Randomized Rumour Spreading

joint work with Benjamin Doerr, Spyros Angelopoulos, Nikolaos Fountoulakis, Ariel Levavi, and Konstantinos Panagiotou

Anna Huber

Durham University (soon to be: Derby University)

Graph Theory and Interactions, July 2013

Graph Theory and Interactions...

Graph Theory and Interactions...

... between neighbouring vertices!

For example:

- Broadcasting updates in distributed databases.

Durham University

For example:

- Broadcasting updates in distributed databases.

Durham University

For example:

- Broadcasting updates in distributed databases.
- Mathematics of infectious diseases.

Randomized Algorithms

Durham
University

Randomized Algorithms

- Efficient and simple.

Durham
University

Randomized Algorithms

- Efficient and simple.
- Robust.

Randomized Algorithms

- Efficient and simple.
- Robust.
- Good average properties.

Randomized Algorithms

- Efficient and simple.
- Robust.
- Good average properties.

Examples:

- Monte Carlo algorithms.

Randomized Algorithms

- Efficient and simple.
- Robust.
- Good average properties.

Examples:

- Monte Carlo algorithms.
- Random walk.

Randomized Algorithms

- Efficient and simple.
- Robust.
- Good average properties.

Examples:

- Monte Carlo algorithms.
- Random walk.
- Randomized rumour spreading.

