Hop-optimal Networks in the Weak Sensor Model

Miguel Mosteiro

Department of Computer Science, Rutgers University

Joint work with Martín Farach-Colton and Rohan Fernandes

A sensor node

Capabilities

- processing
- sensing
- communication

Limitations

- range
- memory
- life cycle

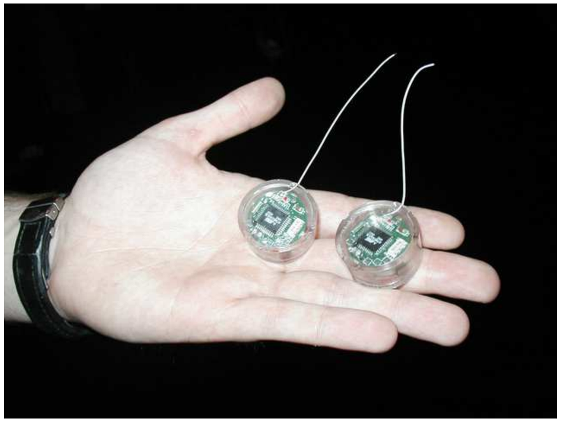

University of California, Berkeley and Intel Berkeley Research Lab.

Deborah Estrin, UCLA, holds a sensor node.

Sensor network

Constraints:

- weak sensors.
- geometric random distribution.

Sensor network

Constraints:

- weak sensors.
- geometric random distribution.

Question:

How to organize such a network optimally?

Sensor network

Constraints:

- weak sensors.
- geometric random distribution.

Question:

How to organize such a network optimally?

Our result:

Optimal-Network Bootstrapping

Our results

Geometric properties

- There exists a hop-optimal subgraph of a random geometric graph with useful properties for weak sensors (details to follow).

Network bootstrapping

- Polylogarithmic localized algorithm to build the network modelled by such a graph within the Weak Sensor Model.

Byproduct

- Fast maximal independent set (MIS) distributed algorithm with contention resolution.

All *with high probability.*

This talk

- Problem details.
  - **The Weak Sensor Model.**
  - Optimization criteria.
  - Random geometric graph model.
- Related work.
- Our results.
- Future work.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- No initial infrastructure.
- Radio tx on a shared channel.
- Binary channel-status: tx|other.
- Discrete tx power range.

- One channel of communication.
- Non-simultaneous rx and tx.
- No position information.
- No synchronicity.
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- Radio tx on a shared channel.
- Binary channel-status: tx|other.
- Discrete tx power range.

- One channel of communication.
- Non-simultaneous rx and tx.
- No position information.
- No synchronicity.
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- Binary channel-status: tx|other.
- Discrete tx power range.

- One channel of communication.
- Non-simultaneous rx and tx.
- No position information.
- No synchronicity.
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- **Binary channel-status: tx|other.**
- Discrete tx power range.
- One channel of communication.
- Non-simultaneous rx and tx.
- No position information.
- No synchronicity.
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- **Binary channel-status: tx|other.**
- **Discrete tx power range.**

- One channel of communication.
- Non-simultaneous rx and tx.
- No position information.
- No synchronicity.
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- **Binary channel-status: tx|other.**
- **Discrete tx power range.**
- **One channel of communication.**
- Non-simultaneous rx and tx.
- No position information.
- No synchronicity.
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- **Binary channel-status: tx|other.**
- **Discrete tx power range.**
- **One channel of communication.**
- **Non-simultaneous rx and tx.**
- No position information.
- No synchronicity.
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- **Binary channel-status: tx|other.**
- **Discrete tx power range.**
- **One channel of communication.**
- **Non-simultaneous rx and tx.**
- **No position information.**
- No synchronicity.
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- **Binary channel-status: tx|other.**
- **Discrete tx power range.**
- **One channel of communication.**
- **Non-simultaneous rx and tx.**
- **No position information.**
- **No synchronicity.**
- Adversarial wake-up schedule.
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- **Binary channel-status: tx|other.**
- **Discrete tx power range.**
- **One channel of communication.**
- **Non-simultaneous rx and tx.**
- **No position information.**
- **No synchronicity.**
- **Adversarial wake-up schedule.**
- No global controller.

tx = transmission.

rx = reception.

The Weak Sensor Model

- **Limited memory size.**
- **Limited life cycle.**
- **Limited range.**
- **No initial infrastructure.**
- **Radio tx on a shared channel.**
- **Binary channel-status: tx|other.**
- **Discrete tx power range.**
- **One channel of communication.**
- **Non-simultaneous rx and tx.**
- **No position information.**
- **No synchronicity.**
- **Adversarial wake-up schedule.**
- **No global controller.**

tx = transmission.

rx = reception.

This talk

- Problem details.
  - The Weak Sensor Model.
  - **Optimization criteria.**
  - Random geometric graph model.
- Related work.
- Our results.
- Future work.

Optimization criteria

Maximize life cycle subject to the Weak Sensor Model constraints.

$$communication\ cost \sim dist^{\alpha} \cdot count$$

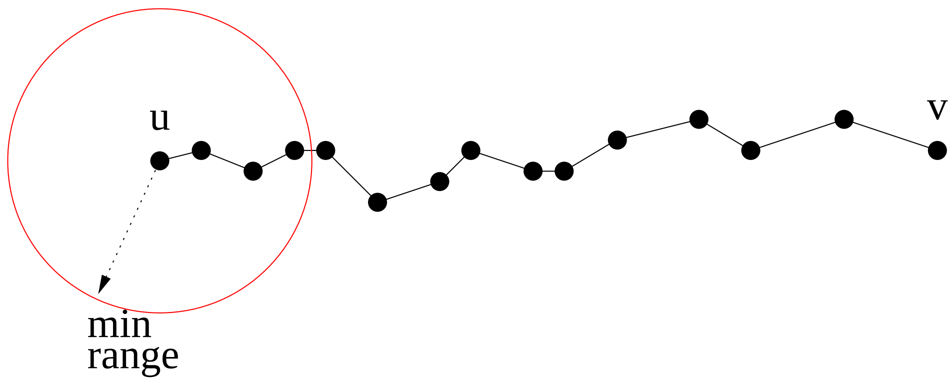

Transmission count due to contention resolution!

fewer hops $\Rightarrow$ ***less energy***

This talk

- Problem details.
  - The Weak Sensor Model.
  - Optimization criteria.
  - **Random geometric graph model.**
- Related work.
- Our results.
- Future work.

Random geometric graph model

Random distributions in $\mathbb{R}^2$:

need to understand geometric properties such as:

- connectivity
- path length
- coverage

$\mathcal{G}_{n,r,\ell}$

- $[0,\ell]^2$
- $\ell \to \infty$
- Structural properties depend on relation among r, n and ℓ.

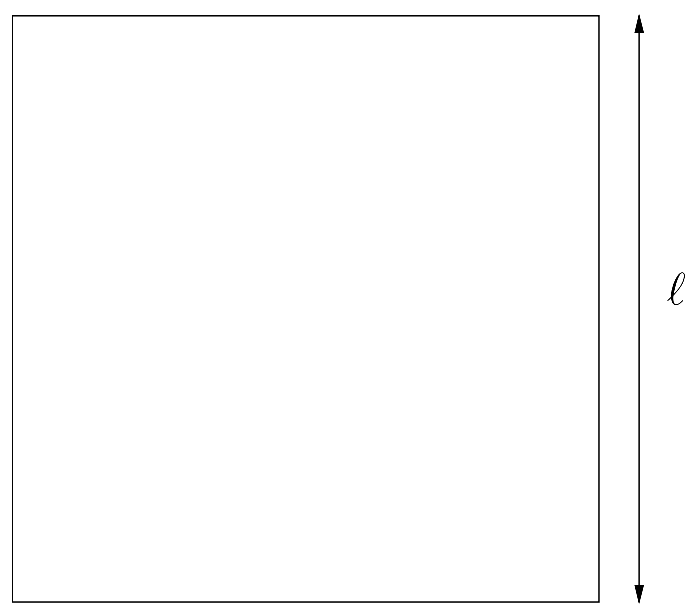

Random geometric graph model

Random distributions in $\mathbb{R}^2$:

need to understand geometric properties such as:

– connectivity

– path length

– coverage

$\mathcal{G}_{n,r,\ell}$

- $[0, \ell]^2$
- $\ell \to \infty$
- Structural properties depend on relation among r, n and ℓ.

