Tight Bounds on Information Spreading in Sparse Mobile Networks

Alberto Pettarin Andrea Pietracaprina Geppino Pucci

Università di Padova, Padova, Italy

Eli Upfal

Brown University, Providence, RI

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. . .

Mobile Networks: a Closer Look

Mobile networks are distributed systems

- dynamic: topology changes over time…
- … but not too fast: mobility speed $<<$ transmission speed

Mobile Networks: a Closer Look

Mobile networks are distributed systems

- dynamic: topology changes over time...
- ... but not too fast: mobility speed $<<$ transmission speed
- no infrastructure $\Rightarrow$ wireless, multi-hop communications

Mobile Networks: a Closer Look

Mobile networks are distributed systems

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

Mobile Networks: a Closer Look

Mobile networks are distributed systems

- dynamic: topology changes over time...
- ... but not too fast: mobility speed $<<$ transmission speed
- no infrastructure $\Rightarrow$ wireless, multi-hop communications
- energy constraints $\Rightarrow$ 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 $<<$ transmission speed
- no infrastructure $\Rightarrow$ wireless, multi-hop communications
- energy constraints $\Rightarrow$ 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]
  - Frog Model: infective RWs on $\mathbb{Z}^d$
  - shape of the infected volume
- Dimitriou *et al.* [Disc. App. Math. '06]
  - RWs on generic graphs
  - expected infection time as function of the graph expansion
- Clementi *et al.* [ICALP '09, IPDPS '09] and Peres *et al.* [SODA '11]
  - time to broadcast a message
  - dense scenarios above percolation threshold

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid

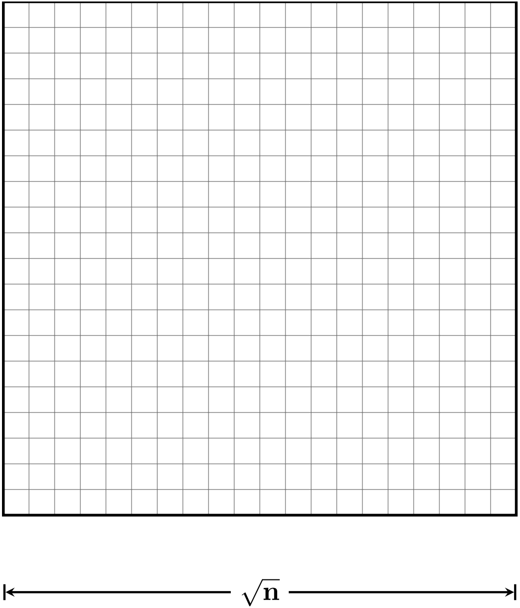

Mobility Model

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

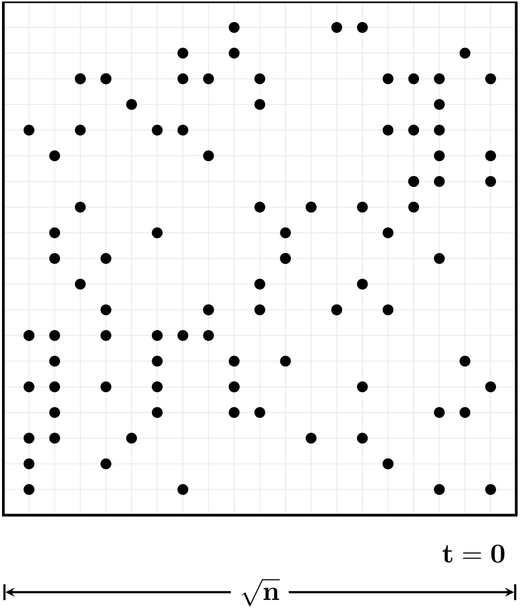

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution
- Independent, simple, discrete-time random walks

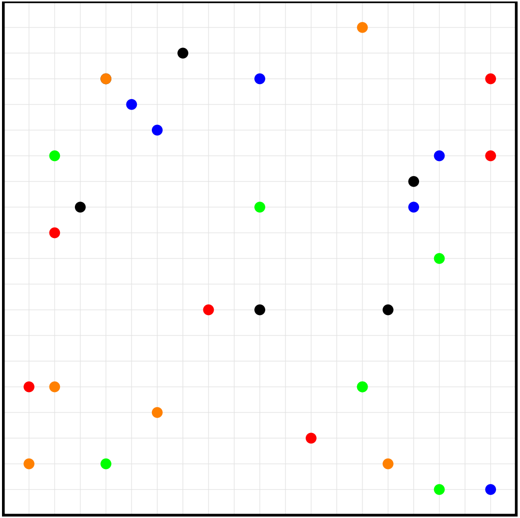

$\mathbf{t = 0}$

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution
- Independent, simple, discrete-time random walks

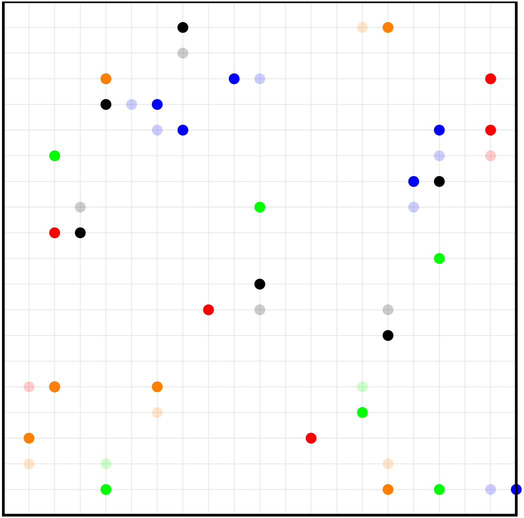

$\mathbf{t = 1}$

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution
- Independent, simple, discrete-time random walks

