Available online at www.sciencedirect.com
ScienceDirect
Electronic Notes in Discrete Mathematics 34 (2009) 243–247
Electronic Notes in DISCRETE MATHEMATICS
www.elsevier.com/locate/endm

Quasirandom Rumor Spreading on Expanders

Benjamin Doerr [a] Tobias Friedrich [a] Thomas Sauerwald [b]

[a] *Max-Planck-Institut für Informatik, Saarbrücken, Germany*

[b] *International Computer Science Institute, Berkeley, CA, USA*

Abstract

Randomized rumor spreading is an efficient way to distribute information in networks. Recently, a quasirandom version of this protocol has been proposed. It was proven that it works equally well or even better in many settings.

In this work, we exhibit a natural expansion property for networks, which ensures that quasirandom rumor spreading informs all nodes of the network in logarithmic time with high probability. This expansion property is satisfied, among others, by many expander graphs, random regular graphs, and Erdős-Rényi random graphs.

1 Introduction

Randomized rumor spreading or *random phone call protocols* are simple randomized epidemic algorithms designed to distribute a piece of information in a network. Many build on the simple approach that informed nodes call random neighbors and make them informed (also called *push model*). Since most of the existing literature regards the push model, so do we in this work. In spite of the simple concept, these broadcasting algorithms succeed in distributing information extremely fast. In contrast to many natural deterministic approaches, they are also highly robust against transmission failures.

Such algorithms have been applied successfully both in the context where a single news has to be distributed from one processor to all others (cf. [6]),

1571-0653/$ – see front matter
doi:10.1016/j.endm.2009.07.040

and in the one where news may be injected at various nodes at different times. The latter problem occurs when maintaining data integrity in a distributed databases, e.g., name servers in large corporate networks [1, 8]. For a more extensive, but still concise discussion of various central aspects of this area, we refer to the paper by Karp, Schindelhauer, Shenker, and Vöcking [7].

1.1 Quasirandom Rumor Spreading

Rumor spreading protocols often assume that all nodes have access to a central clock. The protocols then proceed in rounds, in each of which each node independent from the others can perform certain actions. In the classical randomized rumor spreading protocol, in each round each node contacts a neighbor chosen independently and uniformly at random. The latter node becomes informed if the first was.

In [2], the authors proposed a quasirandom version of randomized rumor spreading. It assumes that each node has a (cyclic) list of its neighbors. Except that each node starts at a random position in the list, this list describes the order in which the node contacts its neighbors. We make no particular assumption on the structure of the list. This allows to use any list that is already present to technically organize the communication of one node with its neighbors. In such, this protocol is rather simpler than the classical, fully random model.

Surprisingly, even though the amount of independent randomness is greatly reduced, similar or even better results could be shown. For the graph being an n–vertex hypercube, $\mathcal{O}(\log n)$ rounds suffice to inform with probability $1-\frac{1}{n}$ all nodes. For random graphs $G \in \mathcal{G}(n,p)$, $p \geqslant (\ln(n)+\omega(1))/n$, $\mathcal{O}(\log n)$ rounds suffice with probability $1 - o(1)$. Similar results are known for the classical model [5], except that for random graphs one requires $p \geqslant (1+\varepsilon)\ln(n)/n$, $\varepsilon > 0$ constant.

These theoretical results are complemented by an experimental investigation [3], which observes that the quasirandom model typically needs less time than the fully random one. For example, we have a runtime reduction by more than 10% for the 12-dimensional hypercube.

1.2 Our Results

In this paper, we greatly expand the first results of [2]. We exhibit a natural expansion property that guarantees that quasirandom rumor spreading succeeds in $\mathcal{O}(\log n)$ iterations.

Our expansion property roughly speaking requires that small sets of vertices have many neighbors, that for large sets of vertices the external vertices have many neighbors in the set, and finally that the vertex degrees are of similar order (see Definition 3.1 for the details).

These requirements are fulfilled by random graphs $\mathcal{G}(n,p)$, where p can be as small as $(\ln(n) + \omega(1))/n$, (regular) expander graphs and random regular graphs. Hence our result also subsumes (and improves in terms of the failure probability) the result on random graphs in [2].

For all these graphs, we show that with probability $1 - n^{-\gamma}$, where γ can be an arbitrary constant, the quasirandom rumor spreading model succeeds in informing all vertices from a single initially informed one in $\mathcal{O}(\log n)$ rounds. This result holds independent of how the cyclic lists look like.

2 Precise Model and Preliminaries

Our aim is spreading a rumor in an undirected graph $G = (V, E)$ with a quasirandom model. Let $n = |V|$ denote the number of vertices. In the quasirandom model, each vertex $v \in V$ is equipped with a cyclic permutation $\pi_v \colon \Gamma(v) \to \Gamma(v)$ of its neighbors $\Gamma(v)$. This can also be seen as a list of its neighbors.

At the start of the protocol each vertex v chooses a first neighbor i_v uniformly at random from $\Gamma(v)$. This is the neighbor it contacts at time $t = 1$. In each following time step $t = 2, 3, \ldots$, the vertex v contacts a vertex $\pi_v^{t-1}(i_v)$. For the quasirandom push model the result of one vertex contacting another one is as follows. If v was informed at time $t-1$, then $\pi_v^t(i_v)$ becomes informed at time t (provided it was not already informed).

