Optimizing Multi-copy Two-hop Routing in Mobile Social Networks

Huanyang Zheng*, Yunsheng Wang†, and Jie Wu*
* Department of CIS, Temple University, USA
†Department of Computer Science, Kettering University, USA

Presenter: Cong Liu

Introduction

Mobile social networks:

- Opportunistic contacts.
- Intermittent connectivity.

Two-hop routing:

- Uses local network information.
- Achieves a high delivery ratio through mobility.
- Each message copy will be forwarded at most twice, resulting in the advantage of the bounded resource (e.g., energy and buffer) consumption.

Introduction

Opportunistic two-hop routing (single-copy case)

- Link weights indicate average delay.

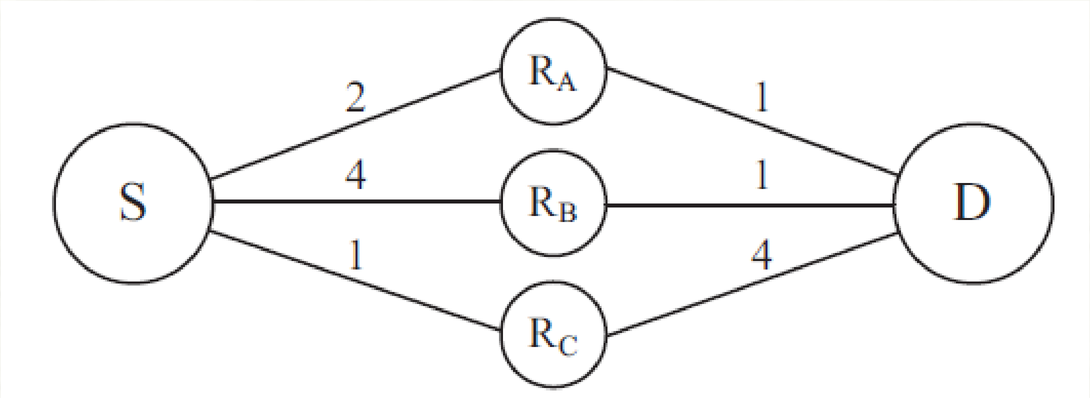

- Forward the message to the first encountered node?
  - ❑ Most likely is $R_{C.}$
  - ❑ Bad decision, since the delay of R_C-D is large.
  - ❑ Wait for S-R_A-D is better (2+1<4)

Introduction

Opportunistic two-hop routing (single-copy case)

- Link weights indicate average delay.

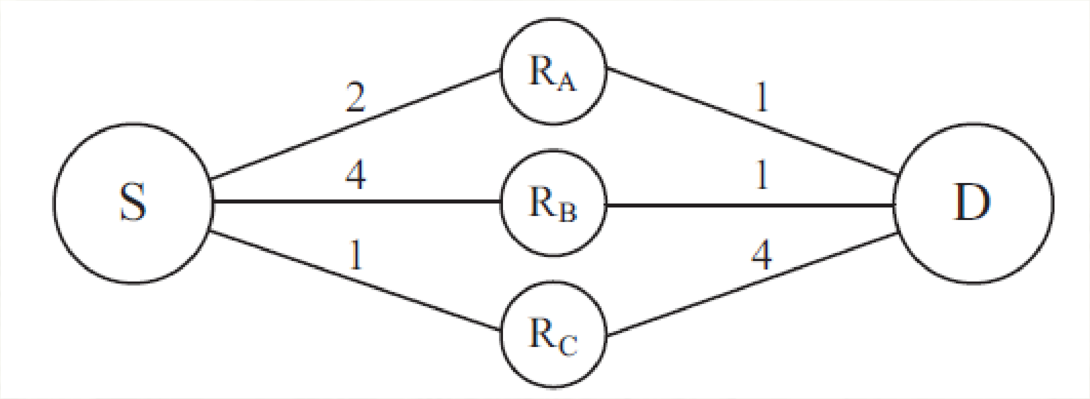

- Always shortest path routing (S-R_A-D) ?
  - ❑ Also bad, when opportunistically meeting R_B.
  - ❑ The delay of R_B-D is smaller than S-R_A-D.

Introduction

Opportunistic two-hop routing (single-copy case)

- Link weights indicate average delay.

