Compounding of Gossip Graphs

Guillaume Fertin
LaBRI-U.M.R. CNRS 5800, Université Bordeaux I, 351 Cours de la Libération, F-33405 Talence Cedex, France

Roger Labahn
Universität Rostock, FB Mathematik, D-18051 Rostock, Germany

Gossiping refers to the following task: In a group of individuals connected by a communication network, every node has a piece of information and needs to transmit it to all the nodes in the network. The networks are modeled by graphs, where the vertices represent the nodes, and the edges, the communication links. In this paper, we concentrate on *minimum gossip graphs* of even order, that is, graphs able to achieve gossiping in minimum time and with a minimum number of links. More precisely, we derive upper bounds for their number of edges from a compounding method, the *k-way split method*, previously introduced for broadcasting by Farley [Networks 9 (1979), 313–332]. We show that this method can be applied to gossiping in some cases and that this generalizes some compounding methods for gossip graphs given in [5]. We also show that, when applicable, this method gives the best-known upper bounds on the size of minimum gossip graphs in most cases, either improving or matching them. Notably, we present for the first time two families of regular gossip graphs of order n and of degree $\lceil \log_2(n) \rceil - 3$ and $\lceil \log_2(n) \rceil - 4$, respectively. We also give some lower bounds on the number of edges of gossip graphs which improve the ones given by Fertin [5]. Moreover, we show that the above compounding method also applies for *minimum linear gossip graphs* (or MLGGs) of even order, which corresponds to a variant of gossiping where the time of information transmission between two nodes depends on the amount of information exchanged. We also prove that this gives the best-known upper bounds for $G_{\beta,\tau}(n)$—the size of an MLGG of order n—in most cases. In particular, we derive from this method the exact value of $G_{\beta,\tau}(72)$, which was previously unknown. © 2000 John Wiley & Sons, Inc.

Keywords: gossiping; graphs; networks; communication; compounding method

1. INTRODUCTION

Gossiping is a task of information dissemination in a group of individuals connected by a communication network. In gossiping, every node knows a piece of information and needs to transmit it to everyone else. This is achieved by placing communication calls over the communication lines of the network. Throughout this paper, we will consider a *1-port*, *full-duplex* model, that is:

- A node can communicate with at most one of its neighbors at any given time, and
- When communication takes place between two nodes, the information flows in both directions.

Depending on the cases, we will consider this model to be either *unit cost* or *linear cost*. In the former, a communication between two nodes takes one time unit, while in the latter, the communication time implies a fixed start-up time β and a propagation time τ proportional to the amount of information exchanged. Note that in this case we suppose that every node holds a unique piece of information that cannot be split and that all pieces have the same length 1. Moreover, we suppose that when two nodes communicate each node can send a *message* consisting of one or more pieces of information to the other node.

In both cases (i.e., *unit cost* and *linear cost*), networks will be modeled by undirected graphs, without loops or multiple edges. The vertices will represent the nodes of the network, and the edges, the communication links.

Most of the recent interest in gossiping is due to its importance in the area of network communications and other areas of parallel and distributed computing. A way to study tasks such as gossiping is to find interconnection networks with the minimum resources necessary to gossip in minimum time. This approach is the one we are dealing with in the following:

Knödel [7] proved that the time g_n to gossip in the complete graph K_n of order n under the *unit cost* model is

- $\lceil \log_2(n) \rceil$ for even n, and
- $\lceil \log_2(n) \rceil + 1$ for odd n.

Received January 1999; accepted April 2000
Correspondence to: G. Fertin; e-mail: fertin@labri.u-bordeaux.fr

A *gossip graph* will then denote a graph able to gossip in minimum time. However, it is not necessary to consider the complete graph to get a gossip graph. Hence, we denote by *minimum gossip graph*, or *MGG*, any gossip graph with a minimum number of edges. For any MGG of order n, or MGG_n, this number is denoted by $G(n)$.

Similarly, in the *linear-cost* model, we denote by $g_{\beta,\tau}(n)$ the minimum time to gossip in the complete graph K_n. When n is even, Fraigniaud and Peters [6] proved that $g_{\beta,\tau}(n) = \lceil \log_2(n) \rceil \beta + (n-1)\tau$. A *linear gossip graph* will denote a graph able to gossip in minimum time, while a *minimum linear gossip graph*, or *MLGG*, is a linear gossip graph with a minimum number of edges. For any MLGG of order n, or MLGG_n, this number is denoted by $G_{\beta,\tau}(n)$.

Very few values of $G(n)$ and $G_{\beta,\tau}(n)$ are known in the general case. $G(n)$ is determined for $n = 2^p, n = 2^p - 2$, and $n = 2^p - 4$ [8], while $G_{\beta,\tau}(n)$ is determined for the same values of n, and also $n = 2^p - 6$ [6]. The only known specific values of $G(n)$ are for $1 \le n \le 16$ (except for $n = 13$) and $n = 24$ [5], while the only known specific values of $G_{\beta,\tau}(n)$ are for even n with $2 \le n \le 32$ (except for $n = 22$), $n = 42, n = 44$, and $n = 48$ [6].

Determining precisely the values of $G(n)$ and $G_{\beta,\tau}(n)$ is known to be a hard problem. In this paper, we focus on a general way to get upper bounds for $G(n)$ and $G_{\beta,\tau}(n)$ for even n (which can be extended to odd n in the unit cost model, thanks to some techniques given in [5]). This can be done either by exhibiting some families of graphs which are known to be (linear) gossip graphs for any even n, like Knödel graphs (cf. [5]), or by constructing (linear) gossip graphs from existing (minimum) (linear) gossip graphs. We will concentrate mostly on the second method in this paper.

Section 2 will first focus on the *unit-cost* model. In Section 2.1.1, we will present a general compounding method, very close in its spirit to the *k-way split method*. The *k-way split method* was first introduced by Farley [4] to get upper bounds on the size of minimum broadcast graphs. It was extended and improved in [2], and, more recently, some other compounding methods have been developed [1, 3, 10, 11]. Surprisingly, no comparable study has been undertaken concerning gossiping. In this paper, we will show a compounding method which applies for gossiping. We will also derive from this method some variants which apply for $k = 3, k = 5, k = 9, k = 10$, and $k = 12$. These results prove for the first time the existence of two families of gossip graphs of order n, each of these two families containing an infinite number of graphs and each graph being regular (of degree $\lceil \log_2(n) \rceil - 3$ and $\lceil \log_2(n) \rceil - 4$, respectively). A summary of these general results in the *unit-cost model*, for even $18 \le n \le 128$, is given in Section 2.1.2. However, the results given in Section 2.1 only give upper bounds for the size of an MGG. This is why, in Section 2.2, we will focus on lower bounds for $G(n)$, both in the even and odd cases. These bounds improve, when applicable, the lower bounds given in [5].

Next, we will focus on the *linear-cost* model, where in Section 3.1 we first give a general upper bound for $G_{\beta,\tau}(n)$ which applies for all even n. We also show that the compounding method from the *unit-cost* model also applies. Notably, we show that it allows us to determine $G_{\beta,\tau}(72)$. Moreover, the family of gossip graphs, regular of degree $\lceil \log_2(n) \rceil - 3$, which is given in Section 2 in the *unit-cost* model, also appears to be a family of linear gossip graphs. Section 3.2 finally gives a summary of these general results in the *linear-cost* model for even $18 \le n \le 128$.

2. THE UNIT-COST MODEL

In this section, we present a compounding method to get upper bounds on $G(n)$, then some improvements on the lower bounds for $G(n)$. First, we give a general upper bound, corresponding to a general compounding method. Next, we give some particular methods which are variants of the general one. Finally, Section 2.2 will be devoted to lower bounds on $G(n)$.

