What's after the Internet?

Christophe Diot

christophe.diot@thomson.net

Communication today

- **Multiple communication technologies co-exist**
- **Wireless is everywhere, bandwidth will be soon higher than wired**
- **All devices can communicate**
  - multiple network interfaces (comes for free)
  - tons of cheap memory,
  - maybe soon tons of batteries
- **Communication services available through multiple technologies**

What does not work

- **The Internet**
  - Not robust
  - Not secure
  - Not capable of integrating wireless (TCP)
  - Not capable of integrating mobility (addressing)
  - Infrastructure based (centralized)
- **The Internet remains the best file transfer network**

Infrastructure based services

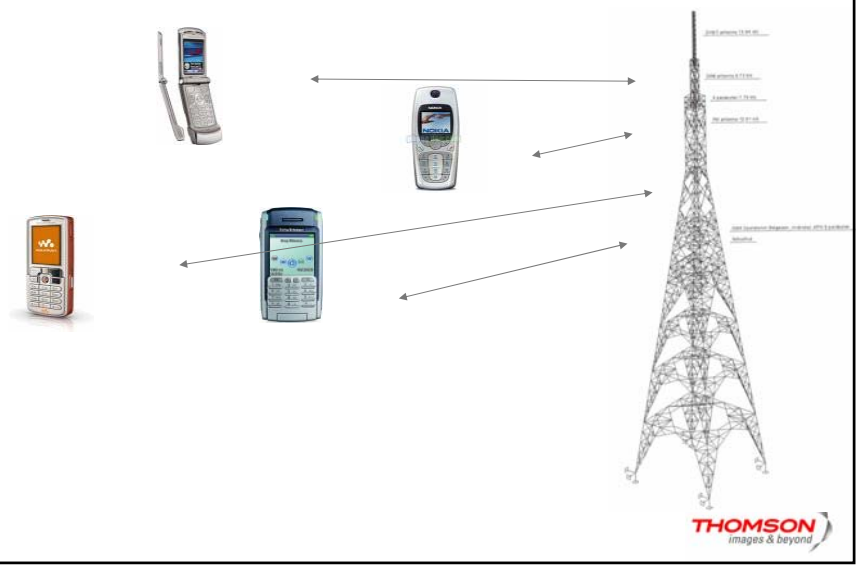

Infrastructure based services

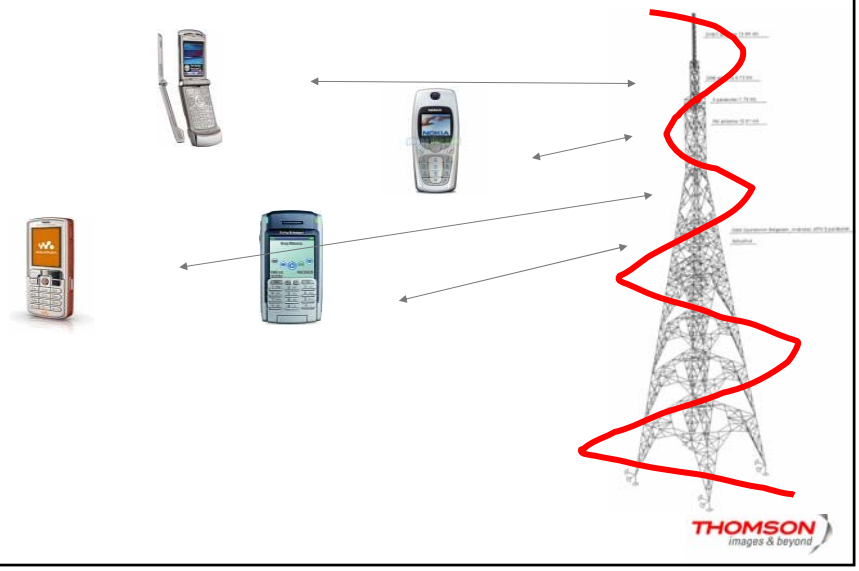

Infrastructure based services

But mobile phones and PDAs are still working….