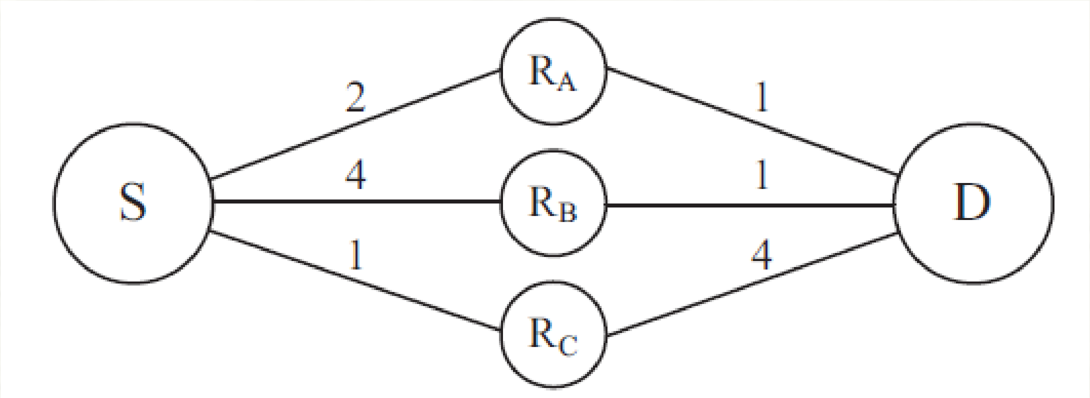

- Forwarding set:
  - The source only forwards its copy to encountered relays in its forwarding set $\{R_A, R_B\}$, ignoring R_C even if it is the next encounter.

Introduction

Opportunistic two-hop routing (multi-copy case)

- Assume the source has 3 copies.

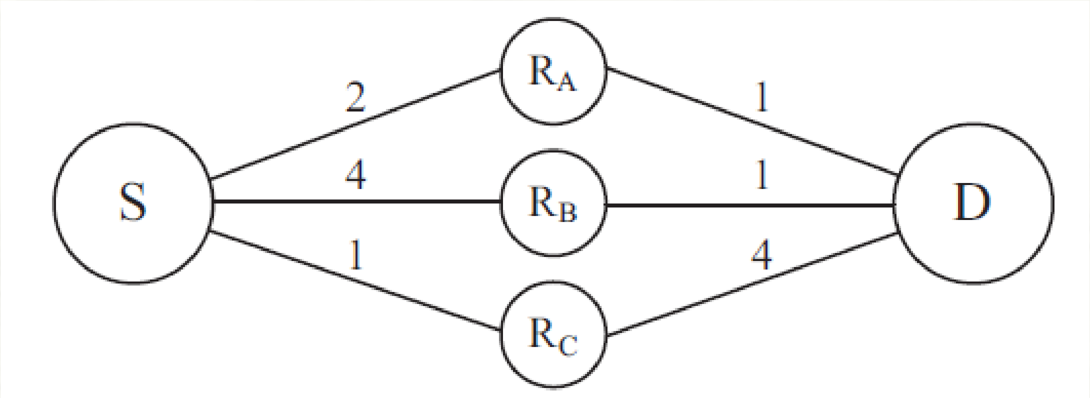

- The forwarding set of the 1st sent copy:
  - Should be{R_A,R_B,R_C}.
  - Enough copies are reserved.
  - Different from the single-copy case.

Introduction

Opportunistic two-hop routing (multi-copy case)

- Assume the source has 2 copies (a complex case).