2.1. Upper Bounds for $G(n)$

2.1.1. Compounding of Gossip Graphs. The method of compounding graphs has been extensively, and is still, used for determining upper bounds on the size of minimum broadcast graphs [1–4, 10, 11]. However, compounding has never been studied in terms of gossiping. Fertin [5] gave some specific compounding methods to get upper bounds for $G(n)$. These have been used as a starting point for our work, the idea then being to find a generalization of the methods exposed in [5]. We soon realized that the underlying idea was none other than the one given by Farley [4] concerning minimum broadcast graphs. However, some parts of the method do not apply for gossiping. Conversely, we can sometimes split our graph into nonequal parts, something which gives, when applicable, even better results than does the general method.

Before introducing the method itself, we need to give the following definition:

Definition 1 (Compoundable Graph). *A* compoundable *gossip graph G of order n is a gossip graph such that there exists a gossip scheme S_G for G having the following property*: *There exists a perfect matching PM_{S_G} with respect to the gossip scheme*, *such that all the edges of PM_{S_G} are used during the same fixed round r and during no other round $r' \neq r$.*

Theorem 1 (Compounding in the Unit-cost Model). *For all k and even n such that there exists a compoundable gossip graph of order $2k$ and of size $G'(2k)$, and such that $\lceil \log_2(nk) \rceil = \lceil \log_2(n) \rceil + \lceil \log_2(k) \rceil$, we* *have*

$$G(nk) \leq k \cdot G(n) + \frac{n}{2} \cdot (G'(2k) - k).$$

Proof. Suppose that we have a compoundable gossip graph G_{2k} of order $2k$ and of size $G'(2k)$. By definition, we know that there exists in G_{2k} a gossip scheme $S_{G_{2k}}$, for which we can find a perfect matching $PM_{S_{G_{2k}}}$ such that every edge $e_i = u_i v_i$ of $PM_{S_{G_{2k}}}$ $(1 \leq e_i \leq k)$ is used exactly once and during the same round r.

Now, let us construct, from G_{2k}, a graph $\mathcal{G}_{nk}$ of order nk, and let us show that this is a gossip graph. The construction is as follows: Replace in G_{2k} every edge $e_i = u_i v_i$ of $PM_{S_{G_{2k}}}$ by a copy G_i of a MGG$_n$, $1 \leq i \leq k$. In each G_i, partition the set V_i of vertices into two subsets $V(u_i)$ and $V(v_i)$, each of cardinality $n/2$ (this is feasible, since n is even). In the following, we make a correspondence between vertex u_i (respectively, v_i) of G_{2k} and the vertex set $V(u_i)$ [respectively, $V(v_i)$]. Indeed, to end the construction, it suffices, for each edge $e' = wx \notin PM_{S_{G_{2k}}}$, to join the two vertex sets $V(w)$ and $V(x)$ by a perfect matching. For a better understanding of the method, we refer to Figure 1, which shows how to construct $\mathcal{G}_{18}$ from an MGG$_6$ (i.e., $k = 3$ and $n = 6$).

Hence, we end up with a graph $\mathcal{G}_{nk}$ of order nk and of size $kG(n) + (n/2) \cdot (G'(2k - k)$ edges. Now, let us prove that $\mathcal{G}_{nk}$ is a gossip graph. For this, we use the following gossip scheme:

1. From round 1 to round $r - 1$, use the scheme $S_{G_{2k}}$;
2. From round r to round $g_n + r - 1$ (i.e., during the g_n following rounds), gossip within each copy G_i of an MGG$_n$, $1 \leq i \leq k$. In other words, we gossip independently in each G_i, with a scheme achieving gossiping in G_i, but with a delay of $r - 1$ rounds;
3. From round $g_n + r$ to round g_{kn}, we use again the scheme $S_{G_{2k}}$, from round $r + 1$, but with a delay of g_n rounds.

We refer to Figure 2 for an example of the method, where $k = 3, n = 6$ and $r = 2$.

We recall that r is the unique round during which the edges of $PM_{S_{G_{2k}}}$ are used in the scheme $S_{G_{2k}}$. Hence, we necessarily have $1 \leq r \leq g_{2k}$. Moreover, we supposed that $\lceil \log_2(nk) \rceil = \lceil \log_2(n) \rceil + \lceil \log_2(k) \rceil$; since n is even by hypothesis, this means that $g_{nk} = g_n + g_{2k} - 1$.

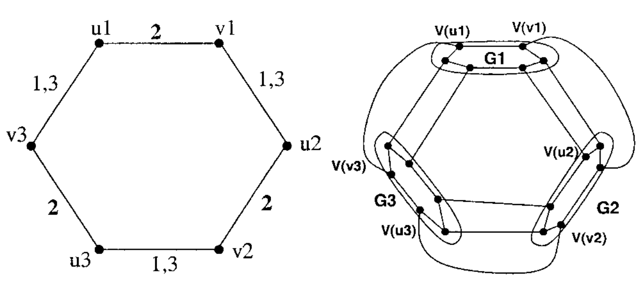

FIG. 1. (a) An MGG$_6$ and a gossip scheme; (b) construction of $\mathcal{G}_{18}$ from (a).

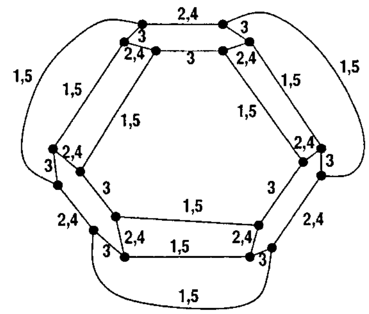

FIG. 2. Gossip scheme in $\mathcal{G}_{18}$.

Let us now prove that, following this scheme, it is possible to achieve gossiping in $\mathcal{G}_{nk}$. First, we note that this is a valid gossip scheme, thanks to the "compoundability" of the gossip graph of order $2k$ and thanks to the constraint $g_{nk} = g_n + g_{2k} - 1$. The former allows us to see each of the MGG$_n$'s as a black box (where the unique round r used to communicate along the edge $u_i v_i$ now takes g_n rounds, after which all the vertices of the MGG$_n$ are informed), while the latter shows that gossiping takes place in minimum time. Moreover, this scheme respects the 1-port model (i.e., no two incident edges communicate during the same round), since $n > 1$.

Following the gossip scheme above, one can see that each vertex of $\mathcal{G}_{nk}$ is able to broadcast its own information to all the other vertices in the graph in minimum time. Moreover, since the gossip scheme is valid, the broadcast of each vertex can be done in parallel, which shows that gossiping is achieved in minimum time. ∎

For some particular cases, the structure of the MGG$_{2k}$ on which we build our compounding is of extreme importance. Indeed, in some cases it is not necessary to take k copies of an MGG$_n$: We can use MGGs of different orders. Propositions 1–5 are based on this particular method.

Proposition 1 (3-way split method). *For all even n_1, n_2, and n_3 such that $\lceil \log_2(n_1 + n_2 + n_3) \rceil = 2 + \lceil \log_2(n_j) \rceil$ $\forall j \in \{1, 2, 3\}$,*

$$G(n_1 + n_2 + n_3) \leq G(n_1) + G(n_2) + G(n_3) + \frac{1}{2} \cdot (n_1 + n_2 + n_3).$$

Proof. Note that if we consider $n_1 = n_2 = n_3 = n$ we get the formula of Theorem 1 in the case $k = 3$. Note also that this proposition is a generalization of a compounding method given in [5], where we had $n_1 = n_3$. Here, we show that we can take three MGGs of distinct orders. In that case, $k = 3$, that is, the MGG$_{2k}$ is the cycle C_6 of order 6. Hence, the perfect matching we will use is necessarily the one where the edges are used at round 2. However, if instead of taking three copies of a MGG$_n$ we take an MGG$_{n_1}$, and MGG$_{n_2}$, and an MGG$_{n_3}$, we show that we still can get a gossip graph. Suppose, w.l.o.g., that $n_1 \geq n_2 \geq n_3$. Let $\alpha = (n_1 + n_2 - n_3)/2$; α

is a strictly positive integer since every n_i is even and since $n_1 \geq n_2 \geq n_3 > 0$. The idea here is to match α vertices among the n_1 of the MGG_{n_1} with as many in the MGG_{n_2}, as shown in Figure 3. Note that this is possible since $0 < \alpha < n_2 \leq n_1$: Indeed, if $\alpha \geq n_2$, that would mean $n_1 \geq n_2 + n_3 \geq 2n_3$, which would imply that $\lceil \log_2(n_1) \rceil > \lceil \log_2(n_3) \rceil$ and would thus violate the condition $\lceil \log_2(n_1 + n_2 + n_3) \rceil = 2 + \lceil \log_2(n_j) \rceil$ $\forall j \in \{1, 2, 3\}$.