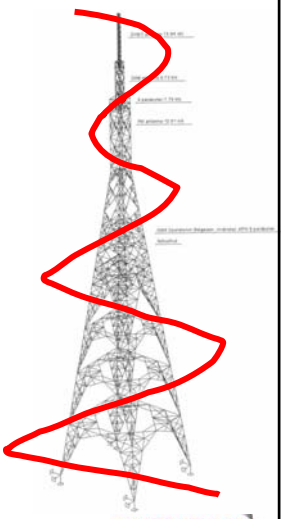

No infrastructure
Every terminal is also a router
Mobility change topology

THOMSON images & beyond

No alternative to global services

Today

OR …

Tomorrow

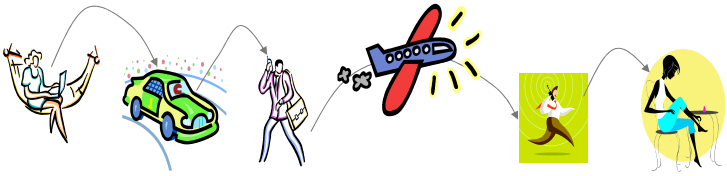

THOMSON images & beyond

A world of content

- **Large diversity**
- **Easy to create and distribute**
  - Internet
- **Access to content remains problematic**
  - A dedicated infrastructure per service/content
  - Not interoperable
  - No infrastructure, no service
  - No dedicated terminal, no content


THOMSON
images & beyond


XPS
THOMSON
images & beyond

Sumary

- **Technology convergence will not happen!**
  - Infrastructures exist and are performant
- **Content must be available**
  - Everywhere and anytime
  - On all terminals
  - With an « acceptable » quality

⇒ **Convergence is happening at the service level**
(triple-play, TV on cell phone, etc.)

A new era for communication

- **New applications, new usages (local & global)**
  - Service personalization (health, tourism)
  - Security (natural disasters)
  - Ease of commercial transactions
  - Legacy applications everywhere
- **New business models**
- **Difficult research problems**
  - Distributed architectures, self-configuration
  - Peer-to-peer search engines
  - Authentication, confidentiality, DRM

A new communication paradigm

- **Pocket Switched Networks**
  - Any network capable device can store and forward content
  - Exchange content using device/user mobility
  - Local peer-to-peer communication
  - Access to legacy services through gateways

THOMSON images & beyond

How would it work?

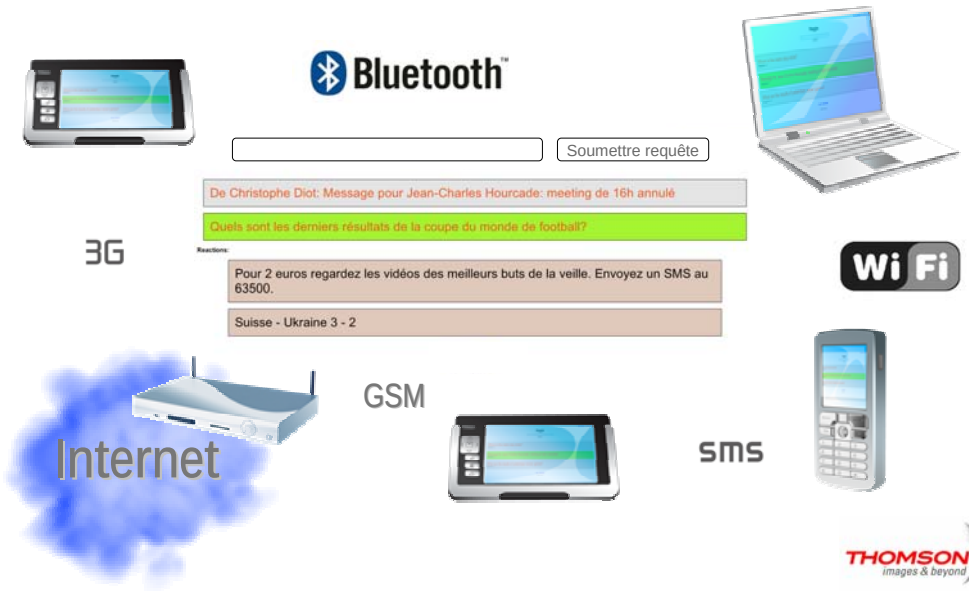

THOMSON images & beyond

Pocket Switched Networks

- **Provide a useful set of services in the absence of a fixed infrastructure and e2e connectivity**
- **Cross devices, cross network technologies**
- **Exploit features of the problem space**
  - node mobility for message delivery
  - build communities
  - distributed, intuitive authentication
- **Integration with legacy systems**
  - email delivery
  - web requests