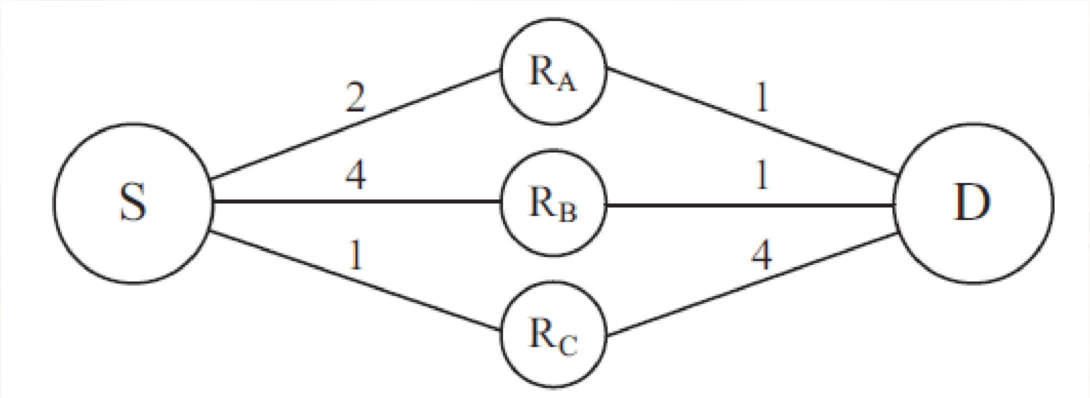

- The forwarding set of the 1st sent copy:
  - ❑ Should be{R_A,R_B,R_C} or {R_A,R_B}?
  - ❑ Not trivial.
  - ❑ The forwarding set of the 2nd sent copy?

Introduction

Opportunistic two-hop routing (multi-copy case)

- Assume the source has 2 copies (a complex case).

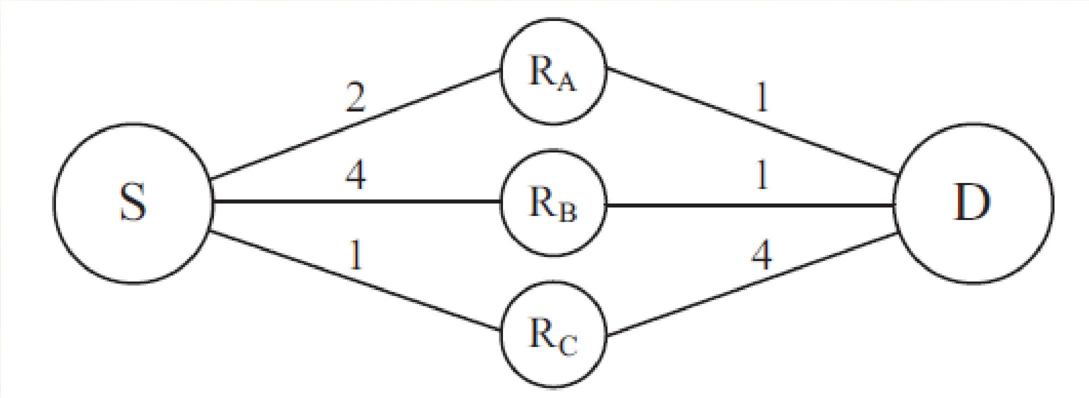

- Suppose the 1st sent copy uses $\{R_A, R_B, R_C\}$:
  - ❑ R_A takes the 1st copy.
  - ❑ R_B takes the 1st copy.
  - ❑ R_C takes the 1st copy.

Introduction

Opportunistic two-hop routing (multi-copy case)

- Assume the source has 2 copies, very complex.

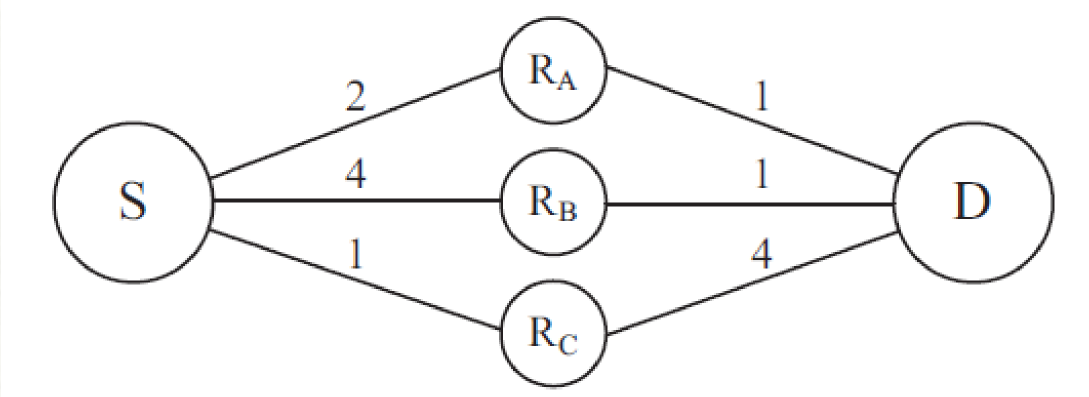

TABLE I. FORWARDING SET OUTLOOK

Forwarding set of the 1^{st} copy (when the source has 2 copies)	Forwarding set of the 2^{nd} copy (when the source has 1 copy)
$\{R_A, R_B, R_C\}$	R_A gets the 1^{st} copy $\Rightarrow \{R_B, R_C\}$
	R_B gets the 1^{st} copy $\Rightarrow \{R_A\}$
	R_C gets the 1^{st} copy $\Rightarrow \{R_A, R_B\}$

Introduction

Very challenging problem:

- To calculate the forwarding set of the current copy, we need to know the delay reduction brought by the remaining copies.

