Trade-Offs for Odd Gossiping
(Extended abstract)

Guillaume Fertin

LaBRI U.M.R. C.N.R.S. 5800, Université Bordeaux I
351 Cours de la Libération, F33405 Talence Cedex
`fertin@labri.u-bordeaux.fr`

1 Introduction

Communication networks are more and more used to realize communications between groups of users. Among the applications, one can think of video-conferencing, distributed data bases, or parallel computing. Different types of communications can appear in these networks, such as one-to-all (broadcast), one-to-many (multicast), or all-to-all (gossip) communications. In the latter, 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 along the links of the network. Here, we assume a *store-and-forward, 1-port, full-duplex* model in which each communication involves two nodes and the single communication link that connects them, each node communicates with at most one other node at any given time, and information can flow simultaneously in both directions along a link. Each node starts with a message of length 1, and messages can be concatenated and sent as a single communication.

Gossiping has been widely studied in the *unit cost* model, where a communication takes a fixed amount of time, independent of the number of messages transmitted. We assume in this paper a more realistic model : the *linear cost* model, in which the time to send a message is proportional to the amount of information exchanged. More precisely, the time to exchange a message of length L is $\beta + L\tau$, where β is the *start-up* time to initiate a call between a pair of nodes and τ is the *propagation* time of a message of length 1 along a link (in other words, $\frac{1}{\tau}$ is the bandwidth of the network). If the two nodes involved in a call send messages of different lengths, then the time for both nodes to complete the call is determined by the length of the longer message.

Moreover, a linear cost model can be either *synchronous* or *asynchronous.* In the synchronous linear cost model, a gossip algorithm consists of a sequence of *rounds* of simultaneous pairwise communications. 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. In the asynchronous linear cost model, a call can start as soon as both nodes are ready to communicate. Thus, a pair of nodes can start communicating while calls between other pairs are in progress.

In the fully-connected network with n nodes (which can be modelled by the complete graph K_n), we denote by $g_{\beta,\tau}(n)$ the minimum time needed to achieve gossiping. Clearly, $g_{\beta,\tau}(n)$ depends on β, τ and n. For a fixed n, it is then possible to reduce $g_{\beta,\tau}(n)$, depending on the relative values of β and τ (which themselves depend on the network). Indeed, if $\beta >> \tau$, then it is interesting to reduce the number of rounds of the gossip algorithm (that is, the factor associated to β). Conversely, when $\tau >> \beta$, it is better to reduce the factor associated to τ. When we are in an intermediate situation, a deeper study needs to be undertaken.

In the following, we call *step* a message of length 1. If two nodes u and v exchange respectively L_u and L_v pieces of information during a call, then this communication will take a time $t_{u,v} = \beta + \max\{L_u, L_v\}\tau$. Let $L = \max\{L_u, L_v\}$. Hence the *number of steps* of the call between u and v refers to L.

Moreover, in the synchronous case, for a given round r of the algorithm, the *number of steps of* r, s_r, is the maximum of all the number of steps of the calls which take place during round r. In other words, $s_r = \max\{L_u,\ u \text{ communicating during round } r\}$. Hence, in the synchronous case, the time necessary to achieve a round r is equal to $\beta + s_r\tau$. Finally, the *number of steps of a gossip algorithm* with R_n rounds, S_n, is defined as follows : $S_n = \sum_{r=1}^{R_n} s_r$.

In the asynchronous case, we only define the *number of steps of a gossip algorithm* : if an asynchronous algorithm takes a time equal to $R_n\beta + S_n\tau$, then S_n denotes the number of steps of such an algorithm.

In the following, we will always consider that we are given the fully-connected network with n nodes (i.e., the complete graph K_n). When n is even, Fraigniaud and Peters [FP94] have proved the following : $g_{\beta,\tau}(n) = \lceil \log_2(n) \rceil \beta + (n-1)\tau$ for any $\beta > 0$ and $\tau > 0$. We see that the factors associated to both β and τ are minimized here. Moreover, they have given a synchronous optimal-time algorithm in that case ; in other words, asynchronicity does not help in the case n even.

