The Freeze-Tag Problem: Theoretical and Experimental Investigations

Esther M. Arkin* Michael A. Bender † Sándor P. Fekete‡ Joseph S. B. Mitchell*
Martin Skutella § Marcelo O. Sztainberg*

Abstract

An optimization problem that naturally arises in the study of "swarm robotics" is to wake up a set of "asleep" robots, starting with only one "awake" robot. One robot can only awaken another when they are in the same location. As soon as a robot is awake, it assists in waking up other robots. The goal is to compute an optimal *awakening schedule* such that all robots are awake by time t^*, for the smallest possible value of t^*.

We consider both scenarios on graphs and in geometric environments. In the graph setting, robots sleep at vertices and there is a length function on the edges. An awake robot can travel from vertex to vertex along edges, and the length of an edge determines the time it takes to travel from one vertex to the other.

While this problem bears some resemblance to problems from various areas in combinatorial optimization such as routing, broadcasting, scheduling and covering, its algorithmic characteristics are surprisingly different.

This talk is based on two companion papers investigating theoretical[1] and experimental[2] questions associated with the freeze-tag problem (FTP).

We prove that the FTP is NP-hard, even for the special case of star graphs. We also establish hardness of approximation, showing that it is NP-hard to obtain an approximation factor better than 5/3, even for graphs of bounded degree. These lower bounds are complemented with several algorithmic results. We present a simple online algorithm that is $O(\log \Delta)$-competitive for graphs with maximum degree Δ. Other results include algorithms that require substantially more sophistication and development of new techniques: (1) The natural greedy strategy on star graphs has a worst-case performance of 7/3, which is tight. (2) There exists a PTAS for star graphs. (3) For the problem on ultrametrics, there is a polynomial-time approximation algorithm with performance ratio $2^{O(\sqrt{\log \log n})}$. (4) There is a PTAS, running in nearly linear time, for geometrically embedded instances (e.g., Euclidean distances in any fixed dimension).

We also perform a suite of experiments designed to investigate the effectiveness of some heuristic algorithms under various design options. These experiments show the relative comparison among four methods when applied to various forms of geometric data as well as more general metric spaces (with datasets from the TSPLIB).

1 Introduction

The following natural problem arises in the study of *swarm robotics.* Consider a set of n *robots*, modeled as points in some metric space (e.g., vertices of an edge-weighted graph). Initially, there is one "awake" robot and all others are *asleep*, i.e., in stand-by mode. Our goal is to "wake up" all of the robots as quickly as possible. In order to wake up a sleeping robot, an awake robot must go to its location; once awake, this new robot is available to assist in awaking other robots. The objective is to minimize the "makespan", which is the time when the last robot is awakened.

This problem is reminiscent of the children's game of "freeze-tag", in which the person who is "it" tags a player, who becomes "frozen" until another player (who is not "it" and not "frozen") tags him to unfreeze him. Our problem arises in the situation in which there are a large number, n, of players that are frozen, and one (not "it") unfrozen player, whose goal it is to unfreeze the rest of the players as quickly as possible. Of course, as soon as a player gets unfrozen, s/he is available to assist in unfreezing other frozen players, so there is a cascading effect. Due to the similarity with this game, we have dubbed our problem the *freeze-tag problem* (FTP).

Other applications of the FTP arise in the context of distributing data (or some other commodity), where physical proximity is required to allow producing a copy. This may be because wireless communication is too costly in terms of bandwidth or because there is too much of a security risk. How does one propagate the data to the entire set of participants in the most efficient manner?

In this paper, we give the first algorithmic and experimental results on this problem, which arises naturally as a hybrid of various problems from the areas of broadcasting, routing, scheduling, and network design. As in *broadcasting problems*, the goal is to dis-

*Dept. of Appl. Math. and Statistics, SUNY Stony Brook, NY 11794-3600, {estie, jsbm, mos}@ams.sunysb.edu.

†Dept. of Computer Science, SUNY Stony Brook, NY 11794-4400, bender@cs.sunysb.edu.

‡Abt. für Mathematische Optimierung, TU Braunschweig, D-38106 Braunschweig, Germany, s.fekete@tu-bs.de.

§Institut für Mathematik, TU Berlin, 10623 Berlin, Germany, skutella@math.tu-berlin.de.

[1] See "The Freeze-Tag Problem: How to Wake Up a Swarm of Robots," by Arkin, Bender, Fekete, Mitchell, Skutella, to appear SODA'02.

