Pocket Switched Networks:
Real-world Mobility and its Consequences for Opportunistic Forwarding

Jon Crowcroft, Pan Hui, Augustin Chaintreau,
Joint work with
Christophe Diot, James Scott, Richard Gass,

Intel Research Cambridge
University of Cambridge

Intel Research

Outline

- Motivation and context
- Experiments
- Results
- Analysis of forwarding algorithms
- Consequences on mobile networking

The world is NOT connected!

- Users move between heterogeneous connectivity islands
- End-to-end is not always possible
  - One or both ends may be disconnected
  - Internet “routing” is a bad idea
- Device should make network decisions
  - Shall I send by email, infrared or Bluetooth?

No alternative to the Internet

Today

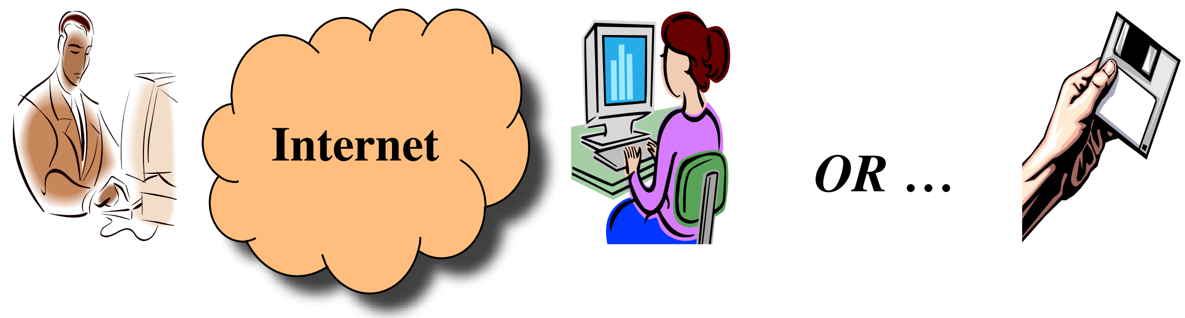

Tomorrow

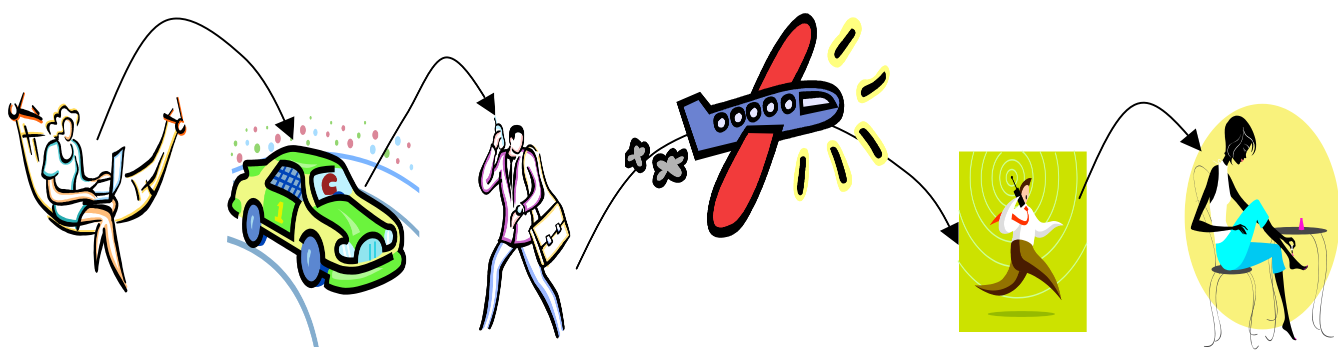

Intel Research

Pocket networking

- A packet can reach destination using network connectivity or user mobility
- Mobility increases capacity.

  [Grossglauser and Tse 2001]

Intel Research

State of the art

- Most efforts try to hack Internet legacy applications so that they work in Delay Tolerant Environments
  - MANET
  - DTN (even if DTN is more general by definition)
- Real “ad-hoc” approaches:
  - Zebranet, Lapnet, Cyberpostman

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

Applications

- Asynchronous, local messaging
- Automatic address book or calendar updates
- Ad-hoc google
- File sharing, bulletin board
- Commercial transactions
- Alerting, tracking or finding people

Outline

- Motivation and context
- Experiments
- Results
- Analysis of forwarding algorithms
- Consequences on mobile networking

Three independent experiments

- In Cambridge
  - Capture mobile users interaction.
- Traces from Wifi network :
  - Dartmouth and UCSD

User Population	Intel	Cambridge	UCSD	Dartmouth
Device	iMote	iMote	PDA	Laptop/PDA
Network type	Bluetooth	Bluetooth	WiFi	WiFi
Duration (days)	3	5	77	114
Granularity (seconds)	120	120	20	300
Nodes participating	141	238	261	6648
Number of contacts	3,984	8,856	175,105	4,058,284

iMote data sets

- Easy to carry devices
- Scan other devices every 2mns
  - Unsync feature
- log data to flash memory for each contact
  - MAC address, start time, end time
- 2 experiments
  - 20 motes, 3 days, 3,984 contacts, IRC employee
  - 20 motes, 5 days, 8,856 contacts, CAM students