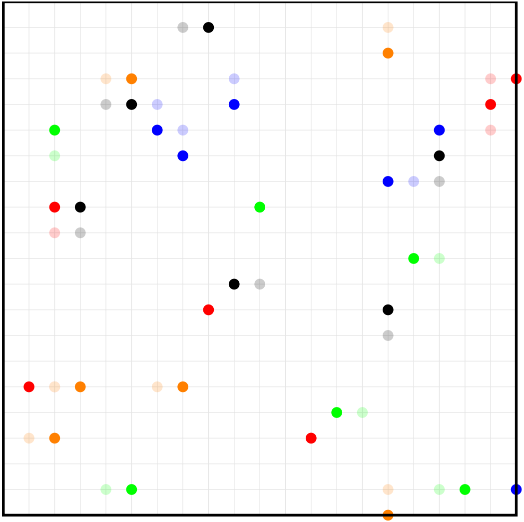

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution
- Independent, simple, discrete-time random walks

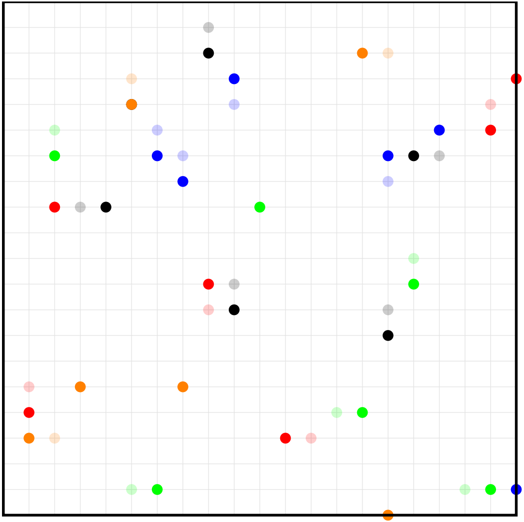

$\mathbf{t = 3}$

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution
- Independent, simple, discrete-time random walks

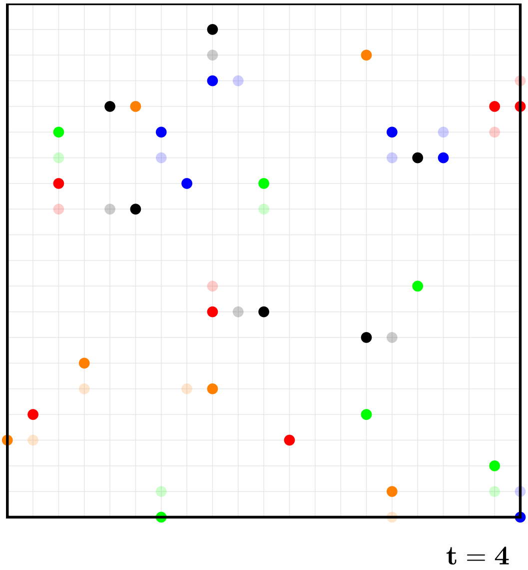

Mobility Model

- $\sqrt{n} \times \sqrt{n}$ 2D grid
- $k = O(n)$ mobile agents
- Initial positions $\equiv$ stationary distribution
- Independent, simple, discrete-time random walks

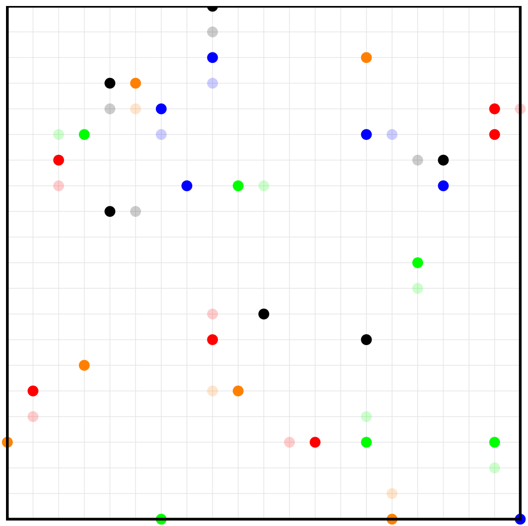

$\mathbf{t = 5}$

Communication Model

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

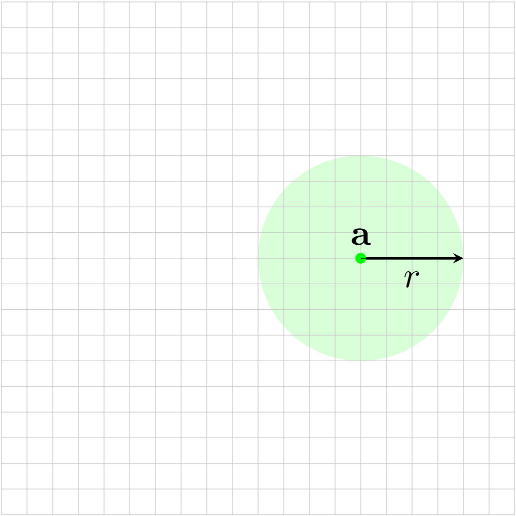

Communication Model

- Each agent has transmission radius $r \geqslant 0$
- $\text{dist}(a, b) \leqslant r \ \Rightarrow \ a \leftrightarrow b$

