Alea Meeting, Munich, February 2016

Lecture 2: Rumor Spreading (2)

Benjamin Doerr, LIX, École Polytechnique, Paris-Saclay

Outline:
Reminder last lecture and an answer to Kosta’s question
Definition: Rumor spreading time
Path arguments, degree-diameter bound, trees and grids
Hypercubes
Preferential attachment graphs

Reminder Last Lecture

- Randomized rumor spreading (RRS):
  - round-based process in a graph $G = (V, E)$
  - starts with one node informed (knows the rumor)
  - in each round, each informed node calls a random neighbor and informs it (if it wasn't already)

- Main result last lecture: With probability $1 - o(1)$, RRS informs all vertices of a complete graph on n vertices in $\log_2 n + \ln n + o(\log n)$ rounds

- Same result holds for $G(n, p)$ random graphs when $p = \omega((\log n)/n)$, but not when $p = \Theta((\log n)/n)$
[FountoulakisHuberPanagiotou'10, PanagiotouPerezSauerwaldSun'15]

- Same result holds for random regular graphs $G(n, d)$ when $d = \omega(1)$, but not for constant d
[PanagiotouPerezSauerwaldSun'15, FountoulakisPanagiotou'10]

Kosta's Question

- End of last lecture: r-termination tricky rumor spreading in complete networks with node labels $1, \ldots, n$:
  - In each round, each informed node calls another node.
    - In the first round of activity, node k calls $k+1$
    - in all other rounds, it calls a random node
  - Termination: Nodes stop calling after having called r informed nodes

- Result: This informs all nodes in $\log_2 n + O\big((\log n)/r + (\log n)^{1/2}\big)$ rounds and $O(rn)$ calls.

- Kosta's question: Can you get $\log_2 n + o(\log n)$ rounds using $O(n)$ calls with a more clever protocol?

- Johannes' idea (that convinced the coffee round): Node k calls first $k+1$, then a random node in $[k + (\log n)^{1/2}, k + 2(\log n)^{1/2}]$, then random nodes.

Part 1: Definition Rumor Spreading Time

- Last lecture, we had a slightly informal start into rumor spreading. Before proceeding with lot's of hand-waiving, let's clarify a basic question first:
  - how to define *rumor spreading time* ?

- Definition [rumor spreading times T_v]: Consider a round-based rumor spreading process in a graph $G = (V, E)$ that surely finishes in finite time.
  - For $v \in V$, the T_v denote the number of rounds after which a rumor starting in v for the first time has reached all nodes of G

- Notes:
  - T_v is a random variable taking values in $\mathbb{N} := \{0,1,2, \dots\}$
  - You can extend all that follows to include the time ∞ denoting that a process does not finish, but you'll not gain a lot from it

Towards Defining T_G

- Definition: Let X and Y be random variables. We say that Y stochastically dominates X, written $X \preccurlyeq Y$, if for all $\lambda \in \mathbb{R}$ we have
$$\Pr[Y \leq \lambda] \leq \Pr[X \leq \lambda].$$
  - Note: a very strong sense of "Y is bigger than X"

- Lemma: There is a unique random variable T such that
  1) for all $v \in V$, $T_v \preccurlyeq T$
  2) for any T' satisfying 1) we have $T \preccurlyeq T'$

- Proof: *Go from the probability mass functions to the cumulative distributions, take the minimum, and convert it to a mass function*
  - $F(t) \coloneqq \min\{\Pr[T_v \leq t] \mid v \in V\}$
  - Define T by $\Pr[T = 0] = F(0)$, $\Pr[T = t] = F(t) - F(t-1)$ for $t \geq 1$
  - This T satisfies 1). If some T' satisfies 1), then $F'(t) \coloneqq \Pr[T' \leq t]$ is at most $\Pr[T_v \leq t]$ for all $v \in \mathrm{V}$ by 1), hence $F' \leq F$, and thus $T \preccurlyeq T'$.

Rumor Spreading Time T_G of G

- Definition [D, Friedrich, Sauerwald '14]: The unique T from the lemma is called the rumor spreading time of G and denoted by T_G

- Trivialities:
  - $\Pr[T_G \leq t] \geq p$ is equivalent to saying that regardless of where the rumor starts, after t rounds with probability at least p all vertices are informed.
  - $\Pr[T_G \geq t] \geq p$ is equivalent to saying that there is a vertex $v \in V$ such that the rumor spreading process started in v with probability at least p has not informed all vertices earlier than after t rounds.

- Note: $E[T_G] \geq \max\{E[T_v] | v \in V\}$, but there is no argument for equality

- Not very important open problem: Do we have equality above?
  - Equivalent question: are there "worst" starting points w for the rumor (in the sense that $T_v \preccurlyeq T_w$ for all $v \in V$)?

