Faster Rumor Spreading with Multiple Calls

Konstantinos Panagiotou[1], Ali Pourmiri[2], and Thomas Sauerwald[2]

[1] The Ludwig Maximilian University of Munich, Germany.
kpanagio@math.lmu.de
[2] Max Planck Institute for Informatics, Saarbrücken, Germany.
{pourmiri, sauerwal}@mpi-inf.mpg.de

Abstract. We consider the random phone call model introduced by Demers et al. [1], which is a well-studied model for information dissemination on networks. One basic protocol in this model is the so-called Push protocol which proceeds in synchronous rounds. Starting with a single node which knows of a rumor, every informed node calls a random neighbor and informs it of the rumor in each round. The Push-Pull protocol works similarly, but additionally every uninformed node calls a random neighbor and may learn the rumor from that neighbor.

In one of the first papers in this area, Frieze and Grimmett [2] proved that if the underlying graph is a complete graph with n nodes, then the runtime of the Push protocol is $\log_2 n + \log n \pm o(\log n)$ with high probability[3], where log denotes the natural logarithm. This result was later strengthened by Pittel [4]. For the standard Push-Pull protocol, Karp et al. [3] proved a runtime bound of $\log_3 n + \mathcal{O}(\log\log n)$.

In this work, we are interested by how much we can speed up the spread of the rumor by enabling nodes to make more than one call in each round. We propose a new model where the number of calls of a node u is chosen independently according to a probability distribution R with bounded mean determined at the beginning of the process. We provide both lower and upper bounds on the rumor spreading time depending on statistical properties of R such as the mean or the variance. For instance, we show that if R follows a power law distribution with exponent $\in (2, 3)$, the Push-Pull protocol spreads a rumor in $\Theta(\log\log n)$ rounds.

References

1. A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. 6th Symp. Principles of Distributed Computing (PODC)*, pages 1–12, 1987.
2. A. Frieze and G. Grimmett. The shortest-path problem for graphs with random-arc-lengths. *Discrete Applied Mathematics*, 10:57–77, 1985.
3. R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized Rumor Spreading. In *Proc. 41st Symp. Foundations of Computer Science (FOCS)*, pages 565–574, 2000.
4. B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.

[3] By with high probability we refer to an event which holds with probability $1 - o(1)$ as $n \to \infty$. For simplicity, we sometimes omit the "with high probability" in the introduction.