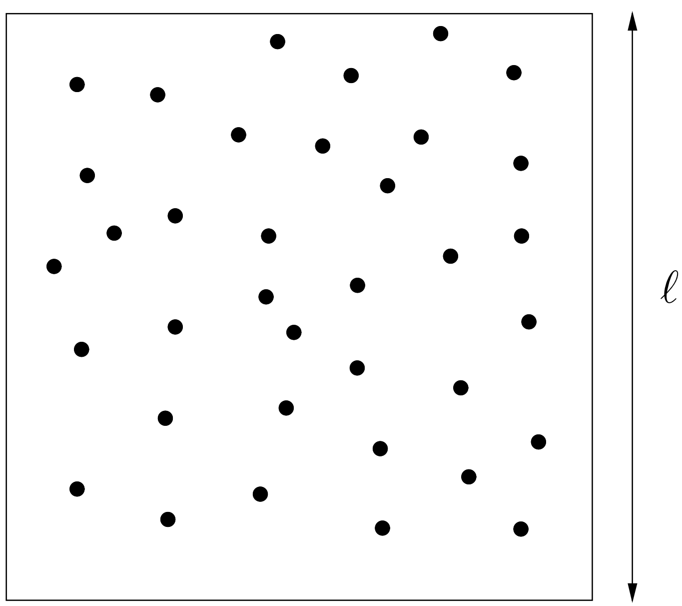

Random geometric graph model

Random distributions in $\mathbb{R}^2$:

need to understand geometric properties such as:

– connectivity

– path length

– coverage

$\mathcal{G}_{n,r,\ell}$

- $[0, \ell]^2$
- $\ell \to \infty$
- Structural properties depend on relation among r, n and ℓ.

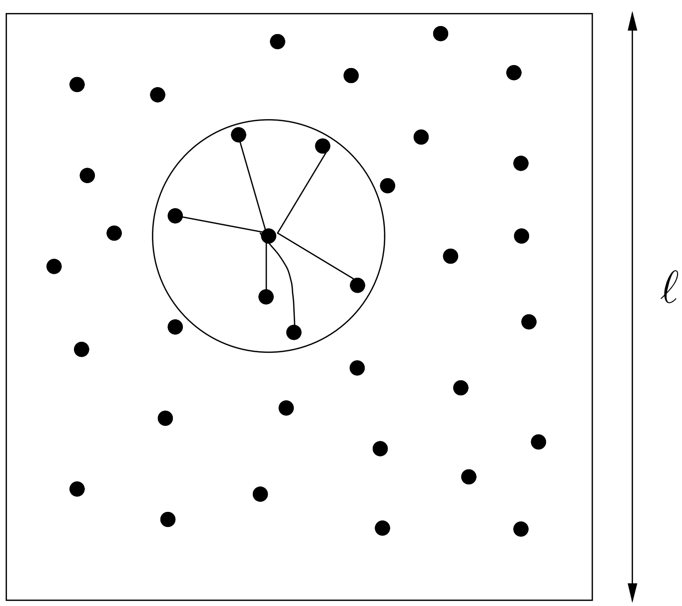

Random geometric graph model

Random distributions in $\mathbb{R}^2$:

need to understand geometric properties such as:

– connectivity

– path length

– coverage

$\mathcal{G}_{n,r,\ell}$

- $[0, \ell]^2$
- $\ell \to \infty$
- Structural properties depend on relation among r, n and ℓ.

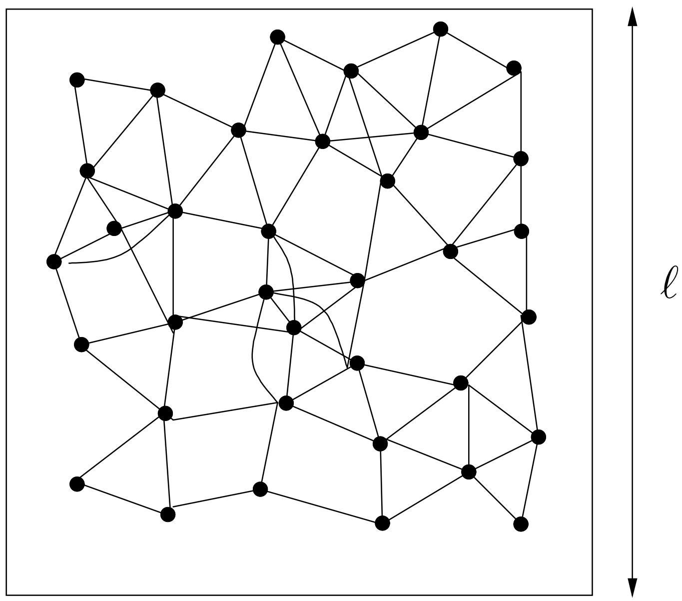

Random geometric graph model

Random distributions in $\mathbb{R}^2$:

need to understand geometric properties such as:

– connectivity

– **path length**

– coverage

$\mathcal{G}_{n,r,\ell}$

- $[0, \ell]^2$
- $\ell \to \infty$
- Structural properties depend on relation among r, n and ℓ.

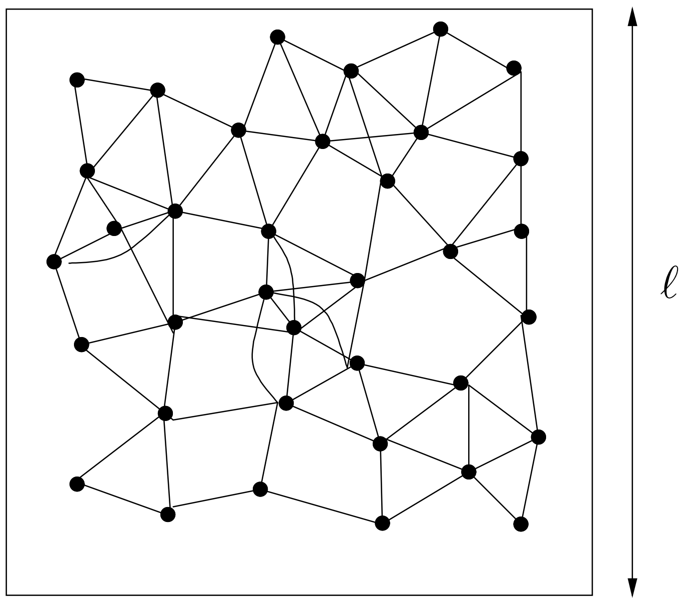

Random geometric graph model

Route-stretch

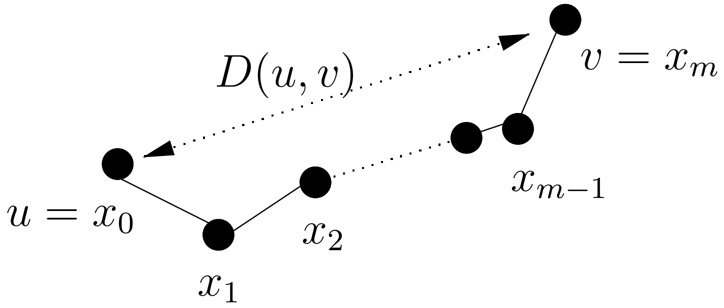

$$\frac{d_{min}(u,v)}{d_{opt}(u,v)} = \frac{\sum_{i=0}^{m-1} D(x_i,x_{i+1})}{D(u,v)}$$

Hop-stretch

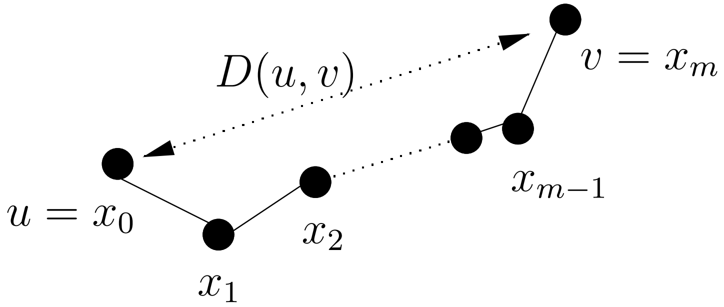

$$\frac{d_{min}(u,v)}{d_{opt}(u,v)} = \frac{m}{D(u,v)/r}$$

This talk

- Problem details.
- **Related work.**
- Our results.
- Future work.

Related work

Geometric properties

- Connectivity in $\mathcal{G}_{n,r}$.

  [Gupta,Kumar,98]

  Graph is connected *a.a.s.* when $\pi r^2 = \frac{\log n + c(n)}{n}$ if $c(n) \to \infty$.

- Threshold properties in $\mathcal{G}_{n,r,\ell}$.