Part 2: Path-Arguments

- Main argument so far:
  - Analyze how many nodes become informed in a round starting with i informed nodes
  - Works well when the graph is highly symmetric (complete graphs, random graphs)

- Now: Use the argument “how long does it take for the rumor to traverse a given path”
  - potential disadvantage: we have to fix a path first and thus ignore the fact that the rumor could also use a different path
  - advantage: since we look at many steps together, we may use concentration on the whole process (instead of only one round)

Path-Arguments: Trivial Lower Bound

- Trivial remark: If the rumor starts in s and v is a node in distance $d(s,v)$ from s, then the rumor cannot reach v earlier than in round $d(s,v)$.
  - $d(s,v)$: smallest length (=number of edges) of a path between s and v

- Consequence: The diameter $diam(G)$ is a lower bound for the rumor spreading time (worst-case over starting vertex) in the graph $G = (V,E)$
  - $diam(G) = \max\{d(u,v) | u,v \in V\}$

Path-Arguments: Upper Bounds

- Lemma: Let $P: x = x_0, x_1, \ldots, x_k = y$ be any path from x to y in G. Assume that the rumor starts in x. Then the first time T_y when y is informed satisfies $\boldsymbol{E\left[T_y\right] \leq \sum_{i=0}^{k-1} \deg x_i}$
  - proof: add the pessimistic waiting times for the events that the rumor moves from x_i to x_{i+1}
  - more precisely: it takes an expected number of $\deg x_i$ rounds for x_i to call x_{i+1}

- Hence the expected time for the rumor to traverse one path is easy, but to say something about the rumor spreading time, we need to pick a path from the source to each vertex y and say something about $T = \max T_y$

Path-Lemma, Degree-Diam. Bound

- **Path-Lemma:**
  - Let G be any graph. Assume that the rumor starts in a vertex x_0.
  - Let $P: x_0, \dots, x_k$ be any path of length k in G.
  - Let $\Delta \coloneqq \max\{\deg x_i \mid i \in [0..k-1]\}$ be the maximum degree of the vertices on P.
  - Let $k' \geq k$. ["safety margin over the expectation"]
  - Then after $T = 2k'\Delta$ rounds, the whole path is informed with probability $1 - \exp(-k'/4)$.
- Proof: Chernoff bound (details next slide)

- **Corollary (degree-diameter bound):** $k' = \max\{diam(G), 8\ln(n)\}$ gives that all vertices are informed after $\boldsymbol{T = O(\Delta \max\{diam(G), \log n\})}$ rounds with probability $1 - n^{-1}$
- Proof: Take a shortest path from the rumor source to each vertex. This is traversed in T rounds with prob. $1 - n^{-2}$. Union bound

Proof of the Path Lemma

- **Path lemma:** Let G be any graph. Let $P: x_0, \dots, x_k$ be a path of length k in G. Let $\Delta := \max\{\deg x_i \mid i \in [0..k-1]\}$ be the maximum degree on P. Assume the rumor starts in x_0. Let $k' \geq k$. Then after $T = 2\Delta k'$ rounds, the whole path is informed with probability at least $1 - \exp(-k'/4)$.

- **Proof:** We analyze the following modified process: In each round, each informed node $x_i, i \in [0..k-1]$, calls node x_{i+1} with probability exactly $1/\Delta$.
- Observation: The modified process clearly is slower in informing x_k.
- For each round t, define a binary random variable X_t as follows. Let $i(t)$ be the maximal $j \in [0..k]$ such that $x_0, \dots, x_j$ are informed at the start of the round.
  - If $i(t) < k$ and $x_{i(t)+1}$ becomes informed in round t, then $X_t := 1$.
  - If $i(t) = k$, then set $X_t := 1$ with probability $1/\Delta$ independently from all other random decisions.
  - In all other cases, set $X_t = 0$.
- Then the X_t are independent and satisfy $\Pr[X_t = 1] = 1/\Delta$.
- Let $X := \sum_{t=1}^{T} X_t$. Then x_k is informed after T rounds if and only if $X \geq k$.
- The multiplicative Chernoff bound shows that this fails with probability at most $\Pr[X < k] = \Pr\left[X < 0.5E[X]\right] \leq \exp(-E[X]/8) = \exp(-k'/4)$.

Application to Trees

- Let G be a k-regular rooted tree of height h, that is, an undirected graph having $n = 1 + k + k^2 + \cdots + k^h$ vertices such that there is one “root” vertex which has k neighbors such that each of them is the root of a r-regular tree of height $h - 1$ (when we delete the original root and all edges incident with it). Assume that you run the randomized rumor spreading protocol in this graph, starting the rumor in the root.

