Tight Bounds on Information Spreading in Sparse Mobile Networks[1]

Alberto Pettarin Andrea Pietracaprina Geppino Pucci
Università di Padova, Padova, Italy

Eli Upfal
Brown University, Providence, RI

---

[1] Presented at PODC 2011

Motivation

Mobile networks are the emerging paradigm of distributed systems

Mobile Devices Networks

Vehicular Networks

Wildlife Surveillance Systems

Field Operations

Mobile Networks: a Closer Look

Mobile networks are distributed systems

- dynamic: topology changes over time…
- … but not too fast: mobility speed $\ll$ transmission speed
- with no infrastructure: wireless, multi-hop communications
- under energy constraints: small transmission radius
- essentially planar

Mobile Networks: a Closer Look

Mobile networks are distributed systems

- dynamic: topology changes over time...
- ... but not too fast: mobility speed $\ll$ transmission speed
- with no infrastructure: wireless, multi-hop communications
- under energy constraints: small transmission radius
- essentially planar

TCS perspective

Analytical characterization of information spreading in mobile networks

Previous Work

- Alves *et al.* [Ann.App.Pr.'02] and Kesten *et al.* [Ann.Pr.'05]
  - shape of the subspace of $\mathbb{Z}^d$ containing "infected" RWs (Frog Model)
- Dimitriou *et al.* [Dis.App.Mat.'06]
  - k agents performing RWs on an n-node graph, bounds on the expected infection time, depending on graph expansion
- Clementi *et al.* [ICALP'09, IPDPS'09]
  - $k = \Theta(n)$ agents on a n-node 2D grid (dense scenario) large maximum speed R and/or large transmission radius r
  - bounds on broadcast time
- Peres *et al.* [SODA'11]
  - Poisson point process in $\mathbb{R}^d$ above percolation threshold agents follow Brownian motion
  - bounds on detection, coverage, broadcast time

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid w/ loops

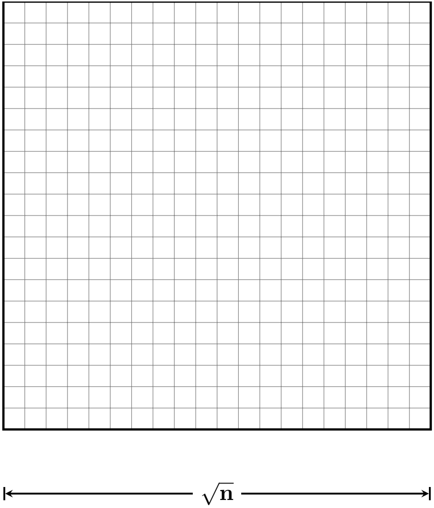

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid w/ loops
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution ($\Rightarrow$ uniform)

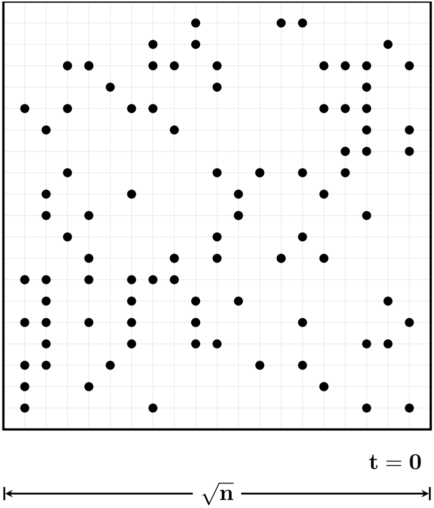

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid w/ loops
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution ($\Rightarrow$ uniform)
- Independent, simple, discrete-time random walks
- $pos_a(t) \equiv$ position of agent a at time $t \in \mathbb{N}$

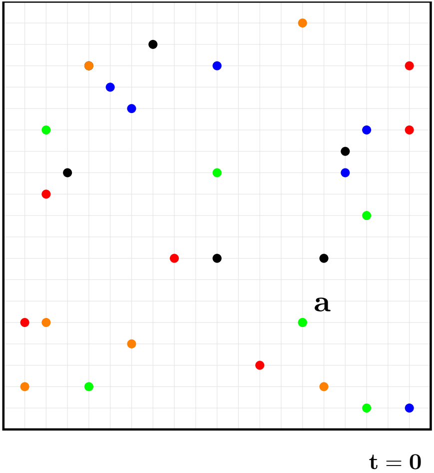

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid w/ loops
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution ($\Rightarrow$ uniform)
- Independent, simple, discrete-time random walks
- $pos_a(t) \equiv$ position of agent a at time $t \in \mathbb{N}$

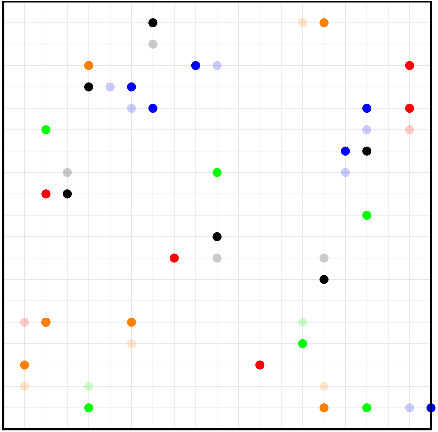

$\mathbf{t = 1}$

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid w/ loops
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution ($\Rightarrow$ uniform)
- Independent, simple, discrete-time random walks
- $pos_a(t) \equiv$ position of agent a at time $t \in \mathbb{N}$