Now there remain $n_1 - \alpha$ vertices from the MGG_{n_1} to match with as many in the MGG_{n_3}. Hence, we must have $n_3 - n_1 + \alpha = n_2 - \alpha$, which is true by definition of α, and $n_1 - \alpha < n_3$. But if we suppose that $n_1 - \alpha \geq n_3$, we would have $n_1 \geq n_2 + n_3 \geq 2n_3$, and we have seen that this cannot occur.

Since the copies of MGG_{n_i} do behave as black boxes as far as gossiping is concerned, and since we have $g_{n_1+n_2+n_3} = g_{n_i} + 2$ $\forall i \in \{1, 2, 3\}$, we still have the property that the whole graph constructed this way is a gossip graph on $n_1 + n_2 + n_3$ vertices. ▪

Proposition 2 (5-way split method).

- *For all even n_1 and n_2 such that* $\lceil \log_2(4n_1 + n_2) \rceil = 3 + \lceil \log_2(n_j) \rceil$ $\forall j \in \{1, 2\}$,

$$G(4n_1 + n_2) \leq 4G(n_1) + G(n_2) + 2n_1 + 2n_2.$$

- *For all even n such that* $\lceil \log_2(5n) \rceil = 3 + \lceil \log_2(n) \rceil$,

$$G(5n) \leq 3G(n) + 4G\left(\frac{n}{2}\right) + 5n.$$

Proof. The first formula of the proposition derives from a similar argument as in Proposition 1. This is done using the MGG_{10} shown in Figure 4(left). Here, we want to replace each edge of the perfect matching corresponding to round 2 by a copy G_i of an MGG_{n_i}, where the n_i $(1 \leq i \leq 5)$ may be pairwise distinct. This is shown in Figure 4(right). In that case, if we suppose that α vertices of G_1 are matched with as many in G_5, we necessarily get the right figure of Figure 4, with $\beta = n_1 - \alpha$. Standard calculations show that we necessarily get the following equalities:

- $n_1 = n_2 = n_4 = n_5$ and
- $2\alpha = n_3$.

Note also that if $n_1 = n_2 = n$, we get the formula of Theorem 1.

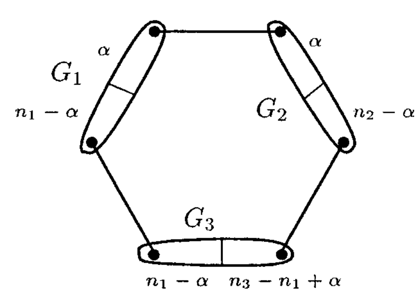

FIG. 3. Compounding using C_6.

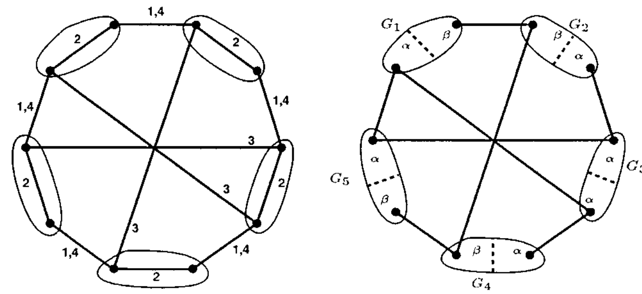

FIG. 4. (Left) An MGG_{10} and (right) a compounding method using it.

The second formula of the proposition relies on a slightly different idea: Suppose that some edges are used during a unique round r, but these edges do not form a perfect matching. In that case, we can still use the same compounding method, but one has to see the "isolated" vertices (i.e., vertices which do not communicate at round r) as a contracted edge. Hence, we replace any isolated vertex by a copy of an $\mathrm{MGG}_{n/2}$, while the edges used in round r are replaced, as before, by a copy of an MGG_n. Then, each edge between an isolated vertex and a nonisolated one will be replaced by a perfect matching between two sets of $n/2$ vertices and an edge between two nonisolated vertices will be replaced, as previously, by a perfect matching between two sets of $n/2$ vertices as well.

In the case where we start from an MGG_{10}, we apply this method with the round $r = 3$, and the method is then shown in Figure 5. The same arguments as in the proof of Theorem 1 show that the graph constructed that way is a gossip graph. Note, though, that we need to have $g_{n/2} \leq g_n$, in order to be able to gossip in any $\mathrm{MGG}_{n/2}$ at worse in the same time as in any copy of an MGG_n. However, we know by definition of g_n that $g_{n/2} \leq g_n$ $\forall n$, even when $n/2$ is odd. ▪

Using the same construction as in the second part of Proposition 2, and relying on the structure and gossip scheme of the gossip graph of order 18 displayed in Figure 6, we can show the following proposition:

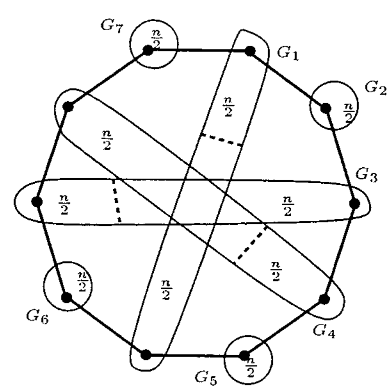

FIG. 5. Another compounding using an MGG_{10}.

Proposition 3 (9-way split method). *For all even n such that $\lceil \log_2(9n) \rceil = 4 + \lceil \log_2(n) \rceil$,*

$$G(9n) \le 7G(n) + 4G\left(\frac{n}{2}\right) + 9n.$$

Proof. The proof relies exactly on the same construction and arguments as in the second part of Proposition 2. The method is displayed in Figure 6(left and right). Note that, as previously, we need $g_{n/2} \le g_n$; but this is always the case, even when $n/2$ is odd. ∎

Proposition 4 (10-way split method). *For all even n_1, n_2, and n_3, let $\alpha = \frac{1}{2} \cdot (n_1 + n_2 - n_3)$, and $\beta = \frac{1}{2} \cdot (n_1 - n_2 + n_3)$. If*

- $\lceil \log_2(6n_1 + 2n_2 + 2n_3) \rceil = 4 + \lceil \log_2(n_j) \rceil \ \forall j \in \{1, 2, 3\}$;
- $g_\alpha \le g_{n_j} \ \forall j \in \{1, 2, 3\}$;
- $g_\beta \le g_{n_j} \ \forall j \in \{1, 2, 3\}$,

$$G(6n_1 + 2n_2 + 2n_3) \le 4G(n_1) + 2G(n_2) + 2G(n_3) + 2G(\alpha) + 2G(\beta) + 6n_1 + 2n_2 + 2n_3.$$

Proof. The method here is similar to the ones given above. Its particularity is that it mixes the two variants of the general method, namely, the matching that we use here is not a perfect matching, and we decide to assign to each isolated vertex (respectively, each edge of the matching) a copy of a MGG_{n_i}, where the n_i may differ.