- To calculate the delay reduction brought by the remaining copies, we need to know the actual relay of the current copy, which is opportunistic.

Introduction

Tradeoff:

- If the forwarding set we selected for the current copy is too small, the subsequent copies will be blocked, losing the advantage of multiple copies.

- On the other hand, if the forwarding set we selected for the current copy is too large, this copy may end up choosing unqualified relays, i.e., this copy is useless.

Outline

- Model
- Insights and Solutions
- Extension
- Evaluation

Model

- Exponential distributed link delay.
  - The parameters for the first and second hops are denoted by λ and μ, respectively.

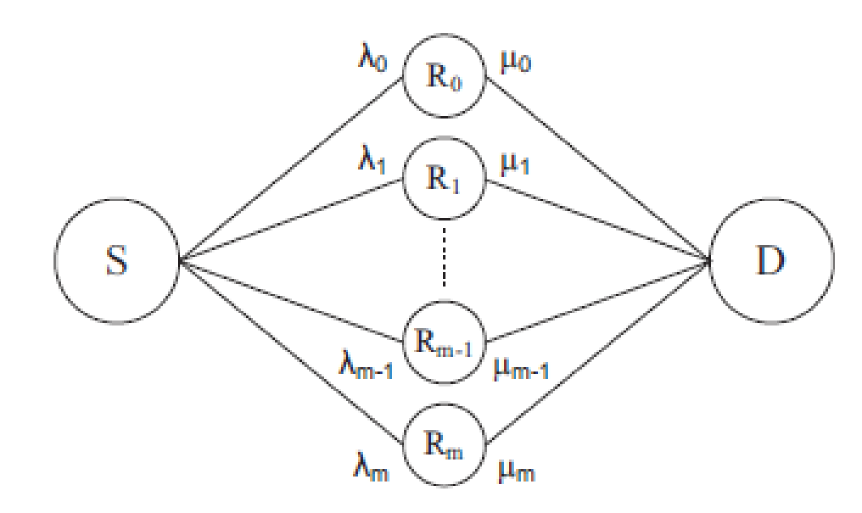

- Direct one-hop routing of S-D?
  - Equivalent to S-R_0-D, where R_0 and D has zero delay (meet infinite-frequently)

Model

- Let F_n denote the current forwarding set, then the expected delay with n copies is

$$E_n = \frac{1}{s_n}[1 + \sum_{i \in F_n} \frac{\lambda_i}{\mu_i}] - \sum_{i \in F_n} \frac{\lambda_i}{s_n} \int_0^{\infty} e^{-\mu_i t} H_{n-1}(t)dt$$

$$s_n = \sum_{i \in F_n} \lambda_i$$

- The former part is the expected delay of the first sent message copy (including the first hop delay and the second hop delay).

- The latter part is the decreased expected delay brought by the remaining n−1 copies.

Insights and Solutions

- For a relay node R_k, we can decide whether R_k is in the forwarding set or not, by comparing
  - The delivery delay of passing a copy to R_k -D path.
  - The delivery delay of not passing a copy (waiting for the other relays).

- This insight means a greedy optimal selection:

Theorem 1: If there are r ($r \gg n \geq 1$) residual relays that have not received a copy, where $\mu_{k_1} > \mu_{k_2} > \ldots > \mu_{k_r}$, then F_n^* satisfies $F_n^* = \{R_{k_1}, R_{k_2}, \ldots, R_{k_j}\}$ for a specified j in $[1, r]$.

Insights and Solutions

- How to deal with the decreased expected delay brought by the remaining copies?

$$E_n = \frac{1}{s_n}[1 + \sum_{i \in F_n} \frac{\lambda_i}{\mu_i}] - \sum_{i \in F_n} \frac{\lambda_i}{s_n} \int_0^{\infty} e^{-\mu_i t} H_{n-1}(t) dt$$

- Key insights:
- The second hop delay of the currently sent copy should have *the same order of magnitude with the delay* reduction brought by the remaining copies.

