

Brief Announcement: Rapid Mixing of Local Dynamics on Graphs

Laurent Massoulié¹ and Rémi Varloot²

- 1 Inria, MSR-Inria Joint Centre, Palaiseau, France
laurent.massoulie@inria.fr
- 2 Inria, MSR-Inria Joint Centre, Palaiseau, France
remi.varloot@inria.fr

Abstract

In peer-to-peer networks, it is desirable that the logical topology of connections between the constituting nodes make a well-connected graph, i.e., a graph with low diameter and high expansion. At the same time, this graph should evolve only through local modifications. These requirements prompt the following question: are there local graph dynamics that i) create a well-connected graph in equilibrium, and ii) converge rapidly to this equilibrium?

In this paper we provide an affirmative answer by exhibiting a local graph dynamic that mixes provably fast. Specifically, for a graph on N nodes, mixing has occurred after each node has performed $O(\text{polylog}(N))$ operations. This is in contrast with previous results, which required at least $\Omega(N \text{ polylog}(N))$ operations per node before the graph had properly mixed.

1998 ACM Subject Classification E.1 Graphs and Networks

Keywords and phrases Markov chains, Mixing time, Dynamic graphs, Local dynamics

Digital Object Identifier 10.4230/LIPIcs.DISC.2017.59

1 Introduction

Peer-to-peer networks perform best if the graph describing their topology is well-connected. Indeed, the diameter of the graph conditions the time required to broadcast information from any one node to the rest of the network. The expansion of the graph conditions the robustness of epidemic algorithms for maintaining content in the system. It also conditions how quickly a random walk over the graph reaches stationarity, and as such determines the performance of various distributed algorithms, e.g., for searching content over the network.

The distributed evolution of such networks, however, can potentially create ill-connected graphs through an unlucky series of node arrivals and departures. This motivates our goal in the present paper: identify local graph dynamics that create a well-connected graph in a short amount of time, i.e., after each node has performed few operations ($O(\text{polylog}(N))$, where N denotes the total number of nodes), and this regardless of how poorly connected the initial graph is.

2 Related Work

Graph models meant to capture properties of real-life networks have been thoroughly studied [7]. Important examples include the Barabási-Albert preferential attachment model, yielding graphs with power law degree distribution [1], and random regular graphs, shown to have the small-world property of social networks (i.e., a small diameter) with high probability [12, 15].



© Laurent Massoulié and Rémi Varloot;
licensed under Creative Commons License CC-BY
31st International Symposium on Distributed Computing (DISC 2017).
Editor: Andréa W. Richa; Article No. 59; pp. 59:1–59:3



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Other papers address the design of dynamics meant to alter the overall structure of the graph in a given way [13] or to minimize the convergence rate [9]. Closer to our present work, these issues have been considered specifically for the construction of p2p networks in [6, 5, 8].

The speed of convergence to equilibrium of graph dynamics has been studied in different contexts. [4] considers non-local dynamics. Local dynamics are considered in [14] and [10] for the synthesis of cladograms and bipartite graphs respectively. Closer to our context, [2] considers the local switch dynamic introduced in [8], and proves that it yields an expander graph after $O(N \text{polylog}(N))$ operations per node. The analysis in [2] is very elaborate, and the stated bound improves upon all previous results on local graph dynamics proposed for peer-to-peer topology maintenance. Nevertheless, this bound is not yet satisfactory, as it still increases quickly (at least linearly) with the system size.

3 Our Contribution

Consider the following setting: fix a vertex set $[N] = \{1, \dots, N\}$, where N is a positive integer, and connect the vertices in $[N]$ as follows. First, add a set of fixed edges $(i, i+1)$ constituting a cycle ($N+1 \equiv 1$). Then, have each node $n \in [N]$ maintains two pointers: a blue pointer to a node $b_n \in [N]$, and a red one to $r_n \in [N]$, such that each node n is the destination of exactly one blue pointer and one red pointer. In essence, b and r constitute permutations over $[N]$. From these, we construct a set of undirected blue edges $\{(n, b_n) : n \in [N]\}$, and a set of undirected red edges $\{(n, r_n) : n \in [N]\}$. The result is a 6-regular graph G over $[N]$ composed of N cycle edges, N blue edges and N red edges.

The dynamic then proceeds as follows. The graph evolves through alternating red and blue phases; during a blue phase, only the blue pointers are modified, while the red pointers are kept fixed. The blue pointers slide along the graph G^r formed by the union of the cycle edges and the undirected edges $\{(n, r_n) : n \in [N]\}$ formed by the red pointers. For the red phases, the roles of the blue and red pointers are swapped.