Randomized Rumour Spreading

Durham University

Randomized Rumour Spreading

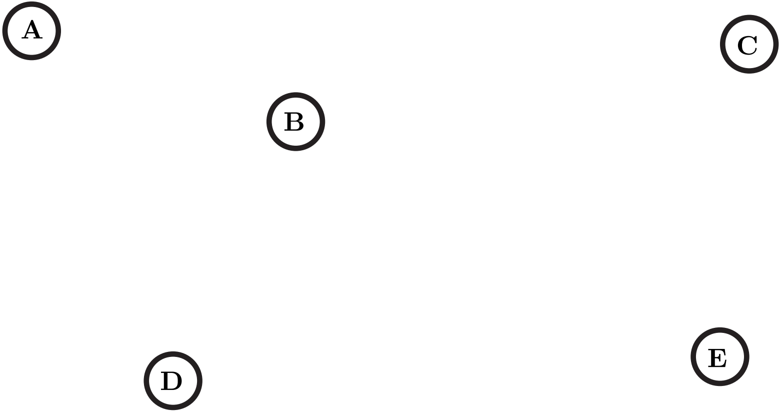

Randomized Rumour Spreading

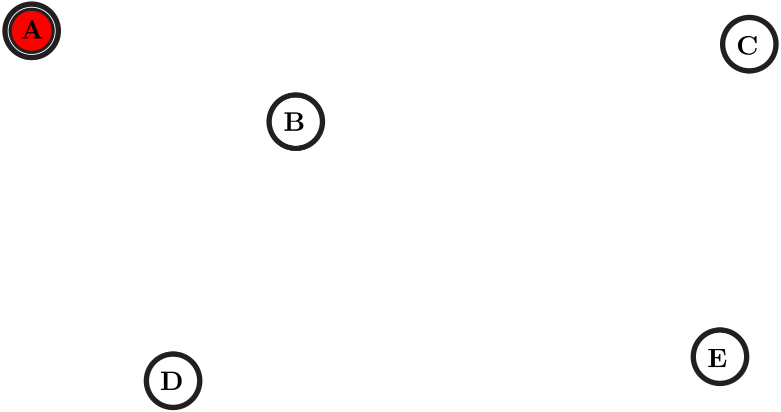

Durham University

Randomized Rumour Spreading

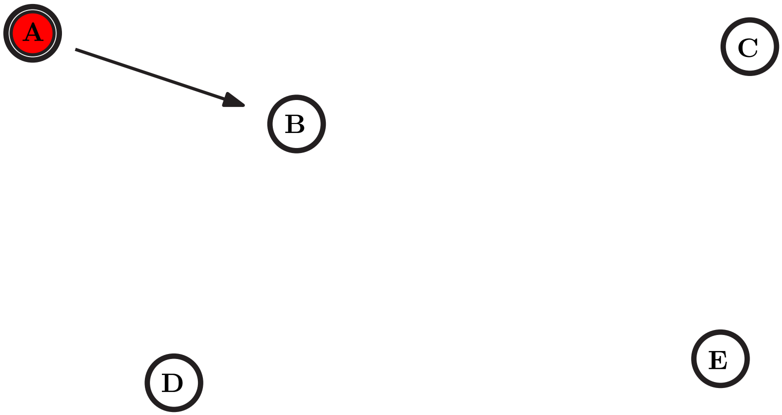

Durham University

Randomized Rumour Spreading

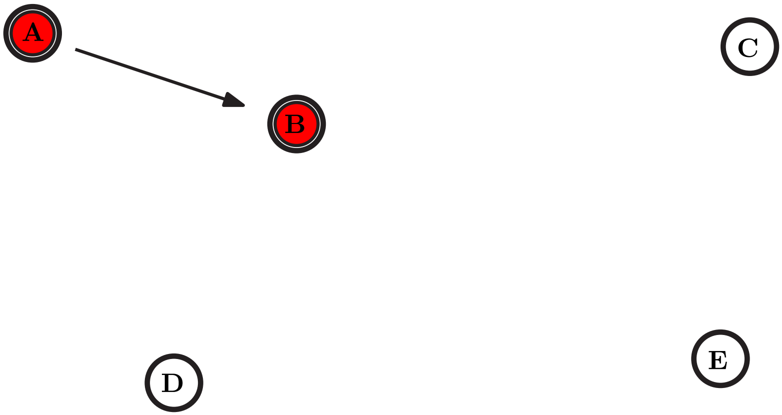

Durham University

Randomized Rumour Spreading

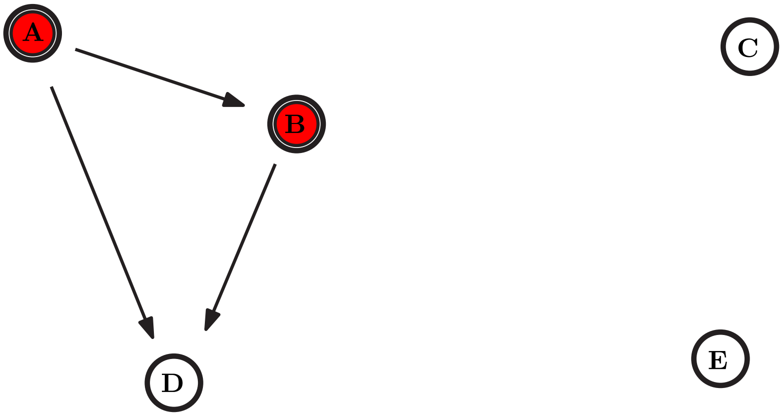

Randomized Rumour Spreading

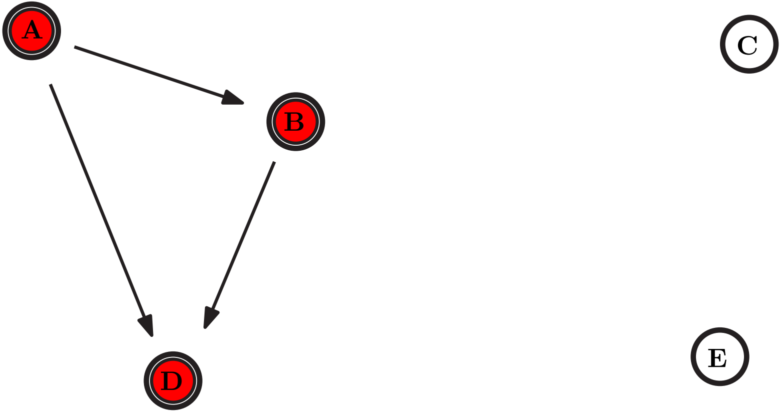

Randomized Rumour Spreading

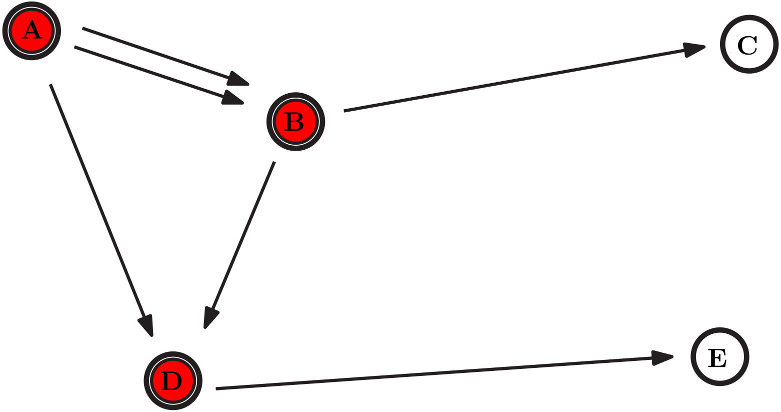

Durham University

Randomized Rumour Spreading

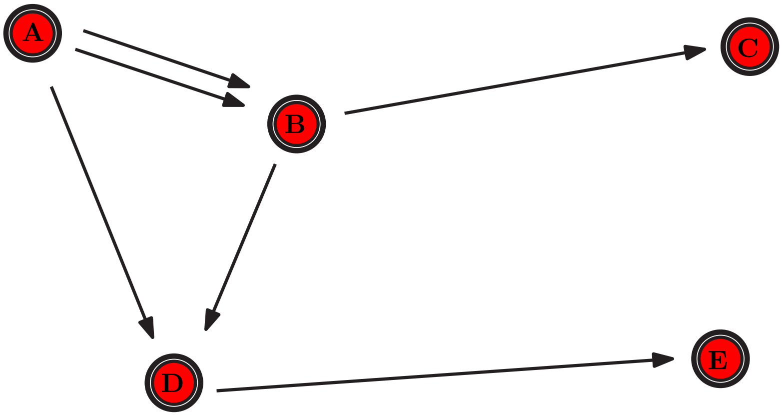

Randomized Rumour Spreading

Question

How long does it take until all n vertices are informed?