  [Muthukrishnan,Pandurangan,03]

  * Physical coverage when $r^2 n \in \Theta(\ell^2)$ *a.a.s.*
  * Graph connectivity when $r^2 n \in \Theta(\ell^2 \ln \ell)$ *a.a.s.*
  * Route stretch of $1 + \frac{\alpha}{2}$ when $r^2 n \in \Omega\left(\frac{1}{\alpha}\ell^2 \ln \ell\right)$*a.a.s.*

- Threshold properties in $\mathcal{G}_{n,r}$.

  [Goel,Rai,Krishnamachari,04]

  All monotone graph properties have a sharp threshold for random geometric graphs.

Related work

Network bootstrapping

- Sensor networks.

  [Sohrabi et al.,00] Flat topology.

  Number of channels function of density.

  [Blough et al., 03] k-neighbors protocol.

  Distance estimation.

  [Song et al., 04] OrdYaoGG structure power spanner.

  Distance estimation, directional antenna.

  All: memory size function of density and no contention resolution in the analysis.

- Bluetooth: scatternet formation.

  [Salonidis et al., 01] One-hop network.

  [Barrière et al. 03] One-hop network, max 32 nodes.

  Other scatternet formation in multi-hop networks are heuristic.

Related work

Related problems

Clustering, dominating set, maximal independent set, leader election, vertex coloring, etc. Most of the solutions assume underlying communication infrastructure.

- MIS

  [Moscibroda,Wattenhofer,04]

  3 channels of communication.

  proof of correctness is broken.

  $\Omega(\log^6 n / \log^2 \log n)$ for one channel.

This talk

- Problem details.
- Related work.
- **Our results.**
  - Disk-cover algorithm.
  - Proof of hop-optimality.
  - Proof of $O(1)$ degree.
- Future work.

Our results

Geometric properties

- There exists a hop-optimal subgraph for any connected random geometric graph, even under a constant-degree assumption.

Network bootstrapping

- $O(\log^3 \ell)$ localized algorithm to build the network modelled by such a graph within the Weak Sensor Model.

Byproduct

- One-channel $O(\log^2 \ell)$ MIS distributed algorithm with contention resolution.

All *with high probability.*

Our results

Geometric properties

We want:

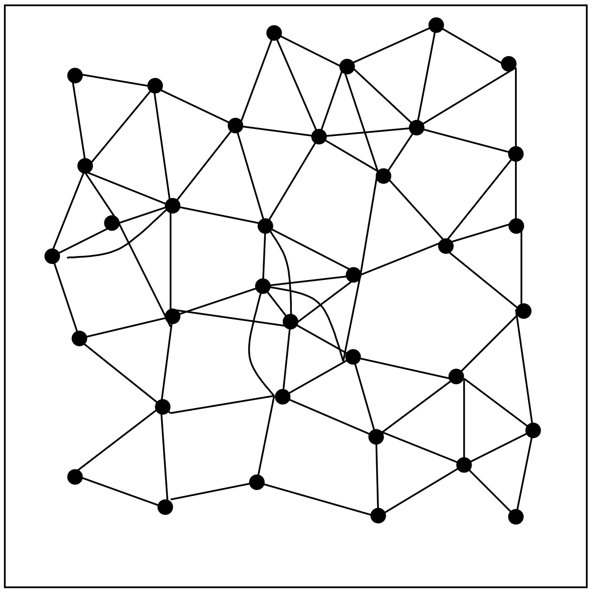

Random geometric graph.

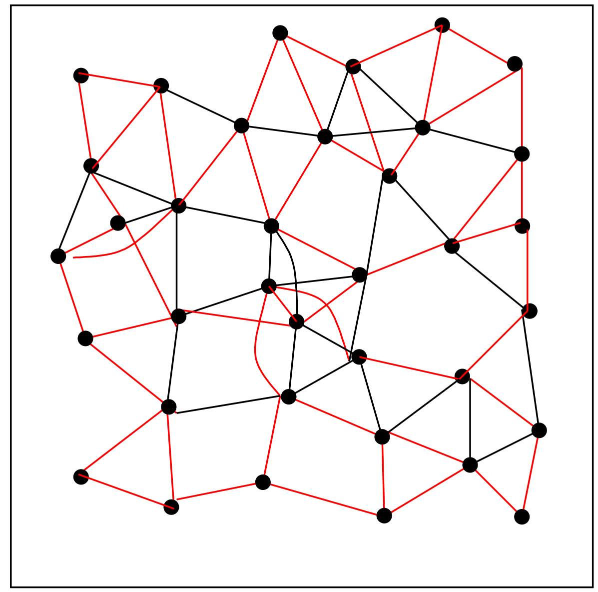

Hop-optimal constant-degree subgraph.

How?

- Define a cover of the rgg with disks.
- Glue all disks using *bridges*.
- Connect all nodes within each disk to its bridge.

Disk-cover algorithm

Given a threshold graph, find an *overlaid graph* as follows:

- Define an MIS with radius $ar/2$ among the nodes ($0 < a < 1$).
- Designate all MIS members as *bridges*.
- Connect all bridges within a distance of r.
- Lay down a disk of radius $r/2$ centered on each bridge.
- Construct a constant-degree spanner within each disk.

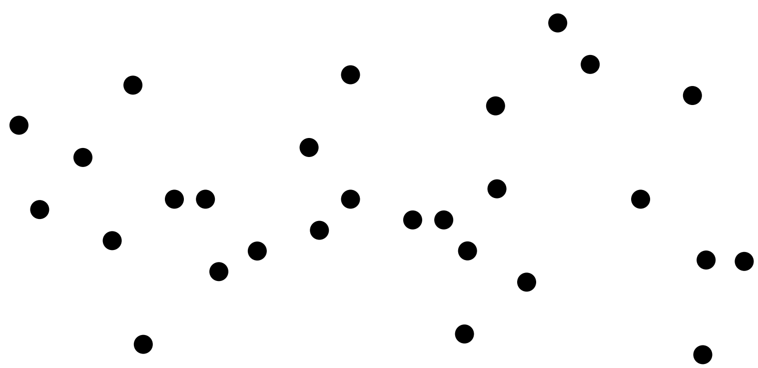

Disk-cover algorithm

Given a threshold graph, find an *overlaid graph* as follows:

- **Define an MIS with radius $ar/2$ among the nodes ($0 < a < 1$).**
- **Designate all MIS members as *bridges*.**
- Connect all bridges within a distance of r.
- Lay down a disk of radius $r/2$ centered on each bridge.
- Construct a constant-degree spanner within each disk.

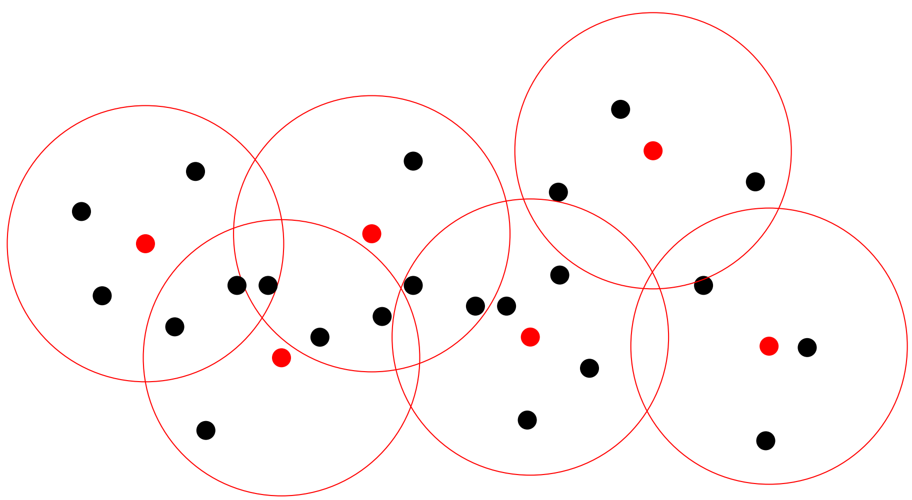

Disk-cover algorithm

Given a threshold graph, find an *overlaid graph* as follows:

- **Define an MIS with radius $ar/2$ among the nodes ($0 < a < 1$).**
- **Designate all MIS members as *bridges*.**
- **Connect all bridges within a distance of r.**
- Lay down a disk of radius $r/2$ centered on each bridge.
- Construct a constant-degree spanner within each disk.

