Seeker-Assisted Information Search in Mobile Clouds

Suzan Bayhan[†], Esa Hyytiä[‡], Jussi Kangasharju[*], and Jörg Ott[‡]

[†]Helsinki Institute for Information Technology (HIIT), Aalto University, Finland

[‡]Aalto University, School of Electrical Engineering, Finland

[*]Department of Computer Science, University of Helsinki, Finland

✉ bayhan@hiit.fi

Aug.12, 2013, Mobile Cloud Computing (MCC'13)

Mobile Cloud

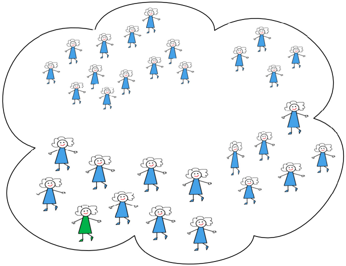

- *Proximate mobile computing entities* may form a mobile cloud.
- Mobile data traffic predicted to 3x by 2017, mobile UGC ↑, mobile storage capacity ↑

Mobile Cloud

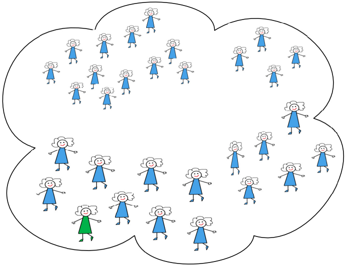

- *Proximate mobile computing entities* may form a mobile cloud.
- Mobile data traffic predicted to 3x by 2017, mobile UGC ↑, mobile storage capacity ↑

Increasing volume of data calls for efficient information search schemes in the mobile cloud

Mobile Cloud Under Intermittent Connectivity

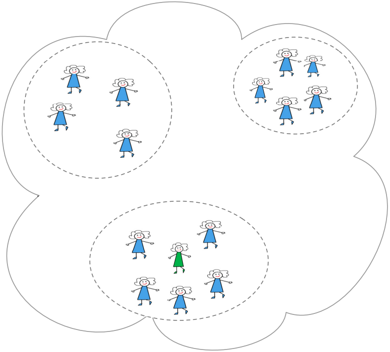

How to form the mobile cloud in this network?

- No infrastructure
- No end-to-end connectivity

Mobile Cloud Under Intermittent Connectivity

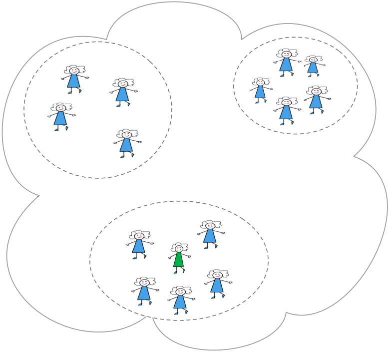

How to form the mobile cloud in this network?

- No infrastructure
- No end-to-end connectivity

But nodes are mobile!

- Exploit mobile nodes as message carriers
- Multiple message copies to increase probability of delivery

Mobile Cloud Under Intermittent Connectivity

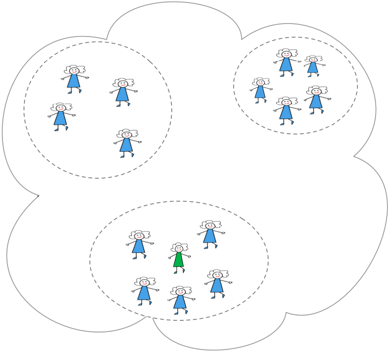

How to form the mobile cloud in this network?

- No infrastructure
- No end-to-end connectivity

But nodes are mobile!

- Exploit mobile nodes as message carriers
- Multiple message copies to increase probability of delivery

Delay/Disruption tolerant networking (DTN)

Outline

1. Background on Delay/Disruption Tolerant Networking (DTN)
   - Information Search in a DTN
2. Seeker-Assisted Search (SAS)
   - Model and Assumptions
   - Performance Analysis
3. Summary and Future Directions