[2] See "An Experimental Investigation of Approximation Algorithms for the Freeze-Tag Problem," by Arkin, Bender, Mitchell, and Sztainberg, manuscript 2001.

seminate information in a network. The FTP also has elements of optimal *routing*, since robots must travel to awaken others or to pass off information. *Scheduling* issues arise in having to assign priorities to sleeping robots due to the limited but growing number of awake robots. One must decide whether to awaken a small nearby cluster to obtain a small number of helpers quickly or whether to awaken a distant but populous cluster to obtain many helpers, but after a longer delay. The relationship to *network design* is given by the fact that finding an optimal schedule can be interpreted as determining a spanning binary tree of minimum depth in a complete weighted graph.

There is an abundance of prior work on the dissemination of data in a graph. Most closely related to the FTP are the *minimum broadcast time problem*, the *multicast problem*, and the related *minimum gossip time problem*. See [4] for a survey; see [1, 6] for recent approximation results. However, that there are significant differences: While the broadcast problem can be solved in polynomial time in tree networks, the FTP turns out to be NP-hard already on the seemingly easy class of weighted stars. And although it is fairly straightforward to obtain an on-line algorithm that is $O(\log n)$-competitive, it is highly nontrivial to obtain an $o(\log n)$ approximation bound, even on some seemingly special graphs, such as trees.

1.1 Summary of Theoretical Results:

(a) We prove that the freeze-tag problem is NP-hard, even for the case of star graphs with an equal number of robots at each vertex. Moreover, there exists a PTAS in this case and we analyze the greedy heuristic, showing a tight performance bound of 7/3.

(b) We give an $o(\log n)$-approximation algorithm for the FTP in ultrametrics [2, 3, 5]. Specifically, we obtain an approximation ratio of $2^{O(\sqrt{\log \log n})}$, based on a series of nontrivial structural results. (Our approximation factor may be novel among existing approximation algorithms; it is a rather unusual term, asymptotically lying between $\log \log n$ and $\log n$.)

(c) We give a simple linear-time on-line algorithm that is $O(\log \Delta)$-competitive for the case of general weighted graphs of maximum degree Δ. On the other hand, we show for the off-line problem that finding a solution within less than 5/3 is NP-hard, even for graphs of maximum degree 5.

(d) We give a PTAS for geometric instances of the FTP in fixed-dimension domains endowed with an L_p metric. In fact, our algorithm runs in near-linear time, $O(n \log n + 2^{poly(1/\varepsilon)})$, with the nonpolynomial dependence on ε showing up only as an additive term in the time complexity.

1.2 Summary of Experimental Results:

(a) We have implemented four main heuristic methods for attacking the FTP:

1. Greedy: when a robot first awakens, it goes to the nearest asleep robot
2. "Bang for the Buck": when a robot first awakens, it selects the next asleep robot to approach based on which of a constant number of sectors around it contains the most "valuable" candidate, where value depends on the number of nearby robots
3. Random Sector Selection: an awakened robot selects the next asleep robot to approach based on randomly choosing among a constant number of sectors
4. Opposite Cone Heuristic: a variant of the Random Sector Selection designed to "mix up" the sector chosen

(b) We have conducted experiments comparing these methods under various choices of the following parameters: (a) whether or not robots lay "claims" to asleep robots, (b) whether or not the assignment of awake robots to asleep robots is "refreshed" at each instant a new robot awakens, and (c) whether or not we perform the heuristic in an "offline" mode.

(c) We have investigated the performance of the methods on a variety of datasets, including uniformly distributed points in the plane, cluster distributions, and rectangular and hexagonal grids. We compare the overall wakeup time to the known lower bound (distance from source to furthest away robot) and find that the methods perform better in practice than the worst-case approximation ratios suggest. We also give comparisons of performance across the four methods, with different choices of parameters, on the various types of data.

Acknowledgements. We thank Tien-Ruey Hsiang and Nenad Jovanovic for many helpful discussions on the topic of this paper.

References

[1] A. Bar-Noy, S. Guha, J. Naor, and B. Schieber. Multicasting in heterogeneous networks. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC-98)*, pages 448–453, New York, May 23–26 1998. ACM Press.

[2] Bartal. On approximating arbitrary metrics by tree metrics. In *STOC: ACM Symposium on Theory of Computing (STOC)*, 1998.

[3] M. Charikar, C. Chekuri, A. Goel, S. Guha, and S. Plotkin. Approximating a finite metric by a small number of tree metrics. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, pages –, 1998.

[4] S. M. Hedetniemi, T. Hedetniemi, and A. L. Liestman. A Survey of Gossiping and Broadcasting in Communication Networks. *NETWORKS*, 18:319–349, 1988.

[5] G. Konjevod, R. Ravi, and F. Salman. On approximating planar metrics by tree metrics, 1997.

[6] R. Ravi. Rapid rumor ramification: Approximating the minimum broadcast time. In S. Goldwasser, editor, *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 202–213, Los Alamitos, CA, USA, Nov. 1994. IEEE Computer Society Press.