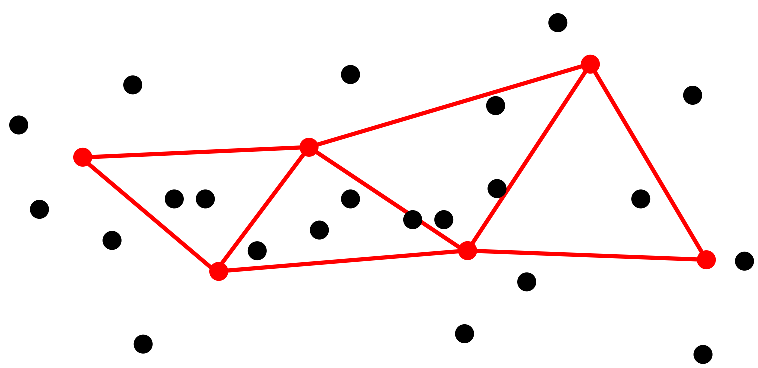

Disk-cover algorithm

Given a threshold graph, find an *overlaid graph* as follows:

- **Define an MIS with radius $ar/2$ among the nodes ($0 < a < 1$).**
- **Designate all MIS members as *bridges*.**
- **Connect all bridges within a distance of r.**
- **Lay down a disk of radius $r/2$ centered on each bridge.**
- Construct a constant-degree spanner within each disk.

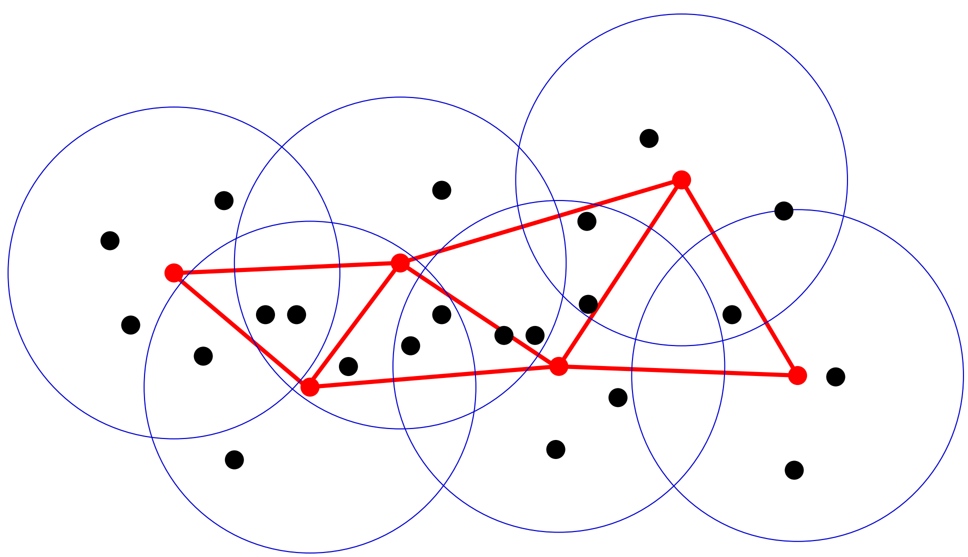

Disk-cover algorithm

Given a threshold graph, find an *overlaid graph* as follows:

- **Define an MIS with radius $ar/2$ among the nodes ($0 < a < 1$).**
- **Designate all MIS members as *bridges*.**
- **Connect all bridges within a distance of r.**
- **Lay down a disk of radius $r/2$ centered on each bridge.**
- **Construct a constant-degree spanner within each disk.**

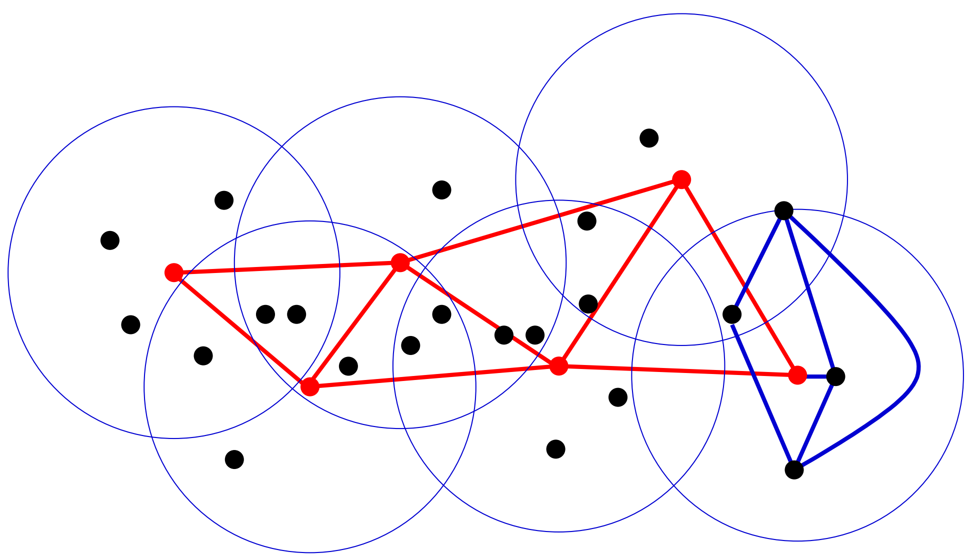

This talk

- Problem details.
- Related work.
- Our results.
  - Disk-cover algorithm.
  - **Proof of hop-optimality.**
  - Proof of $O(1)$ degree.
- Future work.

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u,v)/r)$ *short* edges,

  i.e. edges of length $\leq \frac{1-a}{b}r$ for any constant $b > 1$.

  Proof: If the density of nodes is $\frac{n}{\ell^2} > 6\frac{4+\alpha^2}{\alpha}\left(\frac{b}{1-a}\right)^2\frac{\ln \ell}{r^2}$

  where $r = \theta(\ell^{\epsilon} f(\ell))$,$f(\ell) \in o(\ell^{\gamma})$, $\gamma > 0$, $0 \leq \epsilon < 1$, $0 < \alpha \leq 1$,

  then there is a path in the threshold graph of $\leq \left\lceil \frac{D(u,v)}{r}\frac{b\sqrt{4+\alpha^2}}{1-a} \right\rceil$ *short* edges w.h.p.

  The points are sufficiently dense to guarantee the existence of such a path.

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u, v)/r)$ *short* edges,

  i.e. edges of length $\leq \frac{1-a}{b} r$ for any constant $b > 1$.

- Each short edge is completely covered by one disk.

  Proof: details to follow.

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u,v)/r)$ *short* edges,

  i.e. edges of length $\leq \frac{1-a}{b}r$ for any constant $b > 1$.
- Each short edge is completely covered by one disk.
- The number of nodes in any disk of radius $\Theta(r)$ is $\Theta(\log \ell)$.

  Proof: Chernoff bounds on a uniform distribution.

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u,v)/r)$ *short* edges, i.e. edges of length $\leq \frac{1-a}{b} r$ for any constant $b > 1$.
- Each short edge is completely covered by one disk.
- The number of nodes in any disk of radius $\Theta(r)$ is $\Theta(\log \ell)$.

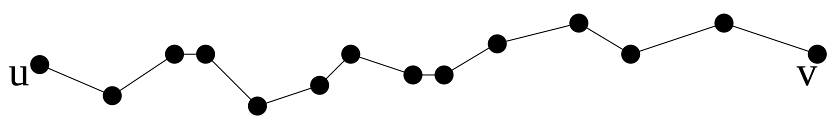

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u,v)/r)$ *short* edges, i.e. edges of length $\leq \frac{1-a}{b}r$ for any constant $b > 1$.
- Each short edge is completely covered by one disk.
- The number of nodes in any disk of radius $\Theta(r)$ is $\Theta(\log \ell)$.

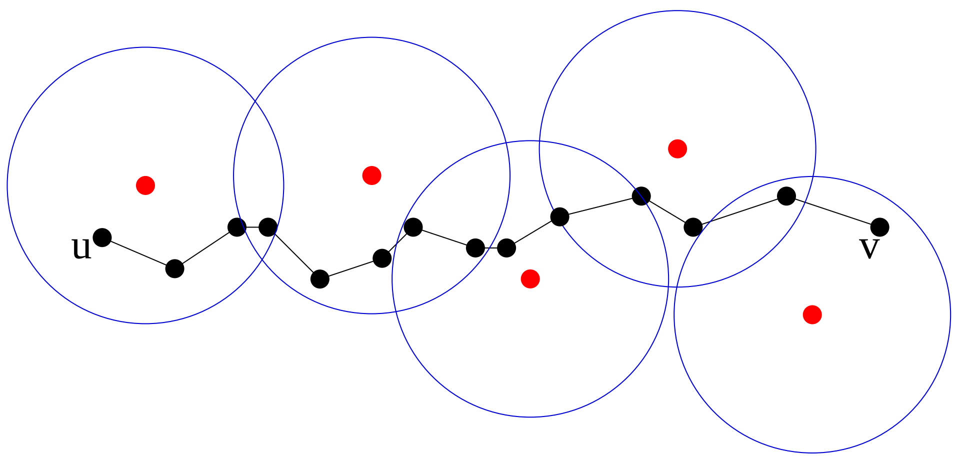

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u,v)/r)$ *short* edges, i.e. edges of length $\leq \frac{1-a}{b} r$ for any constant $b > 1$.
- Each short edge is completely covered by one disk.
- The number of nodes in any disk of radius $\Theta(r)$ is $\Theta(\log \ell)$.