An Example of DTN: Time Tick 0

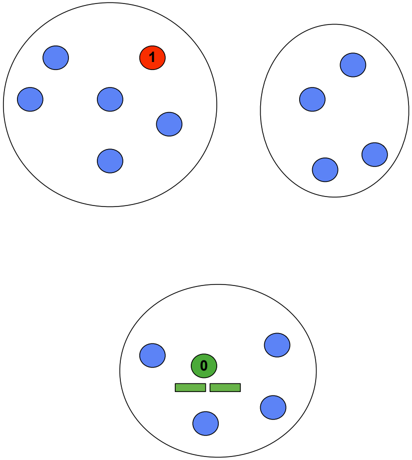

- t=0: 2 copies of the message

An Example of DTN: Time Tick 1

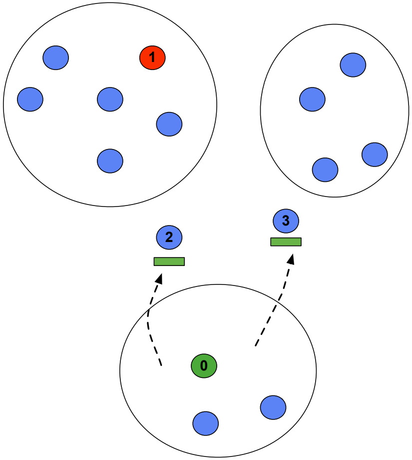

- t=0: 2 copies of the message
- t=1: One copy to 2 and 3

An Example of DTN: Time Tick 2

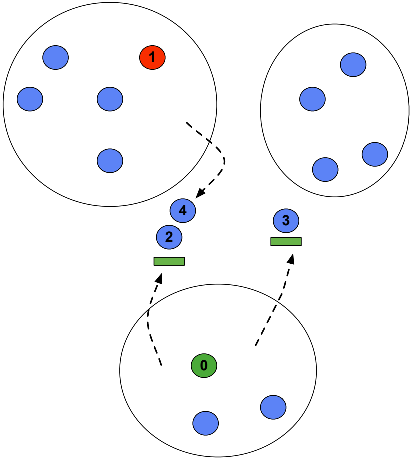

- t=0: 2 copies of the message
- t=1: One copy to 2 and 3
- t=2: 2 meets 4 and 3 moves

An Example of DTN: Time Tick 3

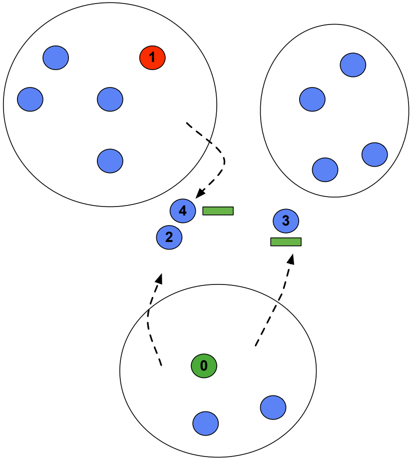

- t=0: 2 copies of the message
- t=1: One copy to 2 and 3
- t=2: 2 meets 4 and 3 moves
- t=3: 2 transmits to 4

An Example of DTN: Time Tick 4

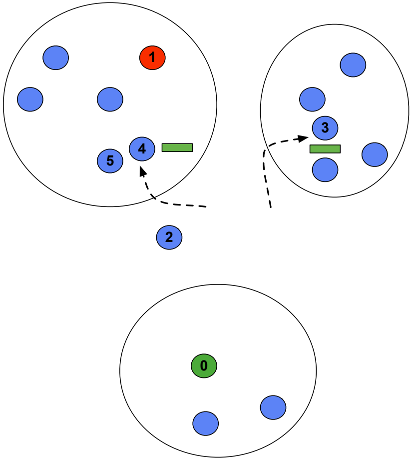

- t=0: 2 copies of the message
- t=1: One copy to 2 and 3
- t=2: 2 meets 4 and 3 moves
- t=3: 2 transmits to 4
- t=4: 4 and 3 carries the message

An Example of DTN: Time Tick 5

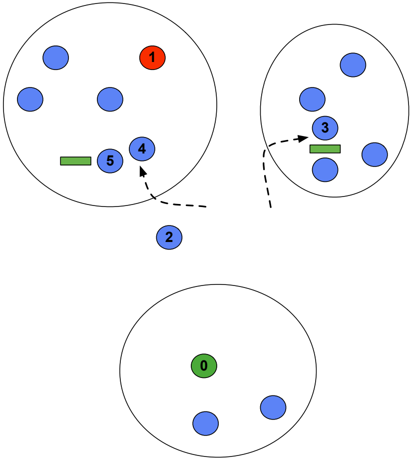

- t=0: 2 copies of the message
- t=1: One copy to 2 and 3
- t=2: 2 meets 4 and 3 moves
- t=3: 2 transmits to 4
- t=4: 4 and 3 carries the message
- t=5: 4 transmits to 5