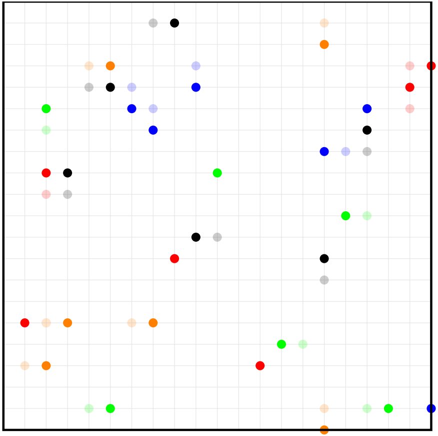

$\mathbf{t = 2}$

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid w/ loops
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution ($\Rightarrow$ uniform)
- Independent, simple, discrete-time random walks
- $pos_a(t) \equiv$ position of agent a at time $t \in \mathbb{N}$

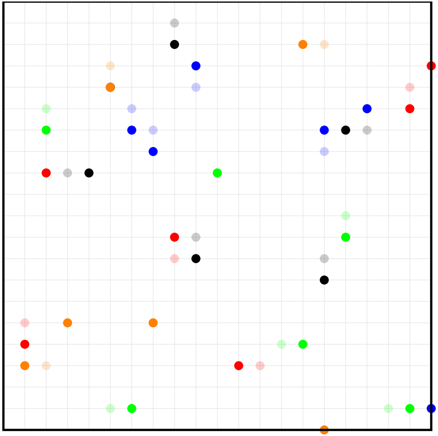

$\mathbf{t = 3}$

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid w/ loops
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution ($\Rightarrow$ uniform)
- Independent, simple, discrete-time random walks
- $pos_a(t) \equiv$ position of agent a at time $t \in \mathbb{N}$

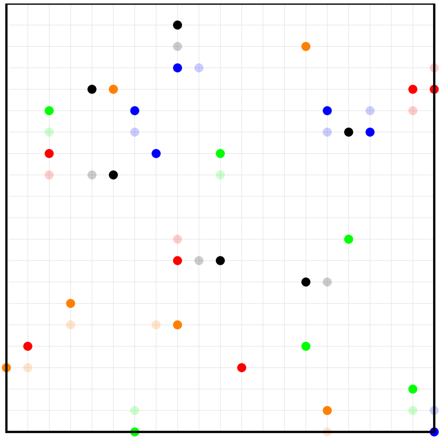

$\mathbf{t = 4}$

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid w/ loops
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution ($\Rightarrow$ uniform)
- Independent, simple, discrete-time random walks
- $pos_a(t) \equiv$ position of agent a at time $t \in \mathbb{N}$