What an iMote looks like

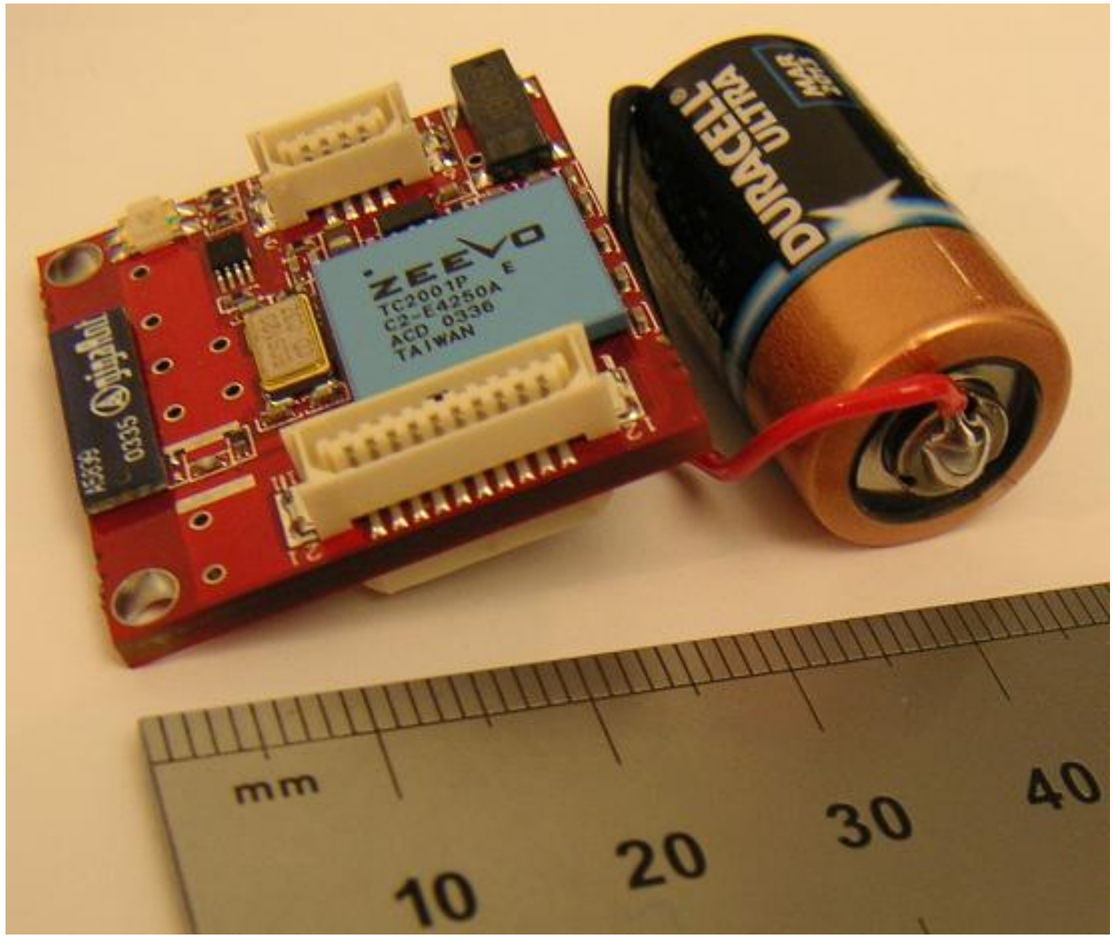

Intel Research

Experimental device

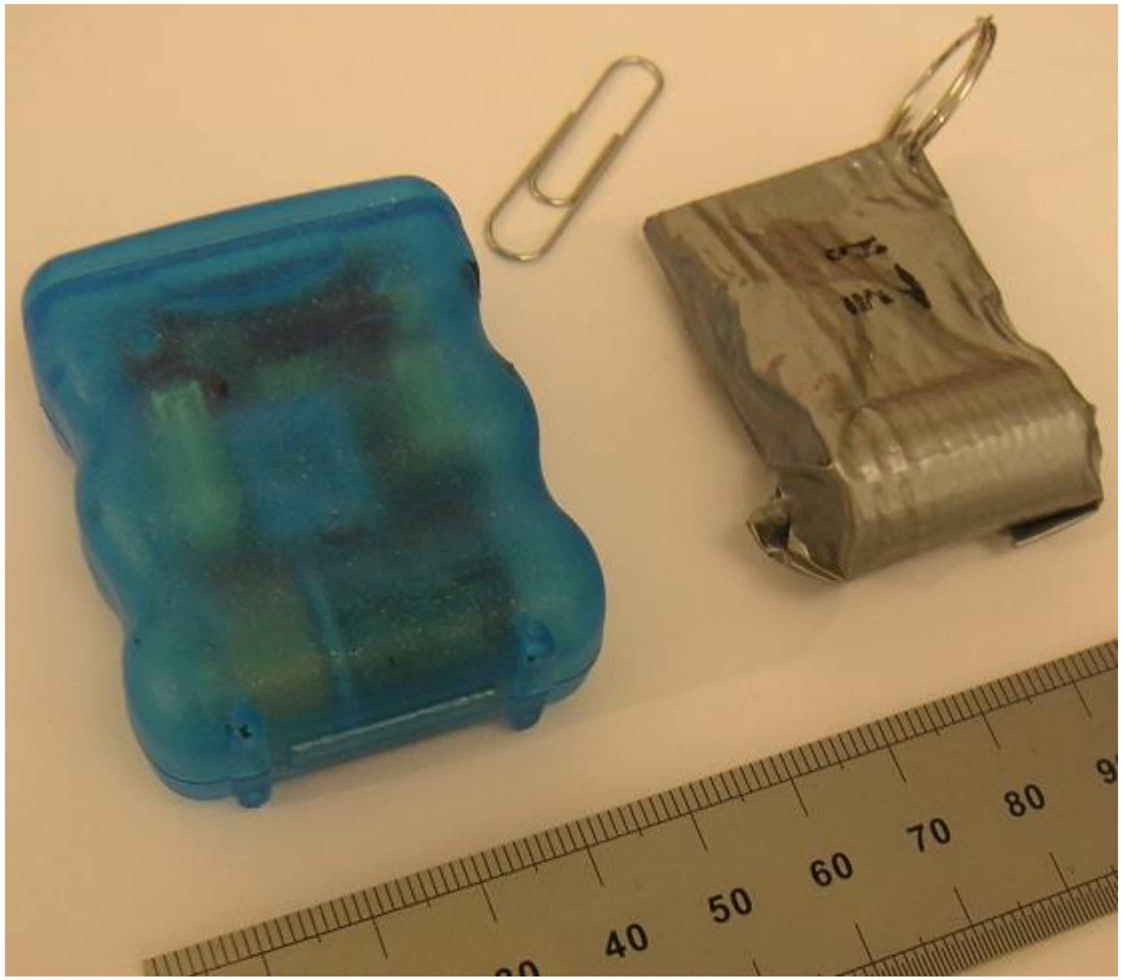

Intel Research

UCSD and Dartmouth Traces

- WiFi access networks
  - Client-based logs of AP (UCSD),
  - SNMP logs from AP (Dartmouth).
- Assumption:
  - Two clients logged on the same AP are in communication range.
- 3 months (UCSD), 4 months (Dartmouth).

Outline

- Motivation and context
- Experiments
- Results
- Analytical analysis
- Consequences on mobile networking

What we measure

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

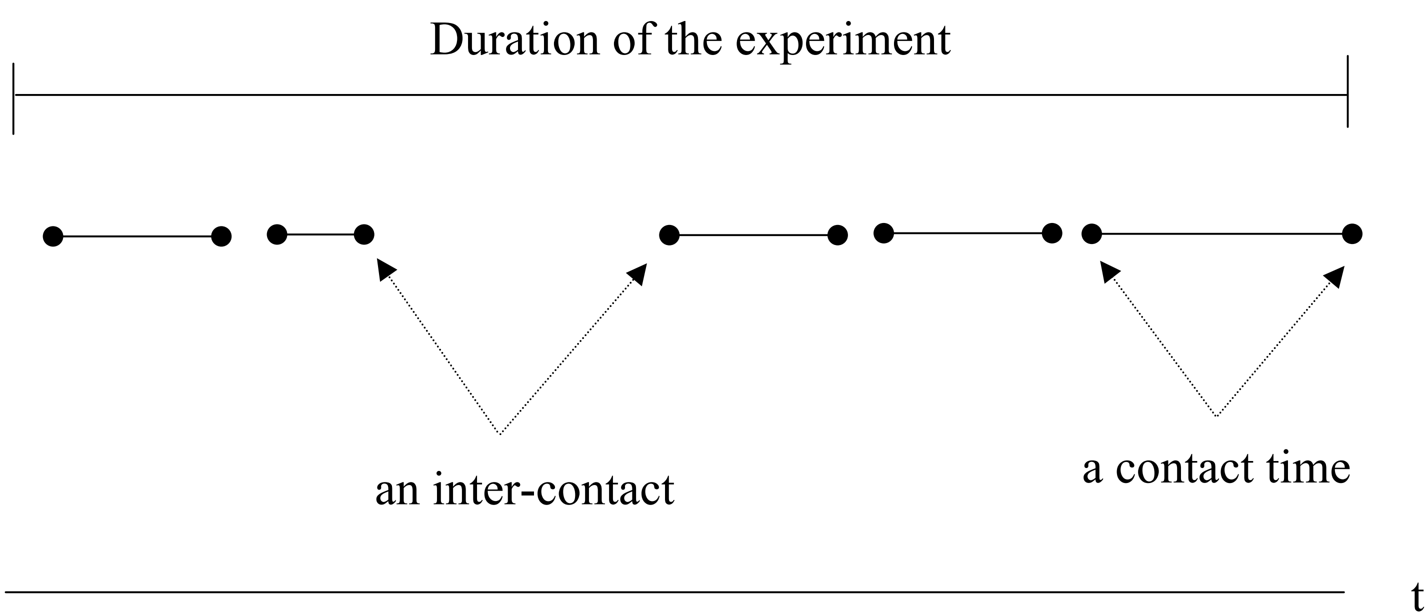

Intel Research

What we measure (cont'd)

- Distribution per event.
  - $\neq$ seen at a random instant in time.
- Plot log-log distributions.
- We aggregate the data of different pairs.

(see the following slides).