Let us take the gossip graph of order 20 and a gossip scheme shown in Figure 7 (note that this graph has been shown to be a gossip graph by [9]). Let us then use the matching given by the edges used in round 3. Note that this is not a perfect matching. Replace each of these edges by an MGG and each of the isolated vertex by a copy of an MGG, where the order of these MGGs may differ. Now we refer to Figure 8 and state the following: Suppose that we decide to match α vertices of G_1 with as many in G_2. In that case, it is easy to see that each matching in the "upper leftmost" cycle will be of size α. Let $\beta = n_1 - \alpha$. Then, each matching of the "upper rightmost" cycle will be of size β. The same occurs for the "lower" cycle, where each matching will be of size γ. Thanks to the constraints on the size of these matchings, we obtain equalities between α (respectively, β, γ) and the n_is, where $n_i = |G_i|$ for $1 \le i \le 3$. Standard calculations then give us the result. ∎

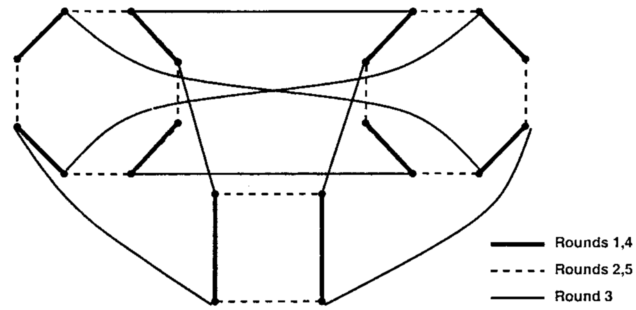

FIG. 7. A gossip graph of order 20 and a gossip scheme.

Proposition 5 (12-way split method). *For all even n_1, n_2, and n_3 such that $\lceil \log_2(4n_1 + 4n_2 + 4n_3) \rceil = 4 + \lceil \log_2(n_j) \rceil \ \forall j \in \{1, 2, 3\}$,*

$$G(4n_1 + 4n_2 + 4n_3) \le 4 \cdot (G(n_1) + G(n_2) + G(n_3)) + 4 \cdot (n_1 + n_2 + n_3).$$

Proof. This relies on the same arguments as, for instance, the proof of Proposition 1. Here, the perfect matching that we use is the one given by edges used in round 3 [cf. for this Fig. 9(left), which shows an MGG_{24} and a gossip scheme]. In that case, if we replace each edge of the perfect matching by an MGG_{n_i}, G_i, where the n_i may differ, and if we partition the set of vertices of G_1 and G_2 as shown in Figure 9(right), we get the following equalities, where $n_i = |G_i|$ for all $1 \le i \le 12$:

- $n_1 = n_6 = n_8 = n_{12} = \alpha + \beta$;
- $n_2 = n_5 = n_9 = n_{11} = \beta + \gamma$;
- $n_3 = n_4 = n_7 = n_{10} = \alpha + \gamma$;

This leads directly to the result. ∎

Thanks to Theorem 1 and Propositions 1–5, we prove for the first time that there exists gossip graphs with $(n/2) \cdot (\lceil \log_2(n) \rceil - 3)$ edges, but also gossip graphs with $(n/2) \cdot (\lceil \log_2(n) \rceil - 4)$ edges, for infinitely many values of n. Indeed, we have the following propositions:

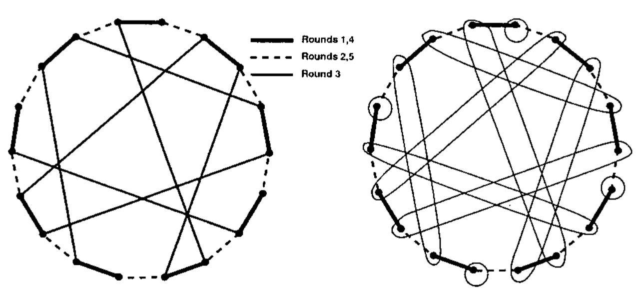

FIG. 6. (Left) A *gossip graph* of order 18 and (right) the compounding method.

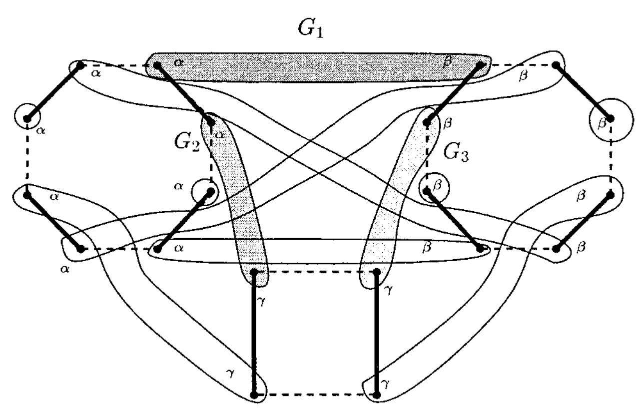

FIG. 8. Compounding using the gossip graph of Figure 7.

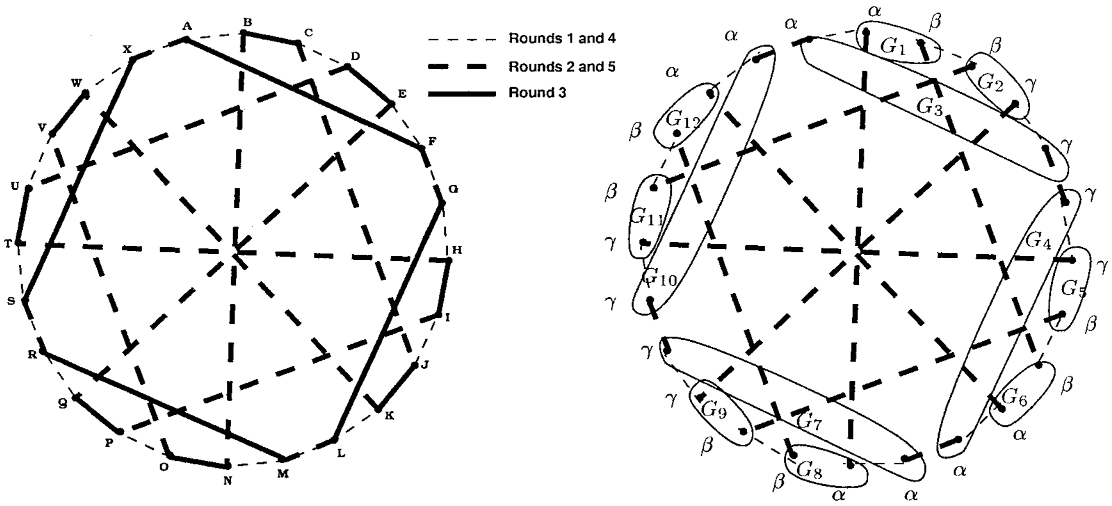

FIG. 9. (Left) An MGG_{24} and gossip scheme and (right) the compounding method.