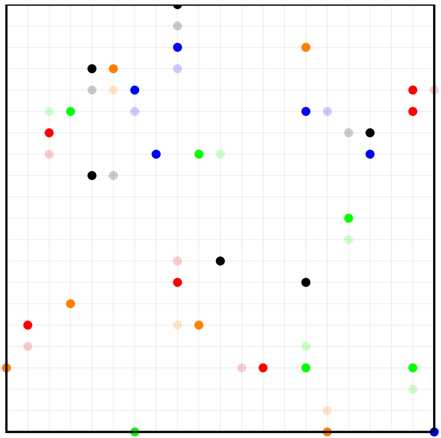

$\mathbf{t = 5}$

Communication Model

- Each agent has transmission radius $r \geqslant 0$

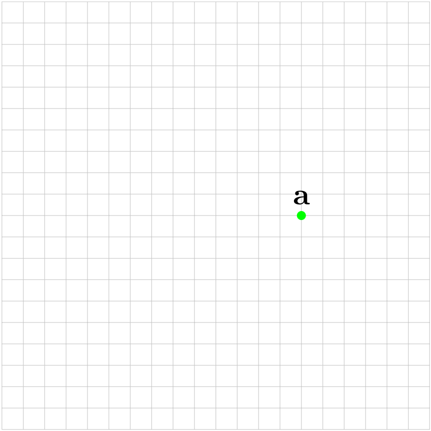

Communication Model

- Each agent has transmission radius $r \geqslant 0$

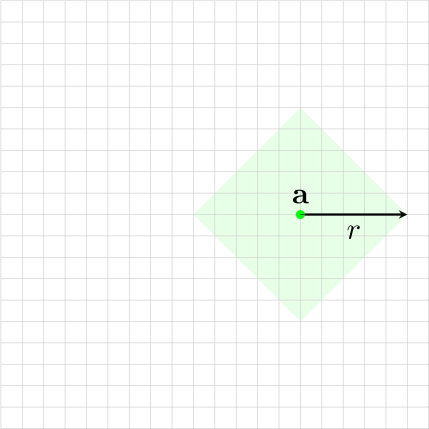

Communication Model

- Each agent has transmission radius $r \geqslant 0$

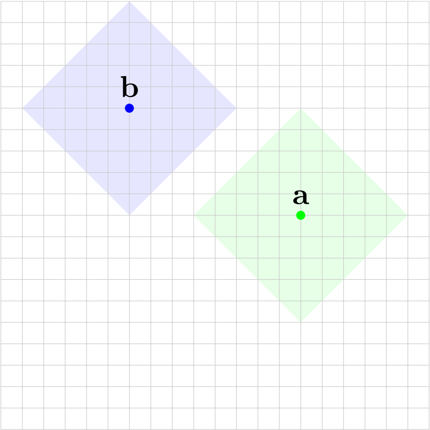

Communication Model

- Each agent has transmission radius $r \geqslant 0$

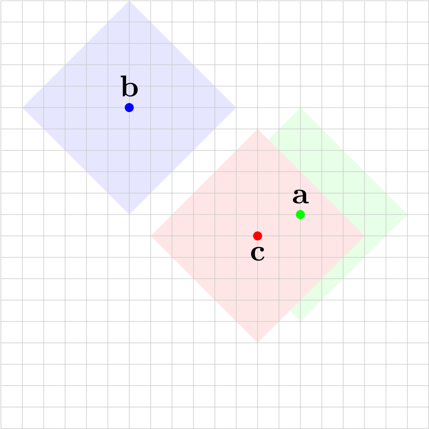

Communication Model

- Each agent has transmission radius $r \geqslant 0$
- Visibility graph $G_t(r)$:
  - vertices $\equiv$ agents

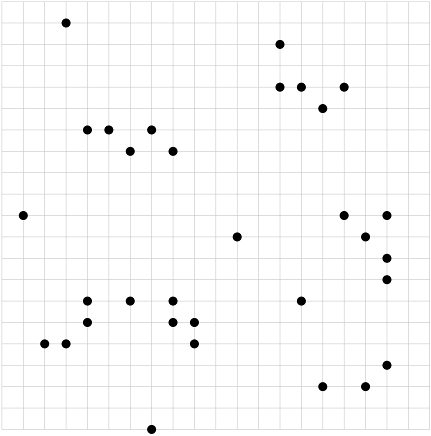

Communication Model

- Each agent has transmission radius $r \geqslant 0$
- Visibility graph $G_t(r)$:
  - vertices $\equiv$ agents
  - edge $\{a, a'\} \in G_t(r) \iff \| pos_a(t) - pos_{a'}(t) \| \leqslant r$
  - each connected component is called “island”

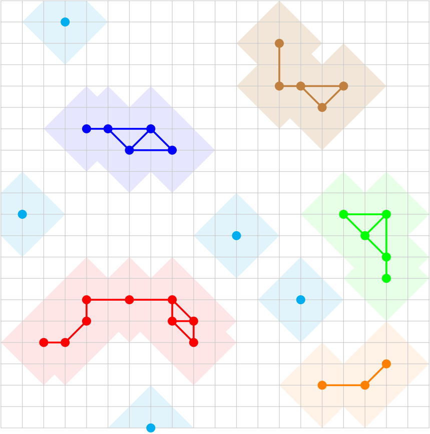

$G_t(r)$

Communication Model

- $M_a(t) \equiv$ messages known by *a* at time *t*
- $M_a(t)$ is non-decreasing (agents don't forget messages)
- On a meeting, agents exchange *all* the messages they know

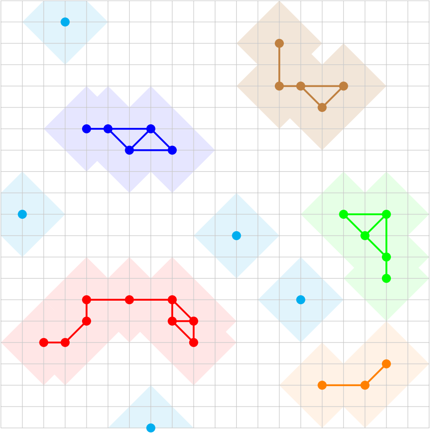

$G_t(r)$

Broadcast Time

Initially, only the *source s* knows the rumor $\mathcal{M}$:

$$M_s(0) = \{\mathcal{M}\} \qquad \text{and} \qquad M_a(0) = \emptyset \quad \forall a \neq s$$

We study the Broadcast Time T_{B} of the system, which is the first time instant when all the agents know the rumor:

$$T_{\mathrm{B}} = \inf\{\, t \geqslant 0 \,:\, \mathcal{M} \in M_a(t) \quad \forall a \,\}$$

Broadcast Time

Initially, only the *source* s knows the rumor $\mathcal{M}$:

$$M_s(0) = \{\mathcal{M}\} \qquad \text{and} \qquad M_a(0) = \emptyset \quad \forall a \neq s$$

We study the Broadcast Time T_{B} of the system, which is the first time instant when all the agents know the rumor:

$$T_{\mathrm{B}} = \inf\{\, t \geqslant 0 \,:\, \mathcal{M} \in M_a(t) \quad \forall a \,\}$$

Remarks

- $T_{\mathrm{B}} \equiv T_{\mathrm{B}}(n, k, r)$
- T_{B} is non-increasing in r: $\qquad r' \geqslant r \Rightarrow T_{\mathrm{B}}(r') \leqslant T_{\mathrm{B}}(r)$
- Broadcast analysis extends to other communication primitives

Our Contribution

Theorem 1 (Upper Bound on T_{B})

Let $r = 0$ (physical meetings). Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{O}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - 1/n^2$.