An Example of DTN: Time Tick 6

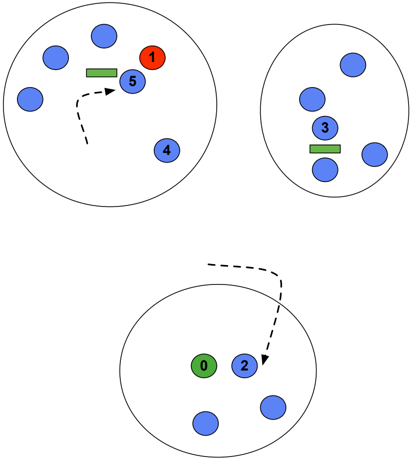

- t=0: 2 copies of the message
- t=1: One copy to 2 and 3
- t=2: 2 meets 4 and 3 moves
- t=3: 2 transmits to 4
- t=4: 4 and 3 carries the message
- t=5: 4 transmits to 5
- t=6: 5 meets 1

An Example of DTN: Time Tick 7

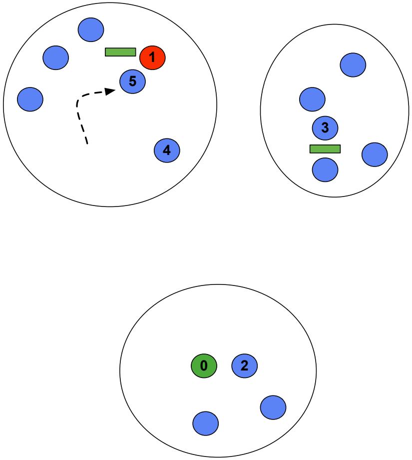

- t=0: 2 copies of the message
- t=1: One copy to 2 and 3
- t=2: 2 meets 4 and 3 moves
- t=3: 2 transmits to 4
- t=4: 4 and 3 carries the message
- t=5: 4 transmits to 5
- t=6: 5 meets 1
- t=7: Transmission is completed

In short

- Communication for weakly-connected nodes is still possible by Delay Tolerant Networking, DTN
- Even if mobile nodes have loose connectivity, they can form a mobile cloud using store-carry-forward approach

Information Search in a DTN

Challenges of search:

- No Google-style database
- High volume of user generated content
- Mystery of the location of the searched content

Easiest way:

- Epidemic: Copy the message to each encountered node
- Direct delivery: Wait till meeting the destination (no replication!)

Information Search in a DTN

Challenges of search:

- No Google-style database
- High volume of user generated content
- Mystery of the location of the searched content

Easiest way:

- Epidemic: Copy the message to each encountered node
- Direct delivery: Wait till meeting the destination (no replication!)

Trade-offs:

- Search completion time
- Search overhead: Replication ratio

Our aim is to design a search scheme considering the tradeoffs

Outline

1. Background on Delay/Disruption Tolerant Networking (DTN)
   - Information Search in a DTN
2. Seeker-Assisted Search (SAS)
   - Model and Assumptions
   - Performance Analysis
3. Summary and Future Directions

Our Model and Assumptions

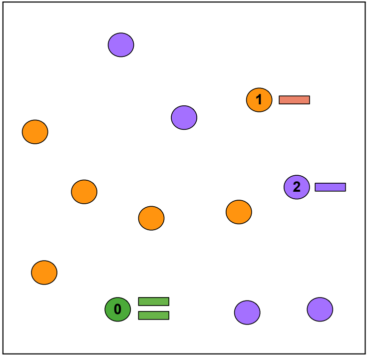

- A DTN-based mobile cloud
- People with similar interests form a community (group of nodes with noticeably higher in-group interactions)
- Community 1 (C_1), Community 2 (C_2)
- Searching node (in C_1)
- Some of the nodes in both communities own the searched content

Our Basic (not-so-unrealistic!) Assumptions

- People search information based on their interests
- The searched item is stored in the same community with higher probability
- People in the same community meet more frequently (i.e., homophily principle)

Five Node Types

1. Searching node: n_s
2. Tagged node (T): A node holding the searched content
3. Seeker node (S): A node holding a copy of the query
4. Tagged seeker node (TS): A node with both the content and the query
5. Passive node (P): Rest of the nodes

A node's type may change upon encounters and the actions triggered by the search scheme

Seeker-Assisted Search (SAS)

Forward Path:

- Seekers assist the searching node by
  - carrying the query and forwarding it
  - replicating the searched content if encounters a tagged node
- Push the query towards *more useful nodes* $\rightarrow$ same community members as the searching node