Proposition 6. *For all $p \geq 7$ and $n' = 24 \cdot (2^{p-5} - 1)$,*

$$G(n') \leq \frac{n'(p-3)}{2}.$$

Proof. Suppose that $n' = 24 \cdot (2^{p-5} - 1)$ with $p \geq 7$. Then, $g_{n'} = p$. In that case, let us apply Theorem 1 where $k = 2^{p-5} - 1$ and $n = 24$. We know this is possible since there exist compoundable gossip graphs of order $2k$ (cf., for instance, [8]). We then have $G(n') \leq k \cdot G(24) + [G(2k) - k]/2 \cdot 24$. Since we know that $G(2k) = G(2^{p-4} - 2) = (p-5) \cdot k$ and $G(24) = 36$, we have $G(n') \leq 36 \cdot k + [(k(p-6)/2] \cdot 24$, that is, $G(n') \leq [24k(p-3)/2]$, where $n' = 24k$. Hence, the result. ∎

It is interesting to note also that there exists an infinity of $(p-4)$-regular gossip graphs. For this, take $n' = 576 = 24 \cdot 24$. In that case, let us apply Theorem 1 where $n = k = 24$. For this, we need to make sure that there exists a compoundable gossip graph of order 48: We then use the 2-way split method to obtain a gossip graph of order 48 such that the perfect matching between the two copies of an MGG_{24} are used in a single round (the first round, for instance). Hence, we get a compoundable gossip graph with $2k = 48$ vertices and 96 edges. Standard calculations then give $G(n') \leq 3n'$, that is, a 6-regular graph where $g_{n'} = 10$. Starting from this graph, and using the compounding method where $k = 2$, we have the following proposition:

Proposition 7. *For all $p \geq 10$ and $n' = 576 \cdot 2^{p-10}$,*

$$G(n') \leq \frac{n'(p-4)}{2}.$$

2.1.2. Summary of the Upper-bounds Results (Unit Cost) Table 1 presents the results given by the k-way split method for even n with $18 \leq n \leq 128$. Note that for the values $n = 2^p, n = 2^p - 2$, and $n = 2^p - 4$ we know by [8] that the result is optimal. For $n = 2^p, G(n) = (pn)/2$, and for $n = 2^p - 2$ and $n = 2^p - 4, G(n) = ((p-1)n)/2$. Note also that $G(24) = 36$ is known to be optimal by [5]. The optimality for $G(n)$ is indicated by an asterisk (*).

The "Comments" column indicates how these bounds have been obtained, and the "Formerly" column indicates the previously known upper bounds on $G(n)$, taken from the results of [5].

The gossip graphs obtained in [9] give better upper bounds than do our k-way split method for $n = 18, 20$, and 22. In particular, these graphs serve as a base for our method, which also helps to improve the following values of n.

Note, finally, that from these upper bounds for even n we can easily derive upper bounds for odd n, thanks to techniques given in [5]. Among others, we note that for all even n and odd k such that $2^p - k < n \leq 2^p$ we have $G(n+k) \leq G(n) + k$.

To completely understand Table 1, it is necessary to introduce the family of Knödel graphs $W_{\Delta,n}$, which appear to be gossip graphs in many cases.

Definition 2 (Knödel graph). *The* Knödel graph [6] *on $n \geq 2$ vertices (n even) and of maximum degree $1 \leq \Delta \leq \lfloor \log_2(n) \rfloor$ is denoted $W_{\Delta,n}$. The vertices of $W_{\Delta,n}$ are the pairs (i,j) with $i = 1,2$ and $0 \leq j \leq (n/2) - 1$. For every $j, 0 \leq j \leq (n/2) - 1$, there is an edge between vertex $(1,j)$ and every vertex $(2, j + 2^k - 1 \bmod n/2)$, for $k = 0, \ldots, \Delta - 1$.*

For $0 \leq k \leq \Delta - 1$, an edge of $W_{\Delta,n}$ which connects a vertex $(1,j)$ to the vertex $(2, j + 2^k - 1 \bmod n/2)$ is said to be *in dimension k*.

From [5], we get the following proposition:

Proposition 8 ([5]). *For all even n not a power of 2, we have*

TABLE 1. Upper bounds for $G(n)(18 \le n \le 128)$.

n	$G(n) \le$	Formerly	Comments
18	25	27	[9]
20	28	30	[9]
22	36	41	[9]
24	36^*	36^*	[5]
26	52	52	$W_{4,26}$
28	56^*	56^*	[5]
30	60^*	60^*	[8]
32	80^*	80^*	[8]
34	64	66	3-way [14–10–10]
36	68	68	[5]
38	74	74	[5]
40	76	76	[5]
42	84	84	$W_{4,42}$
44	88	88	$W_{4,44}$
46	108	108	[5]
48	96	96	[5]
50	125	125	$W_{5,50}$
52	130	130	$W_{5,52}$
54	135	135	$W_{5,54}$
56	140	140	$W_{5,56}$
58	145	145	$W_{5,58}$
60	150^*	150^*	[8]
62	155^*	155^*	[8]
64	192^*	192^*	[8]
66	130	165	3-way [24–24–18]
68	134	170	3-way [24–24–20]
70	143	175	3-way [24–24–22]
72	144	180	3-way [24–24–24]
74	161	185	3-way [26–24–22]
76	170	190	3-way [30–24–22]
78	171	195	3-way [30–24–24]
80	176	200	12-way [8–6–6]
82	197	205	3-way [30–30–22]
84	198	210	3-way [30–30–24]
86	215	215	$W_{5,86}$
88	208	220	12-way [8–8–6]
90	225	225	$W_{5,90}$
92	230	230	$W_{5,92}$
94	267	282	3-way [32–32–30]
96	240	240	[5]
98	294	294	$W_{6,98}$
100	300	300	$W_{6,100}$
102	306	306	$W_{6,102}$
104	312	312	$W_{6,104}$
106	318	318	$W_{6,106}$
108	324	324	$W_{6,108}$
110	330	330	$W_{6,110}$
112	336	336	$W_{6,112}$
114	342	342	$W_{6,114}$
116	348	348	$W_{6,116}$
118	354	354	$W_{6,118}$
120	360	360	$W_{6,120}$
122	366	366	$W_{6,122}$
124	372^*	372^*	[8]
126	378^*	378^*	[8]
128	448^*	448^*	[8]

- $W_{p-2,n}$ *is a gossip graph for any* $2^{p-1}+2 \le n \le 3 \cdot 2^{p-2} - 4$,
- $W_{p-1,n}$ *is a gossip graph for any* $3 \cdot 2^{p-2} - 2 \le n \le 2^p - 2$.

2.2. Lower Bounds for the Unit-cost Model

Lower bounds for $G(n)$ in the *unit-cost* model were previously studied in [5, 8] where the techniques used rely on the analysis of the structure of an MGG_n and mostly on the degrees of its vertices. In this section, following the techniques of [5], some improvements on these lower bounds are given. Before presenting these new results, we need to introduce the following notation: An edge $uv \in E$ is called a $(1,d)-$ *edge* iff $\deg u = 1$ and $\deg v = d$.

Note also that in the following we will abbreviate $p := \lceil \log_2 n \rceil$, that is, n (the number of vertices considered) is as follows: $2^{p-1}+1 \le n \le 2^p$. Then, $g_n = p$ if n is even and $g_n = p+1$ if n is odd.

Provided that we know certain rounds during which a fixed root v communicates with its neighbors, the first standard argument in the following lower-bound proofs is to estimate the maximum number of vertices which eventually can receive v's item: Any neighbor u of v learns this news during the first round, i, the edge vu is used in. From u, this item can then be broadcast to at most 2^{g_n-1} vertices (including u itself) during the remaining $g_n - i$ rounds. But every usage of vu in a later round $j > i$ prevents u from passing v's item to a yet uninformed neighbor which results in losing 2^{g_n-j} vertices from the above amount.

In the following, we will denote by

$$mbt_{g_n}(i, \overline{j_1}, \overline{j_2}, \ldots, \overline{j_l})$$

the maximum number of vertices to which a vertex v in an MGG_n (thus, with gossip time g_n) can broadcast its information via one of its neighbors u, provided that it uses the edge vu during rounds $i, j_1, j_2 \cdots j_l$, where $i < j_q$ for any $1 \le q \le l$. Note that v is counted among the informed vertices in $mbt_{g_n}(i, \overline{j_1}, \overline{j_2}, \ldots, \overline{j_l})$.