Our Contribution

Theorem 1 (Upper Bound on T_{B})

Let $r = 0$ (physical meetings). Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{O}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - 1/n^2$.

Since $T_{\mathrm{B}}(r)$ is non-increasing:

Corollary 1

$T_{\mathrm{B}} = \tilde{O}\left(n/\sqrt{k}\right)$ w.h.p. for any $k \geqslant 2$, $r \geqslant 0$.

Quite surprisingly, this bound is essentially tight (see next slide)

Our Contribution

Theorem 2 (Lower Bound on T_{B})

Let $r \leqslant \frac{1}{8e^3}\sqrt{n/k}$. Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{\Omega}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - (2^{-(k-1)} + 1/n + 2/n^2)$.

Our Contribution

Theorem 2 (Lower Bound on T_{B})

Let $r \leqslant \frac{1}{8e^3}\sqrt{n/k}$. Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{\Omega}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - (2^{-(k-1)} + 1/n + 2/n^2)$.

Peres *et al.* [SODA'11] consider the case $r > r_c \approx \sqrt{n/k}$ (i.e., above percolation threshold)

Our Contribution

Theorem 2 (Lower Bound on T_{B})

Let $r \leqslant \frac{1}{8e^3}\sqrt{n/k}$. Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{\Omega}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - (2^{-(k-1)} + 1/n + 2/n^2)$.

Peres *et al.* [SODA'11] consider the case $r > r_c \approx \sqrt{n/k}$ (i.e., above percolation threshold)

Together with Corollary 1, we have the tight result:

Corollary 2

If $k = \Omega(\log n)$ and $r \leqslant \frac{1}{8e^3}\sqrt{n/k}$, then $T_{\mathrm{B}} = \tilde{\Theta}\left(n/\sqrt{k}\right)$ w.h.p.

Two Key Facts about 2D Random Walks

Lemma (Large Deviation Bound)

Let S be a RW with $pos(0) = v_0$. Then,

$$\Pr\left[\exists 1 \leqslant t \leqslant \ell \quad s.t. \quad \| pos(t) - v_0 \| \geqslant \lambda\sqrt{\ell}\right] \leqslant 2e^{-\lambda^2/2}.$$

Two Key Facts about 2D Random Walks

Lemma (Large Deviation Bound)

Let S be a RW with $pos(0) = v_0$. Then,

$$\Pr\left[\exists 1 \leqslant t \leqslant \ell \quad s.t. \quad \| pos(t) - v_0 \| \geqslant \lambda\sqrt{\ell}\right] \leqslant 2e^{-\lambda^2/2}.$$

Lemma (Visit Probability)

Let S be a RW with $pos(0) = v_0$, and let $d = \| v - v_0 \|$. Then,

$$\Pr\left[S \text{ visits } v \text{ within time } d^2\right] \geqslant \frac{c_1}{\max\{1, \log d\}}.$$

Upper Bound

Theorem 1 (Upper Bound on T_{B})

Let $r = 0$ (physical meetings). Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{O}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - 1/n^2$.

Upper Bound

Theorem 1 (Upper Bound on T_{B})

Let $r = 0$ (physical meetings). Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{O}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - 1/n^2$.

Proof Strategy

- Lower bound the probability that two RWs meet

Upper Bound

Theorem 1 (Upper Bound on T_{B})

Let $r = 0$ (physical meetings). Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{O}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - 1/n^2$.

Proof Strategy

- Lower bound the probability that two RWs meet
- Show that many agents sufficiently close to an informed agent will be informed in a short time interval

Upper Bound

Theorem 1 (Upper Bound on T_{B})

Let $r = 0$ (physical meetings). Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \tilde{O}\left(\frac{n}{\sqrt{k}}\right)$$

with probability $\geqslant 1 - 1/n^2$.

Proof Strategy

- Lower bound the probability that two RWs meet
- Show that many agents sufficiently close to an informed agent will be informed in a short time interval
- Prove that the spreading process proceeds smoothly

Upper Bound: Meeting Probability

Lemma (Meeting Probability)

Let a, b be two RWs with $\| pos_a(0) - pos_b(0) \| = d \geqslant 1$*. Then,*

$$\Pr\left[a, b \text{ meet within time } d^2\right] \geqslant \frac{c_3}{\max\{1, \log d\}}.$$

Upper Bound: Meeting Probability

Lemma (Meeting Probability)

Let a, b be two RWs with $\| pos_a(0) - pos_b(0) \| = d \geqslant 1$. Then,

$$\Pr\left[a, b \textit{ meet within time } d^2\right] \geqslant \frac{c_3}{\max\{1, \log d\}}.$$

Proof

- Coupling with the difference RW $a - b$ and using the Visit Lemma to lower bound the probability of hitting the origin