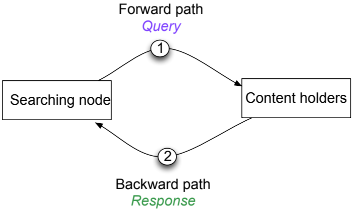

Backward path:

- Response delivered in single hop

An Example Search Scenario

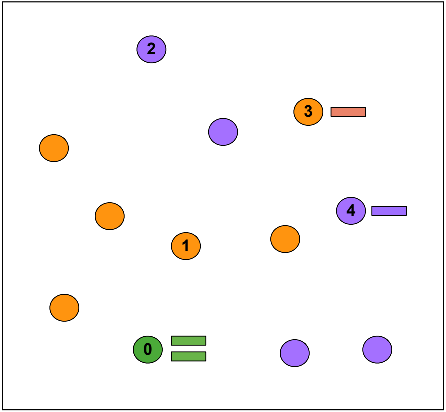

History

$t = 0$: 2 message copies

Passive ⇒ Seeker ⇒ Tagged Seeker ⇒ Completed.

An Example Search Scenario

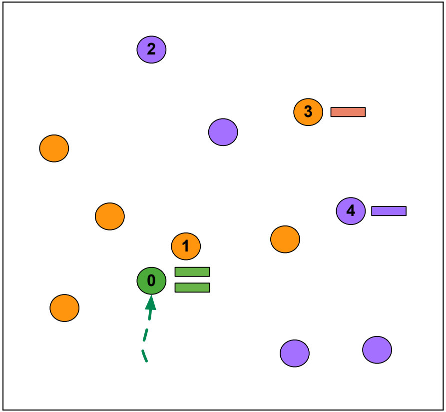

History

$t = 0$: 2 message copies

$t = 1$: n_s meets 1

Passive ⇒ Seeker ⇒ Tagged Seeker ⇒ Completed.

An Example Search Scenario

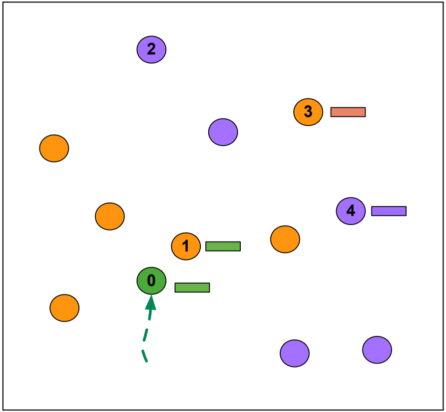

History

$t = 0$: 2 message copies

$t = 1$: n_s forwards to 1

Passive ⇒ **Seeker** ⇒ Tagged Seeker ⇒ Completed.

An Example Search Scenario

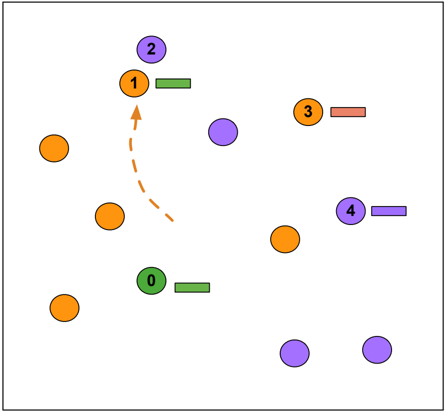

History

$t = 0$: 2 message copies

$t = 1$: n_s forwards to 1

$t = 2$: 1 meets 2

Passive ⇒ **Seeker** ⇒ Tagged Seeker ⇒ Completed.

An Example Search Scenario

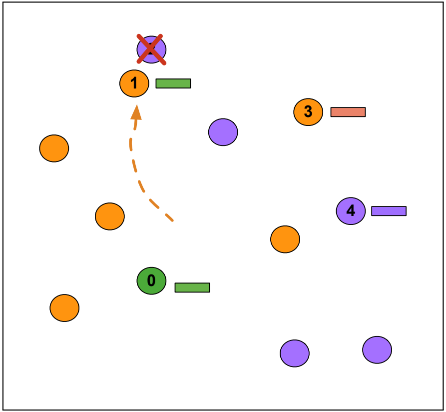

History

$t = 0$: 2 message copies

$t = 1$: n_s forwards to 1

$t = 2$: 1 meets 2, no forwarding

Passive ⇒ **Seeker** ⇒ Tagged Seeker ⇒ Completed.