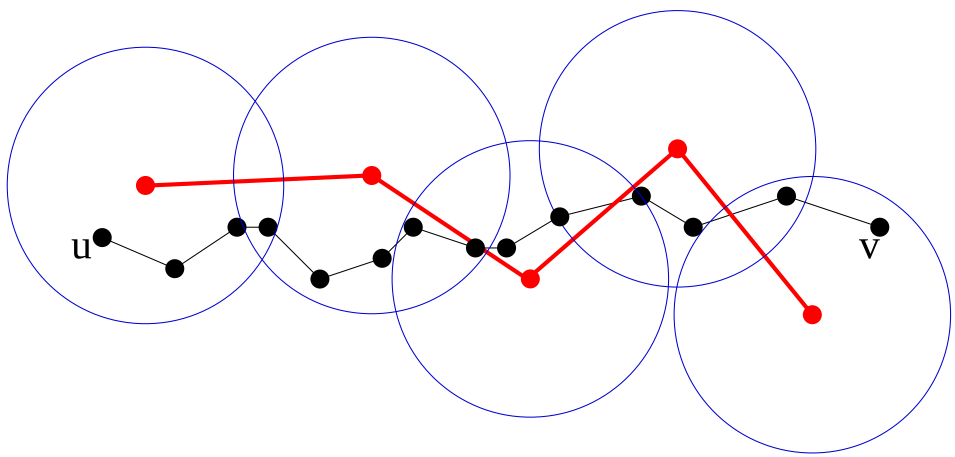

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u,v)/r)$ *short* edges, i.e. edges of length $\leq \frac{1-a}{b}r$ for any constant $b > 1$.
- Each short edge is completely covered by one disk.
- The number of nodes in any disk of radius $\Theta(r)$ is $\Theta(\log \ell)$.

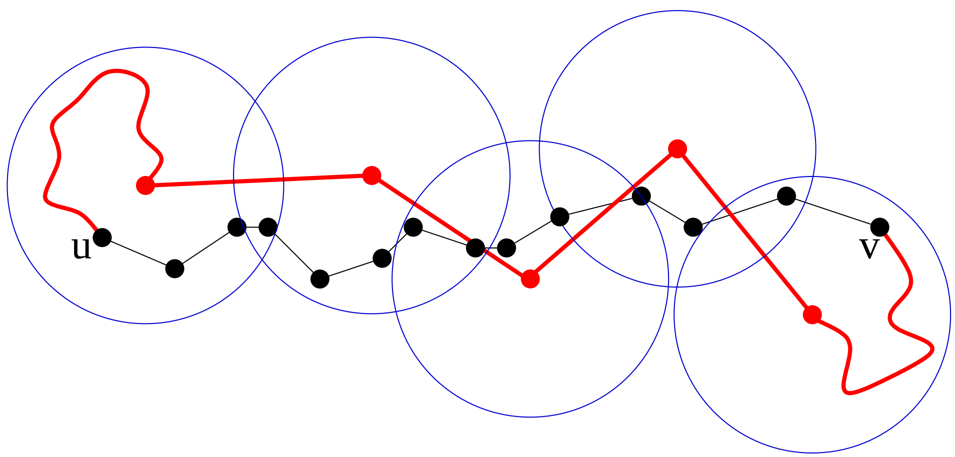

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u,v)/r)$ *short* edges, i.e. edges of length $\leq \frac{1-a}{b}r$ for any constant $b > 1$.
- Each short edge is completely covered by one disk.
- The number of nodes in any disk of radius $\Theta(r)$ is $\Theta(\log \ell)$.

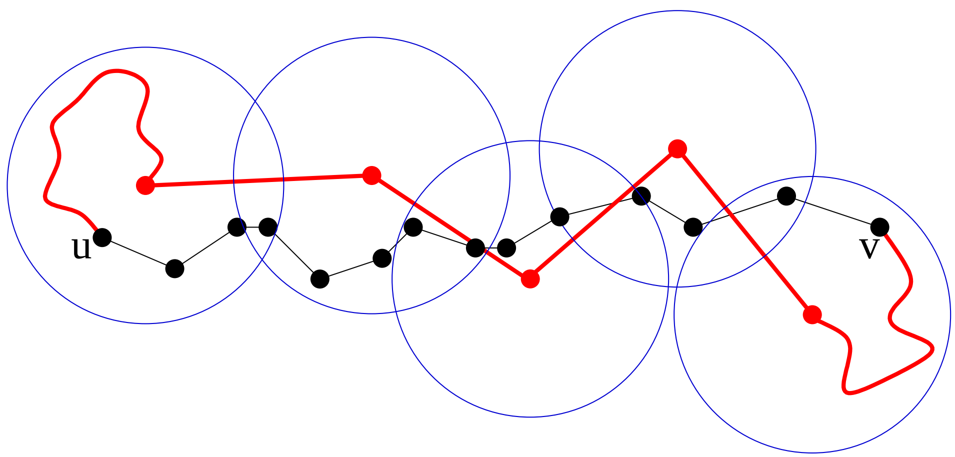

Theorem: $d(u,v) \in O(D(u,v)/r + \log \ell)$ is asymptotically optimal.

Proof of hop-optimality

What is the optimal path between u and v?

Lemmas:

- There is a path in the threshold graph of $O(D(u,v)/r)$ *short* edges,

  i.e. edges of length $\leq \frac{1-a}{b}r$ for any constant $b > 1$.
- **Each short edge is completely covered by one disk.**
- The number of nodes in any disk of radius $\Theta(r)$ is $\Theta(\log \ell)$.

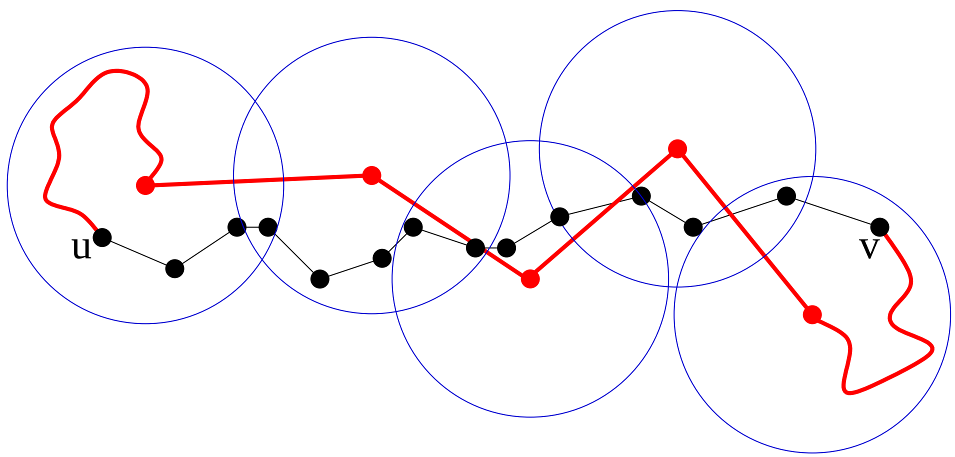

Theorem: $d(u,v) \in O(D(u,v)/r + \log \ell)$ is asymptotically optimal.

Proof of hop-optimality

Each short edge is completely covered by one disk

Length of short edges $\leq (1-a)r/b$ for some constant $b > 1$.

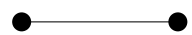

Proof of hop-optimality

Each short edge is completely covered by one disk

Length of short edges $\leq (1-a)r/b$ for some constant $b > 1$.