Randomized Rumour Spreading

Question

How long does it take until all n vertices are informed?

- At least $\log_2 n$.

Randomized Rumour Spreading

Question

How long does it take until all n vertices are informed with high probability (with probability $1 - o(1)$)?

- At least $\log_2 n$.

Durham University

Complete Graphs

Complete Graphs

Theorem (Frieze and Grimmett 1985)

The randomized model on the complete graph informs all n vertices within

$$(1 \pm o(1))(\log_2 n + \ln n)$$

time-steps with high probability.

Durham University

Complete Graphs

Theorem (Pittel 1987)

For every $h \in \omega(1)$ the randomized model on the complete graph informs all n vertices within

$$\log_2 n + \ln n \pm h(n)$$

time-steps with high probability.

Durham University

Hypercubes and Random Graphs

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

Theorem (Feige, Peleg, Raghavan and Upfal 1990)

The randomized model on hypercubes and random graphs $G_{n,p}$ with $p \geq \frac{(1+\epsilon)\ln n}{n}$ informs all n vertices within

$$\Theta(\log n)$$

time-steps with high probability.

Durham University

Leading Constant for Random Graphs

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

Theorem (Fountoulakis, H., Panagiotou, 2009)

The randomized model on the random graph $G_{n,p}$, where $p = \omega\left(\frac{\ln n}{n}\right)$, informs all vertices within

$$(1 \pm o(1))(\log_2 n + \ln n)$$

time-steps with high probability.

Durham University

Leading Constant for Random Graphs

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

Theorem (Fountoulakis, H., Panagiotou, 2009)

The randomized model on the random graph $G_{n,p}$, where $p = \omega\left(\frac{\ln n}{n}\right)$, informs all vertices within

$$(1 \pm o(1))(\log_2 n + \ln n)$$

time-steps with high probability.

- No dependence on p – Same as complete graph!

Durham University

Robustness

Robustness

- Messages successfully sent with constant probability $q \in [0, 1]$.

Robustness

- Messages successfully sent with constant probability $q \in [0, 1]$.
- Independent failures.

Robustness

- Messages successfully sent with constant probability $q \in [0, 1]$.
- Independent failures.
- $q = 1$: As before.
- What will the runtime be for $q \in]0, 1[$?

Durham University

Robustness

Theorem (Fountoulakis, H., Panagiotou, 2010)

The randomized model with transmission success probability q informs all vertices of the random graph $G_{n,p}$, where $p = \omega\left(\frac{\ln n}{n}\right)$, within

$$(1 \pm o(1))(\log_{1+q} n + \frac{1}{q} \ln n)$$

time-steps with high probability.

Durham University

Robustness

Theorem (Fountoulakis, H., Panagiotou, 2010)

The randomized model with transmission success probability q informs all vertices of the random graph $G_{n,p}$, where $p = \omega\left(\frac{\ln n}{n}\right)$, within

$$(1 \pm o(1))(\log_{1+q} n + \frac{1}{q} \ln n)$$

time-steps with high probability.

- Note that $\log_{1+q} n + \frac{1}{q} \ln n < \frac{1}{q}(\log_2 n + \ln n)$.

Durham University

Robustness

Theorem (Fountoulakis, H., Panagiotou, 2010)

The randomized model with transmission success probability q informs all vertices of the random graph $G_{n,p}$, where $p = \omega\left(\frac{\ln n}{n}\right)$, within

$$(1 \pm o(1))(\log_{1+q} n + \frac{1}{q} \ln n)$$

time-steps with high probability.

- Note that $\log_{1+q} n + \frac{1}{q} \ln n < \frac{1}{q}(\log_2 n + \ln n)$.
- No dependence on p – Same as complete graph!

Durham University

Leading Constant for Random Graphs

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

Theorem (Fountoulakis, H., Panagiotou, 2009)

The randomized model on the random graph $G_{n,p}$, where $p = \omega\left(\frac{\ln n}{n}\right)$, informs all vertices within

$$(1 \pm o(1))(\log_2 n + \ln n)$$

time-steps with high probability.

Durham University

How the Rumour Spreads

in $(1 \pm o(1))(\log_2 n + \ln n)$ time-steps:

Durham University

How the Rumour Spreads

in $(1 \pm o(1))(\log_2 n + \ln n)$ time-steps:

- "The $\log_2$-Part": $\sim \log_2 n$ steps.
  Nearly doublings with high probability. εn informed vertices.

Durham University

How the Rumour Spreads

in $(1 \pm o(1))(\log_2 n + \ln n)$ time-steps:

- "The $\log_2$-Part": $\sim \log_2 n$ steps.
  Nearly doublings with high probability. εn informed vertices.
- "few steps huge effect": $(1 - \epsilon)n$ informed vertices.