Insights and Solutions

- The second hop delay of the currently sent copy should have *the same order of magnitude with the delay* reduction brought by the remaining copies.

  - If the former one is the major delay, then we should select more qualified relays into the forwarding set of the current copy, i.e., remove unqualified relays.

  - On the other hand, if the latter one is the major issue, then we should sent out the first copy as soon as possible to take full advantage of subsequent copies.

Insights and Solutions

- Bounded solution:

$$E_n = \frac{1}{s_n}[1 + \sum_{i \in F_n} \frac{\lambda_i}{\mu_i}] - \sum_{i \in F_n} \frac{\lambda_i}{s_n} \int_0^{\infty} e^{-\mu_i t} H_{n-1}(t) dt$$

$$E_n < \frac{1}{s'_n}[\lambda_{max} + \frac{(2\mu_{max} E_1)^{\frac{1}{2^{n-1}}}}{2} \sum_{i \in F_n} \frac{\lambda_i^2}{\mu_i}]$$

$$s'_n = \sum_{i \in F_n} \lambda_i^2$$

- Greedily add the relay that has the smallest relay-destination delay into the forwarding set, until the above upper bound increases.

Extension

- Feature space routing:

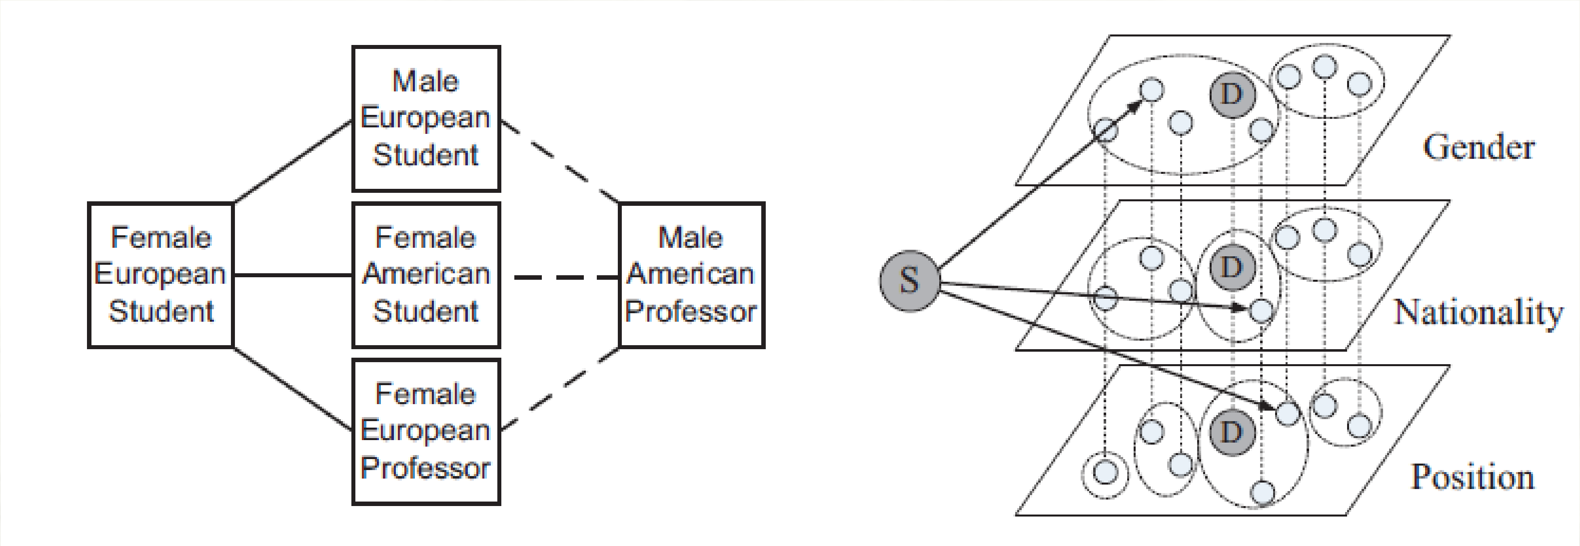

- Use the feature differences of two nodes to estimate their contact frequency.
- Then iteratively apply the two-hop routing.

Evaluation

- Synthetic trace
  - 30 relays between the source and the destination.
  - Uniform distributed contact frequency.
- Intel trace
  - 2-hop connected.
- MIT trace.
- Infocom06 trace.

