Odd Gossiping in the Linear Cost Model

Guillaume Fertin[1] and Joseph G. Peters[2]

[1] LaBRI - U.M.R. CNRS 5800, Université Bordeaux I,
351 Cours de la Libération, F33405 Talence Cedex, France
fertin@labri.u-bordeaux.fr
[2] School of Computing Science, Simon Fraser University,
Burnaby, British Columbia, Canada V5A 1S6
peters@cs.sfu.ca

Abstract. 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 synchronous algorithms that match the lower bounds in [1] for a majority of odd values of n. In *asynchronous* gossiping, a pair of nodes can start communicating while communications between other pairs are in progress. We provide a new short intuitive proof of the asynchronous lower bound in [1].

1 Introduction

Gossiping is an information dissemination problem 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 such as hypercubes and meshes. See [2, 3] 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. Knödel [4] proved that the number of *rounds* of communication necessary to gossip 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 [5, 6]. 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. When messages are long, a *linear cost* model is more realistic since the length of the messages in most gossip algorithms grows exponentially.

In this paper, we assume a *store-and-forward, 1-port, full-duplex* model in which each communication involves two nodes and a single communication link, 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. Messages can be concatenated and sent as a single communication. We assume a *linear cost* model in which the time to send a message of length k is $\beta + k\tau$ 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. 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. A call involving messages of length k can be thought of as a sequence of k *steps* each of which takes time τ.

The constant cost model is *synchronous*; each call takes one time unit, so a gossip algorithm consists of a sequence of *rounds* of simultaneous pairwise calls. Fraigniaud and Peters [7] investigated the structure of minimum-time gossip algorithms using a linear cost model. They established lower and upper bounds on the time to gossip when the number of nodes n is even and showed that minimum-time gossip algorithms for n even are synchronous. They also gave examples to show that the time to gossip can sometimes be reduced when n is odd by allowing pairs of nodes to begin calls at different times.

Peters, Raabe, and Xu [1] studied gossiping with n odd and a linear cost model. They proved a general lower bound of $(\lceil log_2(n) \rceil + 1)\beta + n\tau$ which holds for all 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 $n = 2^k - 1$, they proved that every gossip algorithm with start-up time $(\lceil log_2(n) \rceil + 1)\beta$ must have total propagation time strictly greater than $n\tau$. For the synchronous case, they proved stronger lower bounds and conjectured that their lower bounds are achievable for all odd n. They gave an ad hoc algorithm that achieves the lower bound for $n = 2^k - 1$.

In Sect. 2, we consider synchronous gossiping. Our main result is a method to construct gossip algorithms that achieve the lower bounds in [1]. We show that such algorithms exist for all odd n in the *top half* of any range between two consecutive powers of 2 (i.e., for $3 \times 2^{k-2} + 1 \leq n \leq 2^k - 1$ for any $k \geq 3$), and also for some values in the *bottom half* of each range. Collectively, our results account for 60% of the odd values of n and give support to the conjecture in [1]. In Sect. 3, we will deal with asynchronous gossiping. The proof in [1] that the general lower bound cannot be achieved when $n = 2^k - 1$ is long and complicated. We give a much shorter and more intuitive proof of this result.

2 Synchronous Gossiping

Knödel [4] showed that gossiping in the unit cost model requires $\lceil log_2(n) \rceil + 1$ rounds when n is odd. This lower bound on number of rounds is also valid for the linear cost model in both the synchronous and asynchronous cases. It is also immediate that at least n steps are required 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 [1] proved a lower bound of $(\lceil log_2(n)\rceil+1)\beta+n\tau$ for odd n for 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 focussing on the required number of steps. We take the same approach to upper bounds.

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 and also the required numbers of steps in each of the rounds depends on whether n is in the *bottom half* of the range, $2^{k-1}+1 \le n \le 3\times 2^{k-2}-1$, or the *top half* of the range, $3\times 2^{k-2}+1 \le n \le 2^k-1$.