Durham University

How the Rumour Spreads

in $(1 \pm o(1))(\log_2 n + \ln n)$ time-steps:

- "The $\log_2$-Part": $\sim \log_2 n$ steps.
  Nearly doublings with high probability. εn informed vertices.
- "few steps huge effect": $(1 - \epsilon)n$ informed vertices.
- "The ln-Part": $\sim \ln n$ steps.
  All vertices informed.

Durham University

Proof Idea

Proof Idea

- Show that $G_{n,p}$ has "good" properties with high probability.

Proof Idea

- Show that $G_{n,p}$ has “good” properties with high probability.
- Bound the number of informed vertices after each round (from above and from below) for graphs with “good” properties.

The "good" Properties

The random graph $G_{n,p}$, where $p = \frac{\alpha(n) \ln n}{n}$ and $\alpha \in \omega(1)$, has w. h. p. the following properties. For any subset S of its vertices

The "good" Properties

The random graph $G_{n,p}$, where $p = \frac{\alpha(n) \ln n}{n}$ and $\alpha \in \omega(1)$, has w. h. p. the following properties. For any subset S of its vertices

- if $|S| \geq \frac{n}{\alpha(n)}$, most vertices outside S have $(1 \pm \alpha(n)^{-1/2})p|S|$ neighbours in S,

The "good" Properties

The random graph $G_{n,p}$, where $p = \frac{\alpha(n) \ln n}{n}$ and $\alpha \in \omega(1)$, has w. h. p. the following properties. For any subset S of its vertices

- if $|S| \geq \frac{n}{\alpha(n)}$, most vertices outside S have $(1 \pm \alpha(n)^{-1/2})p|S|$ neighbours in S,
- if $|S| \leq \frac{n}{\alpha(n)}$, most vertices outside S have at most ϵpn neighbours in S,

Durham University

The "good" Properties

The random graph $G_{n,p}$, where $p = \frac{\alpha(n) \ln n}{n}$ and $\alpha \in \omega(1)$, has w. h. p. the following properties. For any subset S of its vertices

- if $|S| \geq \frac{n}{\alpha(n)}$, most vertices outside S have $(1 \pm \alpha(n)^{-1/2})p|S|$ neighbours in S,
- if $|S| \leq \frac{n}{\alpha(n)}$, most vertices outside S have at most ϵpn neighbours in S,
- the cut of S contains $|S|(n - |S|)p \left(1 \pm \sqrt{\frac{8}{\alpha(n)}}\right)$ edges.

Durham University

Why Random Graphs?

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

Why Random Graphs?

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

- "Good" model for "real" networks...???

Why Random Graphs?

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

- "Good" model for "real" networks...???
- p is measure for density

Durham University

Why Random Graphs?

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

- "Good" model for "real" networks...???
- p is measure for density
- Quite regular

Why Random Graphs?

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

- "Good" model for "real" networks...???
- p is measure for density
- Quite regular
- Quite high expansion

Durham University

Why Random Graphs?

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

- "Good" model for "real" networks...???
- p is measure for density
- Quite regular
- Quite high expansion
- Connected $\left(\text{If } p \geq \frac{(1+\varepsilon)\ln n}{n}\right)$

Durham University

Why Random Graphs?

Let $G_{n,p}$ denote the Erdős-Rényi random graph on n vertices with edge probability p.

- "Good" model for "real" networks...???
- p is measure for density
- Quite regular
- Quite high expansion
- Connected $\left(\text{If } p \geq \frac{(1+\varepsilon)\ln n}{n}\right)$
- Effort to design networks with properties of random graphs

Durham University

Why Not Random Graphs?

Why Not Random Graphs?

- “Good” model for “real” networks...???

Durham
University

Why Not Random Graphs?

- "Good" model for "real" networks...???
- maybe not?

Why Not Random Graphs?

- "Good" model for "real" networks...???
- maybe not?
- Too regular for real-world graphs?

Durham University

Why Not Random Graphs?

- "Good" model for "real" networks...???
- maybe not?
- Too regular for real-world graphs?
- Too "static" for real-world graphs?

Durham University

Related Work

Related Work

- Push – and pull! [Demers et al,1988]

Durham
University

Related Work

- Push – and pull! [Demers et al,1988]
- Random regular graphs and expanders [Fountoulakis, Panagiotou 2010]

Durham University

Related Work

- Push – and pull! [Demers et al,1988]
- Random regular graphs and expanders [Fountoulakis, Panagiotou 2010]
- Social networks [Fountoulakis, Panagiotou, Sauerwald / Doerr, Fouz, Friedrich 2012]

Durham University

Related Work

- Push – and pull! [Demers et al,1988]
- Random regular graphs and expanders [Fountoulakis, Panagiotou 2010]
- Social networks [Fountoulakis, Panagiotou, Sauerwald / Doerr, Fouz, Friedrich 2012]
- Dynamic random graph models [Clementi, Crescenzi, Doerr, Fraigniaud, Isopi, Panconesi, Pasquale, Silvestri 2013]

Related Work

- Push – and pull! [Demers et al,1988]
- Random regular graphs and expanders [Fountoulakis, Panagiotou 2010]
- Social networks [Fountoulakis, Panagiotou, Sauerwald / Doerr, Fouz, Friedrich 2012]
- Dynamic random graph models [Clementi, Crescenzi, Doerr, Fraigniaud, Isopi, Panconesi, Pasquale, Silvestri 2013]
- Safe termination criteria, memory of nodes...