Example: a typical pair

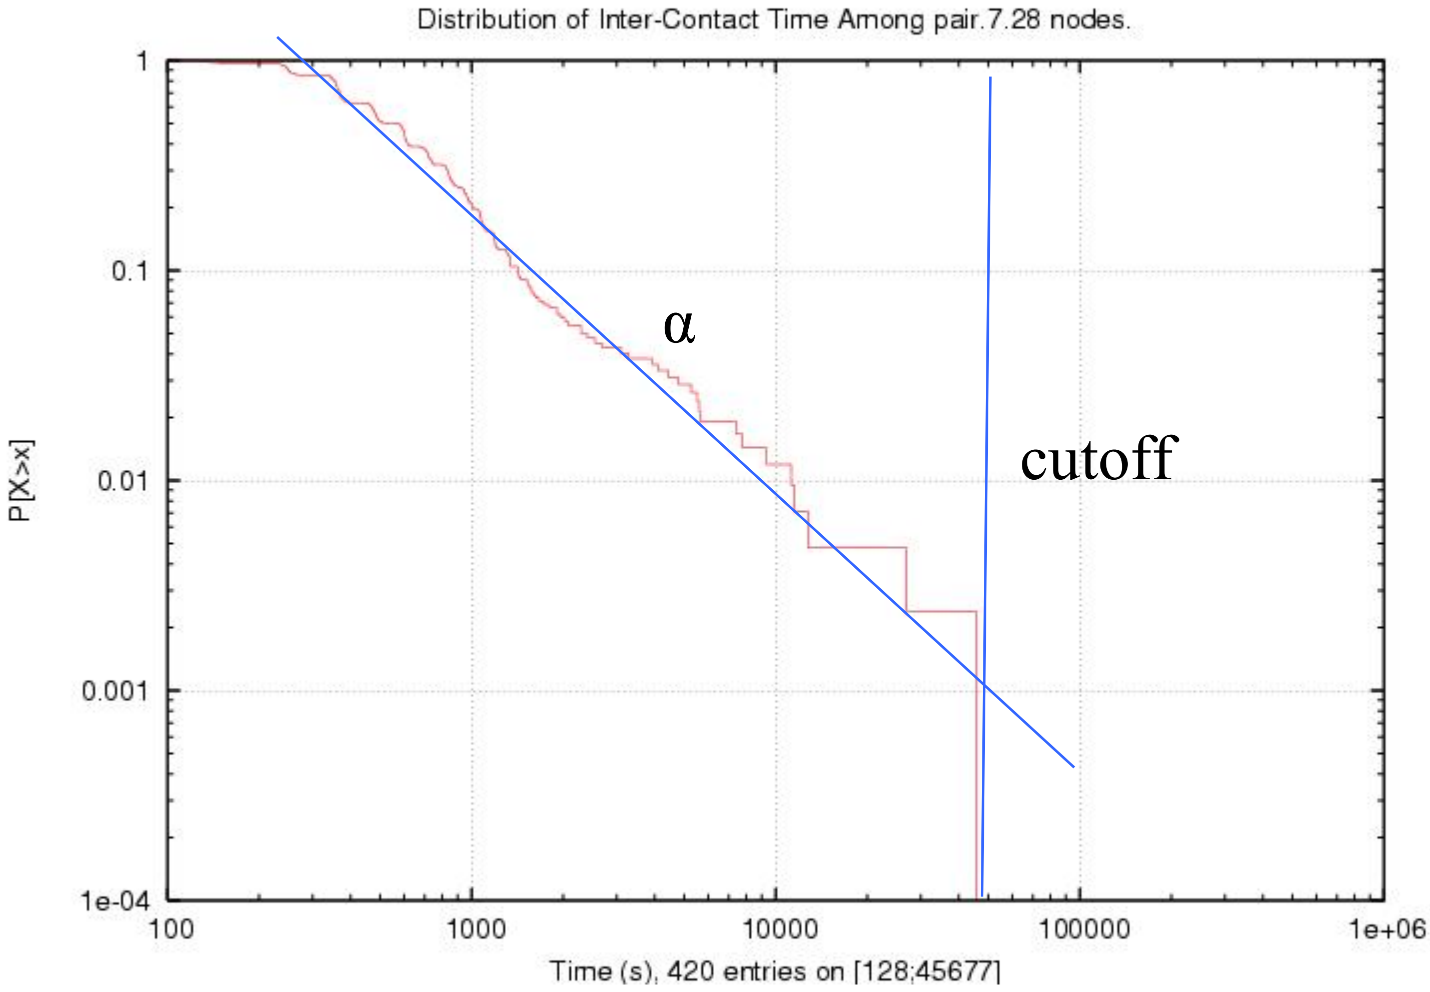

Intel Research

Examples : Other pairs

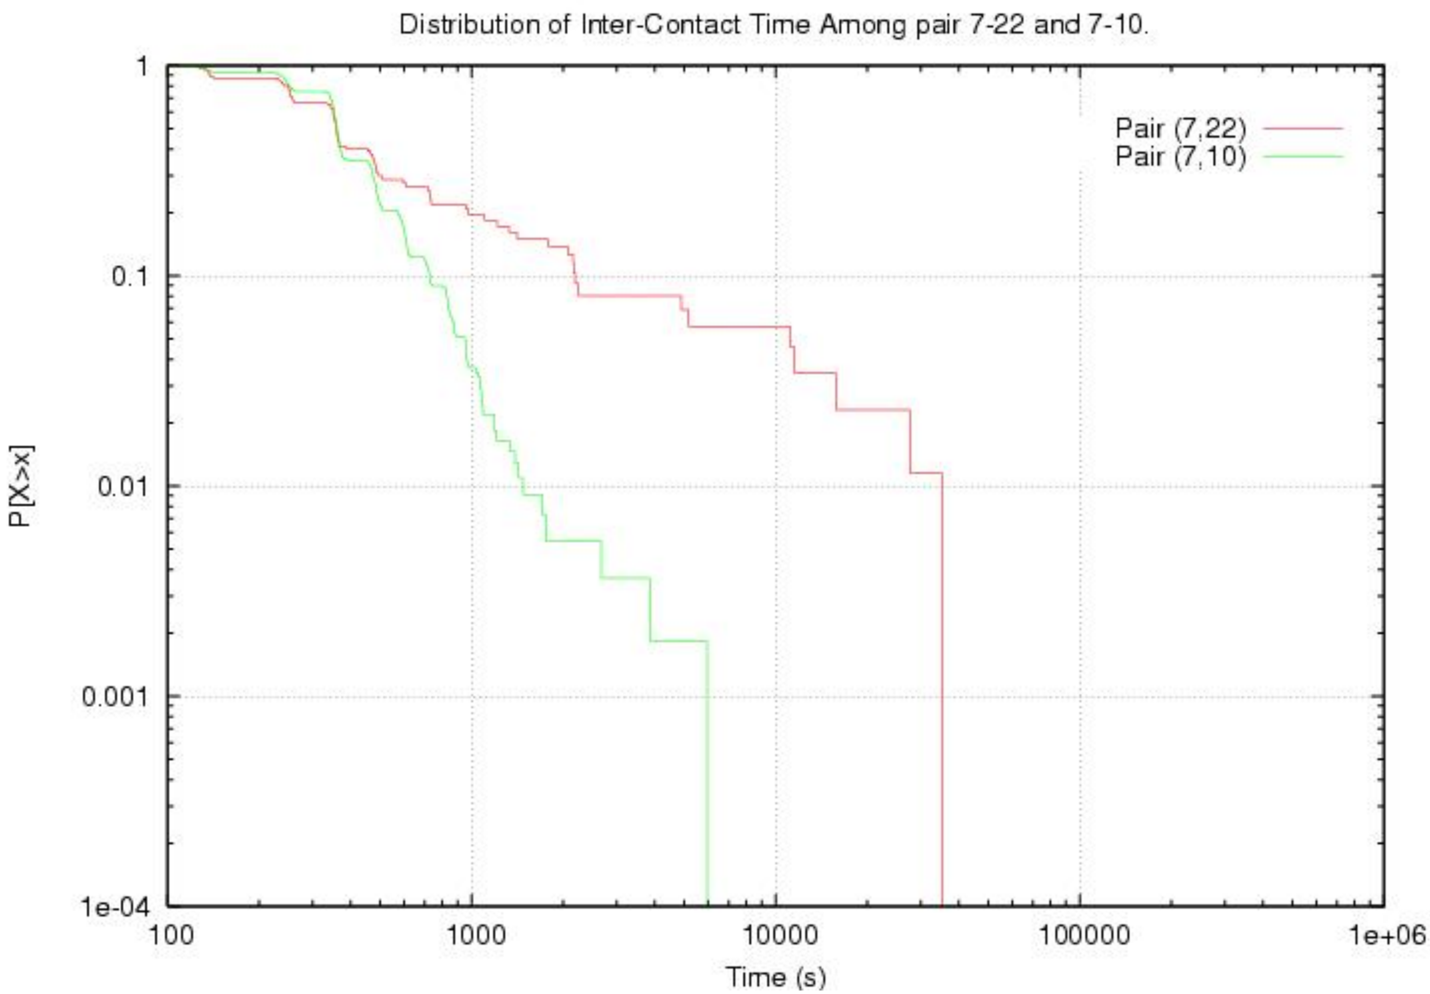

Intel Research

Aggregation (1): for one fixed node

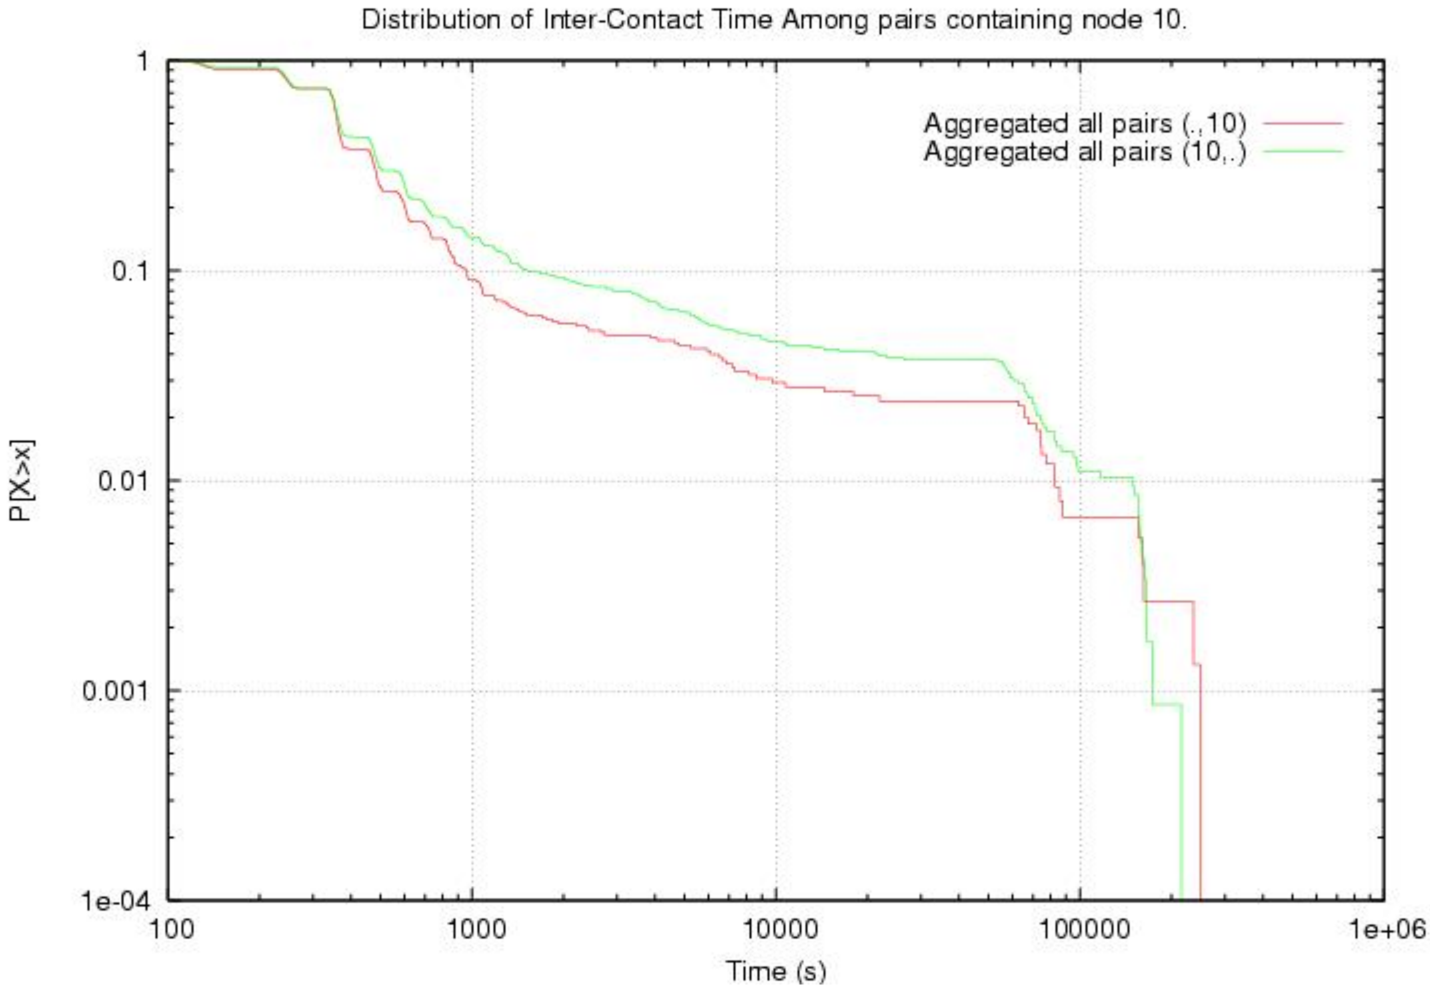

Intel Research