- Degree diameter bound
  - diameter $\Theta(h)$
  - max-degree $\Theta(k)$
  - rumor spreading time $O(k \max\{h, \log n\}) = O(hk \log k)$
  - matching lower bound: just look at how long it takes to inform all leaves (assuming that all the rest is informed)

Application to Grids

- Let $G = (V, E)$ be a d-dimensional grid, that is, $V = [1..k]^d$ and two vertices are neighbors if they differ in exactly one coordinate, and this difference is exactly one.
- Theorem: A rumor starting in an arbitrary vertex reaches all vertices in time $O(d^2 k)$ with probability $1 - o(1)$ [asymptotics for $n := k^d$ tending to infinity]
- Proof: Degree-diameter bound. $\Delta = \Theta(d)$ and $\mathrm{diam}(G) = d(k-1)$, the latter being $\Omega(\log n)$ for all values of d and k
- Comment:
  - For $d = \Theta(1)$, this is tight (diameter is a lower bound).
  - For $k = 2, d = \log_2 n$, this is not tight: We now prove $O(\log n)$ instead of the above $O(\log^2 n)$.
    - "so many paths that one will be much faster than its expectation."

- **Definition:** The d-dimensional *hypercube* is a graph H_d having
  - $V = \{0,1\}^d$ as vertex set (hence $n := |V| = 2^d$), and
  - two vertices are adjacent if they differ in exactly one position.

- Note: $\Delta(G) = d = \log n$
  - distances in H_d: $d(u, v) =$ “number of positions u and v differ in”.
  - diameter (max. distance between vertices): $\mathrm{diam}(G) = d = \log n$,

- Good communication network: Small diameter, relatively few edges, high connectivity (d disjoint paths between any two vertices)

Rumor Spreading in Hypercubes

- The degree-diameter bound gives a rumor spreading time of
$$O(\Delta \max\{\mathrm{diam}(H_d), \log n\}) = O(\log^2 n)$$

- Might be overly pessimistic, because there are many path between any pair of vertices:
  - $d!$ different shortest paths between $(0, \dots, 0)$ and $(1, \dots, 1)$
  - so there might be one path where we are much more lucky than what the expectation tell us.

- **Theorem:** With probability $1 - 1/n$, a rumor started in an arbitrary node of the hypercube has reached all nodes after $O(\log n)$ rounds.
  - beautiful proof (next couple of slides)
  - major open problem to determine the leading constant

Proof: General Stuff

- We assume that the rumor starts in $s = (0, \ldots, 0)$. [symmetry]

- We show that for any $\beta > 0$ there is a $K > 0$ such that after $K \log n$ rounds, the vertex $t = (1, \ldots, 1)$ is informed with probability $1 - n^{-\beta}$
  - similar arguments work for any target t
  - a union bound shows that all vertices are informed w.p. $1 - n^{-\beta+1}$

- Two technical assumptions that do not change how the rumor spreads, but help in the proof
  - all-work assumption: We assume that in each round *every* node calls a random neighbor – if the caller is not informed, nothing happens
  - everything-predefined assumption: We assume that before the process starts, each node has already fixed whom to call in which round

Expansion Phase

- Observation: The rumor quickly moves away from $s = (0, \dots, 0)$, but it is increasingly difficult to argue that the rumor truly approaches the target.

- Plan: Show that you get at least close to the target!
  - for reasons that will become clear later, we show that we get close to any target we want.

- Expansion Lemma: Let $\alpha > 0$. Let $v \in V$. Let $C \geq 2$. After Cd/α rounds, with probability at least $1 - \exp(-\ Cd/8)$ there is an informed vertex w such that $d(v, w) \leq \alpha d$.
  - "in $\Theta(d)$ rounds the rumor reaches v apart from at most the last αd steps (and apart from an $O(n^{-\Omega(1)})$ failure probability"

Proof: Expansion Lemma

- Similar to the analysis how rumors traverse a path.
- Let d_t denote the distance of v to the closest informed vertex after round t.
- Define binary random variables X_t (counting true/artificial progress) as follows
  - if $d_{t-1} > \alpha d$, then $X_t = 1$ if and only if $d_{t-1} > d_t$
  - if $d_{t-1} \leq \alpha d$, then $X_t = 1$ with probability α (independent of everything)
- $\Pr[X_t = 1] \geq \alpha$ for all t
- Note: $X^T \coloneqq \sum_{t=1}^{T} X_t \geq d(s,v) - \alpha d$ if and only if $d_T \leq \alpha d$ (our aim)
- The X_t are not independent, but we have $\Pr[X_t = 1 | X_1 = x_1, \dots, X_{t-1} = x_{t-1}] \geq \alpha$ for all $x_1, \dots, x_{t-1} \in \{0,1\}$. Hence X^T dominates a sum Y^T of T independent random variables that are 1 with probability exactly α (Lemma 1.18 in book chapter).
- For $T = Cd/\alpha$ we have
$$\Pr[X^T \leq d] \leq \Pr[Y^T \leq d] \leq \Pr\left[Y^T \leq \frac{1}{2} E[Y^T]\right] \leq \exp\left(-\frac{E[Y^T]}{8}\right) \leq \exp\left(-\frac{Cd}{8}\right)$$
by the multiplicative Chernoff bound.