Durham University

Related Work

- Push – and pull! [Demers et al,1988]
- Random regular graphs and expanders [Fountoulakis, Panagiotou 2010]
- Social networks [Fountoulakis, Panagiotou, Sauerwald / Doerr, Fouz, Friedrich 2012]
- Dynamic random graph models [Clementi, Crescenzi, Doerr, Fraigniaud, Isopi, Panconesi, Pasquale, Silvestri 2013]
- Safe termination criteria, memory of nodes...

not exhaustive – for more, see e. g.

Durham University

Related Work

- Push – and pull! [Demers et al,1988]
- Random regular graphs and expanders [Fountoulakis, Panagiotou 2010]
- Social networks [Fountoulakis, Panagiotou, Sauerwald / Doerr, Fouz, Friedrich 2012]
- Dynamic random graph models [Clementi, Crescenzi, Doerr, Fraigniaud, Isopi, Panconesi, Pasquale, Silvestri 2013]
- Safe termination criteria, memory of nodes...

not exhaustive – for more, see e. g.

- Dagstuhl Seminar 13042 "Epidemic Algorithms and Processes"

Durham University

Why Quasirandom?

Why Quasirandom?

Deterministic

- Fix runtime and result.
- Guarantees.

Durham University

Why Quasirandom?

Deterministic

- Fix runtime and result.
- Guarantees.

Randomized

- Efficient and simple.
- Robust.
- Good average properties.

Durham University

Why Quasirandom?

Deterministic

- Fix runtime and result.
- Guarantees.

Quasirandom

- Use dependencies.
- Reduce randomness.
- Keep or improve property of the random method.

Randomized

- Efficient and simple.
- Robust.
- Good average properties.