Assumptions (1)

- **Users carry devices with connectivity**
  - Bluetooth, 802.11, Ethernet, cradle, etc.
- **Storage is not a problem**
  - storage density doubling every year
- **Battery is more of a problem**
  - heuristics to determine how scarce battery is
  - scale Haggle operations appropriately
  - allow user control

Assumptions (2)

- Local (e.g. Bluetooth, WiFi) and global (e.g. Internet, GPRS) connectivity are *transfer opportunities*
- Use application-layer data instead of (address, port) for forwarding
- No centralized service available (naming, security, etc.)

THOMSON images & beyond

Challenges

- **Exploit massive aggregate bandwidth**
  - Devices with local connectivity
  - Make use of MBs of local storage
  - Heterogeneous network types
- **Distributed naming**
- **Nodes need to "locate" themselves and their neighbours**
- **Forwarding decision**
- **Security, trust and reputation**
- **Paid to user acceptance, usability, and privacy**

THOMSON images & beyond

Community networks

- **Recently, many applications have relied on communities sharing a common interest/goal:**
  - Overlay networks, VPNs, etc.
  - File sharing P2P networks
  - Ad-hoc networks
- **Communities may be transient (concert attendees) or long-lived (interest groups)**
- **Users may be in multiple communities at any time, and change communities over time**
- **Issues with naming, trustability, security, incentive to cooperate**

Small world forwarding

- **State information are flooded**
- **No routes, no routing**
- **Network Coding ???**
- **Application forwarding based on neighbourhood status and history:**
  - Find the next hop that has the highest probability to deliver
  - Avoid flooding
  - Data aging

Localization

- **Node localization is important**
  - Neighbourhood discovery/community creation
  - Location-aware applications
  - Trust and security
- **Various options**
  - Relative to other nodes
  - Absolute (e.g. GPS)
  - Based on some external service (c.f. PlaceLab use of base stations)

Trust and Security

- **Human in the loop for bootstrapping trust**
- **New paradigms for trust establishment:**
  - history of past interactions
  - incentives based on virtual cash
  - Recommendation schemes
- **Opportunistic forwarding vs. cryptography: combine encryption with network coding**
- **Overall privacy preserving mechanisms**

Usability and User Involvement

- **Usability will require informing the user about network state issues – but how?**
  - Spinning globe, signal strength bars, greyed-out names
- **When is user involvement required?**
  - Most decisions must be made by devices
  - Some decisions may have to be left to user
- **What incentive to cooperate?**

THOMSON images & beyond

Traces driven design of forwarding

- **Bluetooth (iMotes)**
  - Cambridge local
  - Infocom 05
  - Hong-Kong
- **WiFi (Laptops)**
  - Dartmouth
  - San Diego
- **GSM & Bluetooth (cell phones)**
  - Reality Mining (MIT)

THOMSON images & beyond

What we measure

- **For a given pairs of devices:**
  - contact times and inter-contact times.