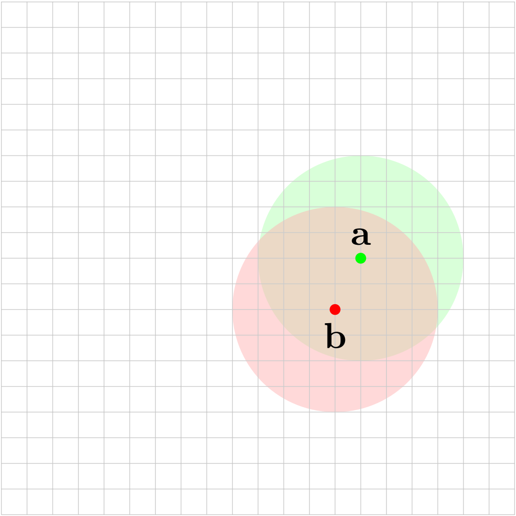

Communication Model

- Each agent has transmission radius $r \geqslant 0$
- $\text{dist}(a, b) \leqslant r \;\Rightarrow\; a \leftrightarrow b$
- $\text{dist}(c, a) > r \;\Rightarrow\; c \not\leftrightarrow a$
- $\text{dist}(c, b) > r \;\Rightarrow\; c \not\leftrightarrow b$

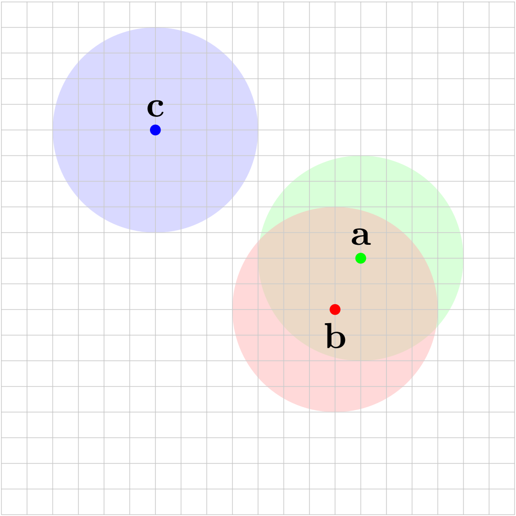

Communication Model

- Each agent has transmission radius $r \geqslant 0$
- $\text{dist}(a, b) \leqslant r \;\Rightarrow\; a \leftrightarrow b$
- $\text{dist}(c, a) > r \;\Rightarrow\; c \not\leftrightarrow a$
- $\text{dist}(c, b) > r \;\Rightarrow\; c \not\leftrightarrow b$
- Reliable transmissions (no faults, no radio interference)

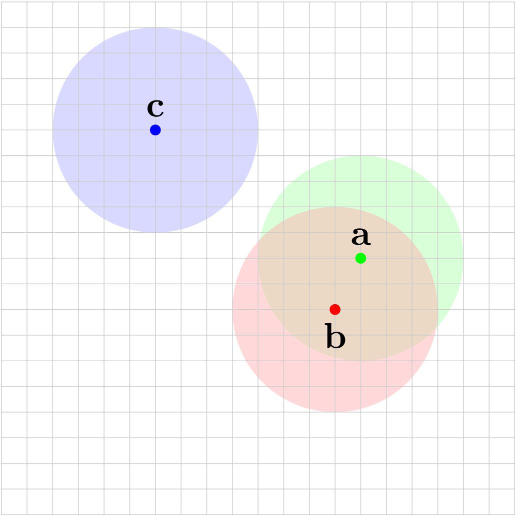

Communication Model

- Visibility graph $G_t(r)$
  - vertex $\equiv$ agent

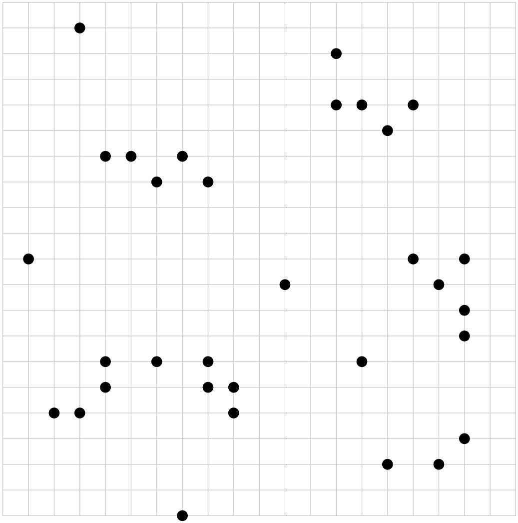

Communication Model

- Visibility graph $G_t(r)$
  - vertex $\equiv$ agent
  - edge $\{a, a'\} \in G_t(r)$

    $\Updownarrow$

    $\text{dist}(a, a') \leqslant r$

    (at time t)