Evaluation

- Two-hop routing algorithms for comparison:
  - Infinite Copies, where the source has infinite copies for two-hop routing. Infinite Copies shows the minimum data delivery delay of two-hop routing algorithms.
  - All New Paths, where the source always forwards one message copy to any inter-meeting relay nodes (if the source has remaining copies).
  - Repeated STRA, the source routes the n copies using single-copy two-hop routing algorithm recursively (the functionality of remaining copies is ignored).

Evaluation

- Other algorithms for comparison:
  - Epidemic, where the nodes continuously replicate and transmit messages to newly discovered contacts that do not already possess a copy. Epidemic represents the minimum data delivery delay of all routing algorithms.
  - (Binary) Spray and Wait, where is composed of a spray phase and wait phase.
  - SimBet where the relays are selected according to similarity and betweenness. Each message holder will give a copy to a inter-meeting relay if this relay does not hold a copy and has shorter feature distance with the destination. Only source holds multiple copies.
  - The feature space routing that is based on Repeated STRA (FSR-RSTRA for short).

Evaluation

- 2-hop routing algorithm in the synthetic trace.
  - ❑ Gain ratio is the delay reduction brought by using one more copy.

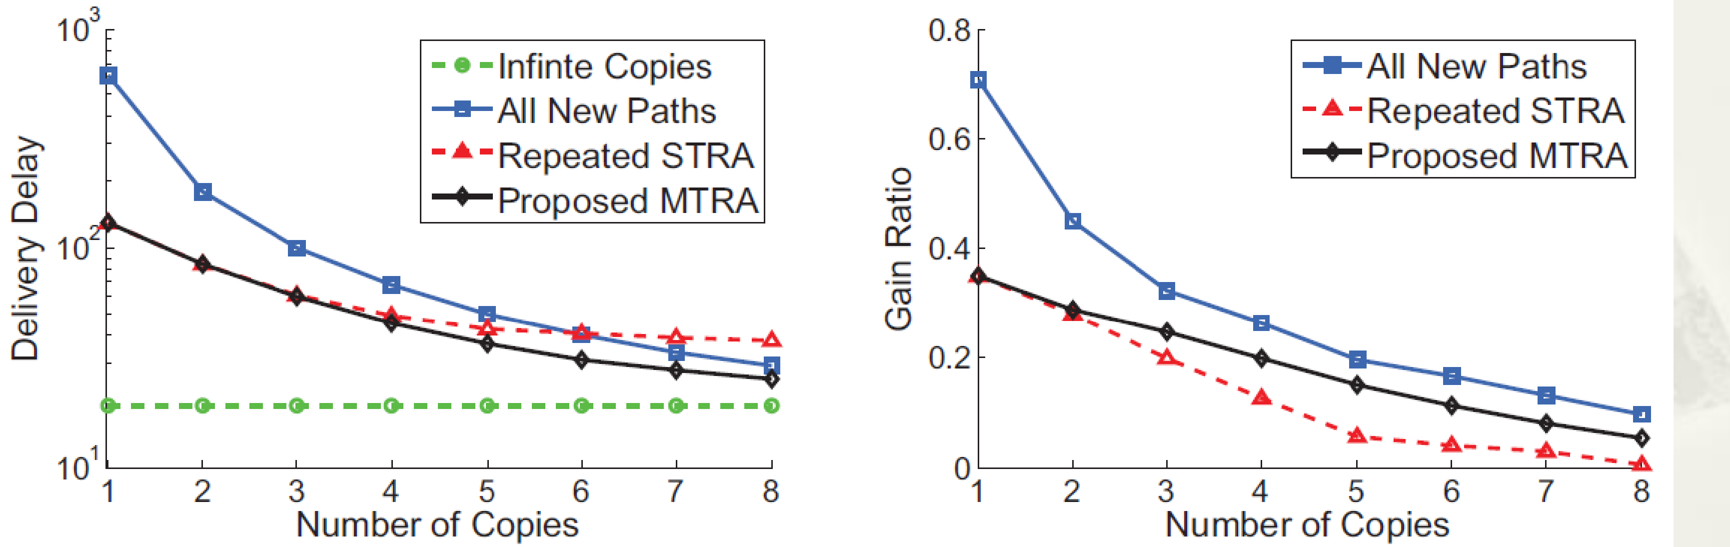

(a) Delivery Delay (synthetic trace)   (b) Gain Ratio (synthetic trace)

Evaluation

- 2-hop routing algorithm in the Intel trace.
  - ❑ Gain ratio is the delay reduction brought by using one more copy.