Upper Bound: Meeting Probability

Lemma (Meeting Probability)

Let a, b be two RWs with $\| pos_a(0) - pos_b(0) \| = d \geqslant 1$. Then,

$$\Pr\left[a, b \text{ meet within time } d^2\right] \geqslant \frac{c_3}{\max\{1, \log d\}}.$$

Proof

- Coupling with the difference RW $a - b$ and using the Visit Lemma to lower bound the probability of hitting the origin

  or

- Direct calculation, via the approximation of $pos_a(t)$ (see [Lawler'91]) and considering the ratio between expected and maximum number of visits to a grid node within time d^2

Upper Bound

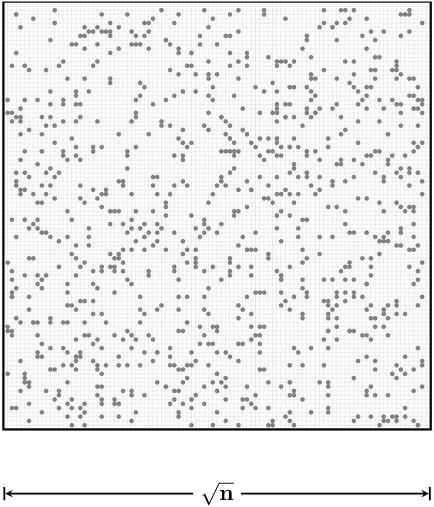

Upper Bound

- Consider the tessellation into cells of side

$$\ell = \Theta\left(\sqrt{(n \log^3 n)/k}\right)$$

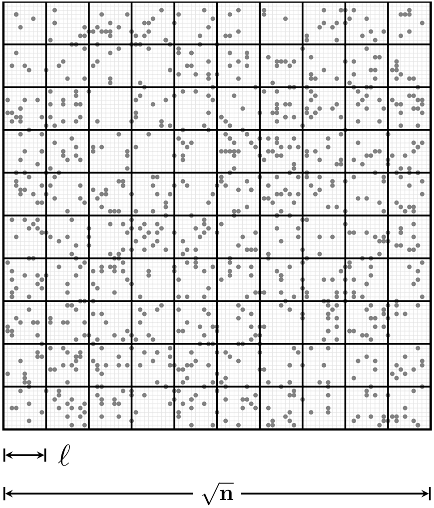

Upper Bound

- Consider the tessellation into cells of side

$$\ell = \Theta\left(\sqrt{(n \log^3 n)/k}\right)$$

- $\Rightarrow$ in each cell, there are always $\Theta\left(\log^3 n\right)$ agents w.h.p.

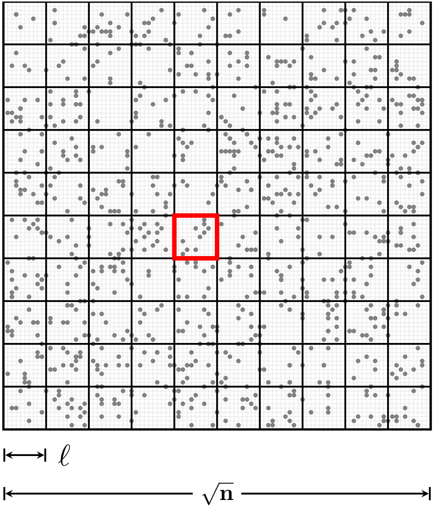

Upper Bound

- $t_Q =$ first time instant s.t. cell Q contains an informed agent
- $T_1 = \Theta\left(\ell^2 \log^4 n\right)$

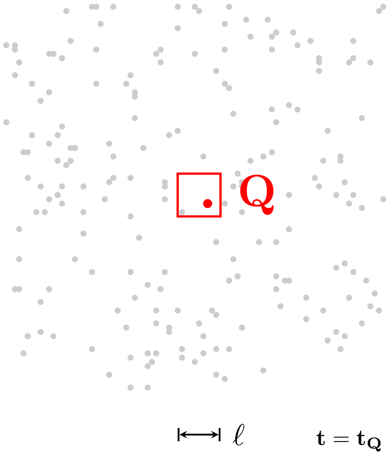

Upper Bound

- $t_Q =$ first time instant s.t. cell Q contains an informed agent
- $T_1 = \Theta\left(\ell^2 \log^4 n\right)$
- by time $\tau_1 = t_Q + T_1$, $\Omega\left(\log^2 n\right)$ agents are informed and at distance $\ell_1 = \tilde{O}(\ell)$ from Q

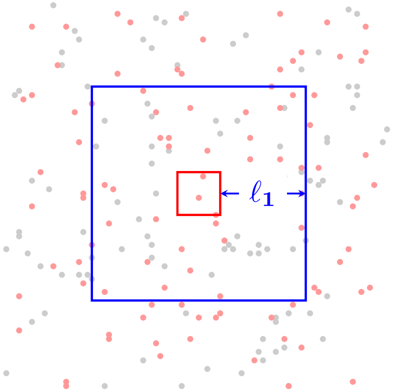

$\mathbf{t} = \tau_1$

Upper Bound

- $T_2 = \tilde{\Theta}\left(\ell_1^2\right) = \tilde{\Theta}\left(\ell^2\right)$
- by time $\tau_2 = \tau_1 + T_2$, each cell adjacent to Q has been reached and $\Omega\left(\log^2 n\right)$ informed agents remain at distance $\ell_2 = \tilde{O}\left(\ell\right)$ from Q during $[\tau_1, \tau_2]$