Backward Phase

- Plan: Do something “dual”: starting in t and going backward in time, spread “uninformedness”

- Recall that we assumed that all nodes call in each round.

- Assume that our target node t is uninformed after some round T.
  - if some node x calls t in round T, then x must be uninformed after round $T-1$
  - iterate this argument to construct a path ending in t such that if the start of the path was informed at some time $T-i$ then t would be informed at time T

- Here we use the all-work and all-predetermined assumptions!

Backward Phase – Some Details

- Lemma: Let T be large. Let $\alpha > 0$. Let $v \in V$. Then with probability at least $1 - \exp(-Cd/8)$ there is a $w \in V$ such that $d(w, v) \le \alpha d$ and if w is informed after round $T - Cd/(1 - \exp(-\alpha))$, then t is informed after round T.

- Proof:
  - For $i = 0,1,2 \ldots$ let d_i be the smallest $d(v, x)$ of a node x having the property that if x is informed at the end of round $T - i$, then t is informed after round T.
    - $d_0 = d(v, t) \le d$
    - if $d_i > \alpha d$, then $\Pr[d_{i+1} = d_i - 1] \ge 1 - (1 - 1/d)^{\alpha d} \ge 1 - \exp(-\alpha)$
  - Use an analogous "artificial progress counting" argument as before
    - $X_i = 1$ if $d_i < d_{i-1}$ and $d_{i-1} > \alpha d$, otherwise independent random bit that is 1 with prob. $1 - \exp(-\alpha)$
    - …

Coupling Phase

- So far: For any $v \in V$, with probability $1 - \exp(\Omega(d))$ [as large as we want]
  - there is an $A_v \in V$ such that $d(v, A_v) \leq \alpha d$ and s informs A_v within $O(d)$ rounds
  - there is a $B_v \in V$ such that $d(v, B_v) \leq \alpha d$ and “the rumor would go from B_v to t in $O(d)$ rounds”

- Remains to do: Get the rumor from A_v to B_v!
  - Problem: Very hard to get the rumor exactly somewhere (we need already d rounds to call a particular neighbor)
  - Solution: Take many v as above, sufficiently far apart, and play this game many times in parallel – once we will be lucky ☺

From A_v to B_v With Small Probability

- Let $B := B(v, 2\alpha d) := \{u \in V |\, d(u,v) \leq 2\alpha d\}$ "$2\alpha d$ ball around v"
- Target: Get the rumor from A_v to B_v, but only using nodes in B
  - needed later to ensure that processes for different v don't interact

- Lemma: The probability that the rumor moves inside B from A_v to B_v in time at most $2\alpha d$, is at least $(2\alpha/e)^{2\alpha d}$.

- Proof: Send the rumor along a direct path with speed one! ($d' := d(A_v, B_v)$)
  - Probability that the rumor moves closer to B_v in every round:
$$\prod_{i=1}^{d'} \frac{i}{d} \geq \prod_{i=1}^{2\alpha d} \frac{i}{d} = \frac{(2\alpha d)!}{d^{2\alpha d}} \geq \frac{(2\alpha d/e)^{2\alpha d}}{d^{2\alpha d}} = \left(\frac{2\alpha}{e}\right)^{2\alpha d}$$
  - Small exercise: Any such path remains in B

Many v's that are far apart

- Target: Find a large set of v's such that the distance of any two is more than $4\alpha d$ – so the $2\alpha d$-balls do not intersect.
- **Lemma**: There is a set $S \subseteq V$ such that $|S| = \exp(d/32) =: m$ and for all $x, y \in S$ with $x \neq y$ we have $d(x, y) \geq d/4$.
- **Proof**: Take a random set!
  - Let $x_1, \dots, x_m$ be random vertices.
  - For $i \neq j$, we have $E\left[d\left(x_i, x_j\right)\right] = d/2$
  - $d\left(x_i, x_j\right)$ is a sum of d independent $\{0,1\}$ random variables
  - Chernoff bound: $p := \Pr\left[d\left(x_i, x_j\right) \leq d/4\right] \leq \exp(-d/16)$
  - Union bound: $\Pr[S \text{ bad}] \leq \sum_{i,j} \Pr\left[d\left(x_i, x_j\right) \leq d/4\right] < m^2 p \leq 1$
    - Consequently, there is such a set S