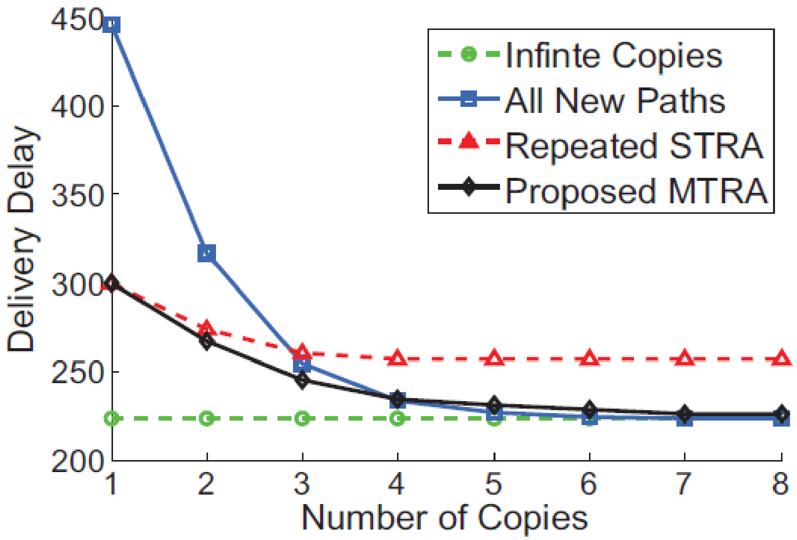

(c) Delivery Delay (Intel trace)

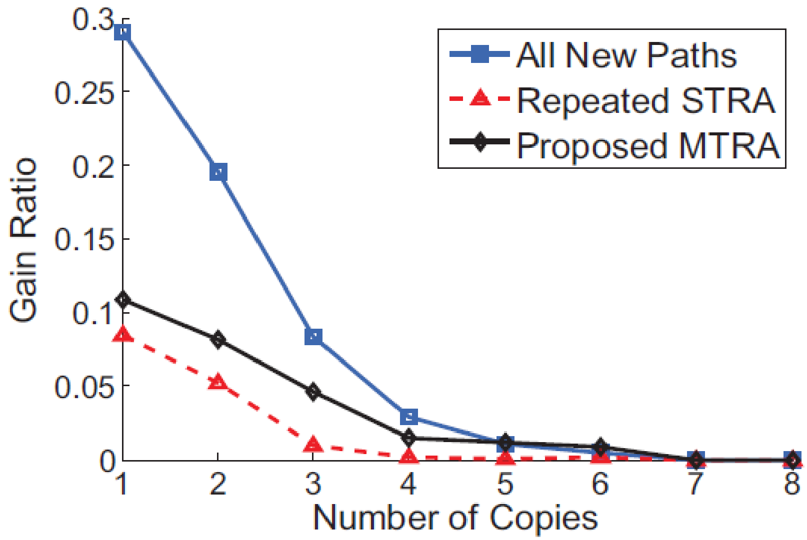

(d) Gain Ratio (Intel trace)