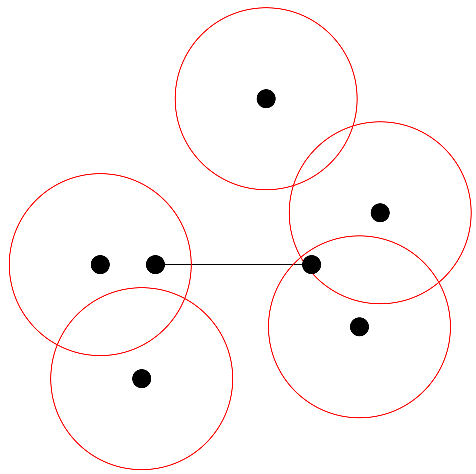

Proof of hop-optimality

Each short edge is completely covered by one disk

Length of short edges $\leq (1-a)r/b$ for some constant $b > 1$.

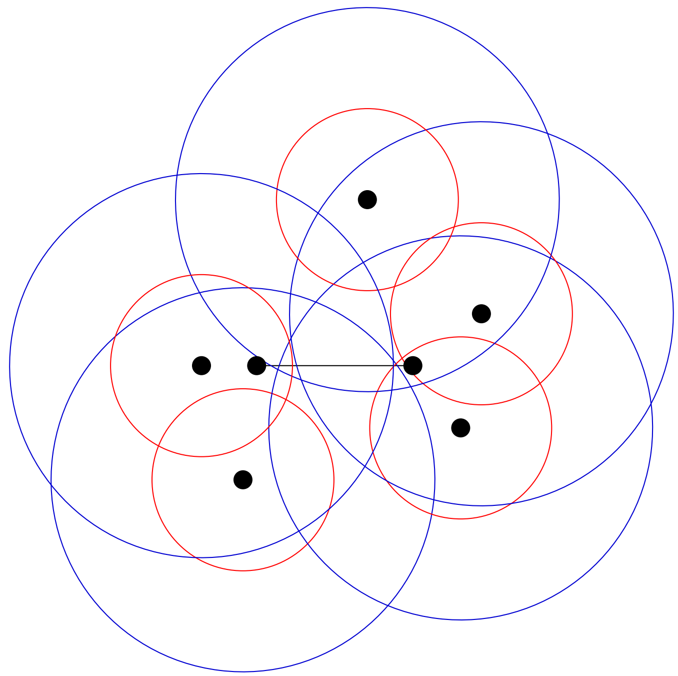

Proof of hop-optimality

Each short edge is completely covered by one disk

Length of short edges $\leq (1-a)r/b$ for some constant $b > 1$.

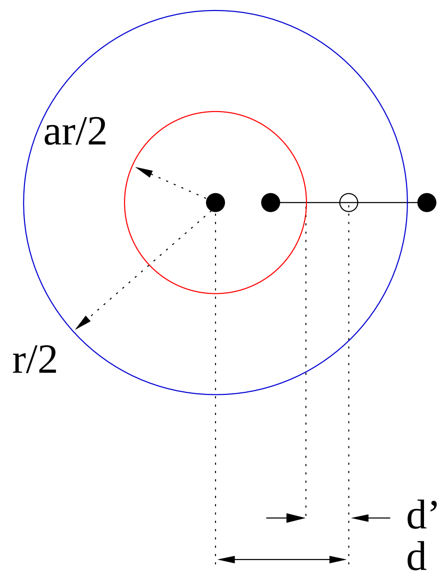

$$d > \frac{r}{2} - \frac{(1-a)r}{2b}$$

$$d' > \frac{r}{2} - \frac{(1-a)r}{2b} - \frac{ar}{2}$$

Proof of hop-optimality

Each short edge is completely covered by one disk

Length of short edges $\leq (1-a)r/b$ for some constant $b > 1$.

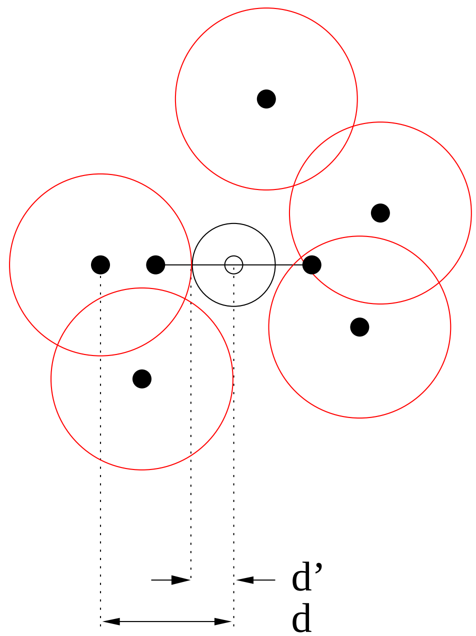

$$d > \frac{r}{2} - \frac{(1-a)r}{2b}$$
$$d' > \frac{r}{2} - \frac{(1-a)r}{2b} - \frac{ar}{2}$$

This talk

- Problem details.
- Related work.
- Our results.
  - Disk-cover algorithm.
  - Proof of hop-optimality.
  - **Proof of $O(1)$ degree.**
- Future work.

Proof of $O(1)$ degree

Bridge nodes

- Bridges are separated by a distance at least $ar/2$ w.h.p.
- Bridges are interconnected within a radius of r w.h.p.

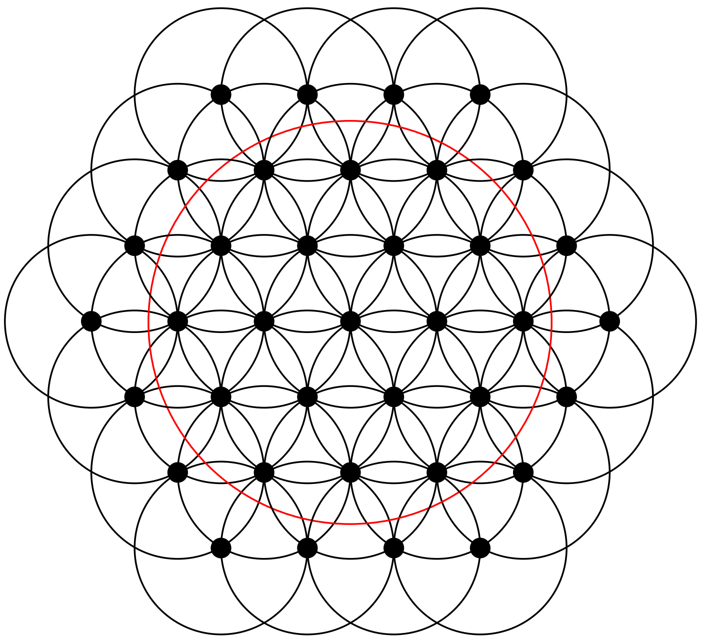

There are $\leq 3\lceil \frac{4}{a\sqrt{3}} \rceil \left(\lceil \frac{4}{a\sqrt{3}} \rceil + 1 \right)$ bridges in any disk of radius r.

Proof of $O(1)$ degree

Non-bridge nodes

- Connected by a constant-degree spanner.
- Covered by a constant number of discs.

Trade-off

Among the length of the optimal path ...

There is a path of $\leq \left\lceil \frac{D(u,v)}{r} \frac{b\sqrt{4+\alpha^2}}{1-a} \right\rceil + O(\log \ell)$ hops w.h.p.

... the maximum degree ...

The degree of any bridge is $\leq 3\lceil \frac{4}{a\sqrt{3}} \rceil \left(\lceil \frac{4}{a\sqrt{3}} \rceil + 1 \right) + 1$ w.h.p.

... and the density ...

The density of nodes is $\frac{n}{\ell^2} > 6 \frac{4+\alpha^2}{\alpha} \left(\frac{b}{1-a} \right)^2 \frac{\ln \ell}{r^2}$.

... where $0 < a < 1$, $b > 1$ and $0 < \alpha \leq 1$.

Longer edges covered $\Longrightarrow$ lower density $\Longrightarrow$ smaller number of hops $\Longrightarrow$ bigger degree.

Our results

Geometric properties

- There exists a hop-optimal subgraph for any connected random geometric graph, even under a constant-degree assumption.

Network bootstrapping

- $O(\log^3 \ell)$ localized algorithm to build the network modelled by such a graph within the Weak Sensor Model.

Byproduct

- One-channel $O(\log^2 \ell)$ MIS distributed algorithm with contention resolution.

All *with high probability.*

This talk

- Problem details.
- Related work.
- Our results.
- **Future work.**

Future work

- Faster network bootstrapping algorithm.
- Lower bounds for MIS for uniform and non-uniform distribution of nodes.
- Extensions of the MIS algorithm to other problems such as coloring.
- Positioning based on local distance estimation.
- Routing in this network.