Proof: Putting Everything Together

- Choose α small enough so that $(2\alpha/e)^{2\alpha} > \exp(-1/32)$ and $\alpha d < d/4$
  - note that $(2\alpha/e)^{2\alpha}$ tends to one for $\alpha \to 0$
- Choose the set S as on the previous slide.
- Apply expansion lemma with C large enough and union bound to show that with probability $1 - n^{-\beta}$ for all $v \in S$ there is an $A_v \in B(v, \alpha d)$ that is informed after $T_1 = O(d)$ rounds
- Apply backward lemma with C large enough and union bound to show that with probability $1 - n^{-\beta}$ for all $v \in S$ there is a $B_v \in B(v, \alpha d)$ such that if B_v is informed after $T_2 := T_1 + 2\alpha d$ rounds, then t is informed after $T_2 + O(d)$ rounds
- Coupling phase: The probability that for no $v \in S$ the rumor goes (inside $B(v, 2\alpha d)$ from A_v to B_v is at most $\left(1 - (2\alpha/e)^{2\alpha d}\right)^{\exp(d/32)} \le \exp\left(-\left((2\alpha/e)^{2\alpha}\exp(1/32)\right)^{d}\right) = \exp\left(-n^{\Theta(1)}\right)$
- Hence apart from a failure prob. of $2n^{-\beta} + \exp(-n^{\Theta(1)})$, t is informed….

Part 4: Rumor Spreading in Social Networks, Real-World Graphs

- "Real-world graph":
  - airports connected by direct flights
  - scientific authors connected by a joint publication
  - Facebook users being "friends"

- Big insight of the last 20 year: Real-world graphs have very special properties!
  - small diameter
  - non-uniform degree distribution:
    - few nodes of high degree: "hubs"
    - many nodes of small (constant) degree
    - power law: number of nodes of degree d is proportional to d^{β} [β a constant, often between 2 and 3]

Preferential Attachment (PA) Graphs

- Barabási, Albert (Science 1999):
  - try to explain why social networks could look like this
  - suggest a model for real-world graphs: preferential attachment (PA)

- Preferential attachment paradigm:
  - networks evolve over time
  - when a new node enters a network, it chooses at random a constant number m of neighbors
  - random choice is not uniform, but gives preference to “popular” nodes
    - probability to attach to node x is proportional to the degree of x

- Once made truly precise (by Bollobás and Riordan (2004)), the PA paradigm defines a cool random graph model (“PA graphs”)
  - Today: One of the most used models for real-world networks

Precise Definition of PA Graphs

- Preferential attachment graph G_m^n
  - n: number of vertices, vertex set $[1..n]$
  - m: density parameter

- The PA graph $\mathrm{G}^n \coloneqq G_m^n$ is recursively defined:
  - G^1: 1 is the single vertex that has m self-loops
  - G^n: Obtained from adding the new vertex n to G^{n-1}
    - one after the other, the new vertex n chooses m neighbors
    - the probability that some vertex x is chosen, is
      - proportional to the current degree of x, if $x \neq n$
      - proportional to "$1 +$ the current degree of x", if $x = n$ (self-loop probability takes into account the current edge starting in n)

Properties of PA Graphs*

- diameter $\Theta(\log n \,/\, \log\log n)$: less than logarithmic despite $\Theta(n)$ edges!
  - $G(n, p)$ with $p = \Theta(1/n)$: far from connected
  - random regular graphs, k-out graphs: diameter $\Theta(\log n)$
- power law degree distribution: For $d \leq n^{1/5}$, the expected number of vertices having degree d is proportional to d^{-3}.
- clustering coefficient = roughly the probability that two neighbors of some node are connected by an edge
  - PA graphs: $\Theta(1/n)$
  - real-world graphs: typically constant ☹

*All statements hold "with high probability" (whp), that is, with prob. $1 - o(1)$

Rumor Spreading in PA Graphs G_m^n

- Chierichetti, Lattanzi, Panconesi (2009):
  - Classic "push" rumor spreading: n^α rounds (α a small constant) with constant probability do not suffice
  - If both informed and uninformed nodes call random neighbors to spread or seek rumors (push-pull protocol), then $O((\log n)^2)$ rounds inform a PA graph G_m^n, $m \geq 2$, whp.

- D, Fouz, Friedrich (2011): In the push-pull protocol, the rumor spreading time is
  - $\Theta(\log n)$ whp
  - $\Theta(\log(n)/\log\log n)$ whp, if contacts are chosen excluding the neighbor contacted in the very previous round (no "double-contacts")
    - Note: Avoiding double-contacts does not improve the $O(\log n)$ times for complete graphs, RGGs, hypercubes, …

Two Questions

- Theorem: Randomized rumor spreading in the push-pull model informs the PA graph G^n (with $m \geq 2$) with high probability in
  - $\Theta(\log n)$ rounds when choosing neighbors uniformly at random
  - $\Theta(\log n / \log \log n)$ rounds without double-contacts

- Two questions:
  - Why do double-contacts matter?
  - What makes PA graphs spread rumors faster than other graphs?