Theorem 2.1 ([1]). *A synchronous gossip algorithm for odd n in the top half which has $\lceil log_2(n)\rceil+1$ rounds requires at least $2n-2^{k-1}-1$ steps where $k=\lceil log_2(n)\rceil$. The numbers of steps in the rounds are 1 2 4 8 ... 2^{k-2} x x respectively where $x=n-2^{k-1}$.*

Theorem 2.2 ([1]). *A synchronous gossip algorithm for odd n in the bottom half which has $\lceil log_2(n)\rceil+1$ rounds requires at least $2^{k-2}-1+2\lceil\frac{n-2^{k-2}}{2}\rceil+\lfloor\frac{n-2^{k-2}}{2}\rfloor$ steps where $k=\lceil log_2(n)\rceil$. The numbers of steps in the first $k-2$ rounds are 1 2 4 8 ... 2^{k-3} respectively. Two of the last three rounds have $y=\lceil\frac{n-2^{k-2}}{2}\rceil$ steps and the other round has $x=\lfloor\frac{n-2^{k-2}}{2}\rfloor$ steps.*

Conjecture 2.1 ([1]). There are synchronous gossip algorithms that achieve the lower bounds of Theorems 2.1 and 2.2 for every odd n.

We note that there can be a trade-off between the number of rounds and the number of steps in a gossip algorithm. If more than $\lceil log_2(n)\rceil+1$ rounds are permitted, then the number of steps can often be reduced. Depending on the relative values of β and τ, the fastest algorithm could have more than $\lceil log_2(n)\rceil+1$ rounds. We do not investigate this trade-off in this paper.

Our algorithms in this section and the next section are based on a property of partial gossip algorithms and the notion of *experts*. We say that a node is an *expert* of a set S if it knows the information of every node in S.

2.1 The Top Half

The main result in this section is an algorithm that achieves the lower bound in Theorem 2.1. This proves Conjecture 2.1 for every odd n in the top half of any range between two consecutive powers of 2. Our result is the following.

Theorem 2.3. *For any odd n in the top half, there is a synchronous gossip algorithm with $\lceil log_2(n)\rceil+1$ rounds and $2n-2^{k-1}-1$ steps, where $k=\lceil \log_2(n)\rceil$.*

As dictated by Theorem 2.1, our algorithm has $k+1$ rounds, where $k=\lceil log_2(n)\rceil$, and the numbers of steps in the rounds are 1 2 4 8 ... 2^{k-2} x x respectively where $x=n-2^{k-1}$. Our algorithm and its proof of correctness are based on the following property which we will prove to be true for every odd $n\ge 3$.

Property 2.1. Let A_n be the following property:
For any odd n such that $\lceil log_2(n) \rceil = k$, there exists a partial gossip algorithm with k rounds and 2^{i-1} steps in each round i such that after k rounds:

1. 2^{k-1} nodes are experts, and
2. the remaining $n - 2^{k-1}$ nodes know at least 2^{k-1} pieces of information.

In order to prove the above Property, let us prove Propositions 2.1 and 2.2, which, together, will prove recursively the correctness of Property 2.1.

Proposition 2.1. *Let $n = 2^k - p$ be an odd number such that $\lceil log_2(n) \rceil = k$ and $n \geq 3$. If A_n is true, then $A_{n'}$ is true for any $n' = 2^{k'} - p$ with $k' \geq k$.*

Proof. Let $n = 2^k - p$ for any fixed k and odd p. Note that $1 \leq p \leq 2^{k-1} - 1$, since $\lceil log_2(n) \rceil = k$. Also note that a gossip algorithm for n nodes has $k + 1$ rounds [4]. We will prove the proposition by induction on k.

Suppose that A_n is true and let $n_1 = 2^{k+1} - p = 2^k + n$. Partition the set S of n_1 nodes into two subsets S_1 and S_2, such that $|S_1| = 2^k$ and $|S_2| = 2^k - p = n$. Gossiping among n_1 nodes requires $k + 2$ rounds. During the first k rounds of a partial gossip algorithm, the nodes of S_1 and the nodes of S_2 communicate independently. The 2^k nodes of S_1 can all be experts of S_1 after k rounds. The $n = 2^k - p$ nodes of S_2 can satisfy property A_n after k rounds by assumption. Round $k + 1$ of the partial gossip algorithm is illustrated in Fig. 1. The 2^{k-1} experts of S_2 exchange all of their information with 2^{k-1} of the experts of S_1 to create 2^k experts of S. The remaining $2^{k-1} - p$ nodes of S_2 exchange information with $2^{k-1} - p$ nodes of S_1. After round $k+1$, these $2^k - 2p$ nodes and the remaining p nodes of S_1 each know at least 2^k pieces of information. Since round $k+1$ takes 2^k steps, all conditions of Property A_{n_1} are satisfied. We complete the induction proof by observing that A_3 is true. □