Durham University

Quasirandom Rumour Spreading

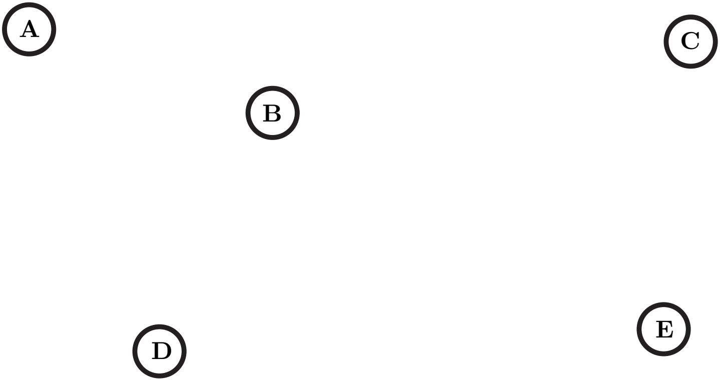

Quasirandom Rumour Spreading

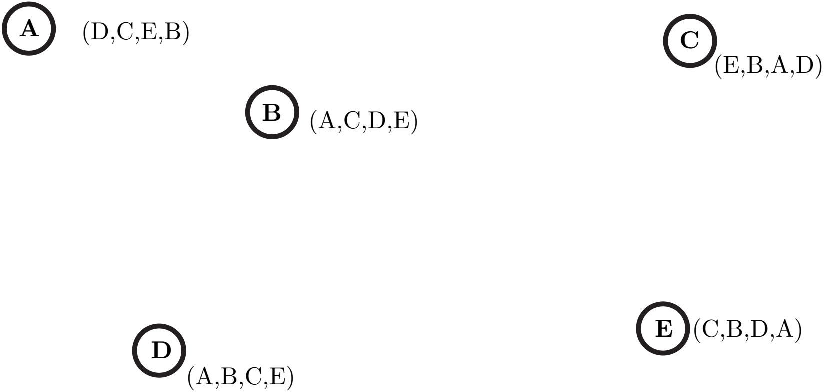

Quasirandom Rumour Spreading

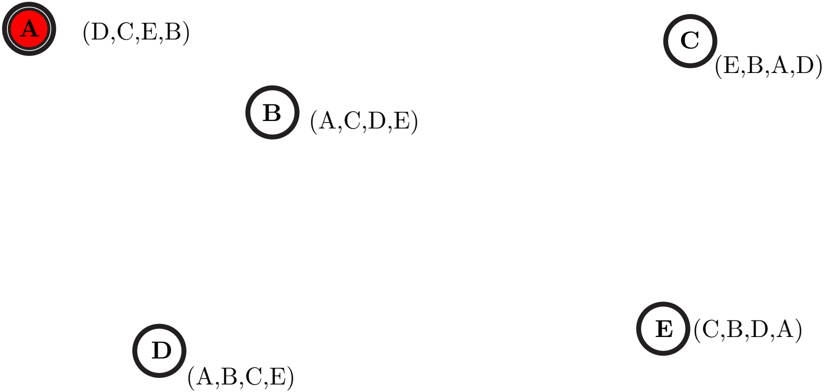

Durham University

Quasirandom Rumour Spreading

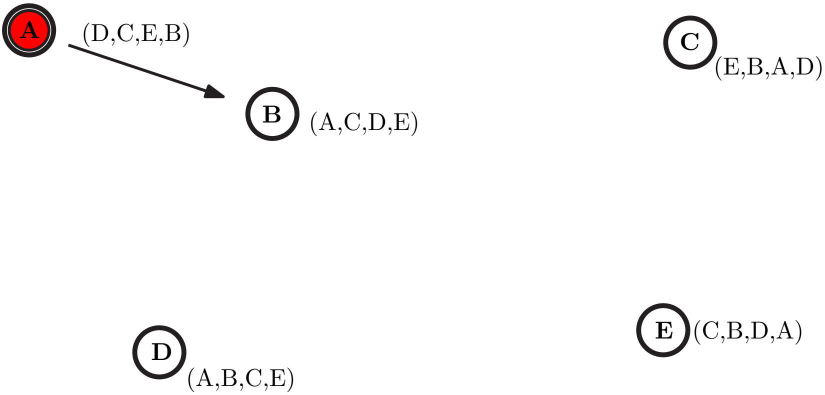

Durham University

Quasirandom Rumour Spreading

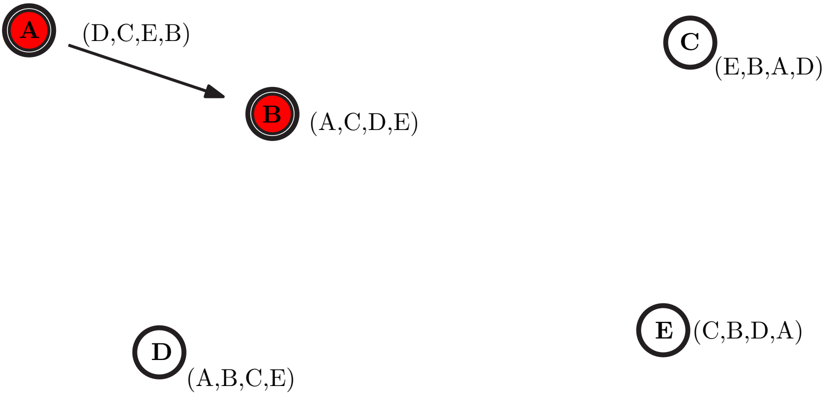

Quasirandom Rumour Spreading

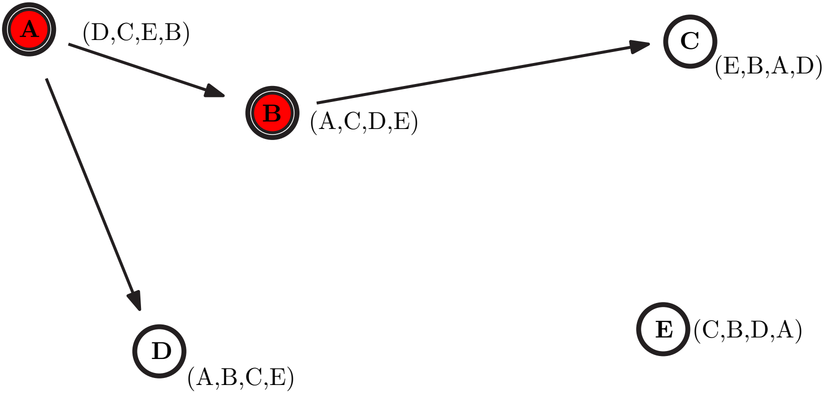

Quasirandom Rumour Spreading

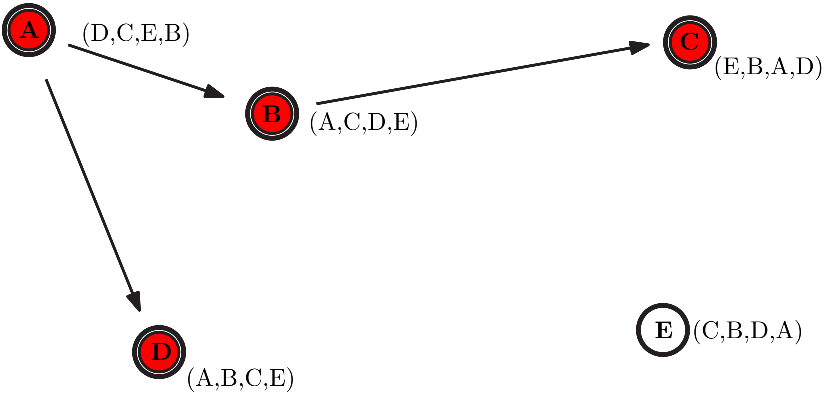

Quasirandom Rumour Spreading

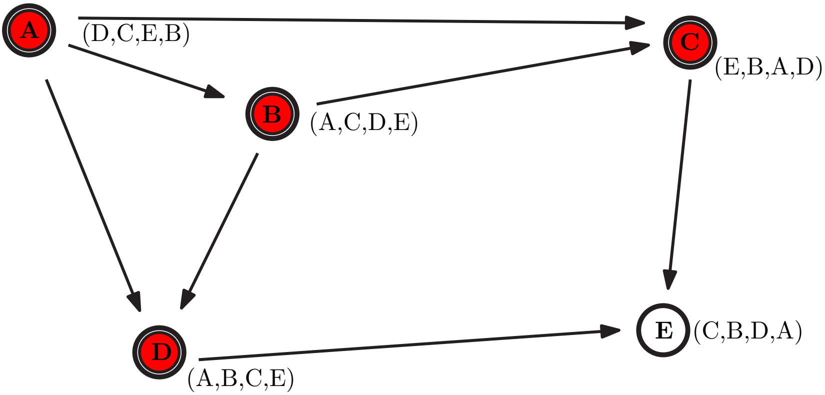

Durham University

Quasirandom Rumour Spreading

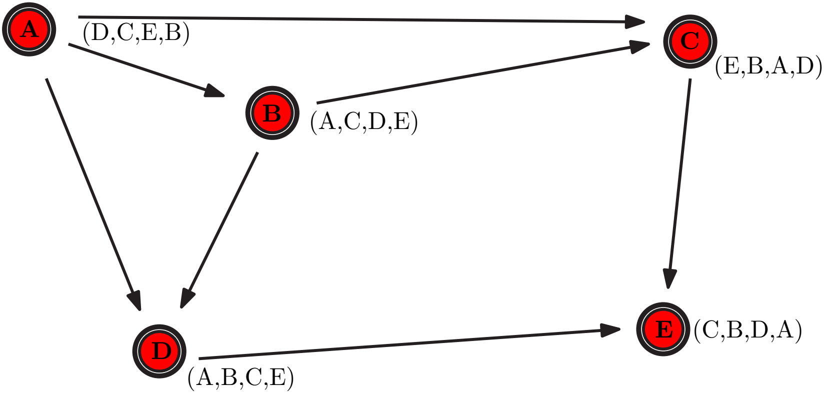

Durham University

Quasirandom Rumour Spreading

Question

How long does it take until all n vertices are informed?

Quasirandom Rumour Spreading

Question

How long does it take until all n vertices are informed?

- At least $\log_2 n$.

Quasirandom Rumour Spreading

Question

How long does it take until all n vertices are informed?

- At least $\log_2 n$.
- Also bounded from above (by $n - 1$ for complete graphs).

Quasirandom Rumour Spreading

Question

How long does it take until all n vertices are informed with high probability?

- At least $\log_2 n$.
- Also bounded from above (by $n-1$ for complete graphs).

Quasirandom Rumour Spreading

Theorem (Doerr, Friedrich and Sauerwald 2008)