Thus,

$$mbt_{g_n}(i, \overline{j_1}, \overline{j_2}, \ldots, \overline{j_l}) := 1 + 2^{g_n-i} - \sum_{k=1}^{l} 2^{g_n-j_k}.$$

Summation over all edges incident to v yields the following general estimate:

Proposition 9. *Let R^+ include the set of all rounds when an edge incident to a fixed root v is used first, and let R^- be contained in the set of all remaining (nonfirst) usages of those edges. Then, including v itself, the number of vertices finally knowing v's item of information after g_n rounds cannot exceed*

$$mbt_{g_n}(t : t \in R^+, \bar{t} : t \in R^-) := 1 + \sum_{t \in R^+} 2^{g_n-t} - \sum_{t \in R^-} 2^{g_n-t}.$$

Our notation refers to the notion of *minimum broadcast trees* [5] because the above maximum can be attained by broadcasting along the edges of a suitable tree rooted at v. The subscript g_n refers to the total number of rounds available and might be omitted if there is no danger of ambiguity. Very often, we will consider consecutive rounds beginning in round r until round s, where sums of the form $\sum_{t=r}^{s} 2^{g_n-t} = 2^{g_n+1-r} - 2^{g_n-s}$ appear frequently.

The second standard argument uses the fact that the *inverse* of a given scheme, obtained by replacing each label $t = 1, \ldots, g_n$ by $g_n + 1 - t$, still provides complete gossiping. Or, equivalently, one considers the process of accumulating all items in the vertex v instead of broadcasting v's item as above. Thus, having proved a statement about round t (usually 1, 2, 3), this method yields an analog statement about round $g_n + 1 - t$ (usually $g_n, g_n - 1, g_n - 2$, respectively).

As an illustration, we recall the well-known

Proposition 10.

(a) *For all even n, every vertex has to communicate during rounds 1 and g_n.*
(b) *For all odd $n > 2^{p-1} + 1$, every vertex of degree 1 has to communicate during rounds 1 and g_n.*

Proof. Assume there is a vertex v [of degree 1 in (b)] using at most rounds $2, \ldots, g_n$ for calling its neighbors. Then, by Proposition 9, the number of vertices finally knowing the corresponding item cannot exceed

(a) $mbt_p(2, \ldots, g_n) = 1 + 2^{p+1-2} - 2^{p-p} = 2^{p-1} < n$;
(b) $mbt_p(2) = 1 + 2^{p+1-2} = 2^{p-1} + 1 < n$.

Because this would contradict the completeness of the entire information exchange, v has to use round 1 also. Considering the accumulation process or the inverse scheme, we know that it also has to be active during round g_n. ∎

Theorem 2. *$G(n) \geq (nd)/2$ for all even n with $2^p - 3 \cdot 2^{p-d} \leq n \leq 2^p, d \geq 3, p \geq 6, p \geq d + 2$.*

Proof. Recall that in this case $g_n = p$. By Proposition 9, we know that from a vertex of degree at most $d-2$ one can reach at most $\sum_{t=1}^{d-2} 2^{p-t} = 2^p - 4 \cdot 2^{p-d} < n$ vertices, and thus for all vertices $v, \deg v \geq d - 1$. Now, let us consider all edges used in round 1, which—thanks to the 1-port model and Proposition 10(a)—form a perfect matching. Either both incident vertices of such an edge have degree at least d, or, if otherwise, one has degree $d-1$ only, the idea is to show that the other one must have degree at least $d + 1$. Thus, the average degree is at least d, that is, the graph contains at least $\frac{1}{2} \sum_{v \in V} \deg v \geq \frac{nd}{2}$ edges—as asserted.

For the rest of the proof, let v be any vertex of degree $d - 1$, and similar to Proposition 9, let R^+ be the set of all rounds when an edge incident to v is used first and let R^- be the set of all remaining (nonfirst) usages of those edges. Assume there is a round $i \in \{1, \ldots, d - 1\}$ not belonging to R^+. Then, v's item of information can be communicated to at most

$$mbt_p(1, 2, \ldots, i - 1, i + 1, \ldots, d) = 1 + 2^p - 2^{p-i} - 2^{p-d}$$

vertices. This must be at least n, which is still possible, but only in the case $i = d - 1, n = 2^p - 3 \cdot 2^{p-d}, R^+ = \{1, \ldots, d-2, d\}$. Then, considering the accumulation process, we analogously get that v must use the rounds $p, p - 1, \ldots, p - d + 3$ and at least one of $p - d + 2$ or $p - d + 1$. If $d \geq 4$, this includes the rounds p and $p-1$ which cannot belong to R^+ since $p-1 \geq d+1$, that is, $R^- \supseteq \{p, p-1\}$. If $d = 3$, then either $R^- \supseteq \{p, p-1\}$ or $R^- \supseteq \{p, p-2\}$ because $p - 2 \geq 4$. In both cases, v's item cannot reach more than

$$\begin{aligned} &mbt_p(1, \ldots, d - 2, d, \bar{p}, \overline{p - 1}) \\ &\quad = 1 + 2^p - 3 \cdot 2^{p-d} - (2^0 + 2^1) < n \end{aligned}$$

vertices. This is a contradiction, and we thus know that v must use its adjacent edges during all the rounds $1, \ldots, d - 1$ and, for accumulating all items, also during the rounds $p, \ldots, p - d + 2$. Hence, $R^+ = \{1, \ldots, d - 1\}$ and $R^- \supseteq \{p, p - 1, p - 2\}$ because $p - 2 \geq d$.

Assume now, finally, that during round 1 v would communicate with a vertex u of degree at most d. Its item could then be disseminated to at most $mbt_p(2, \underline{\ldots, d})$ vertices via u and to at most $mbt_p(2, \ldots, d - 1, \bar{p}, \overline{p - 1})$ vertices via the other neighbors because at least two of $\{p, p - 1, p - 2\}$ cannot belong to the edge vu. But this adds up to only

$$\begin{aligned} &1 + 2^{p-1} - 2^{p-d} + 1 + 2^{p-1} - 2^{p-d+1} - (2^0 + 2^1) \\ &\quad = 2^p - 3 \cdot 2^{p-d} - 1 < n, \end{aligned}$$

that is, our final assumption was wrong, and $\deg u \geq d+1$ is proved. ∎

In the remainder of this section, we will always deal with MGG_n for odd $2^{p-1} + 1 \leq n \leq 2^p - 1$; in other words, $g_n = p+1$ and the *mbt*-function is always mbt_{p+1}.

Proposition 11. *For all odd n such that $2^p - 2^{p-d} + 1 \leq n \leq 2^p - 1$ with $p \geq d + 2 \geq 5$, there is no $(1, q)$-type edge with $q \leq d + 1$ in an MGG_n.*

Proof. Suppose that v and u are neighbors in an MGG_n, where $\deg v = 1$ and $\deg u \leq d + 1$. By Proposition 10, we know that v calls u in rounds 1 and $g_n = p+1$. Hence, u is fully informed already before round $p + 1$, but this cannot happen before round $\lceil \log_2 n \rceil = p$, that is, along an incident edge e, u calls one of its neighbors other than v in round p. None of them can be called earlier than in round 2. By Proposition 9, we can bound the number of vertices other than v, which eventually get u's item of information, namely, if p is the first usage of e, then it is at most

$$mbt_{p+1}(2, \ldots, d, p) = 1 + 2^p - 2^{p+1-d} + 2^{p+1-p} < n - 1;$$

otherwise, it is at most

$$mbt_{p+1}(2, \ldots, d + 1, \bar{p}) = 1 + 2^p - 2^{p-d} - 2^{p+1-p} < n - 1.$$

Note that $p \notin \{2, \ldots, d+1\}$ because of the assumption $p \geq d + 2$. In any case, u's information will not be disseminated through the entire network. ∎

Proposition 12 ([5]). *For all odd n such that $2^p - 2^{p-d} + 1 \le n \le 2^p - 1$ with $p \ge d + 2 \ge 5$, any vertex v of degree 2 has to communicate in rounds 1 and $p + 1$ in an MGG_n.*

Proof. Suppose that v starts to broadcast in round 2 only. Then, via the two incident edges, its item can reach at most $mbt_{p+1}(2)$ and $mbt_{p+1}(3)$ vertices, respectively, whereby v itself is counted twice. But

$$\begin{aligned} mbt_{p+1}(2) + mbt_{p+1}(3) - 1 &= 2 + 2^{p-1} + 2^{p-2} - 1 \\ &= 2^p - 2^{p-2} + 1 < n \end{aligned}$$

because $p - d \le p - 3$. Since this is not sufficient, one assertion is proved. To see the corresponding result for round $p + 1$, consider the inverse scheme. ∎

Proposition 13. *Let n be odd with $2^p - 2^{p-d} + 1 \le n \le 2^p - 1$, $p \ge d + 2 \ge 5$. In any MGG_n, if in round 1 there is a call between two vertices v and u both of degree 2, then in round 2, they have to call their other neighbors, v' and u', respectively, which both have to be of degree at least $d + 1$.*

Proof.