When n is odd, Knödel [Knö75] showed that gossiping in the unit cost model requires $\lceil \log_2(n) \rceil + 1$ rounds. This lower bound on the number of rounds is also valid for the linear cost model in both the synchronous and asynchronous cases. It is also immediate that a gossip algorithm takes at least n steps because each node needs to acquire $n-1$ pieces of information, and at least one node is idle at any given time. This gives a lower bound of $\max\{(\lceil \log_2(n) \rceil + 1)\beta, n\tau\}$. Peters, Raabe, and Xu [PRX96] proved a lower bound of $(\lceil \log_2(n) \rceil + 1)\beta + n\tau$ for odd n for both the synchronous and asynchronous cases. The bound is achievable in the asynchronous case for some odd values of n, but for $n = 2^k - 1$, every gossip algorithm requires time strictly greater than $(\lceil \log_2(n) \rceil + 1)\beta + n\tau$ (cf. [PRX96, FP98b]). 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 focussing on the required number of steps. They also conjectured that these lower bounds are achievable for all odd n, which has been proved to be true in [FP98b].

Here, we study, as suggested in [FP94, FP98b] the possible trade-offs between the number of rounds R_n and the number of steps S_n for a (synchronous or asynchronous) gossip algorithm in a n-node network, n being odd. As said before, depending on the relative values of β and τ, these trade-offs are of interest in order to reduce the value of $g_{\beta,\tau}(n)$. We then study, for any integer $r \geq 0$, the function $S_n(r)$ corresponding to the number of steps which are necessary for a gossip algorithm (in a n-node network) with $R_{min} + r$ rounds, where $R_{min} = \lceil \log_2(n) \rceil + 1$. We study such a function in both the synchronous and asynchronous models.

Similarly, we study, for any integer $s \geq 0$, the function $R_n(s)$, corresponding to the optimal number of rounds necessary for a gossip algorithm (in a n-node network) which takes exactly $S_{min} + s$ steps, where $S_{min} = n$. We study such a function in the synchronous model only.

2 Synchronous Model

2.1 Relaxing the condition on R_n

We discuss the value of $S_n(r)$, the optimal number of steps of a gossip algorithm among n nodes taking $R_n = R_{min} + r$ rounds, where $R_{min} = \lceil \log_2(n) \rceil + 1$. We give lower and upper bounds for $S_n(r)$, and prove that these bounds are tight in several cases (cf. Table 1). One of the main results is the following Theorem, which gives a general lower bound for $S_n(r)$.

Theorem 1 *For any odd n with $\lceil \log_2(n) \rceil = k$, let $R_n = k + 1 + r$ for any arbitrary r. For any $\alpha \geq 2$, let p and q be such that $n - 2^{k-\alpha+1} = (\alpha + r - 1) \cdot p + q$, with $0 \leq q \leq \alpha + r - 2$). Then :*

- *If $q = 0$, for all $n \geq 2^{k-\alpha} \cdot (\alpha + r + 1)$, we have $S_n(r) \geq 2^{k-\alpha+1} + (\alpha + r) \cdot p - 1$;*
- *If $q \neq 0$, for all $n \geq 2^{k-\alpha} \cdot (\alpha + r + 1) - ((\alpha + r - 1) - q)$, we have $S_n(r) \geq 2^{k-\alpha+1} + (\alpha + r) \cdot p + q$.*

Sketch of Proof: The idea here is to give a sequence σ of $k + 1 + r$ sets of steps s_i $(1 \leq i \leq k + 1 + r)$, each $|s_i|$ being to the number of steps used in round i. A necessary condition for σ to be valid for gossiping is the Basic Premise defined below. The Step Decreasing Property, also described below, implies that $\sum_{i=1}^{k+r+1} |s_i|$ is a lower bound for $S_n(r)$.

- The *Basic Premise* : for any round $1 \leq i \leq k + r + 1$, the sum of the number of steps in all the other rounds (that is, excluding round i) must be at least equal to $n - 1$. This comes from the fact that since n is odd, there is at least one idle vertex in round i. Hence we must have, for any sequence σ, and for any $1 \leq i \leq k + r + 1$, $\sum_{j=1, j \neq i}^{k+1+r} |s_j| \geq n - 1$.
- The *Step Decreasing Property* : for any round $1 \leq i \leq k + r + 1$, any decrease of 1 step in s_i implies an increase of at least one step in one or several s_j, with $j \neq i$.

Hence, if the Basic Premise and the Step Decreasing Property hold for the sequences σ we give below, then $S_n(r) \geq \sum_{i=1}^{k+r+1} |s_i|$.

