

## 1. Introduction

Researchers have proposed novel Internet applications for social network information beyond “traditional” social network sites, e.g., search, distributed computation, and security.

Users' **privacy concerns** are typically overlooked when considering such novel applications.

We present a **case study** of the **performance of social network routing** in opportunistic networks, before and after considering users' privacy preferences.

### Research Questions

- How might we **model the privacy concerns** which users have?
- What is the **impact on performance** on taking these concerns into account?

### Opportunistic Networks

Mobile devices, such as mobile phones, are commonly carried around by people during their daily lives.

Messages may be **directly exchanged between devices when they are in physical proximity** to each other in an ad-hoc manner.

In this way, an **opportunistic network** may be formed, where people send messages to each other via intermediary devices utilising a **disconnected store-and-forward** architecture.





Image credits:  
<http://www.flickr.com/photos/cobalt>  
<http://flickr.com/photos/mujitra>

### Social Network Routing

How can we route messages in an opportunistic network? Flooding messages to all encountered devices would be costly, draining battery life rapidly. Messages should be **selectively forwarded** during encounters.

Making the assumption that **encounters between mobile devices are more likely to occur between people in the same social network**, messages may be forwarded selectively only within the sender's social network.



## 2. Privacy threats

**Possible privacy threats** to participants in an opportunistic network which employs social network routing include [1]:

- Message content intercepted.
- Social network information used to inform social network routing decisions leaked.
- Messages traced progressing through the network, to infer communication patterns.
- Locations of participants inferred from the messages which their mobile devices carry. This may be in absolute terms (*Alice is at the supermarket*), or relative terms (*Alice and Bob were in the same location this afternoon*).

A privacy-conscious user **may choose to disable** their device's opportunistic network features at various times due to these privacy concerns. How might doing so affect **performance**?

### References

- [1] I. Parris, G. Bigwood and T. Henderson. Privacy-enhanced social network routing in opportunistic networks. In *Proc. of the IEEE International Workshop on Security and Social Networking (SESOC 2010)*, Mannheim, Germany, Mar 2010.
- [2] F. Ben Abdesslem, I. Parris, and T. Henderson. Mobile experience sampling: Reaching the parts of Facebook other methods cannot reach. In *Proc. of the Privacy and Usability Methods Pow-Wow (PUMP)*, Dundee, UK, Sep 2010. <http://scone.cs.st-andrews.ac.uk/pump2010/papers/benabdesslem.pdf>.
- [3] N. Eagle, A. S. Pentland, and D. Lazer. Inferring friendship network structure by using mobile phone data. In *Proc. of the National Academy of Sciences*, 106(36):15274–15278, Aug 2009.

## 3. Modelling users' privacy concerns

Developing a privacy model requires data on users' privacy behaviours, but collecting such data is not straightforward. Building an experimental, large-scale opportunistic network is impractical.

Therefore, we measure privacy behaviour with a **smaller-scale user study**, where we measure **location sharing preferences** of 40 *Facebook* users [2].

Category	Size	Location sharing choice		
		Nobody	Friends	Everyone
Open	17.5%	3.0%	9.1%	87.9%
Social	57.5%	7.1%	75.6%	17.3%
Closed	17.5%	59.1%	31.6%	9.4%
Variable	7.5%	31.7%	34.2%	34.2%

Table 1: Location-sharing behaviour on Facebook, by participant category.

We segment the participants into **categories** according to their privacy behaviour – i.e., according to their responses to prompted questions. We use these statistics (Table 1) as a dataset-independent **privacy model** – indicating how likely a node is to exchange messages with other nodes in a given encounter.

## 4. Performance results

To evaluate the performance impact of the privacy model on social network routing, we performed trace-driven simulation with two real-world datasets: our LocShare dataset [2], and the Reality Mining dataset collected at MIT [3].