Aggregation (2) : among iMotes

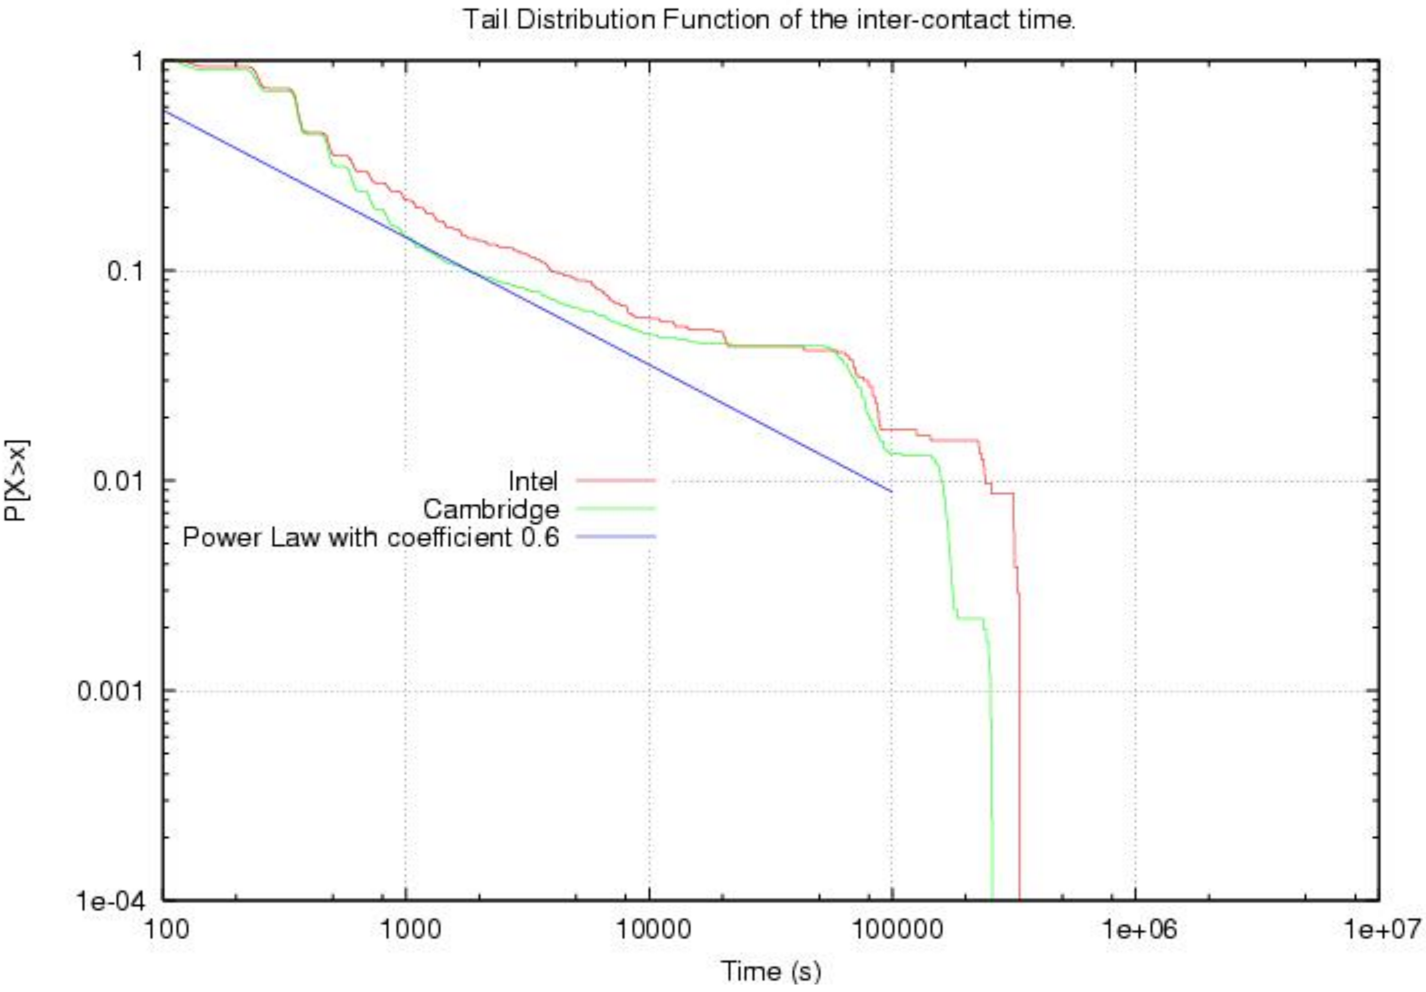

Intel Research

Summary

- Some heterogeneity among iMotes.
- Inter-contact distributions seem to follow a power law on [2mn; 1day].

- What about other nodes ? Campus WiFi experiments ? the time of the day ?