Thank you

Localized algorithm

For each node i in parallel

Run the MIS algorithm with range $ar/2$

If $i \in$ MIS

Designate i as a bridge

Connect to neighboring bridges by broadcasting ID with range r

Lay down a disk of radius $r/2$ centered in i by broadcasting with range $r/2$

Connect with disk neighbors forming a constant-degree spanner.

MIS algorithm

Initialize a counter to 0.

Repeat

Broadcast the counter with probability $1/\delta_1 \log \ell$.

Else

If a counter was received and $|counterreceived - counter| \leq \lceil \delta_2 \log \ell \rceil$ then

Set counter to $-\lceil \delta_2 \log \ell \rceil$.

If an MIS member ID was received then stop.

If this node has ever transmitted, increase the counter.

If the counter has reached $\lceil \delta_3 \log^2 \ell \rceil$ then

This node declares itself an MIS member.

Repeat

Broadcast the ID with probability $1/\delta_4$

Spanner construction algorithm

Bridge nodes

Asign local index to non-bridge nodes upon request.

Non-bridge nodes

Obtain a local index from the bridge.

Connect to current neighbors to form a butterfly network.

Handle new neighbor arrivals.

Hop-stretch

A local spanner of small diameter: Hamilton-expander

- If G is δ-regular $\rightarrow \lambda_0 = \delta$ and $\lambda_{n-1} \geq -\delta$
- [AlMi85] If G is δ-regular $\rightarrow Diameter(G) \leq 2\sqrt{2\delta/\delta - \lambda_1}\log n$.
- [Al86][Fr03] Random δ-regular graphs $\rightarrow \lambda_1 \leq 2\sqrt{\delta - 1} + \epsilon$ for any $\epsilon > 0$ w.h.p.
- [Fr03] Same result for multigraphs composed of $\delta/2$ random Hamilton cycles with probability $O(1 - 1/n^\gamma)$ where $\gamma = \lceil\sqrt{\delta - 1}\rceil - 1$.

If G is a multigraph on n nodes composed of $\delta/2$ random Hamilton cycles:

$$Diameter(G) \in O(\log n) \text{ with probability } O(1 - 1/n^\gamma)\text{ , }\gamma = \lceil\sqrt{\delta - 1}\rceil - 1$$

But, within a given disk, there are $O(\log \ell)$ nodes, then:

$$Diameter(Hamilton - expander) \in O(\log\log \ell) \text{ with probability } O(1 - 1/\log^\gamma \ell)$$

Hamilton-expander algorithm

Bridge nodes

Initialize an index to 0.

Repeat

If an index request is received then

Increase index.

Send the current index for δ_6 steps with probability $1/\beta_3$.

Non-bridge nodes

Phase 1: Ordering the nodes locally using the bridge

Initialize a counter to 0.

Repeat

With probability $1/\beta_3$ request a new index from the bridge.

If not requesting and an index is received then stop.

Increase the counter.

Wait for $(\delta_6 \log^3 n - counter + (index - 1)\delta_7 \log^2 n)$ steps.

Hamilton-expander algorithm

Non-bridge nodes

Phase 2: Joining the Hamilton-expander

Choose d nodes at random in the index range $[1, index - 1]]$.

For τ_1 steps, request the ID's of the chosen nodes and their successors.

Repeat

If an ID is received then update linked list.

If all answers were received then stop.

Phase 3: Handling insertion requests

Repeat

If an ID request is received then

Broadcast the ID for τ_2 steps with probability $1/\beta_4$.

Proof of hop-optimality

A path of $O(D(u,v)/r + \log \ell)$ hops is asymptotically optimal

$D(u,v)/r$ is a lower bound of the length of an optimal path.

In a δ-regular graph:

$$Pr(d(u,v) < c \log n) \leq \frac{1}{n-1} \sum_{i=0}^{c \log n - 2} \delta(\delta-1)^i \in O(n^{-\gamma})$$

Thus, in $G(n,r,\ell)$, where $r^2 n = k\ell^2 \ln \ell$, $r = \theta(\ell^\epsilon f(\ell))$,$f(\ell) \in o(\ell^\gamma)$, $\gamma > 0$, $0 \leq \epsilon < 1$.

$$d(u,v) \in \Omega(\log \ell) \text{ w.h.p.}$$

Hence, $(D(u,v)/r + \log \ell)/2$ is a lower bound of the length of such a path.

Proof of hop-optimality

The number of nodes in any disk of radius $\Theta(r)$ is $\Theta(\log \ell)$.

Consider $G(n, r, \ell)$, where $r^2 n = k\ell^2 \ln \ell$, $r = \theta(\ell^\epsilon f(\ell))$,$f(\ell) \in o(\ell^\gamma)$, $\gamma > 0$, $0 \leq \epsilon < 1$.

Consider a circle of radius βr for any constant $\beta > 0$.

The probability of falling in the circle is $\pi\beta^2 r^2/\ell^2$.

Using Chernoff bounds and the parameter conditions:

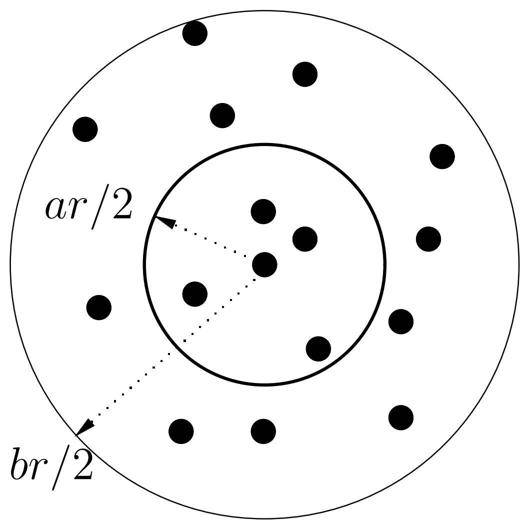

$$Pr(X \geq (1+\epsilon)\pi\beta^2 k \ln \ell) \leq \ell^{-\frac{\epsilon^2 \pi \beta^2 k}{3}}$$

$$Pr(X \leq (1-\epsilon)\pi\beta^2 k \ln \ell) \leq \ell^{-\frac{\epsilon^2 \pi \beta^2 k}{2}}$$

Proof of hop-optimality

There is a path in the threshold graph of $O(D(u,v)/r)$ short edges

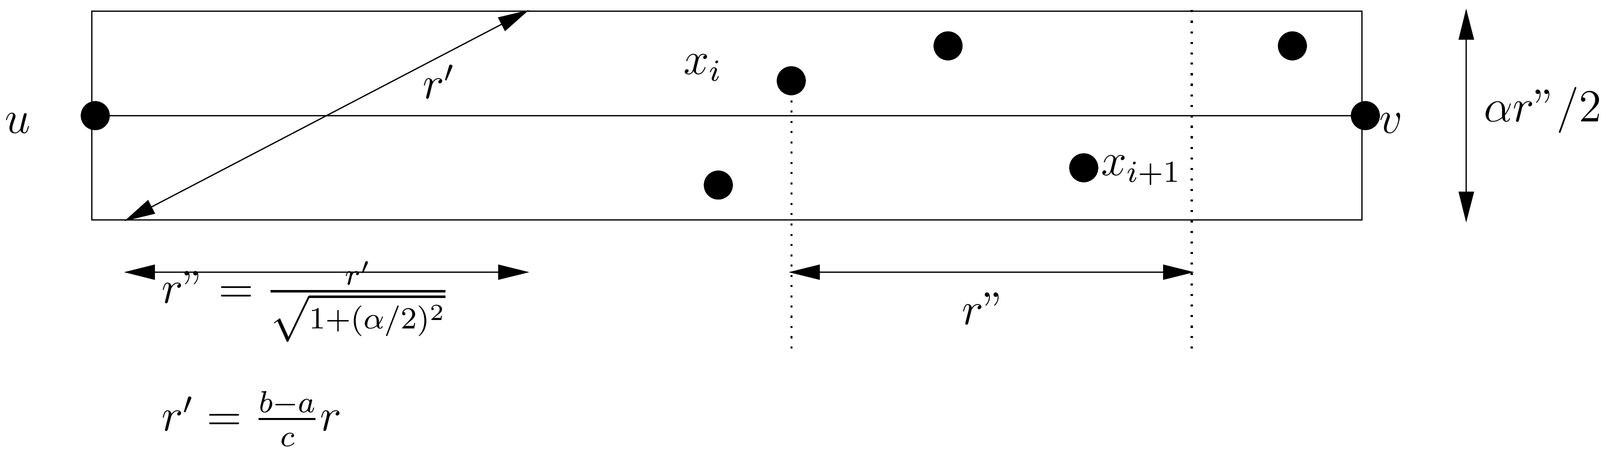

$0 < \alpha \leq 1$

Path definition: for any node x_i

- The node x_{i+1} lies in the strip.
- $D_h(x_i, x_{i+1}) \leq r''$.
- The horizontal distance $D_h(x_{i+1}, v)$ is minimized.