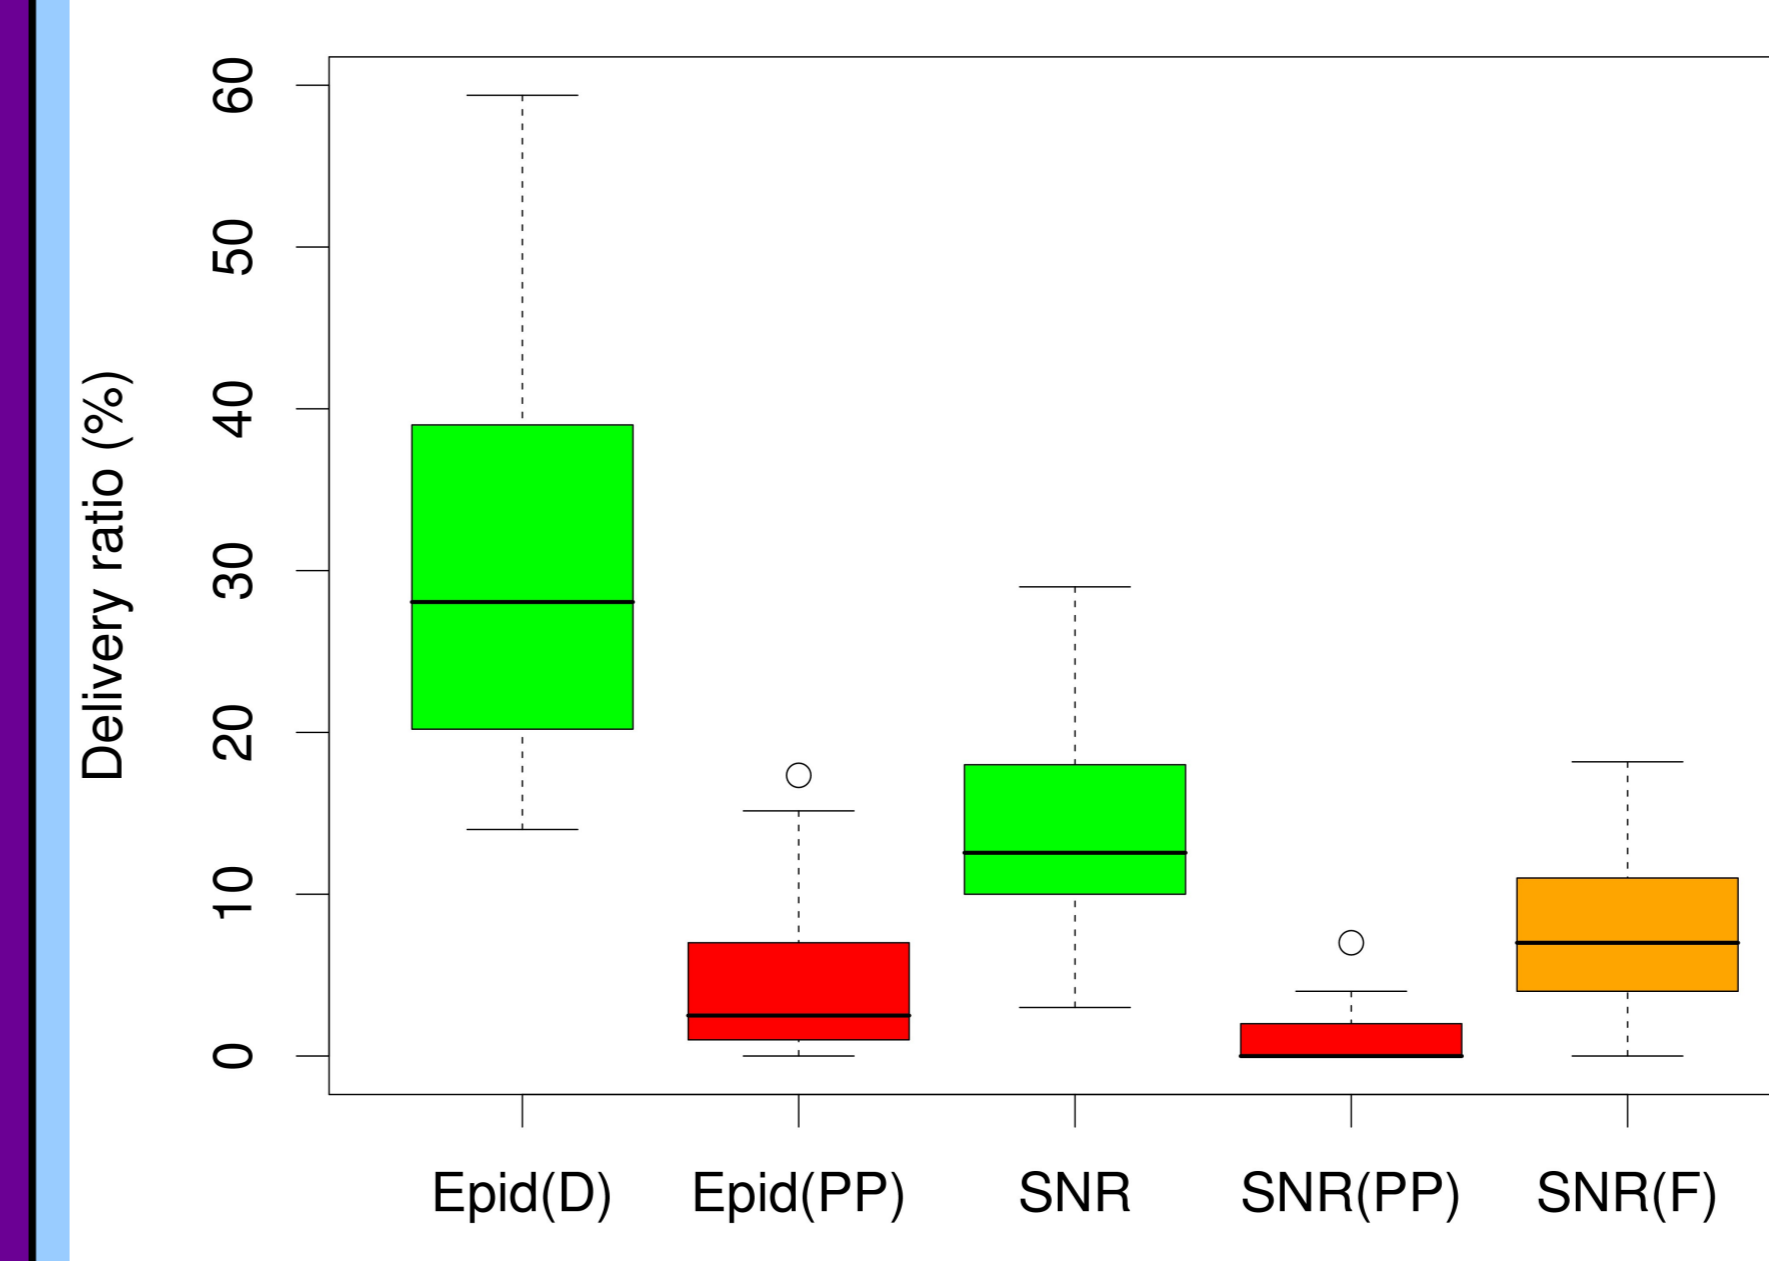
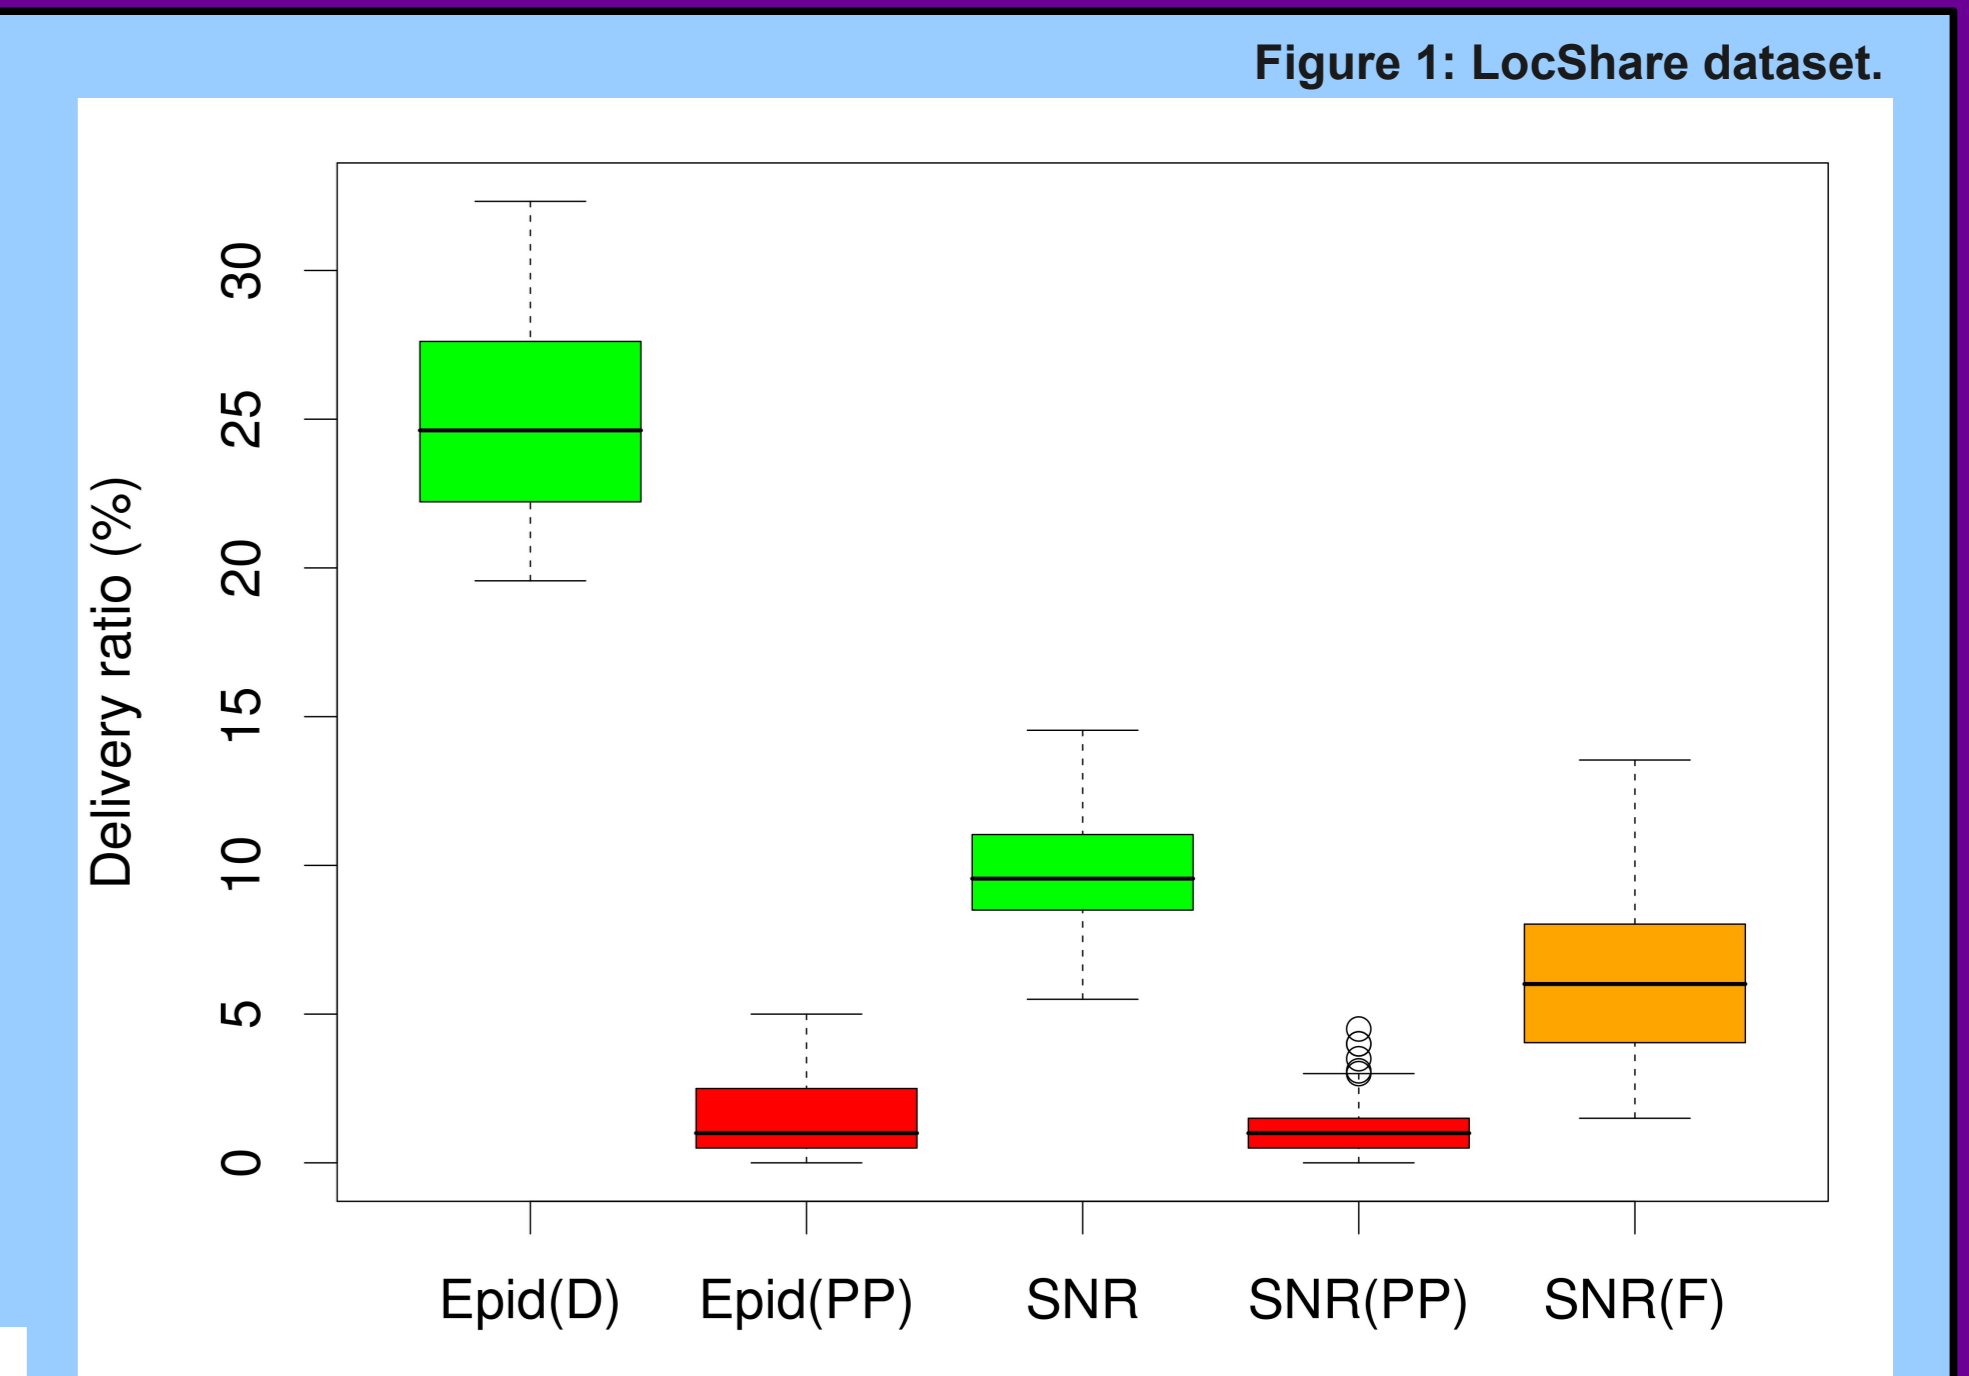


Figure 2: Reality Mining dataset.

We find that:

(1) Taking account of privacy concerns (orange boxes) leads to a **significant reduction in routing performance** compared to standard schemes (green boxes).

(2) The situation is compounded if we assume a stronger privacy model (red boxes). For SNR (PP) and the Reality Mining dataset, the **median delivery ratio falls from 12.6% to zero**.

## 5. Discussion

Taking users' privacy concerns into account may lead to dramatically lower performance for social network routing schemes. What does this mean for designers of future systems?

We present social network routing as a **cautionary tale**. **Failure to consider privacy concerns in early stages of design** may hinder the very novel Internet applications which the social network information had been intended to facilitate.