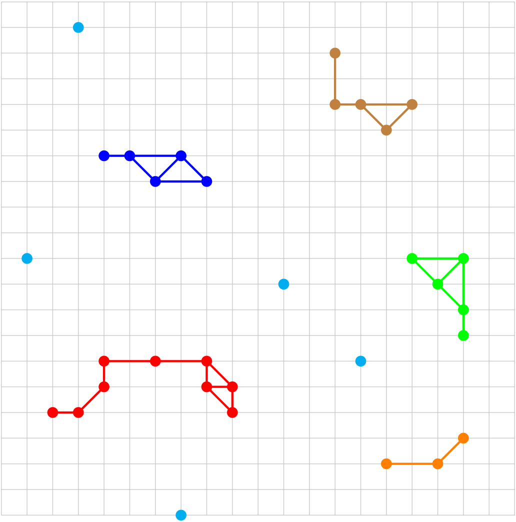

$G_t(r)$

Communication Model

- Visibility graph $G_t(r)$
  - vertex $\equiv$ agent
  - edge $\{a, a'\} \in G_t(r)$

    $\Updownarrow$

    $\text{dist}(a, a') \leqslant r$

    (at time t)
  - each connected component is called "island"

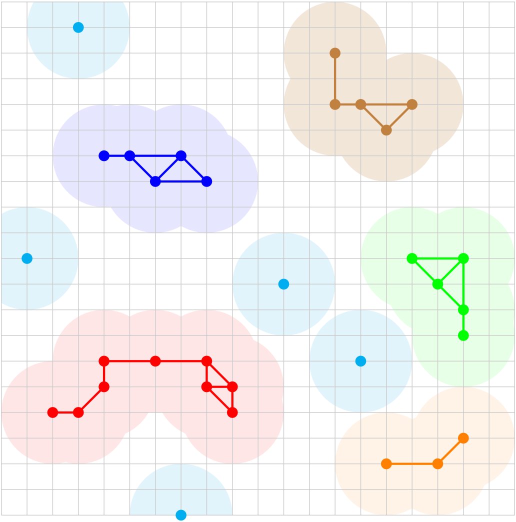

$G_t(r)$

Broadcast Time

At $t = 0$, only one agent is informed of a rumor $\mathcal{M}$, while all the other agents are uninformed.

Broadcast Time

At $t = 0$, only one agent is informed of a rumor $\mathcal{M}$, while all the other agents are uninformed.

An uninformed agent becomes informed when it comes at distance $\leqslant r$ from an informed agent (meeting).

Broadcast Time

At $t = 0$, only one agent is informed of a rumor $\mathcal{M}$, while all the other agents are uninformed.

An uninformed agent becomes informed when it comes at distance $\leqslant r$ from an informed agent (meeting).

We study the Broadcast Time T_{B}, which is the first time instant when all the agents know $\mathcal{M}$.

Broadcast Time

At $t = 0$, only one agent is informed of a rumor $\mathcal{M}$, while all the other agents are uninformed.

An uninformed agent becomes informed when it comes at distance $\leqslant r$ from an informed agent (meeting).

We study the Broadcast Time T_{B}, which is the first time instant when all the agents know $\mathcal{M}$.

Remarks

- $T_{\mathrm{B}} \equiv T_{\mathrm{B}}(n, k, r)$

Broadcast Time

At $t = 0$, only one agent is informed of a rumor $\mathcal{M}$, while all the other agents are uninformed.

An uninformed agent becomes informed when it comes at distance $\leqslant r$ from an informed agent (meeting).

We study the Broadcast Time T_{B}, which is the first time instant when all the agents know $\mathcal{M}$.

Remarks

- $T_{\mathrm{B}} \equiv T_{\mathrm{B}}(n, k, r)$
- T_{B} is non-increasing in r: $r' \geqslant r \Rightarrow T_{\mathrm{B}}(r') \leqslant T_{\mathrm{B}}(r)$

Broadcast Time

At $t = 0$, only one agent is informed of a rumor $\mathcal{M}$, while all the other agents are uninformed.

An uninformed agent becomes informed when it comes at distance $\leqslant r$ from an informed agent (meeting).

We study the Broadcast Time T_{B}, which is the first time instant when all the agents know $\mathcal{M}$.

Remarks

- $T_{\mathrm{B}} \equiv T_{\mathrm{B}}(n, k, r)$
- T_{B} is non-increasing in r: $\quad r' \geqslant r \Rightarrow T_{\mathrm{B}}(r') \leqslant T_{\mathrm{B}}(r)$
- Broadcast analysis extends to gossip, multicast, etc.

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

Note that the percolation threshold ($\Rightarrow$ giant connected component) is $r_c \approx \sqrt{n/k}$.

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

Note that the percolation threshold ($\Rightarrow$ giant connected component) is $r_c \approx \sqrt{n/k}$.

Peres *et al.* consider the case $r > r_c$ and prove that the time to hit the giant connected component is logarithmic.

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

Note that the percolation threshold ($\Rightarrow$ giant connected component) is $r_c \approx \sqrt{n/k}$.

Peres *et al.* consider the case $r > r_c$ and prove that the time to hit the giant connected component is logarithmic.

If $G_t(r)$ is connected, the results of Clementi *et al.* hold.