In that case, we decide to choose a sequence σ which is as follows : during a certain number of rounds (precisely, during the first $k - \alpha + 1$ rounds), we use the maximum number of steps, that is 2^{i-1} steps for each round $1 \leq i \leq k - \alpha + 1$. Then, each of the the remaining $\alpha + r$ rounds will contain approximately the same number of steps. More precisely, we choose an integer $\alpha \geq 2$, and define p and q to be such that $n - 2^{k-\alpha+1} = (\alpha + r - 1) \cdot p + q$ (with $0 \leq q < \alpha + r - 1$). That is, $p = \lfloor \frac{n - 2^{k-\alpha+1}}{\alpha + r - 1} \rfloor = \frac{n - 2^{k-\alpha+1} - q}{\alpha + r - 1}$. Now, depending on the value of q, we have :

- $q = 0$. In that case, we consider the sequence : $\underbrace{1\ 2\ 4 \ldots 2^{k-\alpha}}_{k-\alpha+1\ Rounds}\ \underbrace{p\ p\ p \ldots p}_{\alpha+r\ Rounds}$

- $q \neq 0$. In that case, we consider the sequence :

$$\underbrace{1\ 2\ 4 \ldots 2^{k-\alpha}}_{k-\alpha+1\ Rounds}\ \underbrace{p\ p \ldots p\ \overbrace{(p+1)\ (p+1) \ldots (p+1)}^{q+1\ Rounds}}_{\alpha+r\ Rounds}$$

In each of the cases, the Basic Premise and Step Decreasing Property both hold for the values of n we consider. Hence, the results given above for $q = 0$ and $q \neq 0$ prove the Theorem. □

A summary of the results in the synchronous case is given in Table 1 below. In each row, for the values of r and n which are given respectively in the first and second column, we give the known lower and upper bound on $S_n(r)$. We assume in these tables that $k = \lceil \log_2(n) \rceil$.

$S_n(r)$ - Synchronous case				
s	**n**	**Lower Bound**	**Upper Bound**	**Optimality**
0	$\forall\, n$ in the Bottom Half	$n + \lfloor \frac{n-2^{k-2}}{2} \rfloor$	$n + \lfloor \frac{n-2^{k-2}}{2} \rfloor$	Yes [FP98b]
0	$\forall\, n$ in the Top Half	$2n - 2^{k-1} - 1$	$2n - 2^{k-1} - 1$	Yes [FP98b]
1	$\forall\, n$ in the Bottom Half	Formulae of Theorem 1	$5 \cdot \lceil \frac{n}{4} \rceil + 4$	
1	$\forall\, n$ in the Top Half	Formulae of Theorem 1	$5 \cdot \lceil \frac{n}{4} \rceil + 5$	
1	$n = 2^k - 1$	$5 \cdot 2^{k-2} - 2$	$5 \cdot 2^{k-2} - 2$	Yes
s	$\forall\, n$	Formulae of Theorem 1		
$n - (k+1)$	$\forall\, n$	n	n	Yes

Table 1: Summary of the results for $S_n(r)$

2.2 Relaxing the condition on S_n

Here, we study the function $R_n(s)$ (for any positive integer s), corresponding to the optimal number of rounds for a gossip algorithm in a n-node network which uses $S_n = S_{min} + s$ steps, where $S_{min} = n$. We give lower and upper bounds for $R_n(s)$, which turn out to be tight in several cases (cf. Table 2). The main result is given in Theorem 2 below.

Theorem 2 *For any* $s \geq 1$*, let* $p = \lfloor \log_2(s+1) \rfloor$ *and* $m = n + s - 2^{p+1} + 1$*. Then we have* $R_n(s) \geq (p+1) + \lceil \frac{m}{s+1} \rceil$.

Proof : The Basic Premise must hold for any chosen steps sequence. Here, we are going to give the "best" steps sequence in order to minimize the number of rounds, such that this steps sequence σ satisfies the Basic Premise.

First, for any steps sequence with $R_n(s)$ rounds, we must have $\sum_{i=1, i \neq j}^{R_n(s)} |s_i| \geq n - 1$ for any $1 \leq j \leq R_n(s)$. Since we suppose $S_n = n + s = \sum_{j=1}^{R_n(s)} |s_j|$, we conclude that for any $1 \leq j \leq R_n(s)$, $|s_j| \leq s + 1$. Hence we will choose a steps sequence σ with a maximum number of s_j such that $|s_j| = s+1$. However, we know that for any $1 \leq j \leq R_n(0)$, $|s_j| \leq 2^{j-1}$ (because the number of informed vertices can at most double at each round). Hence, the first $p + 1$ steps in σ will be $1 \quad 2 \quad 4 \ldots 2^p$, where p is such that $2^p \leq s + 1 < 2^{p+1}$, that is $p = \lfloor \log_2(s + 1) \rfloor$.