Proposition 2.2. *Let $n = 2^k - p$ be an odd number such that $\lceil log_2(n) \rceil = k$ and $n \geq 3$. If A_n is true, then:*
(a) *A_{2n-1} is true, and*
(b) *if $p \neq 1$, then A_{2n+1} is true.*

Proof. (a) Suppose that A_n is true and consider a set S of $2n-1$ nodes. Partition S into three subsets S_1, S_2, and S_3 such that $|S_1| = n$, $|S_2| = 2^{k-1}$, and $|S_3| = 2^{k-1} - (p+1)$. Note that $k+2$ rounds are required to gossip among $2n-1$ nodes. During the first $k - 1$ rounds of a partial gossip algorithm, the nodes of S_1, S_2, and S_3 communicate within their own subsets. Since $|S_1| = n$, the nodes of S_1 can satisfy A_n by assumption. After $k - 1$ rounds, all nodes of S_2 and S_3 can be experts of their respective subsets because the number of nodes in each set is at most 2^{k-1} and is even. The communications during rounds k and $k + 1$ of the partial gossip algorithm are shown in Fig. 2. It is not difficult to verify that S satisfies all conditions of Property A_{2n-1} after round $k + 1$.

(b) The proof that A_n implies A_{2n+1} when $p \neq 1$ is similar to the proof of part (a). Partition the set S of $2n + 1$ nodes into three subsets S_1, S_2, and S_3

