Discrete Applied Mathematics 247 (2018) 53–69

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

Optimal odd gossiping

Guillaume Fertin [a], Joseph G. Peters [b,*]

[a] LS2N, UMR CNRS 6004, Université de Nantes, Nantes, France
[b] School of Computing Science, Simon Fraser University, Burnaby, British, Columbia, Canada

A R T I C L E I N F O

Article history:
Received 18 July 2017
Received in revised form 4 March 2018
Accepted 10 March 2018
Available online 4 April 2018

Keywords:
Gossiping
Information dissemination
Communication networks

A B S T R A C T

In the *gossiping* problem, each node in a network starts with a unique piece of information and must acquire the information of all other nodes using two-way communications between pairs of nodes. In this paper we investigate gossiping in n-node networks with n odd. We use a *linear cost model* in which the cost of communication is proportional to the amount of information transmitted. In *synchronous* gossiping, the pairwise communications are organized into *rounds*, and all communications in a round start at the same time. We present optimal synchronous gossip algorithms for all odd values of n, proving the truth of a conjecture by Fertin, Peters, Raabe, and Xu. Central to the construction of the algorithms is a doubly-inductive proof about properties of optimal gossip algorithms.

1. Introduction

Gossiping is an information dissemination process in which each node of a communication network has a piece of information that must be acquired by all the other nodes. Information is communicated between pairs of nodes using two-way communications or *calls* along the communication links of the network. Gossiping is a well-studied problem. There are many papers describing algorithms that minimize the gossip time on various interconnection networks (e.g. hypercubes, meshes, Cayley graphs), using a variety of switching models (store-and-forward, circuit-switched, cut-through) and cost models (unit cost, linear cost), as well as various models of fault-tolerance. There are also papers describing methods to construct gossip graphs that minimize the resources (usually the number of communication links) needed to allow minimum time gossiping, again for various models. See [7,9–12] for surveys of these results.

There has been less study of the minimum time needed to gossip when the topology of the interconnection network does not restrict the communication patterns. In an early paper, Knödel [14] proved that the number of *rounds* of communication necessary to gossip among n nodes is $\lceil \log_2 n \rceil$ when n is even, and $\lceil \log_2 n \rceil + 1$ when n is odd. He also proved sufficiency by describing gossip algorithms that meet the lower bounds on numbers of rounds. The *half-duplex* version of this problem, in which communication links can only be used in one direction at any given time, has also been studied [3,15]. All of these papers assume a *unit cost* model in which a communication takes one time unit independent of the amount of information being transmitted.

There is considerable current interest in *gossip-based algorithms*. First introduced by Demers et al. [2], these are distributed algorithms that attempt to mimic the ways that epidemics spread. The algorithms consist of pairwise communications between nodes (peers) in which one node contacts a neighbour (usually chosen randomly) and exchanges some information. Originally devised for distributed database replication, they are also used for information dissemination, overlay network management, distributed consensus, aggregation, and resource management. They differ from the algorithms that we study

* Corresponding author.
E-mail addresses: guillaume.fertin@univ-nantes.fr (G. Fertin), peters@cs.sfu.ca (J.G. Peters).

https://doi.org/10.1016/j.dam.2018.03.033

in this paper mainly in that they are distributed, randomized, and are focused on the rate of convergence to a result rather than achieving a specific deterministic result. Kermarrec and van Steen [13] give a good overview of this area, and other papers in the same volume [1] examine the issues in more detail including a framework for classifying problems by Fernandess et al. [4].

In this paper, we study the structure of minimum-time gossip algorithms when the topology of the interconnection network does not restrict the communication patterns using a classical *store-and-forward, 1-port, full-duplex* model. In this model, each communication involves two nodes and a communication link that connects them, and information can flow simultaneously in both directions along a link. Each node starts with a piece of information of length 1. Information can be combined into longer messages and sent in a single communication. We will use a *linear cost* model in which the time to send a message of length ℓ is $\beta + \ell\tau$ where β is the time for the leading edge of a message to propagate along a link from sender to receiver, and $\frac{1}{\tau}$ is the data rate of the link. It will be convenient to think of a call involving messages of length ℓ as a start-up period that takes time β followed by a sequence of ℓ *steps* each of which takes time τ. If a call involves messages of different lengths, then the time for both nodes to complete the communication is determined by the length of the longer message.