Inter-contact with External nodes

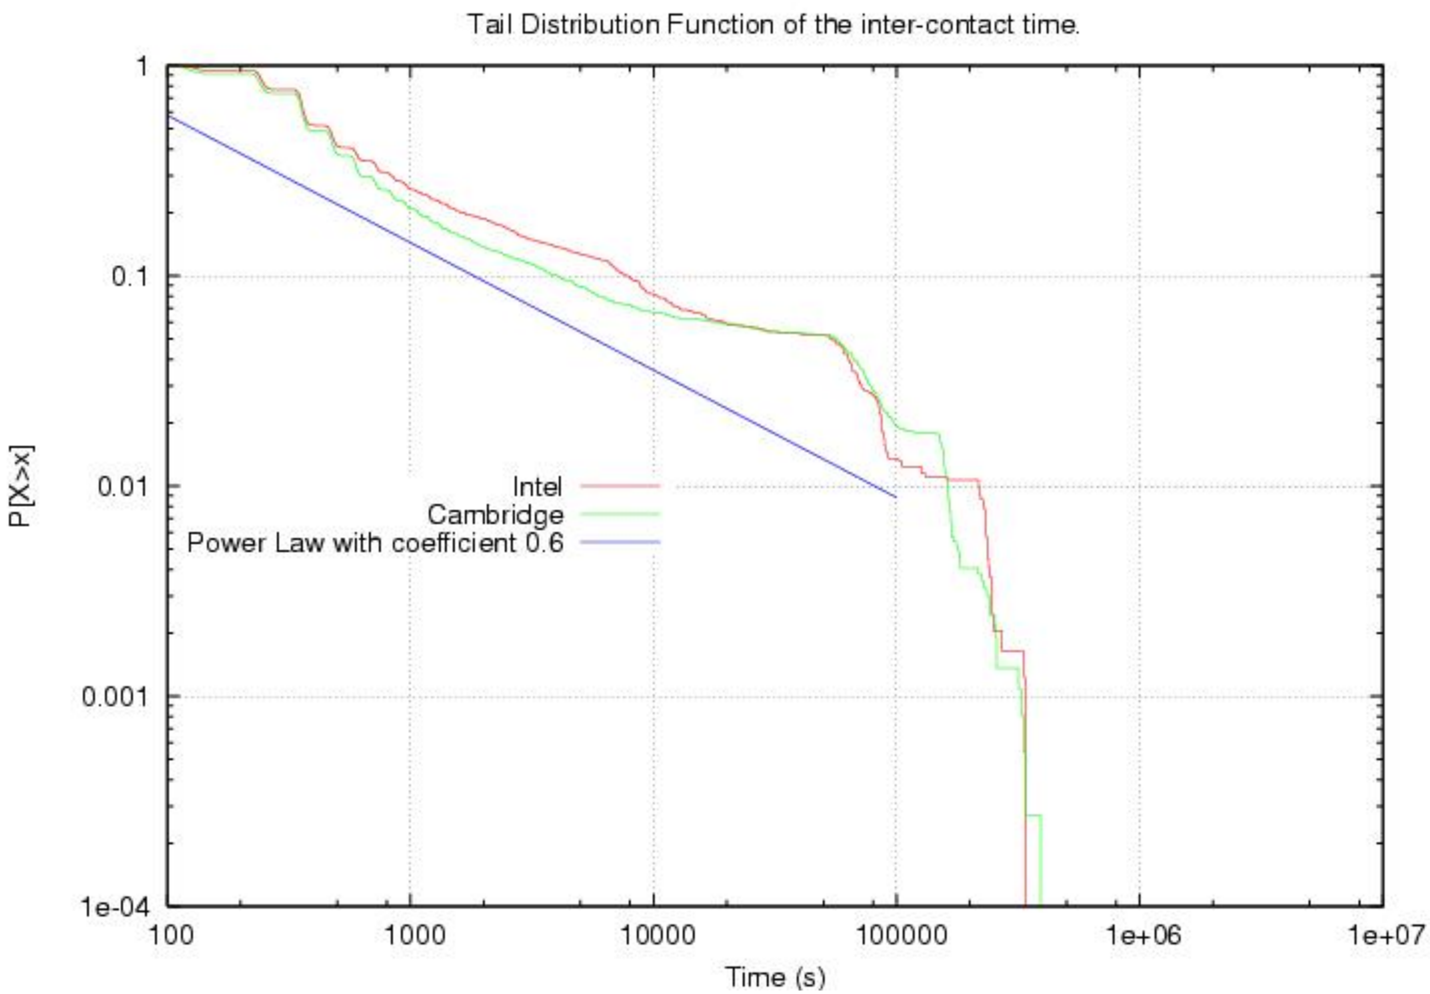

IntelResearch

Inter-contact time for WiFi traces

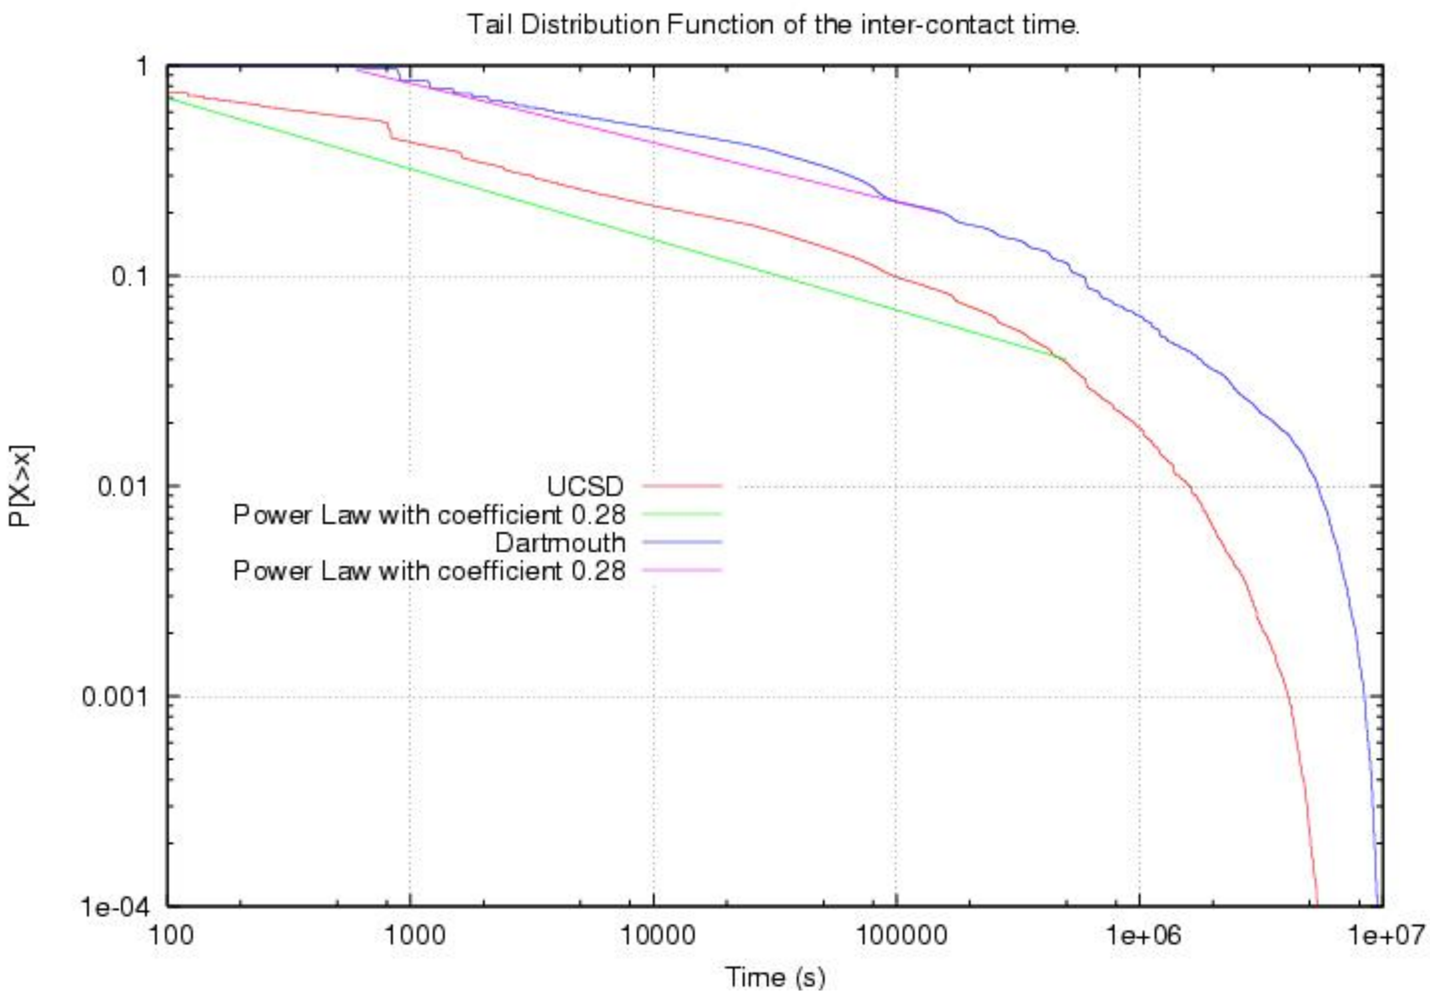

Intel Research

Inter-contact time during the day

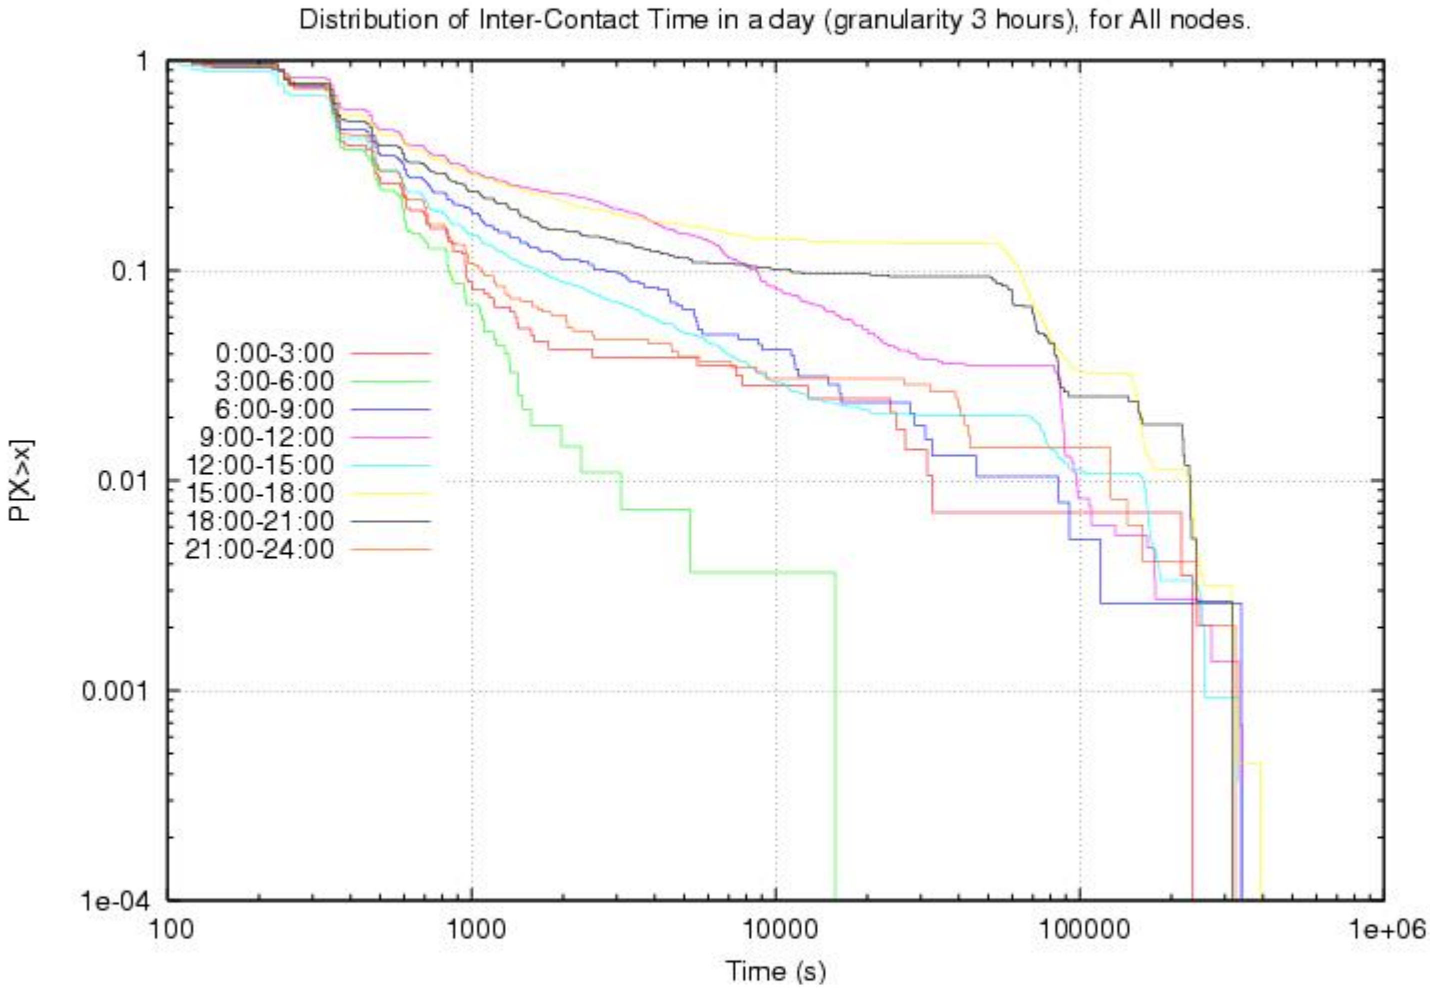

Intel Research

Inter-contact time during the day

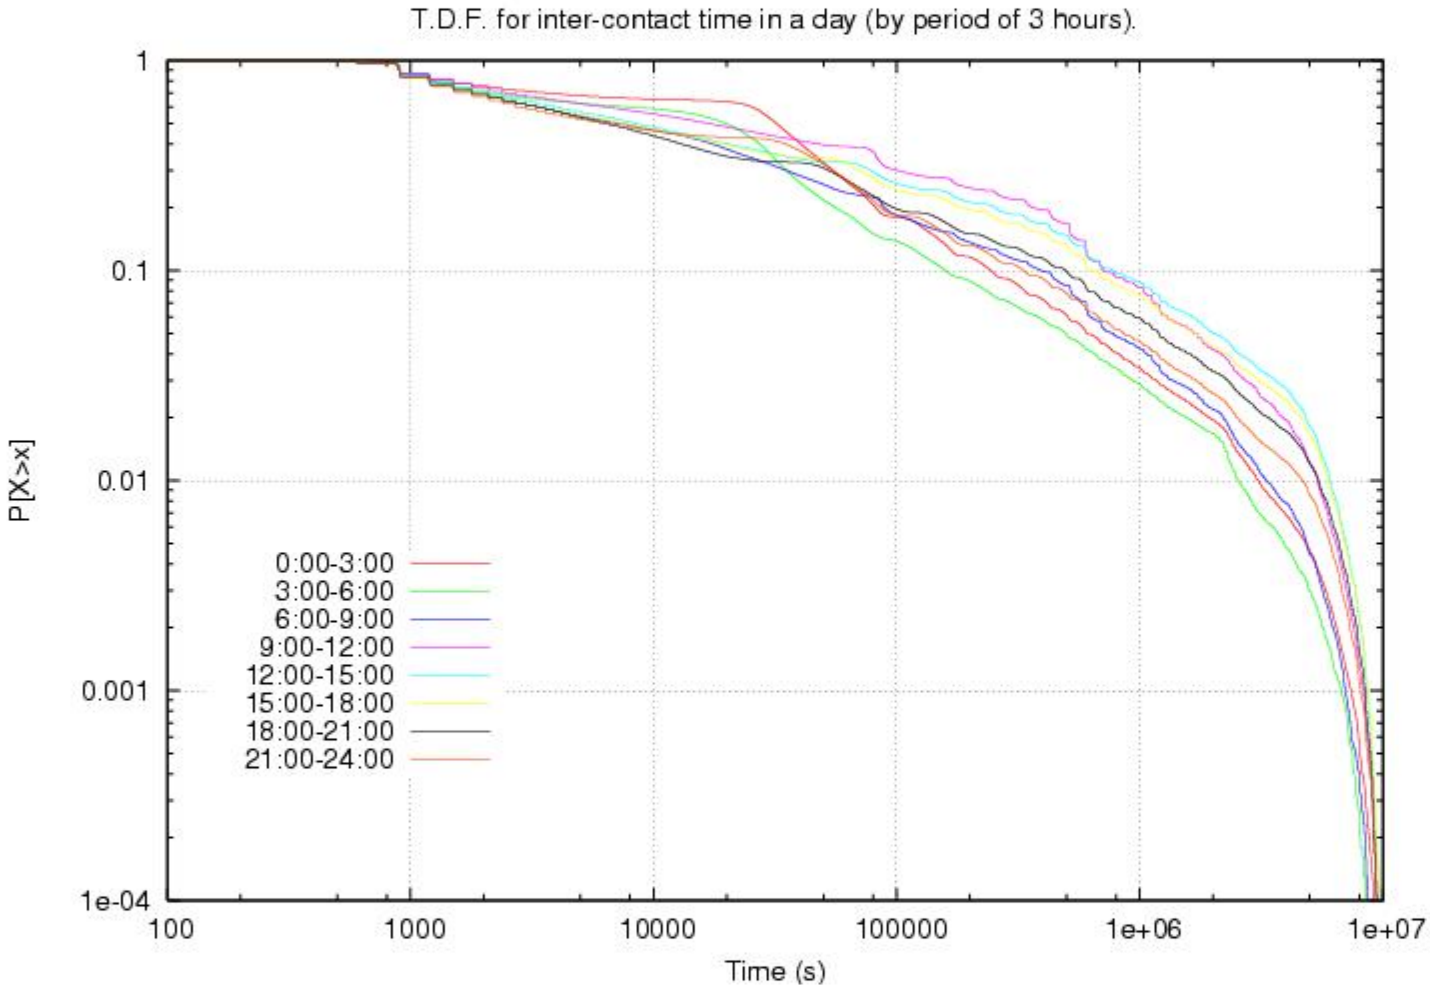

Intel Research

Summary of observations

- Inter-contact time follows an approximate power-law shape in all experiments.
- $\alpha < 1$ most of the time (very heavily tailed).
- Variation of parameter with the time of day, or among pairs.

Outline

- Motivation and context
- Experiments
- Results
- Analysis of forwarding algorithms
- Consequences on mobile networking

Problem

- Given that all data set exhibit approximate power law shape of the inter-contact time distribution:
  - Would a purely opportunistic point-to-point forwarding algorithm converge (i.e. guarantee bounded transmission delays) ?