Evaluation

- Feature space routing

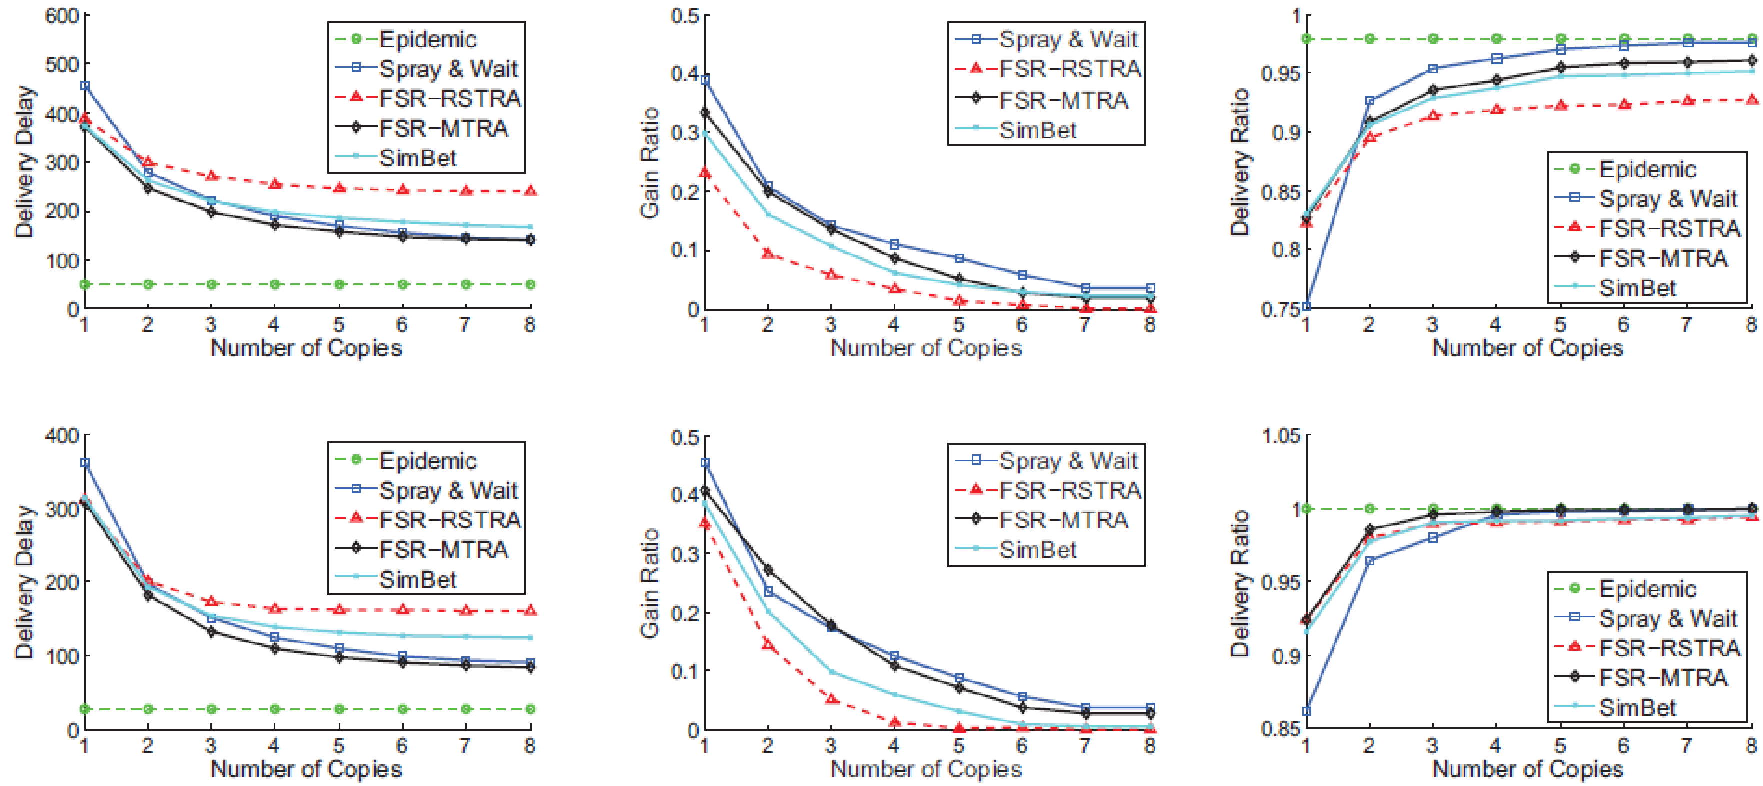

(a) Delivery Delay (b) Gain Ratio (c) Delivery Ratio

Fig. 7. The feature space routing in MIT trace (top) and Infocom 2006 trace (bottom). Here, Epidemic and Spray & Wait are not related to the feature information, while FSR-RSTRA, FSR-MTRA, and SimBet are feature-based routing algorithms.

Conclusion

- A multi-copy two-hop routing algorithm (MTRA) is proposed with a performance bound.
- All the forwarding sets for the n copies can be efficiently determined with a time complexity of O(mlogm+nm), where m is the number of available relays.
- MTRA can be further applied to a feature space routing scheme, where the contact frequencies are estimated by feature distances.
- Simulation results show competitive performances of the proposed algorithms, which fully utilize the opportunistic nature of MSNs.

The End

Questions & Answer