An Example Search Scenario

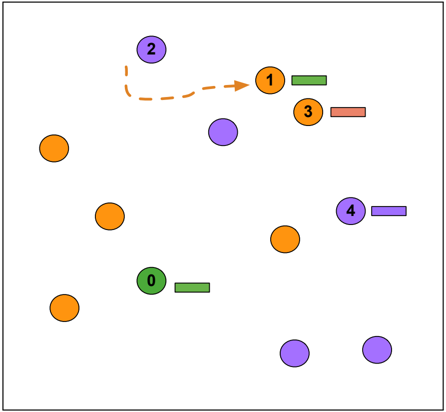

History

$t = 0$: 2 message copies

$t = 1$: n_s forwards to 1

$t = 2$: 1 meets 2, no forwarding

$t = 3$: 1 meets 3

Passive ⇒ **Seeker** ⇒ Tagged Seeker ⇒ Completed.

An Example Search Scenario

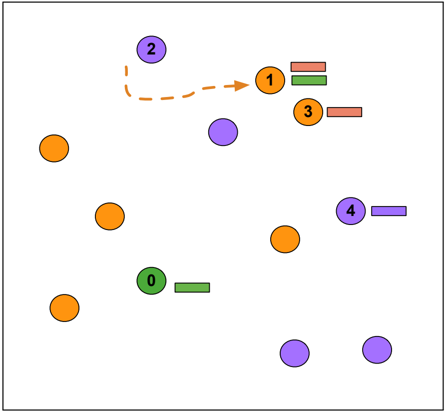

History

$t = 0$: 2 message copies

$t = 1$: n_s forwards to 1

$t = 2$: 1 meets 2, no forwarding

$t = 3$: 1 meets 3, replicates the content

Passive ⇒ Seeker ⇒ **Tagged Seeker** ⇒ Completed.

An Example Search Scenario

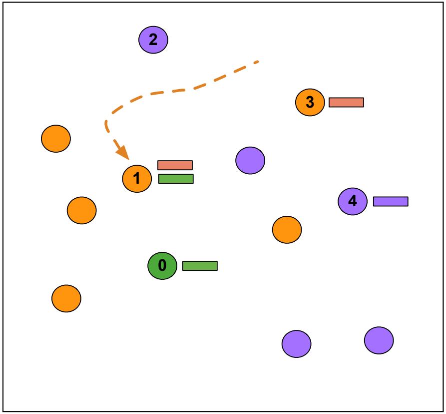

History

$t = 0$: 2 message copies

$t = 1$: n_s forwards to 1

$t = 2$: 1 meets 2, no forwarding

$t = 3$: 1 meets 3, replicates the content

$t = 4$: 1 carries the content and the query

Passive ⇒ Seeker ⇒ **Tagged Seeker** ⇒ Completed.

An Example Search Scenario

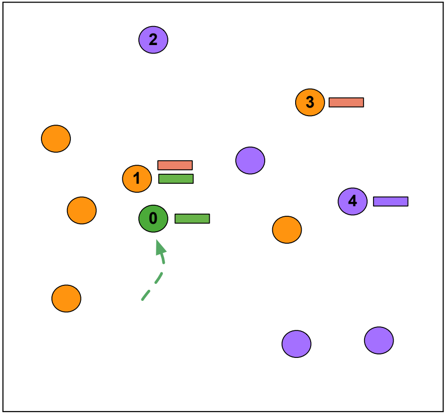

History
$t = 0$: 2 message copies
$t = 1$: n_s forwards to 1
$t = 2$: 1 meets 2, no forwarding
$t = 3$: 1 meets 3, replicates the content
$t = 4$: 1 carries the content and the query
$t = 5$: 1 meets n_s

Passive ⇒ Seeker ⇒ **Tagged Seeker** ⇒ Completed

An Example Search Scenario

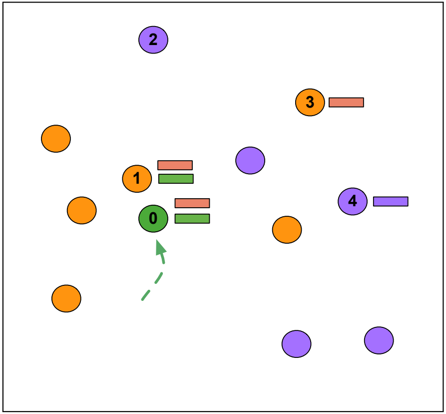

History
$t = 0$: 2 message copies
$t = 1$: n_s forwards to 1
$t = 2$: 1 meets 2, no forwarding
$t = 3$: 1 meets 3, replicates the content
$t = 4$: 1 carries the content and the query
$t = 5$: Search ends