  - Under what conditions ?

Forwarding algorithms

- Based on opportunities, and “Stateless” :
  - Decision does not depend on the nodes you meet.
- Between two extreme relaying strategies :
  - Wait-and-forward.
  - Flooding.
- Upper and Lower bounds on bandwidth:
  - Short contact time.
  - Full contact time (best case, treated here).

Two-hop relaying strategy

- Grossglauser & Tse (2001) :

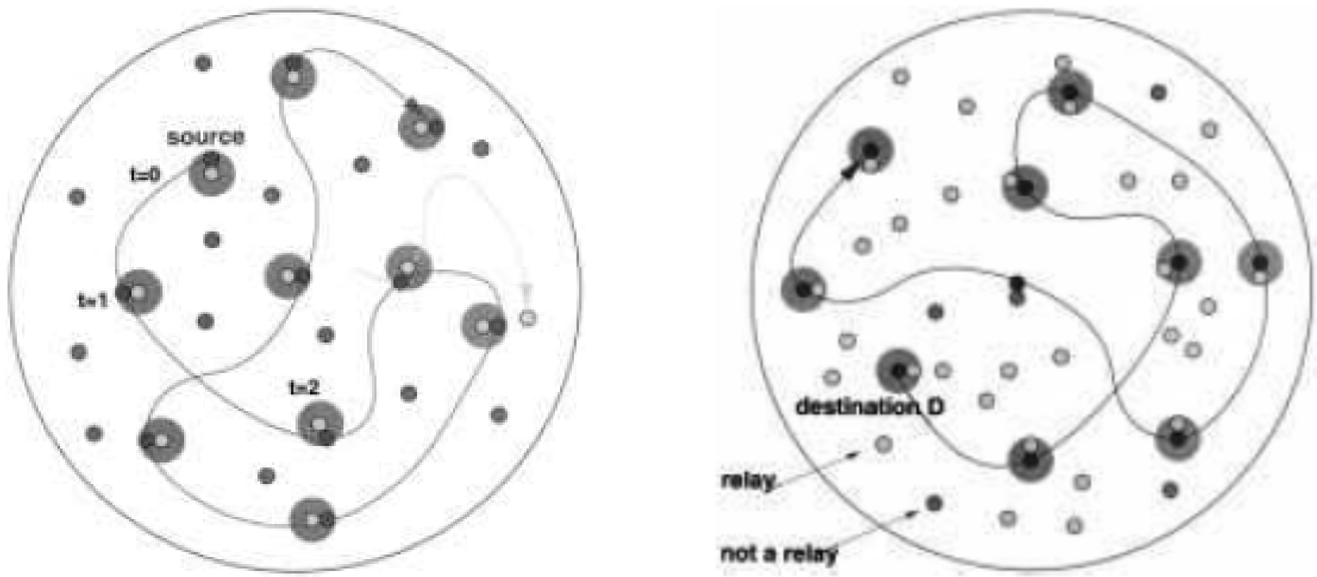

  – Maximizes capacity of dense ad-hoc networks.
  – Authors assume nodes location i.i.d. uniform.

Intel Research

Our assumptions on Mobility

- Homogeneity
  - Inter-contact for every pairs follows power law.

$$P[X \geq t] = t^{\alpha}$$

  - No cut-off bound.
- Independence
  - In “time”: contacts are renewal instants.
  - In “space”: pairs are independent.

Two-hop: stability/instability

- $\alpha > 2$

  The two hop relaying algorithm converges, and it achieves a finite expected delay.

- $\alpha < 2$