Then, we have to make sure that $\sum_{j=1}^{R_n(s)} |s_j| = n + s$. Indeed, we need the following sequence σ : $1 \ 2 \ 4 \ldots 2^p \quad x \quad (s + 1) \quad (s + 1) \ldots (s + 1)$, where x is the rest of the division between $m = n + s - (2^{p+1} - 1)$ and $s + 1$, that is $0 \leq x < s + 1$. However, if $x = 0$, then we can save one round by not including it in σ. In both cases (that is $x = 0$ as well as $x > 0$), we see that σ is the shortest sequence, in terms of rounds, that we could get. Now let us distinguish the two cases :

- $x = 0$. In that case, this means that the smallest number of rounds we could get, say R', satisfies : $2^{p+1} - 1 + (R' - (p + 1)) \cdot (s + 1) = n + s$. From this, since $R_n(s) \geq R'$, we conclude $R_n(s) \geq (p + 1) + \frac{m}{s+1}$.
- $x > 0$. In that case, R' satisfies : $2^{p+1} - 1 + x + (R' - (p + 2)) \cdot (s + 1) = n + s$. Since $R_n(s) \geq R'$, we conclude $R_n(s) \geq (p + 1) + \lceil \frac{m}{s+1} \rceil$.

Collectively, the two cases $x = 0$ and $x > 0$ imply the result. □

A summary of the results in the synchronous case is given in Table 2 below. As for Table 1, we assume here that $k = \lceil \log_2(n) \rceil$.

$R_n(s)$ - Synchronous case				
s	**n**	**Lower Bound**	**Upper Bound**	**Optimality**
0	$\forall\ n$	n	n	Yes
1	$\forall\ n$	$\frac{n+3}{2}$		
s	$\forall\ n$	$p + 1 + \lceil \frac{n+s-2^{p+1}+1}{s+1} \rceil$ where $p = \lfloor \log_2(s + 1) \rfloor$ (Theorem 2)		
$2^{k-2} - 1$	$n = 2^k - 1$	$k + 2$	$k + 2$	Yes
$\lfloor \frac{n-2^{k-2}}{2} \rfloor$	$\forall\ n$ in the Bottom Half	$k + 1$	$k + 1$	Yes
$n - 2^{k-1} - 1$	$\forall\ n$ in the Top Half	$k + 1$	$k + 1$	Yes

Table 2: Summary of the results for $R_n(s)$

3 Asynchronous Case

In the asynchronous case, only the function $R_n(s)$ has been studied. The results we have obtained are gathered in Property 1 below.

Property 1 $R_n(0) \leq \begin{cases} 2k-1 & \textit{for all } n = 2^k - 1 \textit{ with } k \geq 2 \textit{ ;} \\ 2k-2 & \textit{for all } n = 3 \cdot 2^{k-2} - 1 \textit{ with } k \geq 3 \textit{ ;} \\ 2k-3 & \textit{for all } n = 5 \cdot 2^{k-3} - 1 \textit{ with } k \geq 4. \end{cases}$

4 Conclusion

This is the first attempt to understand the trade-offs between the number of steps and the number of rounds for a gossip algorithm among n nodes, n being odd, in the linear cost model. Several quantitative results can be derived from this study ; in particular, it is possible to compare the asynchronous model to the synchronous one. Indeed, thanks to our results, we can show that when $S_n = S_{min} = n$, an asynchronous gossip algorithm is asymptotically at least $\frac{n}{2\log_2 n}$ faster, in terms of rounds, than a synchronous one.

References

[Fer] G. Fertin. Trade-offs for odd gossiping. In *Proc. 6th International Colloquium on Structural Information and Communication Complexity (SIROCCO '99)*. To appear.

[FP94] P. Fraigniaud and J.G. Peters. Minimum linear gossip graphs and maximal linear (Δ, k)-gossip graphs. Technical Report CMPT TR 94-06, Simon Fraser University, Burnaby, B.C., 1994. Available at `http://fas.sfu.ca/pub/cs/techreports/1994/`.

[FP98a] G. Fertin and J.G. Peters. Odd gossiping in the linear cost model. *Proc. Workshop on Communication - 23rd Int. Symp. on Mathematical Foundations of Computer Science (MFCS), Brno, Czech Republic*, 1998.

[FP98b] G. Fertin and J.G. Peters. Optimal odd gossiping. Technical Report CMPT TR 1998-24, Simon Fraser University, Burnaby, B.C., December 1998. Submitted for publication.

[Knö75] W. Knödel. New gossips and telephones. *Discrete Mathematics*, 13:95, 1975.

[PRX96] J.G. Peters, L. Raabe, and C. Xu. Odd gossiping. 1996. Manuscript.