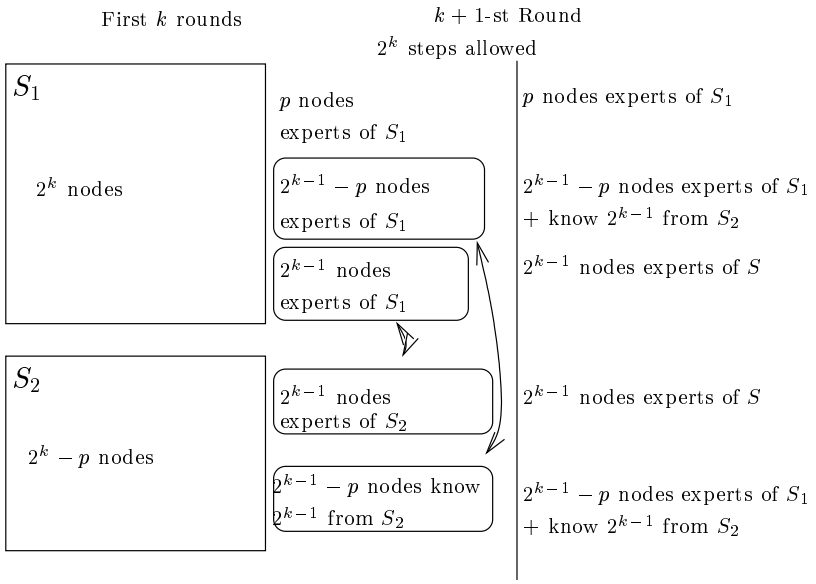

Fig. 1. $A_{n'}$ is true if A_n is true

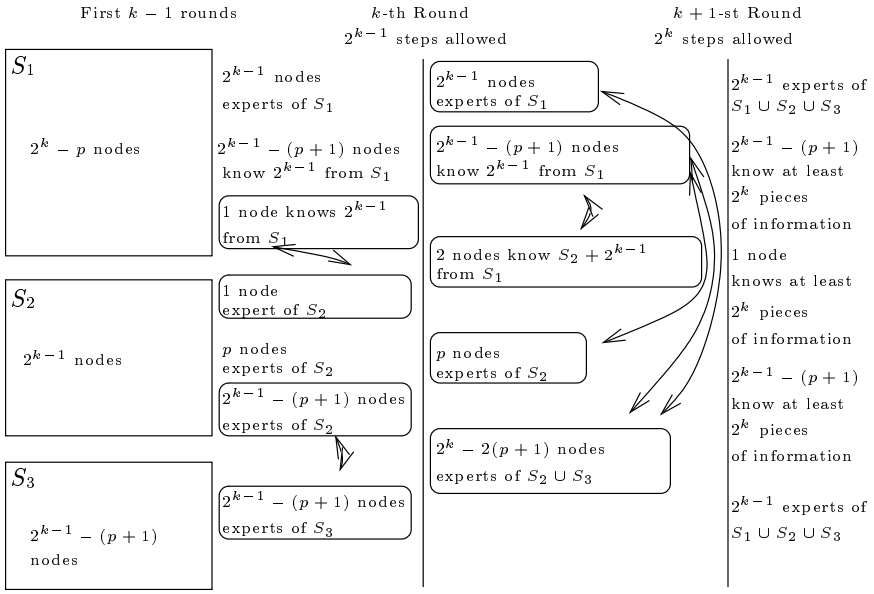

Fig. 2. A_{2n-1} is true if A_n is true

such that $|S_1| = n$, $|S_2| = 2^{k-1}$, and $|S_3| = 2^{k-1} - (p-1)$. Gossiping among $2n+1$ nodes requires $k+2$ rounds. As in part (a), the nodes of S_1, S_2, and S_3 communicate within their own subsets during the first $k-1$ rounds of a partial gossip algorithm. After $k-1$ rounds, the nodes of S_1 can satisfy A_n by assumption, while all nodes of S_2 and S_3 can be experts of their respective sets. The communications during rounds k and $k+1$ of the partial gossip algorithm are shown in Fig. 3 and prove that S satisfies all conditions of Property A_{2n+1} after round $k+1$. The reason for the condition $p \neq 1$ can be seen in the diagram: one node of S_2 must communicate with a node of S_1, so we need $|S_3| < |S_2|$. □

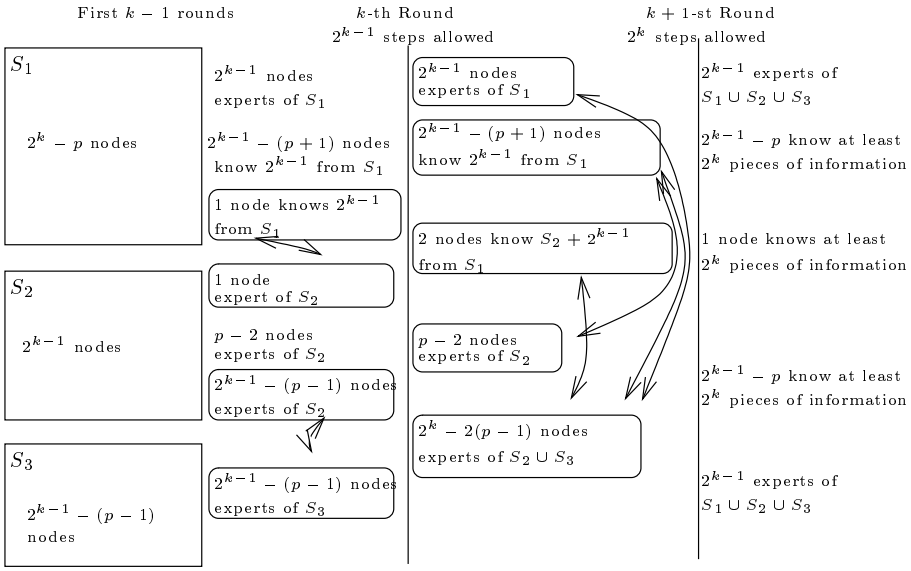

Fig. 3. A_{2n+1} is true if A_n is true (with $p \neq 1$)

Proof of Theorem 2.3. First we establish that Property A_n is true for every odd $n \geq 3$. The case A_3 is true by inspection. Suppose that A_n is true for every odd n with $\lceil \log_2(n) \rceil = k$ for some k. That is, A_n is true for $2^{k-1} < n < 2^k$. Then A_n is true for $n = 2^{k+1} - 1$ by Proposition 2.1, and for every other odd n with $\lceil \log_2(n) \rceil = k+1$ by Proposition 2.2. Hence A_n holds for every odd $n \geq 3$.