With Double-Contacts…

- Critical situation:
  - A pair of uninformed neighboring nodes, each having a constant number of outside neighbors

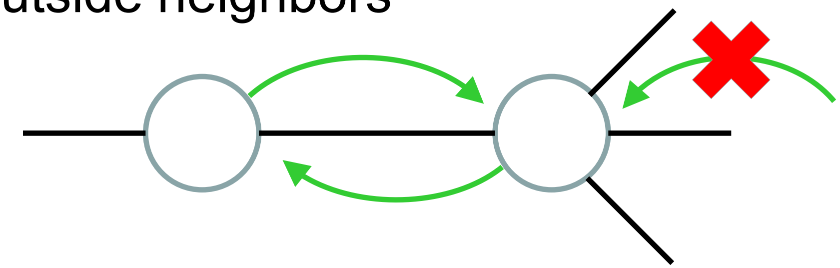

- With constant probability, the following happens in one round:
  - the two nodes of the pair call each other
  - all their neighbors call someone outside the pair
  - → hence the situation remains critical (pair uninformed)

- Problem: Initially, there are $\Theta(n)$ such critical situations in a PA graph. Since each is solved with constant probability in one round, $\Omega(\log n)$ rounds are necessary

Without Double-Contacts

- The uninformed pair is not critical anymore, because the two nodes cannot call each other twice in a row ☺

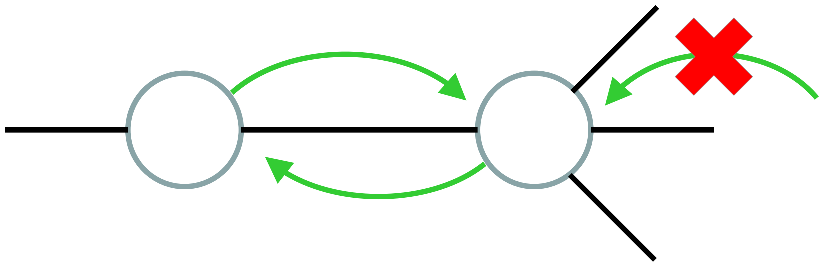

- Remaining critical situations: Uninformed cycles having a constant number of outside neighbors in total.
  - Again, each round, with constant probability the situation remains critical (cycle uninformed)

- No problem! There are only $O(\exp((\log n)^{3/4}))$ such critical situations initially in a PA graph. If each is solved with constant probability, we need $O\left(\log(\exp((\log n)^{3/4}))\right) = O\left((\log n)^{3/4}\right)$ rounds to solve them all ☺

Why are PA Graphs Faster?

- Large- and small-degree nodes:
  - hub: node with degree $(\log n)^3$ or greater
  - poor node: node with degree exactly m (as small as possible)

- Observation: Poor nodes convey rumors fast!
  - Let a and b be neighbors of a poor node x
  - If a is informed, the expected time for x to pull the rumor from a is less than m
  - After that, it takes another less than m rounds (in expectation) for x to push the news to b

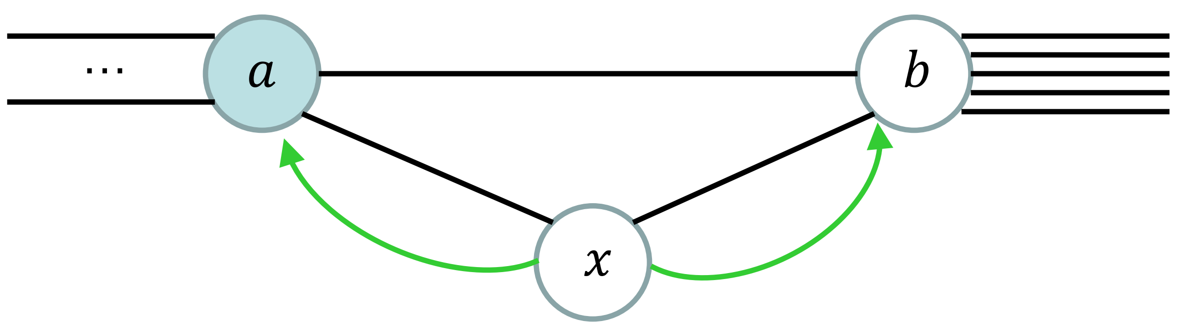

Why Are PA Graphs Faster (2)?

- Large- and small-degree nodes:
  - hub: node with degree $(\log n)^3$ or greater
  - poor node: node with degree exactly m (as small as possible)

- Observation: Poor nodes convey rumors fast!
  - Let a and b be neighbors of a poor node x
  - If a is informed, the expected time for x to pull the rumor from a is less than m
  - After that, it takes another less than m rounds (in expectation) for x to push the news to b

- Key lemma: Between any two hubs, there is a path of length $O(\log n \;/\; \log\log n)$ with every second node a poor node.

- Key lemma + observation + some extra arguments: If one hub is informed, after $O(\log n \;/\; \log\log n)$ rounds all hubs are.