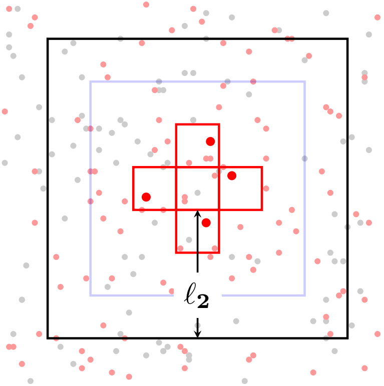

$\mathbf{t} = \tau_2$

Upper Bound

Repeating the above argument, all the cells will be reached by time

$$T^* = \frac{2\sqrt{n}}{\ell}(T_1 + T_2) = \tilde{O}\left(\frac{n}{\sqrt{k}}\right),$$

since $\ell = \tilde{\Theta}\left(\sqrt{n/k}\right)$ and T_1, $T_2 = \tilde{\Theta}\left(\ell^2\right)$.

Upper Bound

Repeating the above argument, all the cells will be reached by time

$$T^* = \frac{2\sqrt{n}}{\ell}(T_1 + T_2) = \tilde{O}\left(\frac{n}{\sqrt{k}}\right),$$

since $\ell = \tilde{\Theta}\left(\sqrt{n/k}\right)$ and T_1, $T_2 = \tilde{\Theta}\left(\ell^2\right)$.

The remaining uninformed agents will be informed within additional $\tilde{\Theta}\left(\ell^2\right)$ steps, since they will be in *some* cell at time T^*.

Upper Bound

Repeating the above argument, all the cells will be reached by time

$$T^* = \frac{2\sqrt{n}}{\ell}(T_1 + T_2) = \tilde{O}\left(\frac{n}{\sqrt{k}}\right),$$

since $\ell = \tilde{\Theta}\left(\sqrt{n/k}\right)$ and $T_1, T_2 = \tilde{\Theta}\left(\ell^2\right)$.

The remaining uninformed agents will be informed within additional $\tilde{\Theta}\left(\ell^2\right)$ steps, since they will be in *some* cell at time T^*.

Hence, the broadcast completes within time

$$T_{\mathrm{B}} \leqslant T^* + \tilde{\Theta}\left(\ell^2\right) = \tilde{O}\left(\frac{n}{\sqrt{k}}\right).$$

Lower Bound

Theorem 2 (Lower Bound on T_{B})

Let $r \leqslant \frac{1}{8e^3}\sqrt{n/k}$. Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \Omega\left(\frac{n}{\sqrt{k}\log^2 n}\right)$$

with probability $\geqslant 1 - (2^{-(k-1)} + 1/n + 2/n^2)$.

Lower Bound

Theorem 2 (Lower Bound on T_{B})

Let $r \leqslant \frac{1}{8e^3}\sqrt{n/k}$. Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \Omega\left(\frac{n}{\sqrt{k}\log^2 n}\right)$$

with probability $\geqslant 1 - (2^{-(k-1)} + 1/n + 2/n^2)$.

Proof Idea

- For a suitable separation parameter $\gamma = \Theta\left(\sqrt{n/k}\right)$, all the islands of $G_t(\gamma)$ have few agents

Lower Bound

Theorem 2 (Lower Bound on T_{B})

Let $r \leqslant \frac{1}{8e^3}\sqrt{n/k}$. Then, for $k \geqslant 2$,

$$T_{\mathrm{B}} = \Omega\left(\frac{n}{\sqrt{k}\log^2 n}\right)$$

with probability $\geqslant 1 - (2^{-(k-1)} + 1/n + 2/n^2)$.

Proof Idea

- For a suitable separation parameter $\gamma = \Theta\left(\sqrt{n/k}\right)$, all the islands of $G_t(\gamma)$ have few agents
- T_{B} is dominated by the time needed to cover the distances between these islands

Lower Bound

By choosing $\gamma = \Theta\left(\sqrt{n/k}\right)$:

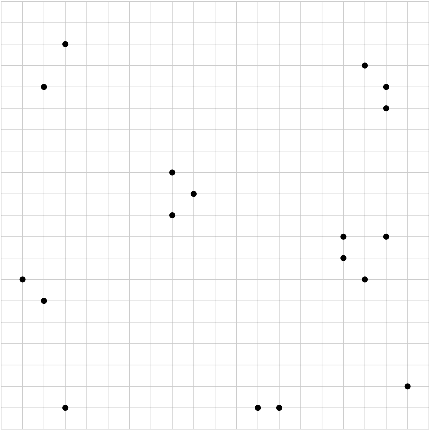

Lower Bound

By choosing $\gamma = \Theta\left(\sqrt{n/k}\right)$:

- at every time instant each island has $\leqslant \log n$ agents

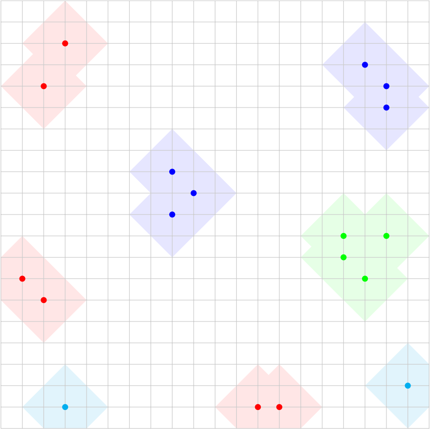

Lower Bound

By choosing $\gamma = \Theta\left(\sqrt{n/k}\right)$:

- at every time instant each island has $\leqslant \log n$ agents

- in a time interval of $\Delta t = \Theta(\gamma^2/\log n)$ steps the rumor cannot spread outside the island, so the informed area cannot “grow” more than $\gamma \log n$

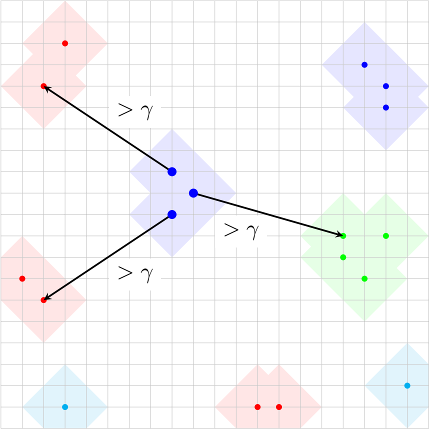

Lower Bound

- At $t = 0$ there is at least an agent at distance $\geqslant \sqrt{n}/4$ from the source

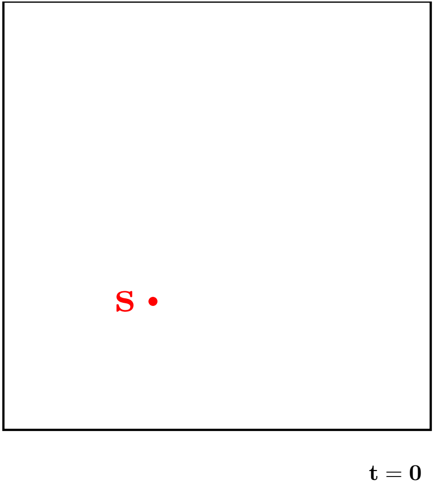

Lower Bound

- At $t = 0$ there is at least an agent at distance $\geqslant \sqrt{n}/4$ from the source

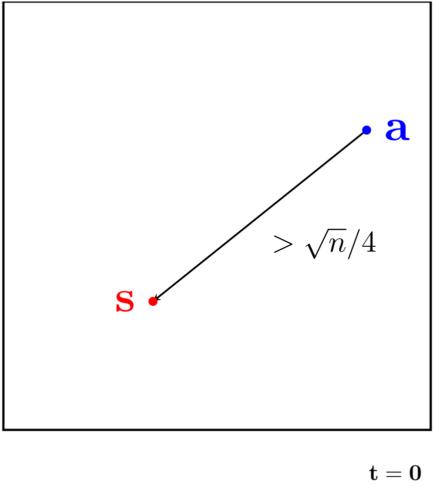

Lower Bound

- At $t = 0$ there is at least an agent at distance $\geqslant \sqrt{n}/4$ from the source
- In a time interval of $\Delta t = \Theta(\gamma^2 / \log n)$ steps the rumor cannot spread outside the island, so the informed area cannot “grow” more than $\gamma \log n$

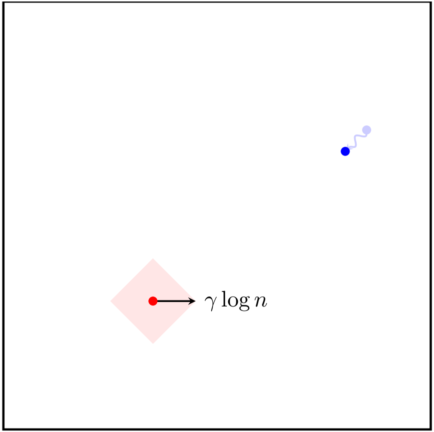

$\mathbf{t = \Delta t}$

Lower Bound

- At $t = 0$ there is at least an agent at distance $\geqslant \sqrt{n}/4$ from the source

- In a time interval of $\Delta t = \Theta(\gamma^2 / \log n)$ steps the rumor cannot spread outside the island, so the informed area cannot "grow" more than $\gamma \log n$

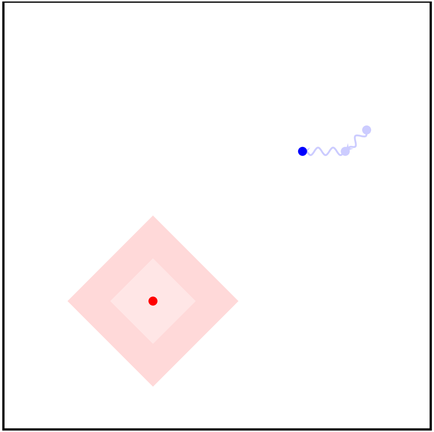

$\mathbf{t = 2\Delta t}$

Lower Bound

- At $t = 0$ there is at least an agent at distance $\geqslant \sqrt{n}/4$ from the source
- In a time interval of $\Delta t = \Theta(\gamma^2/\log n)$ steps the rumor cannot spread outside the island, so the informed area cannot "grow" more than $\gamma \log n$