Now, to prove Theorem 2.3 we use a gossip algorithm that partitions the set S of n nodes into two subsets S_1 and S_2 with $|S_1| = 2^{k-1}$ and $|S_2| = n - 2^{k-1} = x$. Since n is in the top half, $2^{k-2} + 1 \leq |S_2| \leq 2^{k-1} - 1$. During the first $k-1$ rounds, the nodes of the two subsets S_1 and S_2 communicate within their own subsets. After $k-1$ rounds, all nodes of S_1 can be experts of S_1 and the nodes of S_2 can satisfy Property A_x. The communications during rounds k and $k+1$ are shown in Fig. 4. It is not difficult to verify that all nodes of S will be experts

of S after round $k+1$. Rounds k and $k+1$ each use $x = n - 2^{k-1}$ steps, so the total number of steps for the algorithm is $2^{k-1} - 1 + 2 \times x = 2n - 2^{k-1} - 1$. □

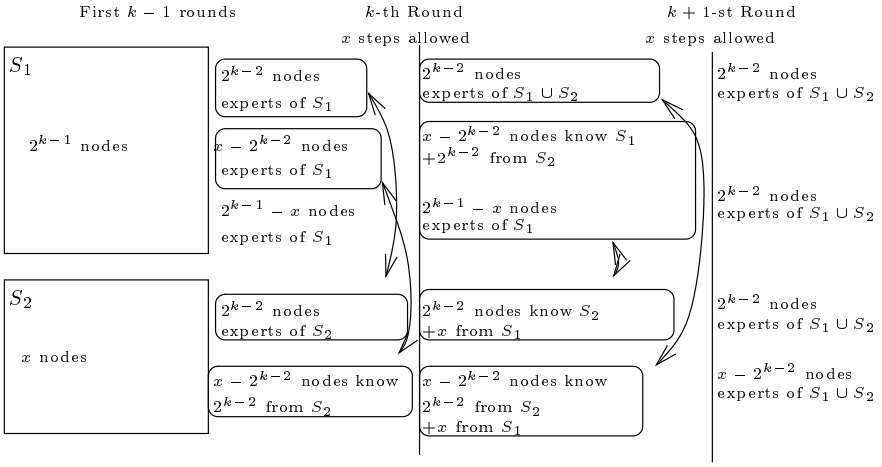

Fig. 4. Gossip algorithm in optimal time in the top half

2.2 The Bottom Half

We did not manage to show the correctness of Conjecture 2.1 for all odd n in the bottom half, that is $2^{k-1}+1 \leq n \leq 3 \times 2^{k-2} - 1$. However, we have shown that the conjecture is true in several cases, which asymptotically account for approximately $\frac{1}{6}$ of the values of odd n in the bottom half. The proofs are rather long but the general ideas are similar to the proofs for the top half; we partition the set of nodes into several subsets, gossip in the subsets independently, and then exchange information among the subsets during the last few rounds using ad hoc methods. We provide here only a summary of our results for the bottom half and refer the reader to a technical report [8] for the details. Table 1 shows the values of odd n for which we know that synchronous gossiping is possible in $\lceil log_2(n) \rceil + 1$ rounds and $2^{k-2} - 1 + 2\lceil \frac{n-2^{k-2}}{2} \rceil + \lfloor \frac{n-2^{k-2}}{2} \rfloor$ steps where $k = \lceil log_2(n) \rceil$. The rightmost column indicates the percentage of values of odd n in each range. We note that these results, together with the results for the top half, confirm the conjecture for all odd $n \leq 39$ and for 60% of the values of odd n asymptotically.