Passive ⇒ Seeker ⇒ Tagged Seeker ⇒ **Completed**

Continuous-time Markov Modeling

In-community and external-community meetings:
Inter-contact times $\sim Exp(\mu_i)$ and $Exp(\mu_x)$

Our model:

- n_i nodes in C_i, $i = 1, 2$
- In community meetings are more frequent, $\mu_i > \mu_x$
- Only $M + 1$ copies of the query (max number of seekers)

Continuous-time Markov Modeling

In-community and external-community meetings:
Inter-contact times $\sim Exp(\mu_i)$ and $Exp(\mu_x)$

Our model:

- n_i nodes in C_i, $i = 1, 2$
- In community meetings are more frequent, $\mu_i > \mu_x$
- Only $M + 1$ copies of the query (max number of seekers)

$$z = (m_1, k_1, mk_1, m_2, k_2, mk_2, c)$$

- m_i: # of tagged nodes in C_i
- k_i: # of seeker nodes in C_i
- mk_i: # of tagged seeker nodes in C_i
- c: # of remaining copies at n_s other than its own copy

Recursive solution for the Markov process

Performance Metrics

- Search completion time:
  What is the average time to meet a tagged node?

- Overhead:
  What is the replication ratio before search completion?

State Transitions

$$z = (m_1, k_1, mk_1, m_2, k_2, mk_2, c)$$

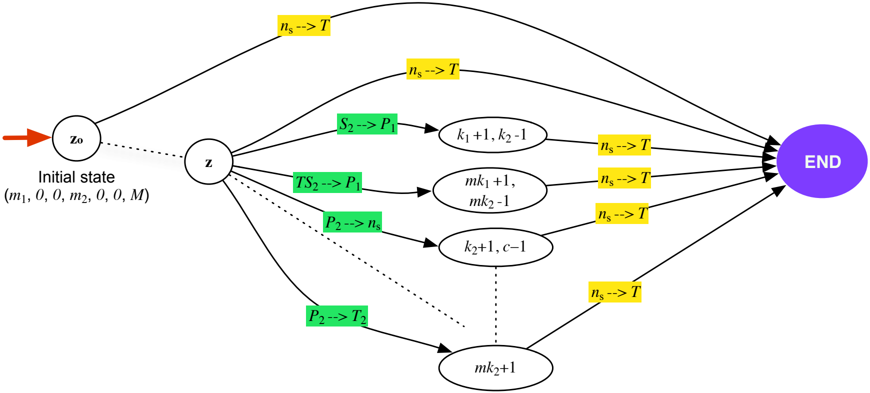

n_s: Searching node, S_i: Seeker (k_i), P_i: Passive
T_i: Tagged, TS_i: Tagged seeker (mk_i), i: Community index

Performance Analysis

We analyze the effect of:

- In-community/external-community meetings
- Location of the searched content, in C_1 or C_2, or both
- Network population

Performance Analysis

We analyze the effect of:

- In-community/external-community meetings
- Location of the searched content, in C_1 or C_2, or both
- Network population

We use the following bounds for comparison:

- Upper bound for performance and overhead
  - Epidemic search
  - No limits on replication, $M = \infty$
- Lower bound
  - Direct delivery
  - Search completes only when n_s meets tagged nodes, i.e., no seekers, $M = 0$

Effect of Average Inter-contact Times (μ_x/μ_i)

Setting: 8 nodes in each community, one content only in C_1

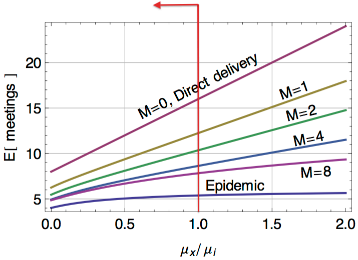

(a) Response time

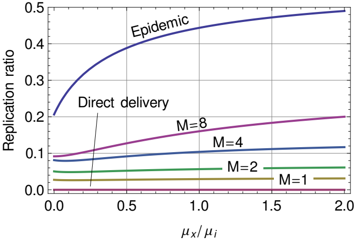

(b) Search overhead

Effect of Average Inter-contact Times (μ_x/μ_i)

Setting: 8 nodes in each community, one content only in C_1

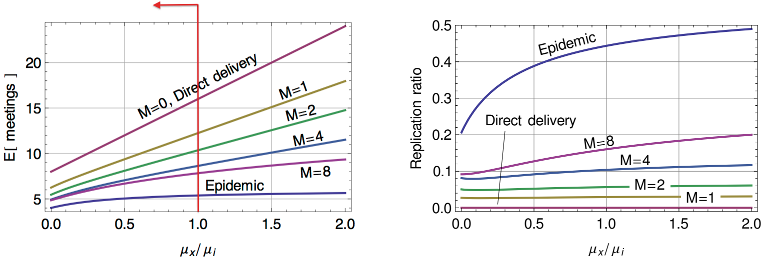

(a) Response time

(b) Search overhead

- With increasing M, the performance of SAS $\rightarrow$ Epidemic

Effect of Average Inter-contact Times (μ_x/μ_i)