Upper Bound: Proof Strategy

To prove the upper bound on T_{B}, we show that:

Upper Bound: Proof Strategy

To prove the upper bound on T_{B}, we show that:

1. **Lemma (Meeting Probability)**

$$\Pr\left[\begin{array}{c}\text{two RWs starting at distance } d \\ \text{meet within } d^2 \text{ steps}\end{array}\right] \geqslant \frac{c_3}{\max\{1, \log d\}}$$

Upper Bound: Proof Strategy

To prove the upper bound on T_{B}, we show that:

1. **Lemma (Meeting Probability)**

$$\Pr\left[\begin{array}{c}\text{two RWs starting at distance } d\\ \text{meet within } d^2 \text{ steps}\end{array}\right] \geqslant \frac{c_3}{\max\{1, \log d\}}$$

2. In a short time interval, an informed agent meets (and informs) many other agents that will stay close to the current position

Upper Bound: Proof Strategy

To prove the upper bound on T_{B}, we show that:

1. **Lemma (Meeting Probability)**

$$\Pr\left[\begin{array}{c}\text{two RWs starting at distance } d \\ \text{meet within } d^2 \text{ steps}\end{array}\right] \geqslant \frac{c_3}{\max\{1, \log d\}}$$

2. In a short time interval, an informed agent meets (and informs) many other agents that will stay close to the current position

3. The spreading process proceeds smoothly

Upper Bound

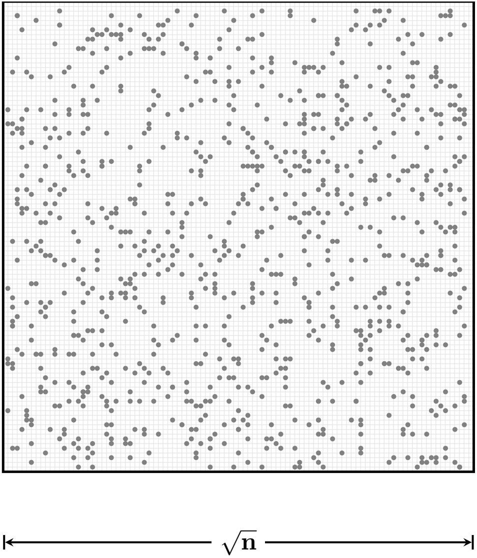

Upper Bound

- Consider the tessellation into cells of side

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

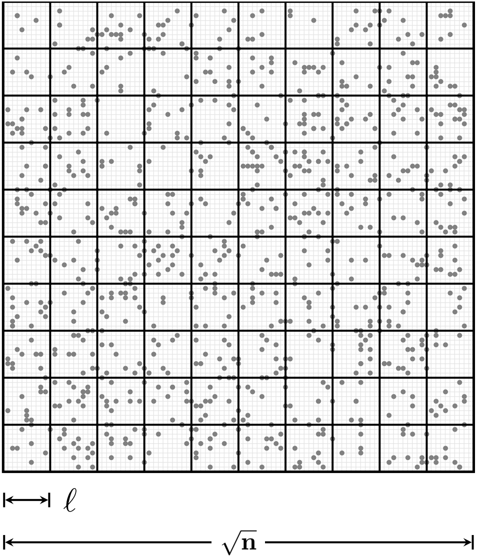

Upper Bound

- 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.

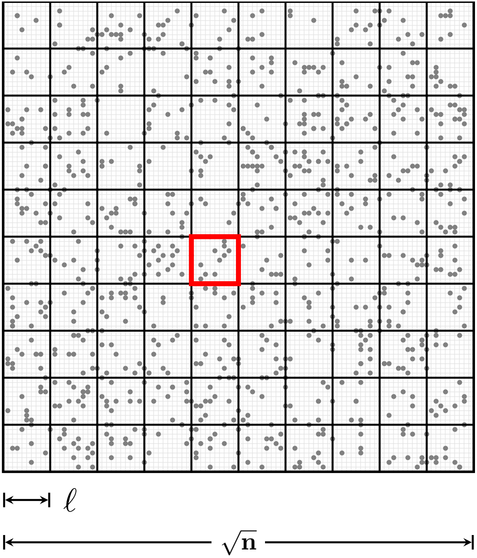

Upper Bound

- t_Q = first time instant s.t. cell Q contains an informed agent

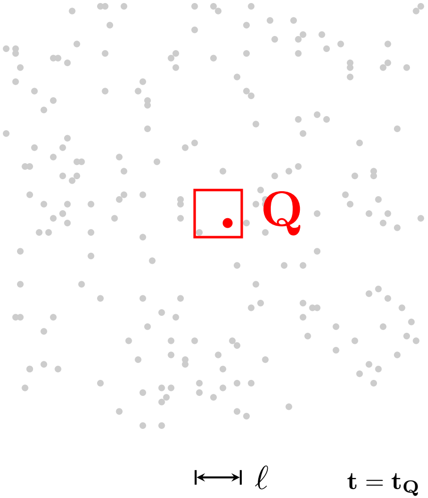

Upper Bound

- $t_Q =$ first time instant s.t. cell Q contains an informed agent
- set $\Delta t = \tilde{\Theta}\left(\ell^2\right)$

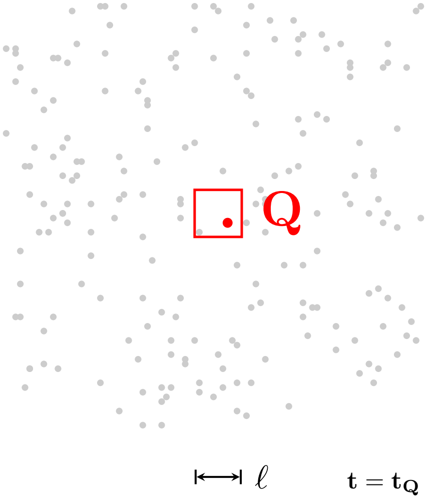

Upper Bound

- $t_Q =$ first time instant s.t. cell Q contains an informed agent
- set $\Delta t = \tilde{\Theta}\left(\ell^2\right)$
- by time $t_Q + \Delta t$,
  - $\Omega\left(\log^2 n\right)$ informed agents remain at distance $\ell_1 = \tilde{O}\left(\ell\right)$ from Q
  - each cell adjacent to Q has been reached