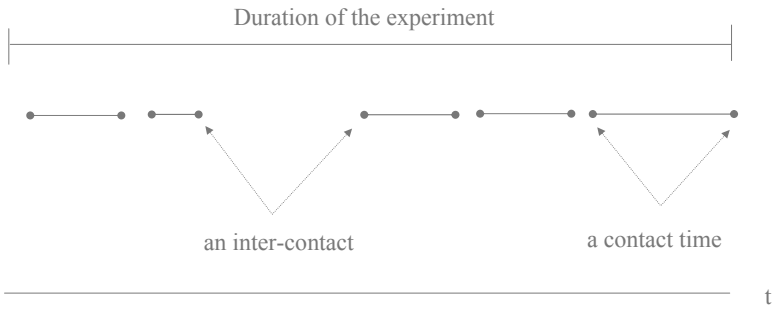

Example: a typical pair

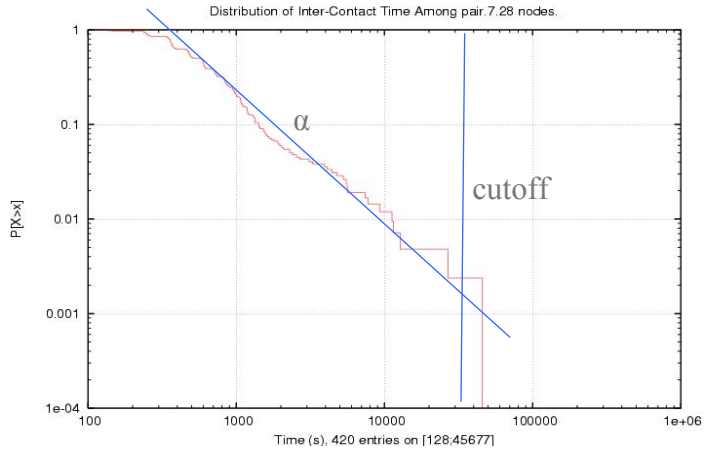

Inter-contact time – iMote experiments

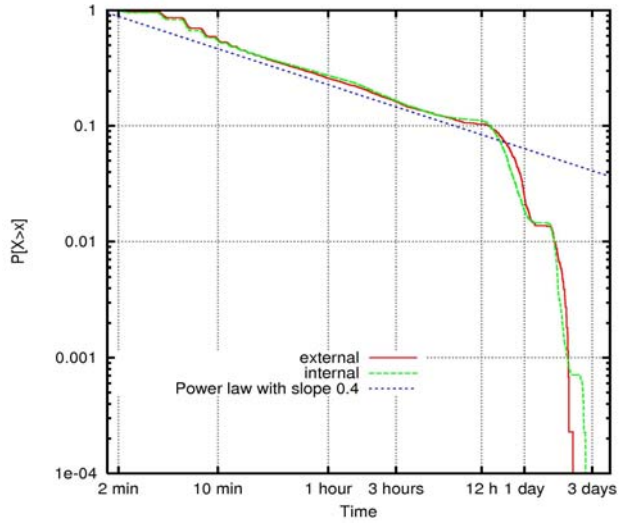

Inter-contact time -- WiFi experiments

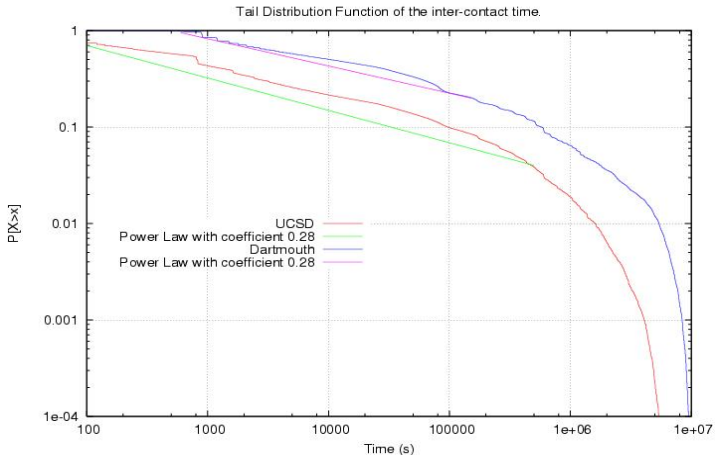

Summary of observations

- Inter-contact time follows an approximate power-law shape in all experiments in the [10mns, 12 hours] range.
- α < 1 most of the time (very heavily tailed).
- Some variability with the time of day, or among pairs.