This model slightly deviates from the description in the introduction, where each vertex chooses the starting point on its list only when it becomes informed. However, the two variants are clearly equivalent and in the proofs it is advantageous to assume that all vertices start contacting their neighbors already when they are uninformed.

Throughout the paper, we use the following notation. For a vertex v of a graph $G = (V, E)$, we denote by $\Gamma(v) := \{u \in V \mid \{u, v\} \in E\}$ the set of its *neighbors* and by $\deg(v) := |\Gamma(v)|$ its *degree*. For any $S \subseteq V$, let $\deg_S(v) := |\Gamma(v) \cap S|$. Let $\delta := \min_{v \in V} \deg(v)$ be the *minimum degree*, $d := 2|E|/n$ be the *average degree*, and $\Delta := \max_{v \in V} \deg(v)$ be the *maximum degree*. All logarithms $\log n$ are natural logarithms to the base e in the following. As we are only interested in the asymptotic behavior, we will sometimes assume that n is sufficiently large.

3 Expanding Graphs

Instead of analyzing specific graphs, we have distilled three simple properties. If these are satisfied, we can prove that quasirandom rumor spreading succeeds in informing all vertices in a logarithmic time. This is independent of which vertex is initially informed and independent of the order of the lists.

Definition 3.1 (expanding graphs) *We call a connected graph expanding if the following properties hold.*

(P1) *For all constants C_α with $0 < C_\alpha \leqslant d/2$ there is a constant $C_\beta \in (0,1)$ such that for any connected subset $S \subseteq V$ with $3 \leqslant |S| \leqslant C_\alpha n/d$, it holds that $|\Gamma(S) \setminus S| \geqslant C_\beta \, d \, |S|$.*

(P2) *There are constants $C_\delta \in (0,1)$ and $C_\omega > 0$ such that for any subset $S \subseteq V$, the number of vertices in S^c which have at least $C_\delta d(|S|/n)$ neighbors in S is at least $|S|^c - \frac{C_\omega n^2}{d|S|}$.*

(P3) $d = \Omega(\Delta)$. *If $d = \omega(\log n)$ then $d = O(\delta)$.*

Some important graph classes satisfy these properties. *Random graphs* in the classical Erdős-Rényi model $\mathcal{G}(n,p)$ are expanding with probability $1 - o(1)$, if $p \geqslant (\log n + \omega(1))/n$. By setting $p = 1$, this includes complete graphs. Let G be a d-regular graph and $\lambda(G) := \max\{|\lambda_2|, |\lambda_n|\}$ be the second largest eigenvalue of its adjaceny matrix. We call G an *expander* if $\lambda(G) \leqslant \min\{d/C, C'\sqrt{d}\}$, where $C > 1, C' > 0$ are arbitrary constants. Naturally, expanders are expanding. Since a random d-regular graph is an expander with probability $1 - o(1)$, it also is expanding with probability $1 - o(1)$.

Our main result is that quasirandom rumor spreading works well on expanding graphs.

Theorem 3.2 *Let $\gamma \geqslant 1$ be a constant. The probability that the quasirandom push model started at an arbitrary vertex of an expanding graph informs all other vertices within $\mathcal{O}(\log n)$ rounds is $1 - \mathcal{O}(n^{-\gamma})$.*

The proof uses the expansion properties roughly speaking as follows. **(P1)** ensures that while only a small number of vertices is informed, these have most of their neighbors uninformed. Consequently, most of the messages sent inform a new vertex. **(P2)** ensured that if many vertices are informed, then most of the remaining ones have many informed neighbors. Consequently, they cannot remain uninformed for too long. However, for all this to work properly, the vertex degrees have to be similar, as prescribed by **(P3)**. More details on the proof can be found in [4].

4 Conclusion and Outlook

In this work, we made significant progress to understand quasirandom rumor spreading. We have shown that on a large class of graphs called expanding graphs, quasirandom rumor spreading succeeds in $\mathcal{O}(\log n)$ rounds. From a broader view-point, this work again shows that a reduced amount of randomness can yield superior algorithms, and that these can still be analyzed in spite of the inherent dependencies.

A question not regarded so far for the quasirandom broadcasting model is how many transmissions are necessary to inform all nodes. Results analyzing how the pure model can be enhanced to reduce the number of transmissions only exist for the fully random model. Hence this seems to be natural direction for future research.

We should note, though, that the quasirandom push model naturally never needs more than $2m = nd$ messages. Hence for really sparse networks, e.g., expander graphs with constant degree d, the quasirandom model without greater optimization immediately yields a protocol that is both time and message efficient.

References

[1] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart, and D. B. Terry. Epidemic algorithms for replicated database maintenance. *Operating Systems Review*, 22:8–32, 1988.

[2] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading. In *Proceedings of SODA 2008*, pp. 773–781, 2008.

[3] B. Doerr, T. Friedrich, M. Künnemann, and T. Sauerwald. Quasirandom rumor spreading: An experimental analysis. In *Proceedings of ALENEX 2009*, pp. 145–153, 2009.

[4] B. Doerr, T. Friedrich, and T. Sauerwald. Quasirandom rumor spreading: Expanders, push vs. pull, and robustness. In *Proceedings of ICALP 2009*. To appear.

[5] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1:447–460, 1990.

[6] S. M. Hedetniemi, S. T. Hedetniemi, and A. L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18:319–349, 1988.

[7] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proceedings of FOCS 2000*, pp. 565–574, 2000.

[8] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proceedings of FOCS 2003*, pp. 482–491, 2003.