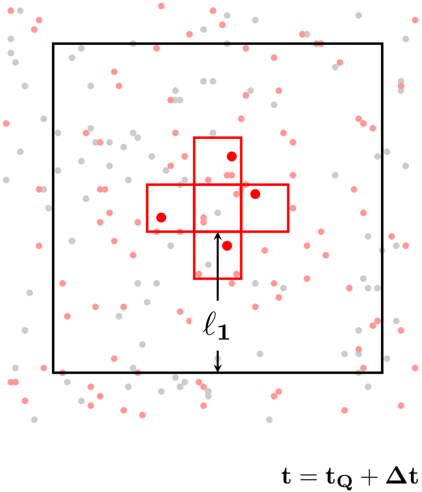

Upper Bound

- by repeating the above argument, at intervals of length Δt ...

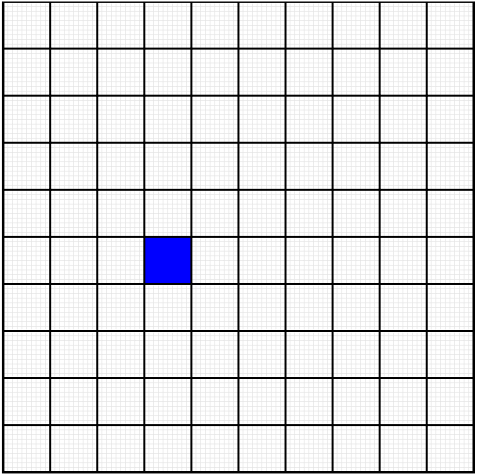

$\mathbf{t = 0}$

Upper Bound

- by repeating the above argument, at intervals of length Δt ...

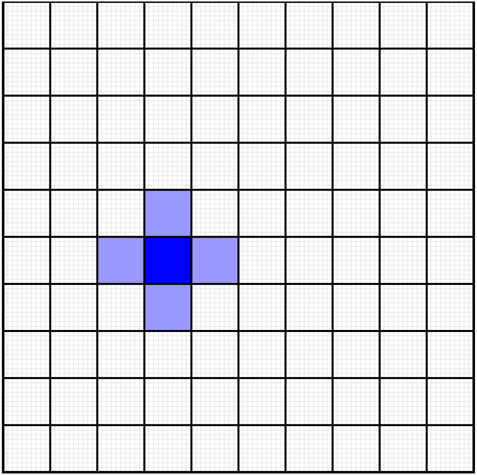

$\mathbf{t = \Delta t}$

Upper Bound

- by repeating the above argument, at intervals of length Δt ...

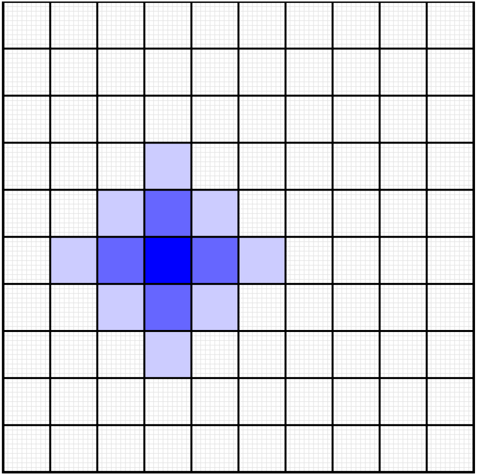

$\mathbf{t = 2\Delta t}$

Upper Bound

- by repeating the above argument, at intervals of length Δt ...
- ... all the cells are reached by time

$$T^* = \frac{2\sqrt{n}}{\ell}\Delta t$$

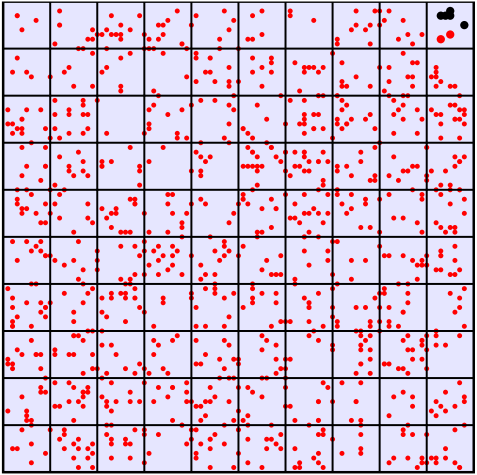

$\mathbf{t} = \mathbf{T}^*$

Upper Bound

- allowing some more time to inform the possible remaining uninformed agents,

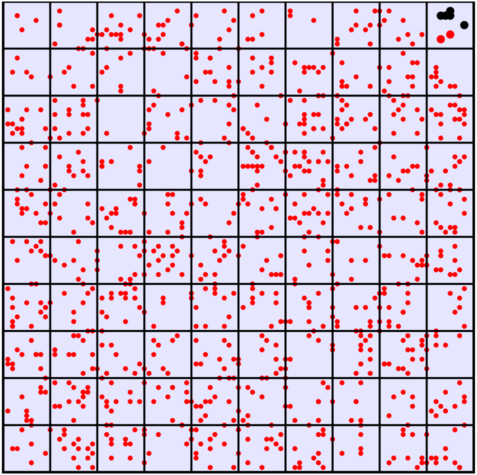

$\mathbf{t} = \mathbf{T}^*$

Upper Bound

- allowing some more time to inform the possible remaining uninformed agents, we can conclude that