  The expected delay grow to infinity with time.

$$f_n = P[X = n]$$

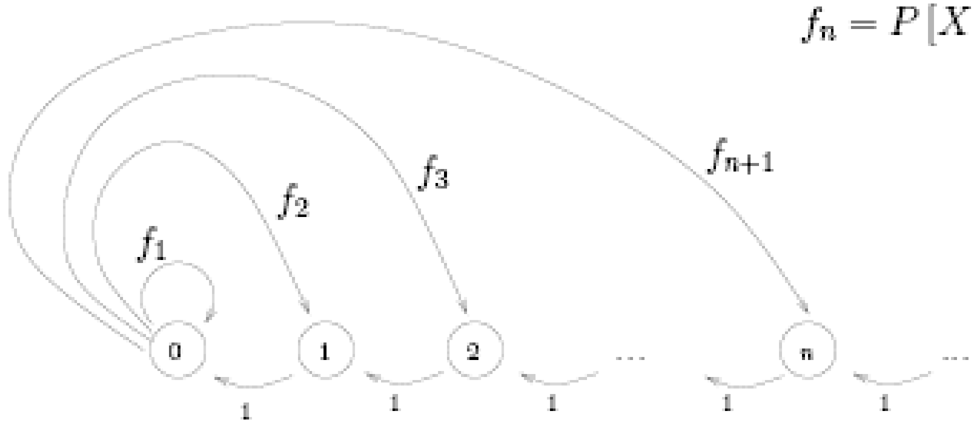

Intel Research

Two-hop: extensions

- Power laws with cut-off:
  - Large expected delay.
- Short contact case:
  - By comparison, all the negative results hold.
  - Convergence for $\alpha > 3$ by Kingman's bound.
  - We believe the same result holds for $\alpha > 2$.

The Impact of redundancy

- The Two-hop strategy is very conservative.
  - What about duplicate packet ? Or epidemics forwarding ?
- This comes to the question:

$$D_1, D_2 \text{ independent, with same law },$$

$$\text{if } \mathbb{E}[D_1] = \mathbb{E}[D_2] = \infty \text{ what is } \mathbb{E}[\min(D_1, D_2)]?$$

Intel Research

Forwarding with redundancy:

- 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).

Forwarding w. redundancy (cont'd)

- Further extensions:
  - The short contact case is open for $1<\alpha<2$.
  - Can we weaken the assumption of independence between pairs ?