Main Tool: BR'04 Definition of Preferential Attachment Model

<u>Equivalent description of the PA model (also Bollobás & Riordan (2004))</u>

- For $m = 1$
  - Choose $2n$ random numbers in $[0,1]$: $x_1, y_1, \dots, x_n, y_n$
  - If $x_i > y_i$, exchange the two values
    - $\Pr(y_i \leq r) = r^2$
  - Sort the (x, y) pairs by increasing y-value; call them again $(x_1, y_1), (x_2, y_2), \dots$
  - For all k, vertex k chooses the unique $i \leq k$ as neighbor which satisfies $y_{i-1} \leq x_k < y_i$
    - Note: x_k is uniform in $[0, y_k]$
- For $m \geq 2$: Generate G_1^{nm} as above, then merge each m consecutive nodes

- Advantage: Many independent random variables, not a sequential process

Some More Results

- Fountoulakis, Panagiotou, Sauerwald (SODA'12):
  - Chung-Lu graphs: variation of the classic Erdős-Rényi random graphs (independent edges) that yields a power-law degree distribution
  - Synchronized randomized rumor spreading in the push-pull model informs all but $o(n)$ nodes of the Chung-Lu graph
    - in $\Theta(\log n)$ rounds, if the power-law exponent $\beta > 3$
    - in $\Theta(\log \log n)$ rounds, if $2 < \beta < 3$
      [no result for $\beta = 3$, the PA exponent]
  - Asynchronous: Nodes call at times triggered by (their private) Poisson clock (with rate 1 → one call per time unit in expectation)
    - $2 < \beta < 3$: most nodes informed after a *constant* number of rounds!

- D, Fouz, Friedrich (SWAT'12):
  - Asynchronous rumor spreading informs most nodes of the PA graph in $O((\log n)^{1/2})$ time [not at all clear is this is sharp]

Summary: Rumor Spreading in Preferential Attachment Graphs

- Theorem: Randomized rumor spreading in the push-pull model informs the PA graph G_m^n (with $m \geq 2$) with high probability in
  - $\Theta(\log n)$ rounds when choosing neighbors uniformly at random
  - $\Theta(\log n / \log\log n)$ rounds without double-contacts
  - asynchronous: most nodes informed after $O((\log n)^{1/2})$ rounds

- Explanation: Interaction between hubs and poor nodes (constant degree)
  - hubs are available to be called
  - poor nodes quickly transport the news from one neighbor to all others

- Difference visible in experiments.

Merci ☺

Part X: Rumor Spreading in Wireless Sensor Networks

- *Wireless sensor network*: spatially distributed system of *sensor nodes* used for collecting data in a large area (often difficult to access)

- Sensor node
  - radio transmitter and receiver: communication with near-by nodes
  - sensors: collect data
  - microcontroller
  - battery: crucial for the life-time of the system

- Typical characteristics: Simple and cheap, so you don't care about setting up a clever network, but you distribute the nodes in a simple fashion (randomly)
  - multi-hop network: communication is via intermediate nodes, so that the communication range can be kept small (saves energy)

Mathematical Model

- *Random geometric graph* (RGG) $G(n,r)$:
  - Nodes: $v_1, \dots, v_n \in [0,1]^2$ be chosen uniformly at random
  - Edges: $\{v_i, v_j\}$ is an edge if and only if $d(v_i, v_j) \le r$

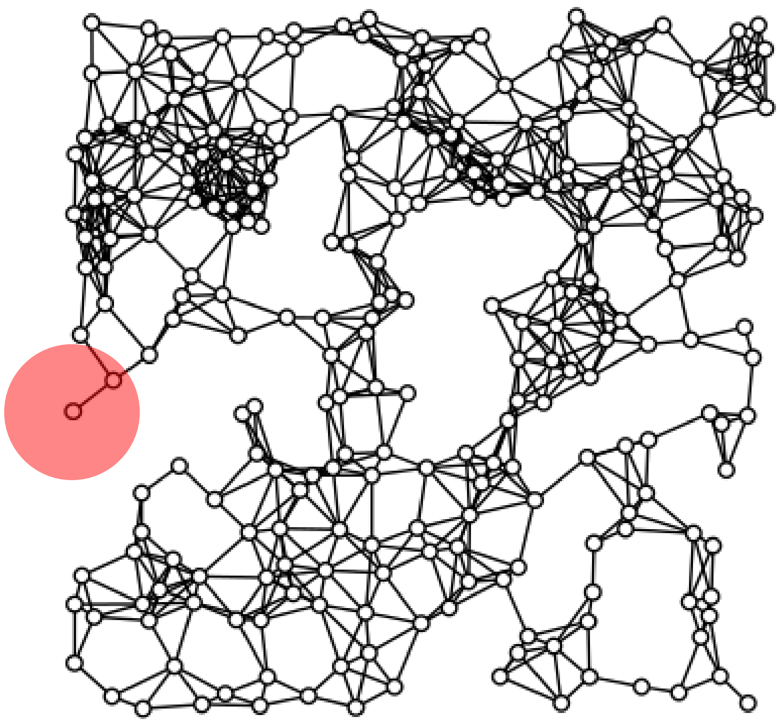

Fig.: A random geometric graph $G(256, 0.1)$ [Wikipedia]