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

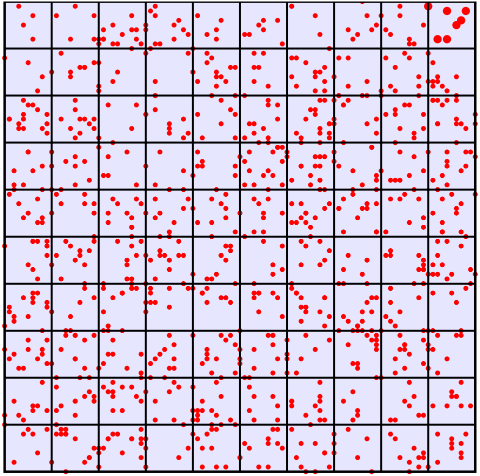

$\mathbf{t = T_B}$

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

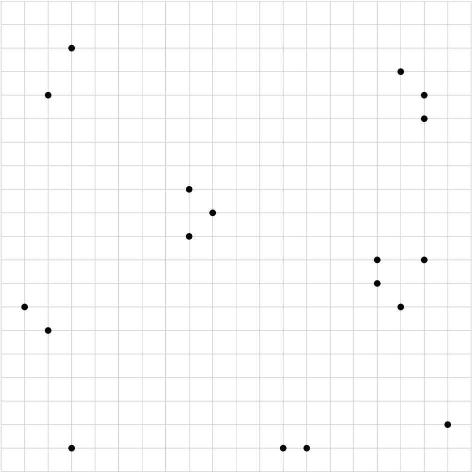

Lower Bound

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

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

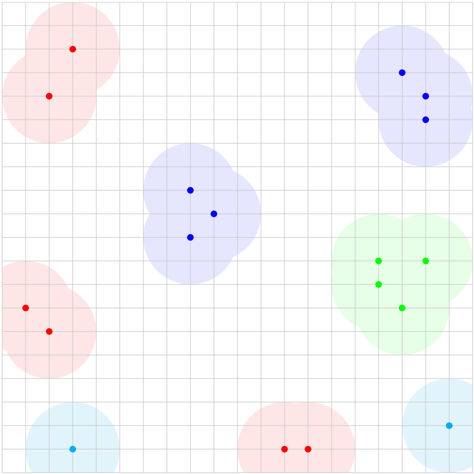

Lower Bound

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

- at every time instant each island has $\leqslant \log n$ agents
- in $\Delta t = \Theta(\gamma^2 / \log n)$ steps the rumor cannot spread outside an island, so the informed area cannot "grow" more than $\gamma \log n$