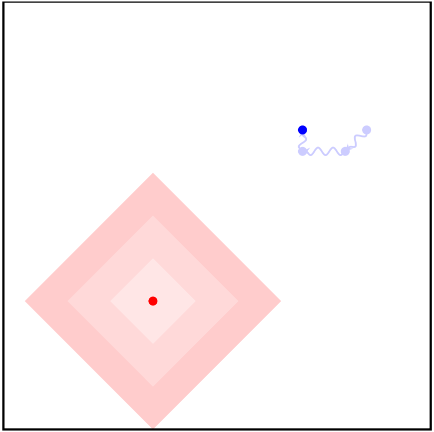

$\mathbf{t = 3\Delta t}$

Lower Bound

Consider an interval of $T = \Theta\left(n/(k \log^2 n)\right)$ steps:

- the informed area cannot cover a distance $> \sqrt{n}/8$

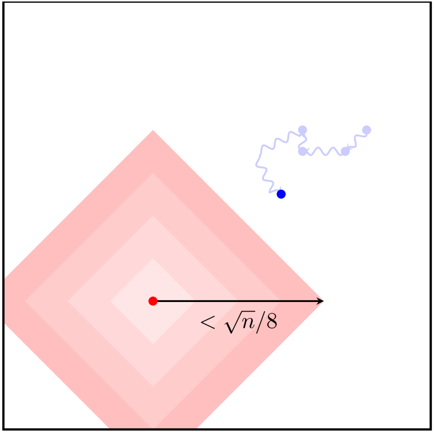

$\mathbf{t = T}$

Lower Bound

Consider an interval of

$T = \Theta\left(n/(k \log^2 n)\right)$ steps:

- the informed area cannot cover a distance $> \sqrt{n}/8$
- the blue agent cannot move towards the informed area more than $2\sqrt{T \log n} < \sqrt{n}/8$, hence it cannot know the rumor at time T

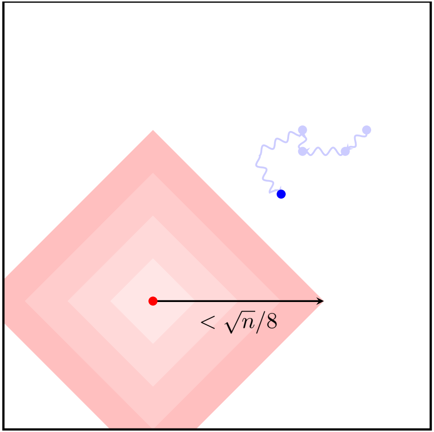

$\mathbf{t = T}$

Our Contribution & Open Problems

- We presented a tight characterization of T_{B} for a sparse system
  - UB: (1) lower bounding the meeting probability of two RWs and (2) showing that the spreading process is smooth
  - LB: showing that T_{B} is dominated by the inter-island distance

Our Contribution & Open Problems

- We presented a tight characterization of T_{B} for a sparse system
  - UB: (1) lower bounding the meeting probability of two RWs and (2) showing that the spreading process is smooth
  - LB: showing that T_{B} is dominated by the inter-island distance

- Our analysis techniques extend to
  - dense scenarios ($k \geqslant n/2$)
  - other communication primitives (gossip, multicast)
  - related models (Frog Model, mobility with jumps, predator-prey)

Our Contribution & Open Problems

- We presented a tight characterization of T_{B} for a sparse system
  - UB: (1) lower bounding the meeting probability of two RWs and (2) showing that the spreading process is smooth
  - LB: showing that T_{B} is dominated by the inter-island distance

- Some open problems
  - Modeling barriers and obstacles
  - More realistic mobility models
  - Tradeoffs between communication complexity and spreading time
  - Tradeoffs between agent's buffer size and spreading time

- Generalization to higher dimensions (Lam *et al.*, SODA'12)

Thank You!

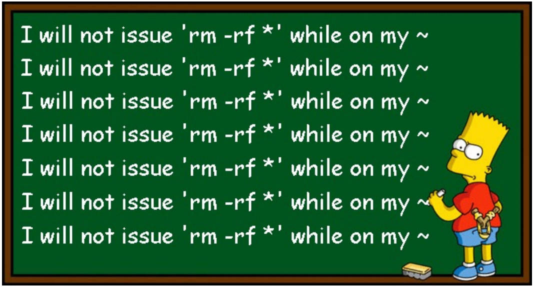
I will not issue 'rm -rf *' while on my ~
I will not issue 'rm -rf *' while on my ~
I will not issue 'rm -rf *' while on my ~
I will not issue 'rm -rf *' while on my ~
I will not issue 'rm -rf *' while on my ~
I will not issue 'rm -rf *' while on my ~
I will not issue 'rm -rf *' while on my ~