Setting: 8 nodes in each community, one content only in C_1

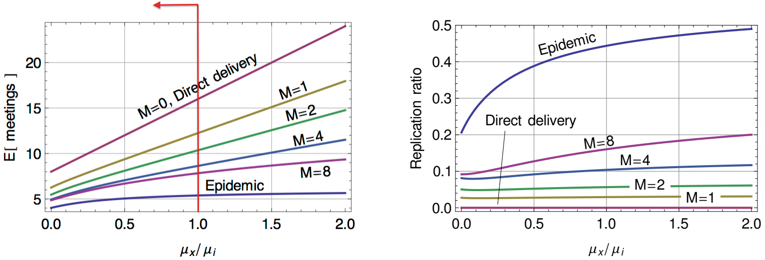

(a) Response time

(b) Search overhead

- With increasing M, the performance of SAS $\rightarrow$ Epidemic
- Epidemic spreads the content to almost 50% while SAS to 10% of nodes for $M = 4$
- As μ_x/μ_i increases, search time $\uparrow$ but cost is almost insensitive

Location of the Content

- Case 1: 2 copies in C_1
- Case 2: 1 copy in C_1, 1 copy in C_2

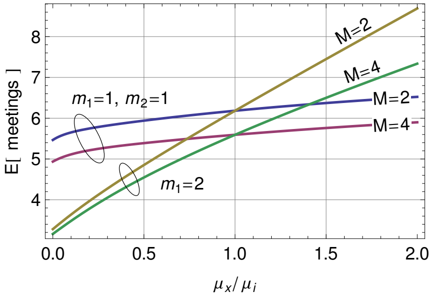

- $\mu_x/\mu_i < 1 \Rightarrow$ Shorter search time for $m_1 = 2$
- As inter-community meetings $\uparrow$, the initial location of the content becomes less important

Mean search time if our same-community assumption does not always hold

What if only C_2 has the content?

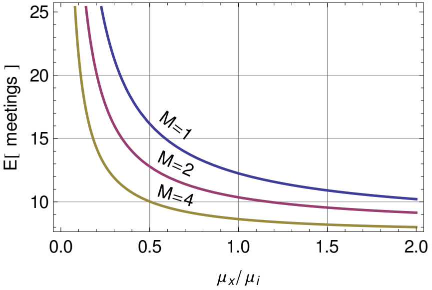

- For low μ_x/μ_i, long time to reach the content
- For larger μ_x/μ_i, search speed increases significantly

Effect of network population

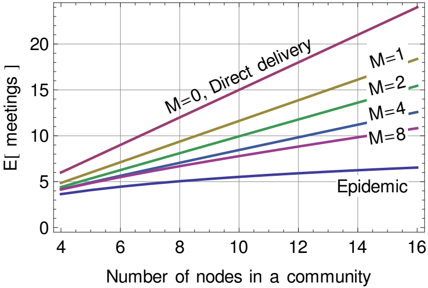

- Response time scales approximately linearly with the increasing network population

Outline

1. Background on Delay/Disruption Tolerant Networking (DTN)
   - Information Search in a DTN
2. Seeker-Assisted Search (SAS)
   - Model and Assumptions
   - Performance Analysis
3. Summary and Future Directions

Summary and Future Directions

- Mobile cloud under intermittent connectivity
- Seeker-Assisted Search (SAS)

Summary and Future Directions

- Mobile cloud under intermittent connectivity
- Seeker-Assisted Search (SAS)
- Future directions:
  - Human mobility and social properties
  - Exploit the relationship between contents and users
  - Simulation-based analysis using real mobility traces

Summary and Future Directions

- Mobile cloud under intermittent connectivity
- Seeker-Assisted Search (SAS)
- Future directions:
  - Human mobility and social properties
  - Exploit the relationship between contents and users
  - Simulation-based analysis using real mobility traces

Thank you!!!