Outline

- Motivation and context
- Experiments
- Results
- Analysis of forwarding algorithms
- Consequences on mobile networking

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.*
- Stateless forwarding does not work
  - *Can we benefit from heterogeneity to forward by communities ?*
  - *Scheme for peer-to-peer information sharing.*

THANK YOU

Tech Report available at:

http://www.cl.cam.ac.uk/TechReports/UCAM-CL-TR-617.html

Jon.Crowcroft@cl.cam.ac.uk, Pan.Hui@cl.cam.ac.uk,
augustin.chaintreau@intel.com

Intel Research

Next steps

- Collect more data
  - More motes
  - Other crowds of users
  - Collect contact time data
- Design algorithms that work
- New mobility models

Contact time distribution

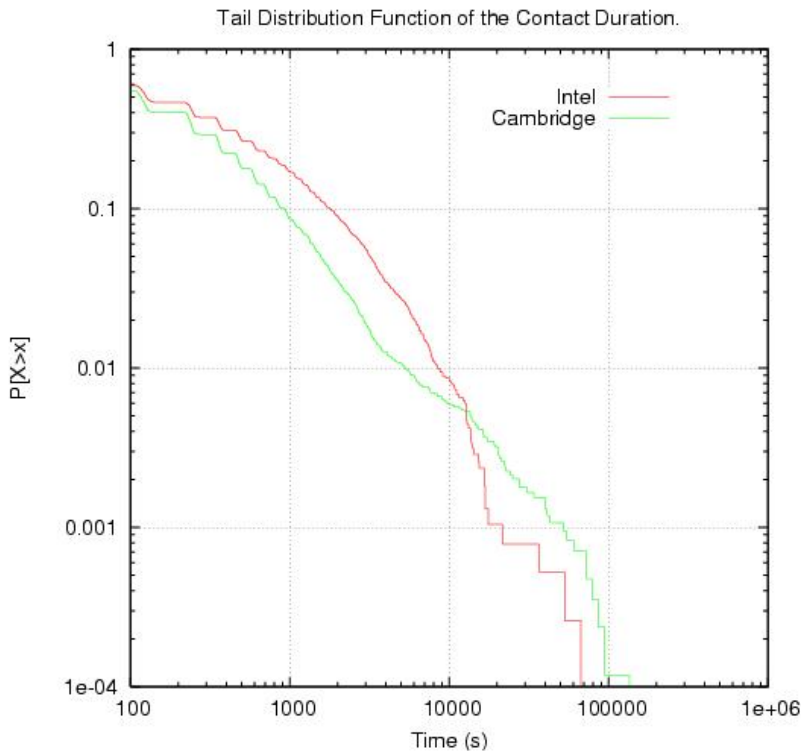

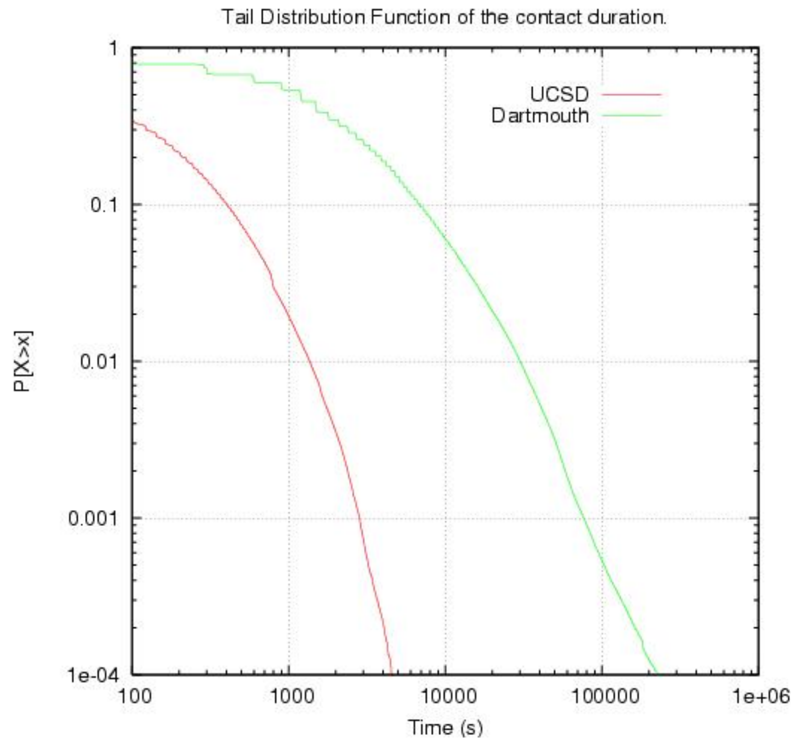

Intel Research

Inter-contact for all pairs

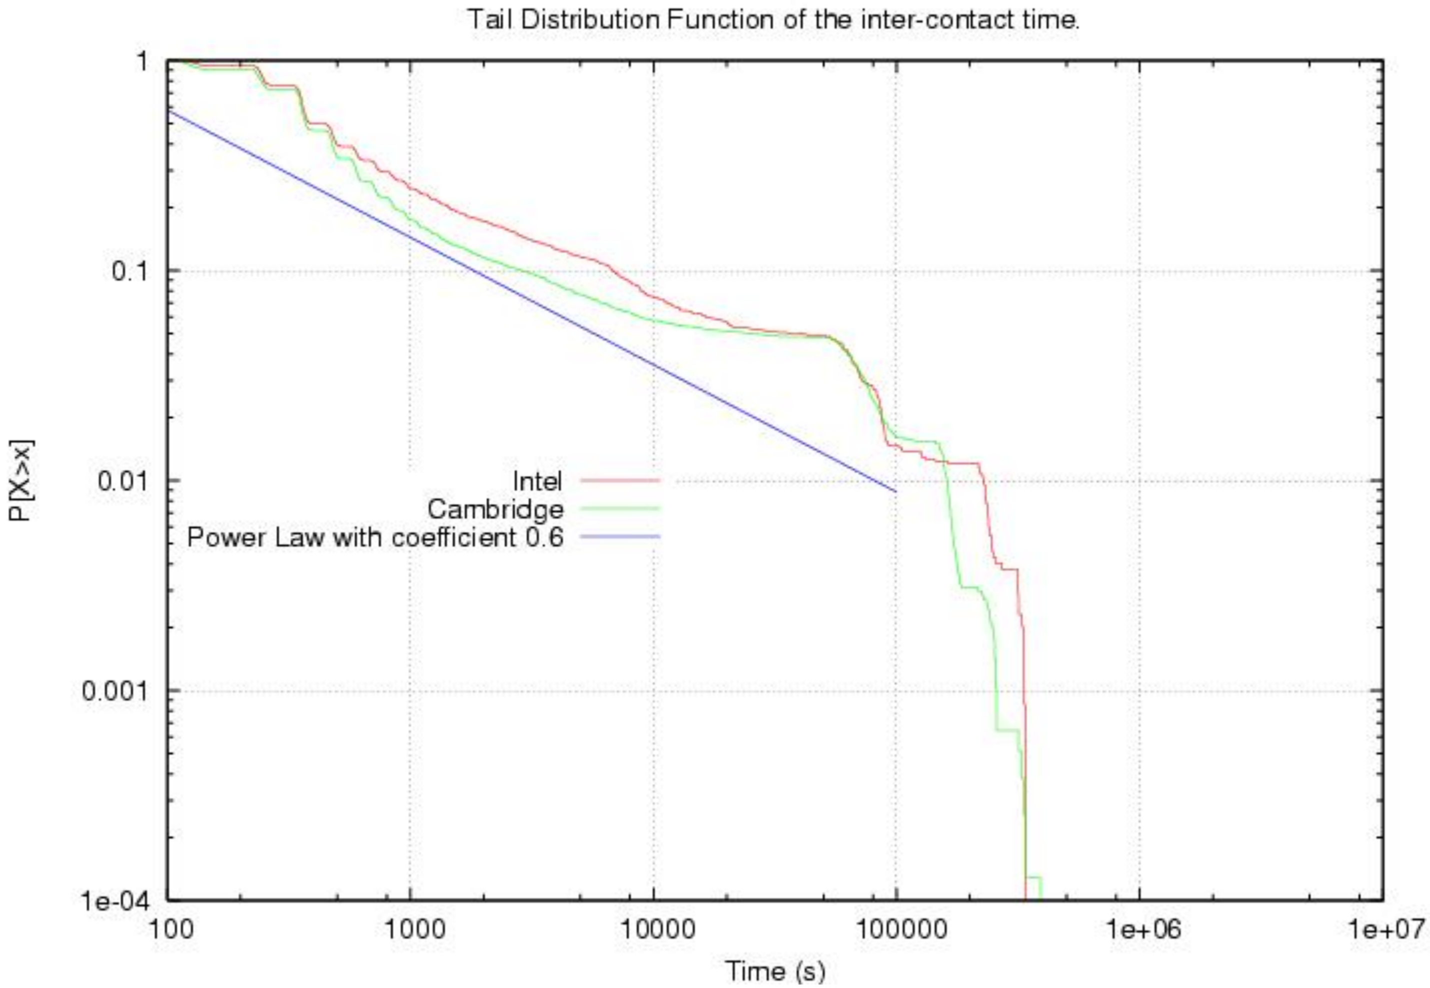

Intel Research