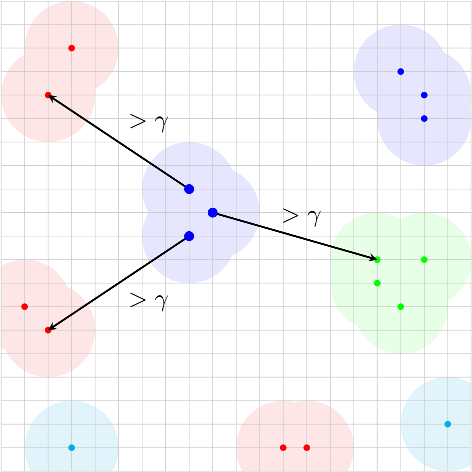

Lower Bound

- At $t = 0$

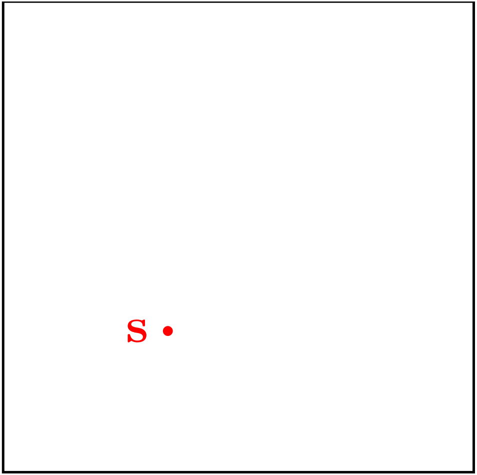

$\mathbf{t = 0}$

Lower Bound

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

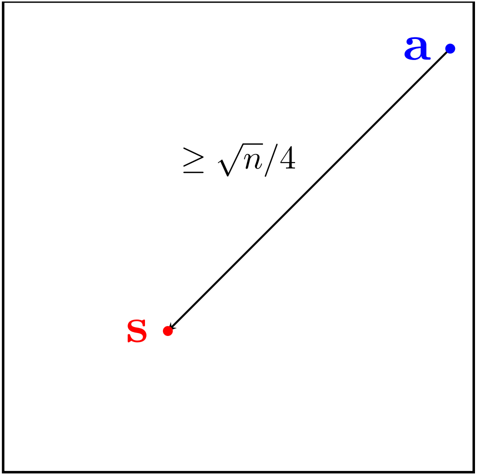

$\mathbf{t = 0}$

Lower Bound

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

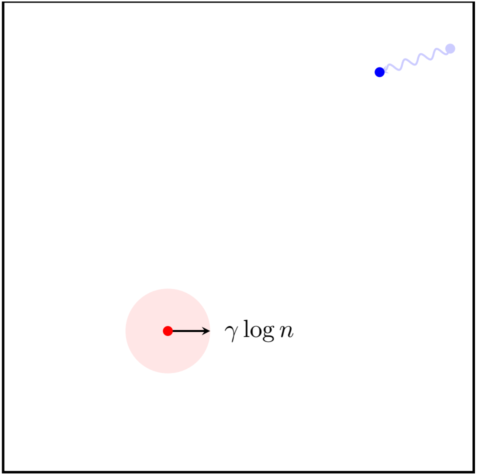

$\mathbf{t = \Delta t}$

Lower Bound

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

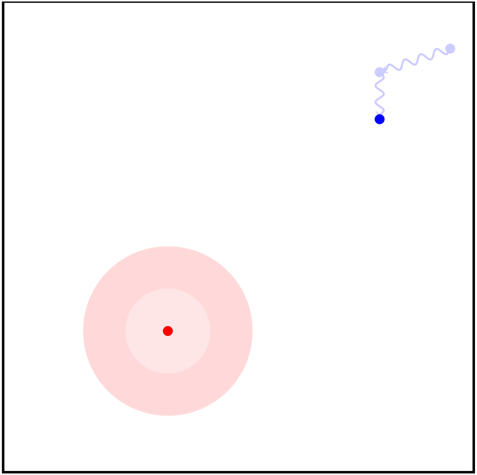

$\mathbf{t = 2\Delta t}$

Lower Bound

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

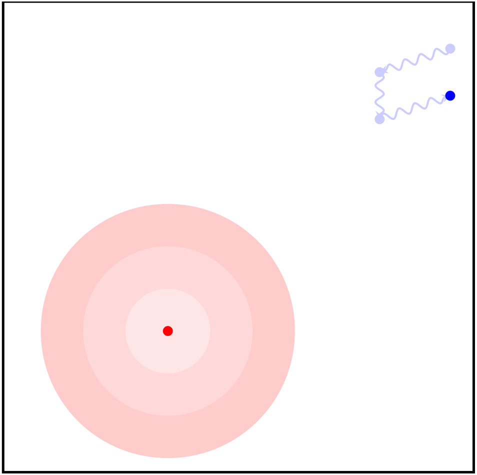

$\mathbf{t = 3\Delta t}$

Lower Bound

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

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

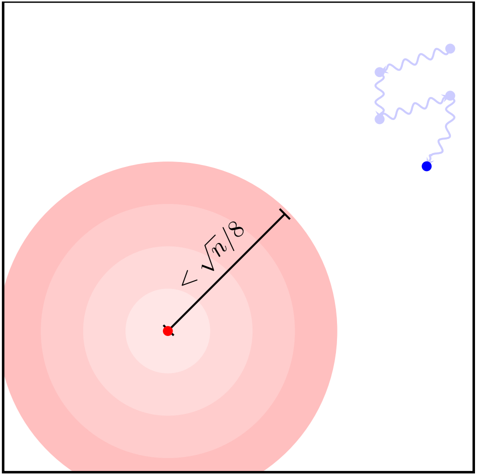

$\mathbf{t = T}$

Lower Bound

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

- the informed area does not cover a distance $> \sqrt{n}/8$
- the blue agent does not move towards the informed area more than $2\sqrt{T \log n} < \sqrt{n}/8$ (large deviation bound)

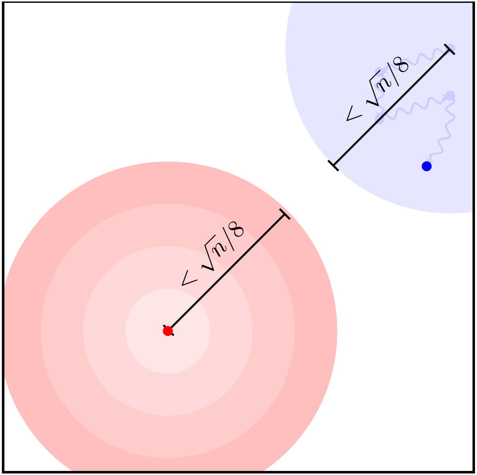

$\mathbf{t = T}$

Lower Bound

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

- the informed area does not cover a distance $> \sqrt{n}/8$
- the blue agent does not move towards the informed area more than $2\sqrt{T \log n} < \sqrt{n}/8$ (large deviation bound)
- $\Rightarrow T_\mathrm{B} > T$

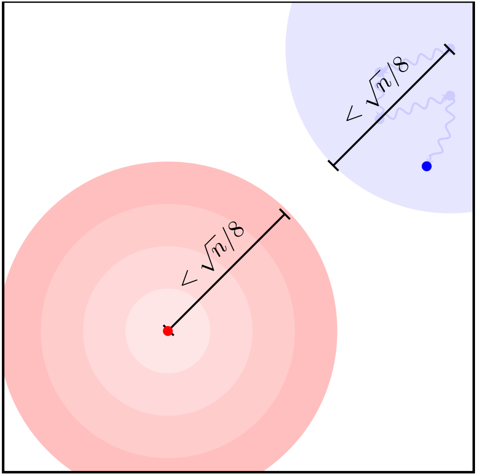

$\mathbf{t = T}$

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
  - 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 (work in progress)
  - Tradeoffs between agents' message buffer size and spreading time (work in progress)
  - More realistic mobility models
- Generalization to higher dimensions (Lam *et al.*, [arXiv:1104.5268])


Thank You!