Table 1. Summary of the results for odd n in the bottom half

Values of n	Gossip algorithm that achieves the lower bound?	Percentage of bottom half
$n = 2^{k-1} + 1$	YES	0
$2^{k-1} + 3 \leq n < \lceil \frac{3 \times 2^{k+1}+1}{11} \rceil$	?	18.18
$\lceil \frac{3 \times 2^{k+1}+1}{11} \rceil \leq n \leq \lceil \frac{11 \times 2^{k-2}}{5} \rceil$	YES	1.82
$\lceil \frac{11 \times 2^{k-2}}{5} \rceil < n < \lceil \frac{7 \times 2^{k-2}-1}{3} \rceil$	?	13.33
$\lceil \frac{7 \times 2^{k-2}-1}{3} \rceil \leq n \leq 5 \times 2^{k-3} - 1$	YES	16.66
$5 \times 2^{k-3} + 1 \leq n \leq 3 \times 2^{k-2} - 5$	?	50
$n = 3 \times 2^{k-2} - 3$	YES	0
$n = 3 \times 2^{k-2} - 1$	YES	0

3 Asynchronous Gossiping

3.1 The Equal Exchange Principle

Any asynchronous gossip algorithm for n nodes with n odd takes time at least $(\lceil log_2(n) \rceil + 1)\beta + n\tau$ [1]. We can derive several properties of gossip algorithms that take time exactly $(\lceil log_2(n) \rceil + 1)\beta + n\tau$.

Property 3.1. No node can be idle during more than one step.

Proof. Suppose some node u is idle during two or more distinct steps. Since u needs $n-1$ steps to acquire the information of the other nodes, the total number of steps will be greater than n. □

Property 3.2. Each node must be idle during at least one step.

Proof. Suppose that some node u is never idle. Since n is odd, there must be at least one idle node during each of the n steps. This means that some other node v must be idle during at least two distinct steps, which contradicts Property 3.1. □

Property 3.3. Two nodes cannot be idle during the same step.

Proof. Each node must be idle during at least one step by Property 3.2 and each node needs $n-1$ steps to acquire the information of the other nodes. Summing over all nodes gives a total requirement of n^2 communication units. Since n is odd, there is at least one idle node during each step. If two nodes are idle during the same step, then the total number of units is at least $n^2 + 1$ and this is not possible in an algorithm with n steps. □

Property 3.4. Each node is idle during exactly one step, and all the idle steps are distinct.

Proof. This follows directly from the other three properties. □

Based on these properties, we get a short proof of the *Equal Exchange Principle* first proved in [1].

Theorem 3.1 (Equal Exchange Principle [1]). *Two nodes exchange the same amount of information when they communicate.*

Proof. Suppose two nodes u and v send different amounts of information to each other during a communication. Then one of these nodes, say u is idle (i.e., not receiving information) during at least one step s while v is busy receiving information from u. Since the number of nodes is odd, and since communications occur between pairs of nodes, there must be another node w which is idle during the same step s. This contradicts Property 3.3. □

3.2 The Case $n = 2^k - 1$

The following theorem shows that the lower bound $(\lceil log_2(n) \rceil + 1)\beta + n\tau$ cannot be achieved by any gossip algorithm when $n = 2^k - 1$. A different proof of this result is given in [1]. The proof that we present here is much shorter and more intuitive.

Theorem 3.2. *Any gossip algorithm for $n = 2^k - 1$ nodes, $k \geq 3$, takes time strictly greater than $(\lceil log_2(n) \rceil + 1)\beta + n\tau$, for all $\beta > 0$, $\tau > 0$.*

Proof. A gossip algorithm can be represented as an $n \times n$ grid. Each row represents a node and each column represents a step of the algorithm. By Property 3.4, each node must be idle during exactly one step, and these idle steps must be distinct. Without loss of generality, we can arrange the idle steps along a diagonal. Figure 5 shows the case $n = 15$ with the idle steps shown in dark gray.

Consider the last two nodes, $u = 2^k - 2$ and $v = 2^k - 1$ (nodes 14 and 15 in Fig. 5). Node v is idle during the first step while all other nodes are busy. When it starts its communication in step 2, it will inherit a delay of $\beta + \tau$ from the node with which it is communicating. So, v can only have k active rounds and the numbers of steps in its rounds must be $1, 2, 4, ..., 2^{k-2}, 2^{k-1} - 1$ as shown in Fig. 5. Any increase in the number of steps in one of the first $k - 1$ rounds would violate the equal exchange principle and any decrease would prevent v from acquiring enough information for its last round. The pattern for node u is the same as for node v by a similar argument.

