data yet
- Our understanding is not quite balanced:
scenarios, node densities, modeling, implement.
- Lack of user and usability studies
- Eat (more of) our own dog food!

A!