Properties of RGG

- Theorem [Penrose (2003)]: If $r = (\pi n)^{-0.5}\,(\ln(n) + \alpha)^{0.5}$, then
$$\lim_{n\to\infty} \Pr[G(n,r) \text{ is connected}] = \exp(-e^{-\alpha})$$
- Theorem [Penrose (2003)]: There is a constant c such that
  - if $r > cn^{-0.5}$, then $G(n,r)$ has a “giant component”, that is, a connected component consisting of $\Theta(n)$ vertices; the proportion of this component tends to one if $r = \alpha n^{-0.5}$ for $\alpha \to \infty$
  - if $r < cn^{-0.5}$, then there is no giant component

- Implications for wireless sensor networks obtained by randomly distribution nodes: 3 regimes
  - disconnected: $r < cn^{-0.5}$ - all nodes can talk to only few others
  - giant component: There a large connected subnetwork, but some nodes are not part of it
  - connected: $r \geq (1+\varepsilon)(\pi n)^{-0.5}\ln^{0.5}(n)$

Rumor Spreading in Well-Connected RGGs (1)

- **Well-connected regime:** $r \geq Cn^{-0.5}\ln^{0.5}(n)$ for C large enough. Typical computational geometry argument shows that G is connected.

- Key approach: Make your life discrete!
  - Partition the unit square $[0,1]^2$ into $\Theta(1/r^2)$ squares of side length $l = r/2\sqrt{2} = \Theta(r)$
  - call two squares adjacent if they touch (vertical, horizontal, diagonal)
  - *vertices in adjacent squares are adjacent in the RGG*

- Claim: Each square S contains a similar number of vertices!
- The number X^S of vertices in S is a sum of n independent binary random variables X_i with $\Pr[X_i = 1] = l^2 = r^2/8 \geq C^2\ln(n)/n$. Hence
  - $E := E[X^S] = n\,r^2/8 \geq C^2\ln(n)$
  - $\Pr\left[|X^S - E[X^S]| \geq 0.25E[X^S]\right] \leq n^{-2}$ when C is sufficiently large.

Rumor Spreading in Well-Connected RGGs (2)

- Properties of well connected RGGs ($r \geq Cn^{-0.5} \ln^{0.5}(n)$ for C large)
  - diameter $\Theta(1/r)$
    - upper bound: the graph of the squares has diameter $\Theta(1/r)$, each square contains at least one vertex, and there is an edge between any two vertices in neighboring squares
    - lower bound follows easily from the geometry
  - all degrees $\Theta(nr^2)$: all vertices in the (usually) 8 neighboring squares are neighbors and all neighbors lie in the (usually 48) squares in distance at most 3; hence $\deg v$ is the sum of a constant number of X^S
- Degree-diameter bound: Rumors spread in time $O(nr)$
- Observation: This bound becomes weaker for larger r
  - if this was true, then a larger communication power would reduce the rumor spreading speed

Rumor Spreading in Well-Connected RGGs (3)

- Better: Two stage argument
  - Call a square informed if it contains one informed vertex. This gives a rumor spreading process similar to the one in 2D grids (time $\Theta(r)$)
  - Once all squares are informed (in the above sense), argue that what happens inside the square is similar to rumor spreading in a complete network (time $\Theta(\log n)$)

- Square process:
  - Since each vertex has a constant fraction of its neighbors in each of the 8 neighbors, it takes expected constant time to call from one square to a given neighbor
  - "follow the path" argument (in grid of square): $\Theta(\max\{r^{-1}, \log n\})$ rounds until each square has an informed node

Rumor Spreading in Well-Connected RGGs (4)

- Making all squares fully informed (all vertices in the square are informed)
  - Each vertex has a probability greater than some constant c to call another vertex in its own square (which then is chosen uniformly at random from the square)
  - Hence ignoring all calls that cross square boundaries, each square is running a rumor spreading process in a complete graph in which each informed vertex only call with probability c.
  - This process is very similar to the usual rumor spreading process in complete graphs, but slowed down by a factor of at most $1/c$. Hence a large $\log(n)$ runtime is enough to have each square fully informed with probability $1 - 1/n^2$.
  - Union bound: All squares fully informed… additional $O(\log n)$ rounds for each square to inform itself

- Theorem: With high probability a well-connected RGGs is such that rumors spread in time $\Theta(diam(G) + \log n)$ with probability $1 - 1/n$.

Summary

- Rumor spreading is an efficient epidemic algorithm to disseminate information in various network topologies
  - hypercubes (as models for man-made communication networks)
  - random geometric graphs (model for wireless sensor networks)
  - preferential attachment graphs (model for social networks)

- Often, we can prove a rumor spreading time of $O(\mathrm{diam}\ G)$, which naturally is asymptotically optimal