Next, consider the $2^{k-1} - 2$ nodes labelled 2 to $2^{k-1} - 1$ (nodes 2 to 7 in Fig. 5). None of these nodes is idle before step $2^{k-1} + 1$. Since the amount of information exchanged during each round i cannot be greater than 2^{i-1}, none of these nodes can start round k later than step 2^{k-1}. Thus, these nodes have at most one round after their idle steps and each node i, $2 \leq i \leq 2^{k-1}$ must exchange exactly $i - 1$ pieces of information with another node during its last round. These communications are indicated by light gray rectangles in Fig. 5. The only available nodes for these exchanges are the $2^{k-1} - 2$ nodes labelled 2^{k-1} to $2^k - 3$ (nodes 8 to 13). Therefore, exactly two of the nodes that are active during the last step must exchange i pieces of information during their

last rounds for each $i = 1, 2, ..., 2^{k-1} - 1$. We can show that this is impossible by examining node $w = 2^{k-1}$ (node 8). If node w has k rounds before its idle step, then it must exchange $2^{k-1} - 1$ pieces of information during a single round after its idle step and this gives three nodes (u, v, and w) exchanging $2^{k-1} - 1$ pieces of information. If node w has $k - 1$ rounds before its idle step, then it can have rounds k and $k + 1$ after its idle step. During round k, node w cannot communicate with any of the nodes 1 through $2^{k-1} - 1$ (nodes 1 to 7) because they must all start round k no later than step 2^{k-1}. Node w cannot communicate with any of nodes $2^{k-1} + 1$ to $2^k - 3$ (nodes 9 to 13) because this would leave three nodes with the same amount of information to exchange during their last rounds (node w, the node with which w communicated in round k, and one of nodes 2 to $2^{k-1} - 1$). □

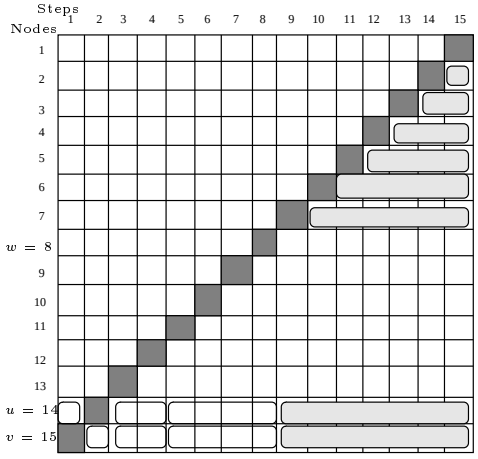

Fig. 5. Asynchronous gossiping with 15 nodes

4 Conclusion

We have shown that synchronous gossiping can be completed in time that matches the lower bounds for approximately 60% of all odd values of n. This provides evidence that the conjecture in [1] is true, but leaves the conjecture open for most of the *bottom halves* of the ranges between consecutive powers of 2. We have also given a simple new method to prove the lower bound on asynchronous gossiping for $n = 2^k - 1$. The extension of this method to other values of n remains open. Finally, the trade-offs between the number of rounds and the number of steps for both synchronous and asynchronous gossiping remain unexplored.

References

1. Peters, J.G. Raabe, L., and Xu, C. Odd Gossiping. Manuscript, 1996.
2. Hedetniemi, S.M., Hedetniemi, S.T., and Liestman, A.L. A survey of gossiping and broadcasting in communication networks. *Networks* 18(1988), 319-349.
3. Fraigniaud, P. and Lazard, E. Methods and problems of communication in usual networks. *Discrete Applied Math.* 53(1994), 79-133.
4. Knödel, W. New Gossips and Telephones. *Discrete Mathematics* **13** (1975) 95.
5. Even, S. and Monien, B. On the number of rounds necessary to disseminate information, in *Proc. 1st ACM Symp. on Parallel Algorithms and Architectures*, Santa Fe, NM, June 1989, 318–327.
6. Krumme, D.W., Cybenko, G., and Venkataraman, K.N. Gossiping in minimal time. *SIAM J. Comput.* **21** (1992) 111–139.
7. Fraigniaud, P. and Peters, J.G. Minimum Linear Gossip Graphs and Maximal Linear (Δ, k)-Gossip Graphs. Technical Report 94-06, School of Computing Science, Simon Fraser Univ., 1994.
8. Fertin, G. and Peters, J.G. Odd Gossiping in the Linear Cost Model. Technical Report, School of Computing Science, Simon Fraser Univ., 1998.