The quasirandom model on hypercubes and random graphs $G_{n,p}$ with $p \geq \frac{(1+\epsilon)\ln n}{n}$ informs all n vertices within

$$\Theta(\log n)$$

time-steps with high probability.

Durham University

Recall:

Let R_n denote the number of rounds needed to inform all n vertices of a complete graph in the random model.

Theorem (Frieze and Grimmett 1985)

$$(1 - o(1))(\log_2 n + \ln n) \leq R_n \leq (1 + o(1))(\log_2 n + \ln n)$$

with high probability.

Durham University

Quasirandom Rumour Spreading

Let Q_n denote the number of rounds needed to inform all n vertices of a complete graph in the quasirandom model.

Theorem (Angelopoulos, Doerr, H., Panagiotou 2009)

$$(1 - o(1))(\log_2 n + \ln n) \leq Q_n \leq (1 + o(1))(\log_2 n + \ln n)$$

with high probability.

Durham University

Recall:

Let R_n denote the number of rounds needed to inform all n vertices of a complete graph in the random model.

Theorem (Pittel 1987)

For every $h \in \omega(1)$ one has with high probability

$$\log_2 n + \ln n - h(n) \leq R_n \leq \log_2 n + \ln n + h(n).$$

Durham University

Quasirandom Rumour Spreading

Let Q_n denote the number of rounds needed to inform all n vertices of a complete graph in the quasirandom model.

Theorem (Fountoulakis, H., SIDMA 2009)

For every $h \in \omega(1)$ one has with high probability

$$\log_2 n + \ln n - 4 \ln \ln n \leq Q_n \leq \log_2 n + \ln n + h(n).$$

Durham University

Robustness

Robustness

- Messages successfully sent with constant probability $q \in [0, 1]$.