The *synchronous* linear cost model that we use in this paper is a generalization of the unit cost model. A synchronous gossip algorithm consists of a sequence of *rounds* of simultaneous pairwise communications, and all calls in a round start at the same time. Calls in a round may end at different times, depending on the lengths of the messages, but no node can start a new call until all nodes are ready to start new calls. Gossiping has also been studied with an *asynchronous* linear cost model [6,8] in which a call can start as soon as both nodes are ready to communicate. This allows a pair of nodes to start communicating while calls between other pairs are in progress. We will restrict attention to synchronous gossip algorithms in this paper. Note that the unit cost model is always synchronous because each call takes one time unit.

Fraigniaud and Peters [8] investigated the structure of minimum-time gossip algorithms using a linear cost model. They proved a lower bound on the time to gossip when the number of nodes n is even and showed that there is a synchronous gossip algorithm that achieves the lower bound for every even n. They also gave examples to show that minimum-time gossip algorithms for some odd values of n must be asynchronous; any synchronous algorithm requires strictly more than minimum time.

Fertin, Peters, Raabe, and Xu [6] studied gossiping with n odd and a linear cost model. They proved a general lower bound of $(\lceil \log_2 n \rceil + 1)\beta + n\tau$ on the time to gossip for any $\beta \geq 0$ and $\tau \geq 0$. This lower bound holds for all odd n for both the synchronous and asynchronous models. The bound is achievable in the asynchronous model for some odd values of n, but they proved that every gossip algorithm for $n = 2^k - 1$, $k \geq 3$, requires time strictly greater than $(\lceil \log_2 n \rceil + 1)\beta + n\tau$. They proved stronger lower bounds for the synchronous model and conjectured that their lower bounds are achievable for all odd n. They gave an ad hoc synchronous algorithm that achieves their lower bound for $n = 2^k - 1$, $k \geq 2$.

In Section 2, we briefly review the lower bounds for synchronous gossiping from [6]. We then analyze Knödel's algorithm using the linear cost model and compare the upper bounds from this analysis to the lower bounds from [6]. As we will see, there is a large gap between the upper and lower bounds. Our main result in this paper is a collection of algorithms that achieves the lower bounds for <u>all</u> odd values of n, thereby establishing the truth of the conjecture in [6]. Our treatment of the synchronous upper bounds is split into two sections. In Section 3, we consider odd values of n that are in the *top half* of any range between two consecutive powers of 2. In Section 4, we consider the *bottom halves* of the ranges. The construction of the algorithms is based on two properties of optimal gossip algorithms which we establish in Section 3 with an interesting doubly-inductive proof.

2. Lower bounds and Knödel's algorithm

At any given time during a gossip algorithm for n odd, at least one node will be *idle* (not involved in a communication) because each call involves a pair of nodes. Based on this fact, Knödel [14] showed that gossiping in the unit cost model requires $\lceil \log_2 n \rceil + 1$ rounds when n is odd. This lower bound on the number of rounds is also valid for the synchronous linear cost model with $\beta \neq 1$ and $\tau = 0$. The only difference in this case is that each round of a gossip algorithm takes time β instead of time 1. By a similar argument, at least n steps are required to gossip when n is odd because each node needs to acquire $n - 1$ pieces of information, and at least one node is idle during each step. This gives a lower bound of $\max\{(\lceil \log_2 n \rceil + 1)\beta, n\tau\}$. Fertin, Peters, Raabe, and Xu [6] proved a lower bound of $(\lceil \log_2 n \rceil + 1)\beta + n\tau$ for odd n which applies to both the synchronous and asynchronous cases. They proved stronger lower bounds for the synchronous case by fixing the number of rounds to be $\lceil \log_2 n \rceil + 1$ and then focusing on the required number of steps. We take the same approach to synchronous upper bounds in this paper.

The pairwise communications of a synchronous gossip algorithm are organized into a sequence of rounds. All calls in a round start at the same time but the nodes may finish communicating at different times during the round depending on the lengths of the messages that they are receiving. All calls begin with a start-up period that takes time β followed by some number of steps. Since all nodes must finish the round before any node can start a new call, the length of the round depends on the number of steps needed to receive the longest message and we can unambiguously refer to the start-up period of a round and the number of steps in the round.

The required number of rounds, $\lceil \log_2 n \rceil + 1$, is the same for every odd n between $2^{k-1} + 1$ and $2^k - 1$, where $k = \lceil \log_2 n \rceil$. The required total number of steps for all rounds and also the required numbers of steps in each of the rounds depend on