(a) v and v' as well as u and u' have to communicate in round 2 and in at least one of the rounds $p + 1$ or p.

Assume that u can start broadcasting its item in round 1 to v and not earlier than in round 3 to u', thus reaching no more than $mbt_{p+1}(2) + mbt_{p+1}(3) = 2^p - 2^{p-2} + 2 < n$ vertices, which is a contradiction. If u and u' would not communicate in one of the rounds $p + 1$ or p, then the above situation appears in the inverse scheme.

Suppose for the rest of the proof that $\deg v' \le d$.

(b) v' has to call one of its neighbors other than v during round p or $p + 1$.

As every vertex, v' becomes fully informed in round p or $p + 1$. Assume now that none of both rounds uses an edge between v' and one of its neighbors other than v. By considering the inverse scheme, we then know that v' can get the items of information from at most $1 + mbt_{p+1}(3)$ vertices via v and u and at most $mbt_{p+1}(3, \ldots, d + 1)$ vertices via its other neighbors. But

$$\begin{aligned} & mbt_{p+1}(3, \ldots, d + 1) + 1 + mbt_{p+1}(3) \\ &= 1 + 2^{p-1} - 2^{p-d} + 1 + 1 + 2^{p-2} \\ &= 2^p - 2^{p-d} - 2^{p-2} + 3 < n \end{aligned}$$

because $p - 2 > 2$, and this is again a contradiction.

Now, v's item can be forwarded from u and v' not earlier than beginning in rounds 2 and 3, respectively. But depending on whether there is a call between v and u in round $p + 1$ or not, Proposition 12, (a) and (b) imply that either u calls u' in round p and v' calls one of its neighbors other than v in round $p + 1$ or vice versa. Consequently, v's item can reach no more than $mbt_{p+1}(2, \overline{p + 1})$ or $mbt_{p+1}(2, \bar{p})$ vertices starting in u and no more than $mbt_{p+1}(3, \ldots, d + 1, \bar{p})$ or $mbt_{p+1}(3, \ldots, d + 1, \overline{p + 1})$ vertices starting in v', respectively. But both

$$mbt_{p+1}(2, \overline{p + 1}) + mbt_{p+1}(3, \ldots, d + 1, \bar{p}) \quad \text{and} \quad mbt_{p+1}(2, \bar{p}) + mbt_{p+1}(3, \ldots, d + 1, \overline{p + 1})$$

equal

$$\begin{aligned} & (1 + 2^{p-1}) + (1 + 2^{p-1} - 2^{p-d}) - (2^0 + 2^1) \\ &= 2^p - 2^{p-d} - 1 < n - 1, \end{aligned}$$

that is, not all vertices besides v itself can be informed. Consequently, $\deg v' > d$, and the proof for $\deg u' > d$ is analog. ∎

Theorem 3. *$G(n) \ge \lceil (5n)/4 \rceil$ for all odd n such that $2^p - 2^{p-d} + 1 \le n \le 2^p - 1$ with $p \ge d + 2 \ge 5$.*

Proof. Let V_1 and V_2 be the set of all vertices of degree 1 and 2, respectively. Moreover, let $V_{\le d+1} := \{x \in V : 3 \le \deg x \le d + 1\}$, and let $V_{\ge d+2,\bullet}$ be the set of vertices of degree at least $d + 1$ which are not adjacent to a vertex of degree 1. Finally, let n_i be the cardinality of the set V_i for any $i \in \{1; 2; \le d + 1; \ge d + 2, \bullet\}$.

In the 1-port model, Proposition 10(b) implies that no two vertices of V_1 can have a common neighbor. By Proposition 11, all neighbors of vertices of V_1 must have degree at least $d + 2$. Hence, there are n_1 vertices besides $V_1 \cup V_2 \cup V_{\le d+1} \cup V_{\ge d+2,\bullet}$, that is, $n = n_1 + n_2 + n_{\le d+1} + n_{\ge d+2,\bullet} + n_1$, which we use in the form

$$n_2 = n - 2n_1 - n_{\le d+1} - n_{\ge d+2,\bullet}. \tag{1}$$

Accordingly, summation of all degrees yields

$$\begin{aligned} 2G(n) \ge{} & n_1 + 2n_2 + 3n_{\le d+1} \\ & + (d + 2)n_{\ge d+2,\bullet} + (d + 2)n_1. \end{aligned} \tag{2}$$

Now, we consider any vertex v of degree 2. By Proposition 12, in round 1, it has to call one of its neighbors u, which itself can neither be of degree only 1 [Proposition 11] nor be adjacent to a vertex of degree 1 [Proposition 10(b)]. Thus, if $\deg u \ge 3$, then $u \in V_{\le d+1} \cup V_{\ge d+2,\bullet}$. If, otherwise, $\deg u = 2$, we are in the situation of Proposition 13 which says that in round 2 v calls its other neighbor v' and $\deg v' \ge d + 1$. Assume there is a vertex w of degree 1 adjacent to v'. By Proposition 10(b), w and v' communicate in rounds 1 and $p + 1$, that is, v' can broadcast its own item to w, at most $mbt_{p+1}(4)$ vertices via v and u and at most $mbt_{p+1}(3, \ldots, p)$ other vertices. But this adds up to at most

$$1 + (1 + 2^{p-1} - 2^1) + (1 + 2^{p-3}) = 2^p - 3 \cdot 2^{p-3} + 1 < n.$$

This contradiction shows that $v' \in V_{\ge d+2,\bullet}$.

Altogether, we showed that assigning to any $v \in V_2$ the uniquely determined first vertex of degree at least 3 it communicates with always leads to vertices of $V_{\le d+1} \cup$

$V_{\geq d+2,\bullet}$, whereby those of $V_{\leq d+1}$ or $V_{\geq d+2,\bullet}$ can appear at most once or twice, respectively. Hence, $n_2 \leq n_{\leq d+1} + 2n_{\geq d+2,\bullet}$ or

$$n_{\leq d+1} \geq n_2 - 2n_{\geq d+2,\bullet}. \tag{3}$$

Putting (3) and (1) into (2) yields

$$2G(n) \geq \tfrac{5}{2}n + (d-2)n_1 + (d - \tfrac{3}{2})n_{\geq d+2,\bullet},$$

and, thus, the assertion because $d \geq 3$. ∎

3. LINEAR COST MODEL

3.1. Compounding Method for the Linear Cost Model

First of all, we give here general upper bounds for $G_{\beta,\tau}(n)$ when n is even, thanks to the use of Knödel graphs. Indeed, the upper bounds from [5] given in Proposition 8 in the *unit-cost* model turn out to be applicable to the *linear-cost* model as well. Moreover, this observation will be useful in the proof of Theorem 4.