Proof of hop-optimality

There is a path in the threshold graph of $O(D(u,v)/r)$ short edges

Assume there is no *hole*.

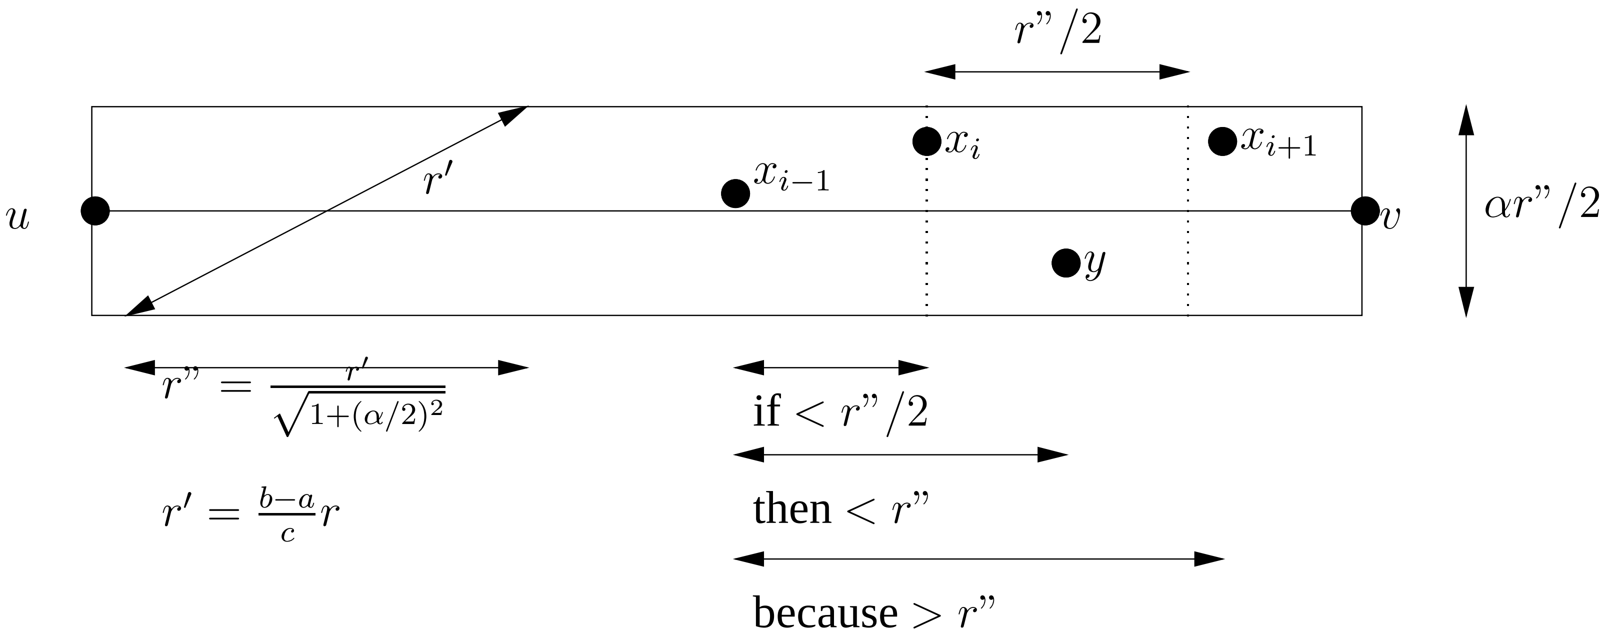

Since $D_h(x_i, x_{i+1}) \geq r''/2$ for $0 \leq i < m$,

$$d(u,v) \leq \left\lceil 2\frac{D_h(u,v)}{r''} \right\rceil = \left\lceil 2\frac{D(u,v)}{r} \frac{b\sqrt{1+(\alpha/2)^2}}{1-a} \right\rceil \in O(D(u,v)/r) \text{ hops.}$$

Proof of hop-optimality

There is no hole within a strip

$G(n, r', \ell)$, where $r'^2 n = k\ell^2 \ln \ell$, $r' = \theta(\ell^\epsilon f(\ell))$,$f(\ell) \in o(\ell^\gamma)$, $\gamma > 0$, $0 \leq \epsilon < 1$.

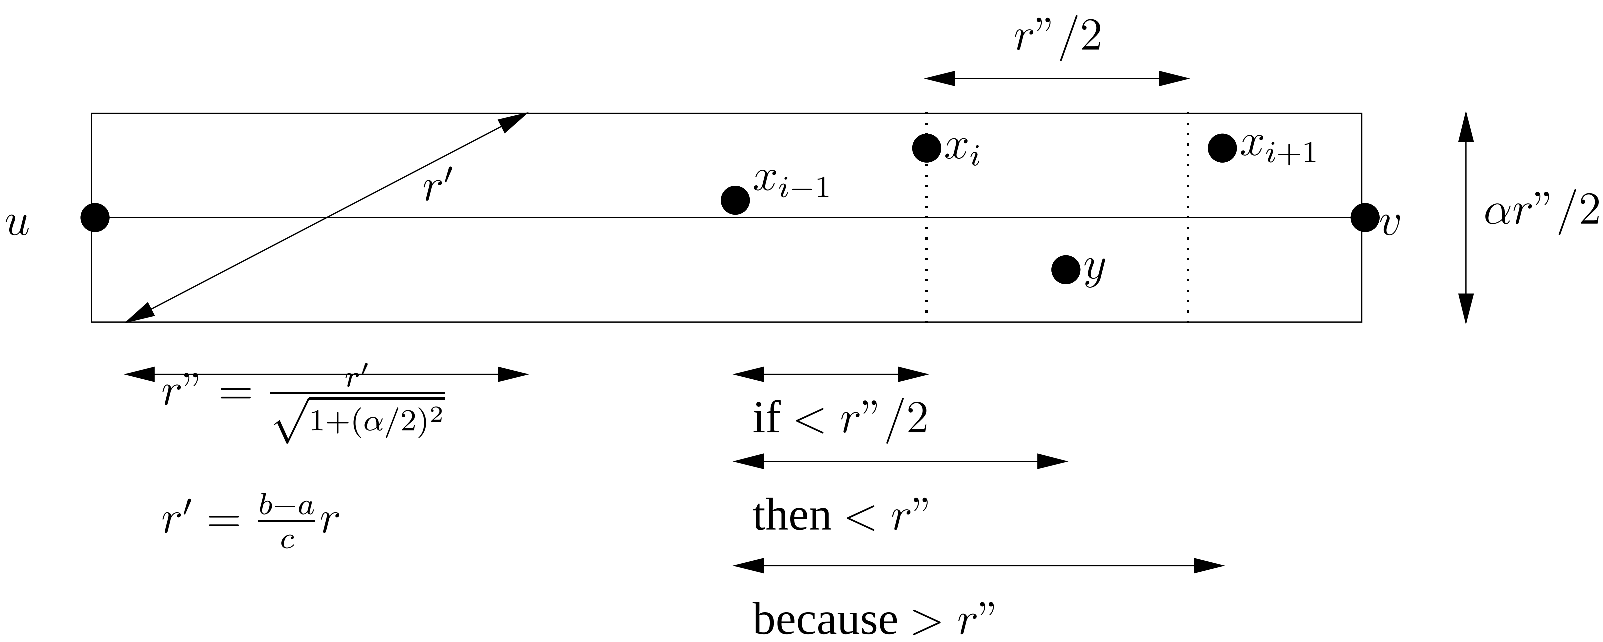

$$Pr[\text{Hole}] \leq \binom{n}{2} n \frac{\alpha r''}{\sqrt{2}\ell} \left(1 - \frac{\alpha r''^2}{4\ell^2}\right)^{n-1}$$
$$\in O(\ell^{-\gamma}) \text{ for } k > 6\frac{4+\alpha^2}{\alpha}$$

Optimization criteria

Maximize life cycle subject to the Weak Sensor Model constraints.

- Minimize transmission power:
  - Polynomial in the distance.
  - Power cannot be adjusted to any number of levels.
  - Not clear how to minimize.

- Minimize transmission count:
  - Transmission count dominated by contention resolution.
  - Each hop in a path requires a new round of contention.
  - Transmission count can be minimized: **the relevant measure is the number of hops**.