Suzan Bayhan

http://www.hiit.fi/u/bayhan

State Transitions from $(m_1, k_1, mk_1, m_2, k_2, mk_2, c)$

Community 1 intra-community meetings					
	n_s	P_1	S_1	T_1	TS_1
P_1	$k_1+1, c-1$	-	-	-	-
S_1	-	-	-	mk_1+1	mk_1+1
T_1	S_{end}	-	mk_1+1	-	-
TS_1	S_{end}	-	mk_1+1	-	-
Community 2 intra-community meetings					
		P_2	S_2	T_2	TS_2
P_2	-	-	-	-	-
S_2	-	-	-	mk_2+1	mk_2+1
T_2	-	-	mk_2+1	-	-
TS_2	-	-	mk_2+1	-	-
Inter-community meetings					
	n_s	P_1	S_1	T_1	TS_1
P_2	$k_2+1, c-1$	-	-	-	-
S_2	-	k_1+1, k_2-1	-	mk_2+1	mk_2+1
T_2	S_{end}	-	mk_2+1	-	-
TS_2	S_{end}	mk_1+1, mk_2-1	mk_1+1	-	-

Attention to P_1 meets S_2 and TS_2 vs. P_2 meets S_1 and TS_1.

Overview of Related Works

Hui *et al.* [1]: File sharing using Osmosis principle

- (Forward path) Epidemic (Backward path) *Osmosis*

[1] Pan Hui, J. Leguay, J. Crowcroft, J. Scott, T.Friedman, and V. Conan. Osmosis in Pocket Switched Networks. In IEEE ChinaCom, 2006.

[2] Mikko Pitkänen, T. Kärkkäinen, J. Greifenberg, and J. Ott. Searching for content in mobile DTNs, IEEE PerCom, 2009.

[3] Jialu Fan, J. Chen, Y. Du, P.Wang, and Y.S.Delque: A socially aware delegation query scheme in delay-tolerant networks. IEEE Trans. on Vehicular Technology, 60(5):2181–2193, 2011.

Overview of Related Works

Hui *et al.* [1]: File sharing using Osmosis principle

- (Forward path) Epidemic (Backward path) *Osmosis*
- We control flooding also in the forward path

---

[1] Pan Hui, J. Leguay, J. Crowcroft, J. Scott, T.Friedman, and V. Conan. Osmosis in Pocket Switched Networks. In IEEE ChinaCom, 2006.

[2] Mikko Pitkänen, T. Kärkkäinen, J. Greifenberg, and J. Ott. Searching for content in mobile DTNs, IEEE PerCom, 2009.

[3] Jialu Fan, J. Chen, Y. Du, P.Wang, and Y.S.Delque: A socially aware delegation query scheme in delay-tolerant networks. IEEE Trans. on Vehicular Technology, 60(5):2181–2193, 2011.

Overview of Related Works

Hui *et al.* [1]: File sharing using Osmosis principle

- (Forward path) Epidemic (Backward path) *Osmosis*
- We control flooding also in the forward path

Pitkänen *et al.* [2]: When to terminate the query?

- Hop-count, TTL, global-response count estimate,

---

[1] Pan Hui, J. Leguay, J. Crowcroft, J. Scott, T.Friedman, and V. Conan. Osmosis in Pocket Switched Networks. In IEEE ChinaCom, 2006.

[2] Mikko Pitkänen, T. Kärkkäinen, J. Greifenberg, and J. Ott. Searching for content in mobile DTNs, IEEE PerCom, 2009.

[3] Jialu Fan, J. Chen, Y. Du, P.Wang, and Y.S.Delque: A socially aware delegation query scheme in delay-tolerant networks. IEEE Trans. on Vehicular Technology, 60(5):2181–2193, 2011.

Overview of Related Works

Hui *et al.* [1]: File sharing using Osmosis principle

- (Forward path) Epidemic (Backward path) *Osmosis*
- We control flooding also in the forward path

Pitkänen *et al.* [2]: When to terminate the query?

- Hop-count, TTL, global-response count estimate,
- No termination policy since we seek the time for getting a response with probability 1

Fan *et al.* [3]: Dynamic geo-community concept

- Our community concept is location-independent

---

[1] Pan Hui, J. Leguay, J. Crowcroft, J. Scott, T.Friedman, and V. Conan. Osmosis in Pocket Switched Networks. In IEEE ChinaCom, 2006.

[2] Mikko Pitkänen, T. Kärkkäinen, J. Greifenberg, and J. Ott. Searching for content in mobile DTNs, IEEE PerCom, 2009.

[3] Jialu Fan, J. Chen, Y. Du, P.Wang, and Y.S.Delque: A socially aware delegation query scheme in delay-tolerant networks. IEEE Trans. on Vehicular Technology, 60(5):2181–2193, 2011.