Formally, the dynamic for the blue edges over G^r is as follows: at each time step, pick an edge (i, j) uniformly at random in G^r , and denote n and m the two nodes in $[N]$ such that $b_n = i$ and $b_m = j$. These two nodes swap their pointers: now $b_n = j$ and $b_m = i$. Notice that b is still a permutation over $[N]$. This dynamic is known in the literature as the interchange process [11, 3]. Our main result is then as follows:

► **Theorem 1.** *Let $T = N \ln(N)^a$ where $a > 8$ is a constant. Then with high probability, after $O(\ln(N))$ alternating phases of length T , the blue and red pointers constitute uniformly and independently distributed permutations of $[N]$.*

► **Corollary 2.** *With the above process, G is an expander with high probability after each node has performed only $O(\text{polylog}(N))$ operations.*

4 Sketch of Proof

For any d -regular graph H over $[N]$, denote $\phi_k(H) = \min_{S \subset [N]: 0 < |S| \leq k} \frac{|\partial S|}{d|S|}$, where ∂S is the set of edges between a node in S and a node not in S . Exploiting properties of the interchange process stated in [3], we show that, if $\phi_{N/2}(G^r) \geq \gamma$ for a well chosen γ , dependent on N , then at the end of the following blue phase, the blue pointers are uniformly distributed with high probability.

We then construct an increasing sequence $(k_t)_{t \in \mathbb{N}}$ such that i) at the end of the t -th blue/red phase, $\phi_{k_t}(G^{b/r}) \geq \gamma$ and ii) there exists $\tau = O(\log(N))$ such that $k_\tau = N/2$. Joining the dots proves the theorem.

Acknowledgements. The authors would like to thank George Giakkoupis for bringing the problem studied in this paper to their attention, and to acknowledge stimulating discussions on the topic with both George Giakkoupis and Marc Lelarge.

References

- 1 Réka Albert and Albert-András Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, page 2002.
- 2 Zeyuan Allen Zhu, Aditya Bhaskara, Silvio Lattanzi, Vahab S. Mirrokni, and Lorenzo Orecchia. Expanders via local edge flips. *CoRR*, abs/1510.07768, 2015.
- 3 N. Berestycki. Mixing times of markov chains: Techniques and examples. *Alea-Latin American Journal of Probability and Mathematical Statistics*, 2016. URL: <http://www.statslab.cam.ac.uk/~beresty/teach/Mixing/mixing3.pdf>.
- 4 Shankar Bhamidi, Guy Bresler, and Allan Sly. Mixing time of exponential random graphs. *The Annals of Applied Probability*, 21(6):2146–2170, 12 2011.
- 5 Colin Cooper, Martin Dyer, and Andrew J. Handley. The flip markov chain and a randomising p2p protocol. In *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing*, PODC’09, pages 141–150, New York, NY, USA, 2009. ACM. doi:10.1145/1582716.1582742.
- 6 Colin Cooper, Ralf Klasing, and Tomasz Radzik. A randomized algorithm for the joining protocol in dynamic distributed networks. Technical Report RR-5376, INRIA, November 2004. URL: <https://hal.inria.fr/inria-00070627>.
- 7 Rick Durrett. *Random Graph Dynamics*. Cambridge University Press, Cambridge, 2007. URL: <http://www.math.cornell.edu/~durrett/RGD/RGD.html>.
- 8 Tomas Feder, Adam Guetz, Milena Mihail, and Amin Saberi. A local switch markov chain on given degree graphs with application in connectivity of peer-to-peer networks. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, FOCS’06, pages 69–76, Washington, DC, USA, 2006. IEEE Computer Society. doi:10.1109/FOCS.2006.5.
- 9 Riko Jacob, Andrea Richa, Christian Scheideler, Stefan Schmid, and Hanjo Täubig. A distributed polylogarithmic time algorithm for self-stabilizing skip graphs. In *Proceedings of the 28th ACM Symposium on Principles of Distributed Computing*, PODC’09, pages 131–140, New York, NY, USA, 2009. ACM. doi:10.1145/1582716.1582741.
- 10 M. Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, dec 1989.
- 11 Johan Jonasson. Mixing times for the interchange process. *Alea-Latin American Journal of Probability and Mathematical Statistics*, 9(2):667–683, 2012.
- 12 Minkyu Kim and Muriel Medard. Robustness in large-scale random networks. *INFOCOM*, 2004.
- 13 Manos Papagelis. Refining social graph connectivity via shortcut edge addition. *ACM Trans. Knowl. Discov. Data*, 10(2):12:1–12:35, oct 2015.
- 14 Jason Schweinsberg. An $o(n^2)$ bound for the relaxation time of a markov chain on cladograms. *Random Struct. Algorithms*, 20(1):59–70, jan 2002.
- 15 Lingsheng Shi and Nicholas Wormald. Models of random regular graphs. In *IN SURVEYS IN COMBINATORICS*, pages 239–298. University Press, 1999.