Impact of α on convergence

- For $\alpha > 2$
  Any stateless algorithm achieves a finite expected delay.
- For $\alpha \geq \frac{m+1}{m}$ and $\#\{\text{nodes}\} \geq 2m$:
  There exist a forwarding algorithm with m copies and a finite expected delay.
- For $\alpha < 1$
  No stateless algorithm (even flooding) achieve a bounded delay (Orey's theorem).

Consequences on mobile networking

- **Mobility models needs to be redesigned**
  - Exponential decay of inter contact is wrong.
  - Mechanisms tested with that model need to be analyzed with new mobility assumptions.
- **Forwarding will not be easy**
  - Are there simple forwarding rules?
  - Can we benefit from heterogeneity to forward by communities ?

THOMSON images & beyond

Mobility increases opportunities

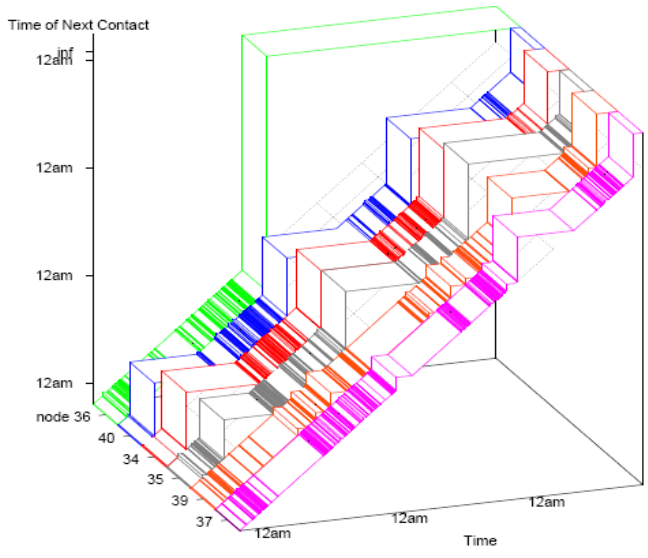

THOMSON images & beyond

Forwarding in PSN

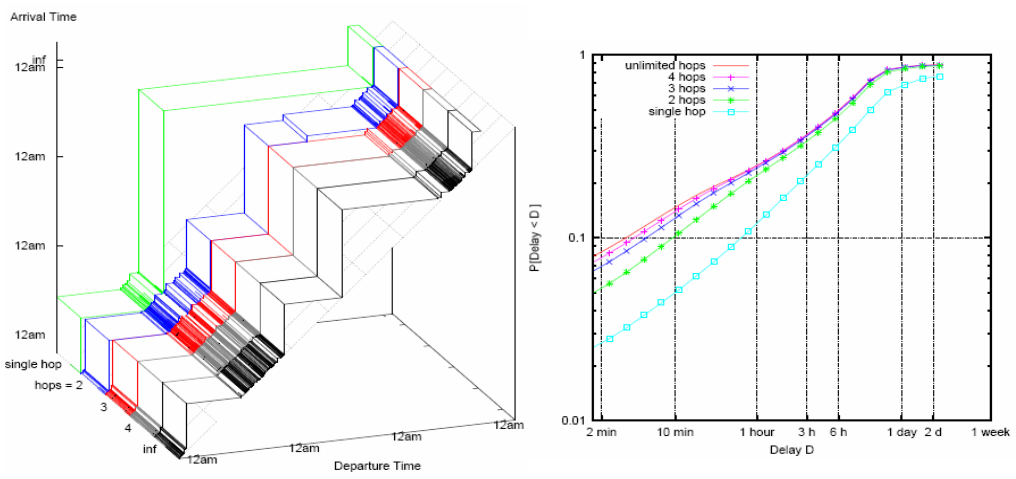

THOMSON images & beyond

On-going work

- **Study the limits of naïve forwarding**
  - 4 hops, no more than 8 copies
- **Are we a small world?**
  - Kleinberg's result would apply
- **Characterization of optimal paths**
- **Communities?**
  - Do they exist
  - Can we use them to forward?
  - How?

THOMSON images & beyond

URLs

- www.haggleproject.net
- www.haggleproject.com
- www.haggleproject.org

- www.cambridge.intel-research.net/haggle/eu/