Robustness

- Messages successfully sent with constant probability $q \in [0, 1]$.
- Independent failures.

Durham University

Robustness

- Messages successfully sent with constant probability $q \in [0, 1]$.
- Independent failures.
- No notifications of success or failure.

Robustness

- Messages successfully sent with constant probability $q \in [0, 1]$.
- Independent failures.
- No notifications of success or failure.
- $q = 1$: As before, quasirandom is as fast as random.
- What about $q \in]0, 1[$?
- Is quasirandom still as fast as random?

Durham University

Robustness

Theorem (Doerr, H., Levavi, ISAAC 2009)

The quasirandom model with transmission success probability q on the complete graph informs all n vertices in

$$(1 + o(1))(\log_{1+q} n + \frac{1}{q} \ln n)$$

time-steps with high probability.

Durham University

Core Questions

- Where do dependencies help?
- How much randomness is necessary?

Durham University

Related Work

Related Work

- Reducing randomness can be costly [Doerr, Fouz 2011]

Related Work

- Reducing randomness can be costly [Doerr, Fouz 2011]
- Asymptotically Optimal Randomized Rumor Spreading [Doerr, Fouz 2011]

Related Work

- Reducing randomness can be costly [Doerr, Fouz 2011]
- Asymptotically Optimal Randomized Rumor Spreading [Doerr, Fouz 2011]
- Randomness-Efficient Rumor Spreading [Guo, Sun 2013]

Durham University

Related Work

- Reducing randomness can be costly [Doerr, Fouz 2011]
- Asymptotically Optimal Randomized Rumor Spreading [Doerr, Fouz 2011]
- Randomness-Efficient Rumor Spreading [Guo, Sun 2013]

not exhaustive – for more, see e. g.

Durham University

Related Work

- Reducing randomness can be costly [Doerr, Fouz 2011]
- Asymptotically Optimal Randomized Rumor Spreading [Doerr, Fouz 2011]
- Randomness-Efficient Rumor Spreading [Guo, Sun 2013]

not exhaustive – for more, see e. g.

- Dagstuhl Seminar 13042 "Epidemic Algorithms and Processes"

Durham University

Summary

- Randomized rumour spreading on sufficiently dense random graphs is as fast and robust as on the complete graph.

Durham University

Summary

- Randomized rumour spreading on sufficiently dense random graphs is as fast and robust as on the complete graph.
- Quasirandom rumour spreading on the complete graph is as fast and robust as randomized rumour spreading.

Open Questions

Open Questions

- Are $\log_2 n + \ln n - h(n)$ time-steps necessary to inform the whole complete graph in the quasirandom model w. h. p.?

Durham University

Open Questions

- Are $\log_2 n + \ln n - h(n)$ time-steps necessary to inform the whole complete graph in the quasirandom model w. h. p.?
- Are there graphs where the quasirandom model is faster?

Durham University

Open Questions

- Are $\log_2 n + \ln n - h(n)$ time-steps necessary to inform the whole complete graph in the quasirandom model w. h. p.?
- Are there graphs where the quasirandom model is faster?
- Quasirandom rumour spreading on random graphs, random regular graphs, expanders, ...

Open Questions

- Are $\log_2 n + \ln n - h(n)$ time-steps necessary to inform the whole complete graph in the quasirandom model w. h. p.?
- Are there graphs where the quasirandom model is faster?
- Quasirandom rumour spreading on random graphs, random regular graphs, expanders, ...
- General result for the robustness of the quasirandom model

Durham University

Open Questions

- Are $\log_2 n + \ln n - h(n)$ time-steps necessary to inform the whole complete graph in the quasirandom model w. h. p.?
- Are there graphs where the quasirandom model is faster?
- Quasirandom rumour spreading on random graphs, random regular graphs, expanders, ...
- General result for the robustness of the quasirandom model
- Other dependency models than lists - generalization of quasirandom rumour spreading

Open Questions

- Are $\log_2 n + \ln n - h(n)$ time-steps necessary to inform the whole complete graph in the quasirandom model w. h. p.?
- Are there graphs where the quasirandom model is faster?
- Quasirandom rumour spreading on random graphs, random regular graphs, expanders, ...
- General result for the robustness of the quasirandom model
- Other dependency models than lists - generalization of quasirandom rumour spreading
- “Good” lists - specialization of quasirandom rumour spreading

Durham University

Open Questions

- Are $\log_2 n + \ln n - h(n)$ time-steps necessary to inform the whole complete graph in the quasirandom model w. h. p.?
- Are there graphs where the quasirandom model is faster?
- Quasirandom rumour spreading on random graphs, random regular graphs, expanders, ...
- General result for the robustness of the quasirandom model
- Other dependency models than lists - generalization of quasirandom rumour spreading
- "Good" lists - specialization of quasirandom rumour spreading

Durham University

Thank you very much for your attention!