Observation 1. *For all even n, we have*

$$G_{\beta,\tau}(n) \leq \begin{cases} \frac{n\cdot(p-2)}{2} & \text{if } 2^{p-1}+2 \leq n \leq 3\cdot 2^{p-2}-4 \\ \frac{n\cdot(p-1)}{2} & \text{if } 3\cdot 2^{p-1}-2 \leq n \leq 2^p - 2. \end{cases}$$

Proof. The first part of this observation is derived from [6], where it was shown that the Knödel graph $W_{p-1,n}$ is a linear gossip graph for any even $2^{p-1}+2 \leq n \leq 2^p - 2$.

The second part of the observation was proved in the *unit-cost* model in [5], where it was shown that the Knödel graph $W_{p-2,n}$ is a gossip graph for any even $2^{p-1}+2 \leq n \leq 3\cdot 2^{p-2}-4$. However, it is easy to see that the proof still holds in the *linear-cost* model. Indeed, the gossip scheme is the following: gossip along edges in dimension $i-1$ for every round $1 \leq i \leq p-2$, then gossip again along dimension 0 during round $p-1$, then along dimension $p-3$ during round p. Such a gossip scheme is illustrated in Figure 10 for the case $p = 5$ and $n = 20$. In this figure, the number(s) in parentheses on the right of the "dim j" line denote the round(s) during which the edges in dimension j communicate.

It is easy to see that, during each round, the vertices respect the properties of an optimal gossip algorithm in the *linear-cost* model, as stated in [6]. Indeed, the first $(p-2)$ rounds take time $t_1 = (p-2)\beta + (2^{p-2}-1)\tau$, while round $p-1$ takes $t_2 = \beta + (2^{p-2}-2)\tau$ and round p takes $t_3 = \beta + n - (2^{p-1}-2)\tau$. Hence, the total gossip time is $t = p\beta + (n-1)\tau$, and $W_{p-2,n}$ is a linear gossip graph for every $2^{p-1}+2 \leq n \leq 3\cdot 2^{p-2}-4$. ∎

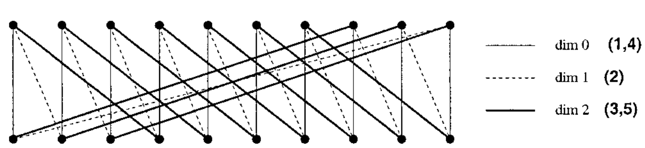

FIG. 10. $W_{3,20}$ and (between parentheses) a gossip scheme.

In [6], Fraigniaud and Peters gave a compounding method to get linear gossip graphs from existing (minimum) linear gossip graphs. However, their method differs from ours, since in their case, they take an MLGG$_k$ and replace each vertex by a copy of an MLGG$_n$ and each edge linking two vertices by a perfect matching between two copies of a MLGG$_n$. Here, we present a method similar in all points to the method proposed in Theorem 1 in the *unit-cost* model, which turns out to give better results than does the compounding method from [6].

Theorem 4 (Compounding in the Linear-cost Model). *For all $k = 2^{p-1}-1$ and even n such that $\lceil \log_2(kn) \rceil = \lceil \log_2(k) \rceil + \lceil \log_2(n) \rceil$, we have*

$$G_{\beta,\tau}(kn) \leq kG_{\beta,\tau}(n) + \frac{n}{2}\cdot(G_{\beta,\tau}(2k) - k).$$

Proof. To prove the theorem, we need, as in the *unit-cost* model, to start from an MLGG$_{2k}$ (i.e., an MLGG of order $2k = 2^p - 2$) which is compoundable. However, by Observation 1, we know that $W_{p-1,2k}$ is an MLGG$_{2k}$. For this, we use the following gossip scheme: Let the vertices communicate along dimension $i-1$ during every round $1 \leq i \leq p-1$ and vertices communicate again along dimension 0 during round p. Hence, $W_{p-1,2k}$ is clearly compoundable, and for our purpose, we will use the perfect matching PM_2 induced by the communications which take place during round 2, that is, the perfect matching corresponding to dimension 1 in the graph. We illustrate this in Figure 11, where $k = 7$. The number(s) in parentheses on the right of the "dim j" line denote the round(s) during which the edges in dimension j communicate.

In that case, let us construct a graph of order nk the same way as in the *unit-cost* model: We replace each edge (and its adjacent vertices) of the perfect matching PM_2 by a copy of an MLGG$_n$. In each of these copies G_i, we split the vertices into two subsets of equal cardinality, $V_{i,1}$ and $V_{i,2}$. Then, each of these $V_{i,j}$ will play the role of what was previously a single vertex in the MLGG$_{2k}$. More precisely, if two vertices u and v were neighbors in $W_{p-1,2k}$, then we add a perfect matching between the corresponding $V_{i,j}$ and $V_{i',j'}$.

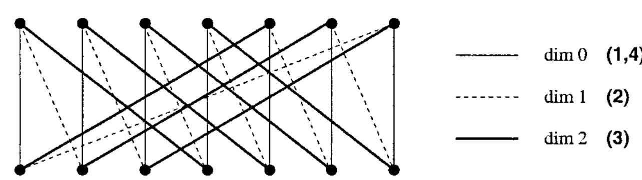

FIG. 11. $W_{3,14}$ and (between parentheses) a gossip scheme.

Now the gossip scheme will be as follows:

1. Gossip along what was previously dimension 0 during the first round;
2. Then, gossip independently in each copy G_i of an MLGG$_n$ during rounds 2 to $\lceil \log_2(n) \rceil + 1$;
3. Finally, gossip from rounds $\lceil \log_2(n) \rceil + 2$ to $\lceil \log_2(k) \rceil + \lceil \log_2(n) \rceil$ the same way as in $W_{p-1,2k}$, but with a delay of $\lceil \log_2(n) \rceil - 1$ rounds.

We then see, by the same argument as in the *unit-cost* model, that every vertex will be informed of all the pieces of information of the other vertices within $\lceil \log_2(kn) \rceil$ rounds. In other words, this proves that $g_{\beta,\tau}(kn) \leq \lceil \log_2(kn) \rceil + f(k,n) \cdot \tau$, where $f(k,n)$ is a function of k and n. However, we still have to show that $f(k,n) \leq kn - 1$:

1. Since each vertex starts with a unique piece of information of length 1, the first round takes time $t_1 = \beta + \tau$.
2. Then, each vertex knows two pieces of information and will gossip independently in its own copy of an MLGG$_n$. Hence, this "internal" gossiping takes time $t_n = \lceil \log_2(n) \rceil \beta + 2(n-1)\tau$.
3. After this, each vertex knows $2n$ pieces of information, and the gossip goes on during $\lceil \log_2(k) \rceil - 1$ more rounds, as it did in $W_{p-1,2k}$. The only difference, as stated above, is that each vertex knows $2n$ pieces of information instead of 1. Moreover, except for the very last round, there is no overlap, that is, every pair of vertices which communicate do not have any piece of information in common. Hence,
   - During the $\lceil \log_2(k) \rceil - 2$ first rounds (among the $\lceil \log_2(k) \rceil - 1$ remaining), the time to exchange information will be $\beta + 2n\tau, \beta + 4n\tau, \beta + 8n\tau, \ldots, \beta + 2^{\lceil \log_2(k) - 2 \rceil} n\tau$.
   - During the very last round, every pair of vertices has n pieces of information in common (since they have already exchanged information at round 1 and since the gossip in each G_i begins just after round 1). Hence, the time needed for the last round is $t_l = \beta + (2^{\lceil \log_2(k) - 1 \rceil} - 1)n\tau$.

Summing all the times needed for each round, we get a total time of $t = t_1 + t_n + (\sum_{i=1}^{\lceil \log_2(k) \rceil - 2} \beta + 2^i n\tau) + t_l$. Standard calculations then give us $t = (\lceil \log_2(n) \rceil + \lceil \log_2(k) \rceil)\beta + (n \cdot (2^{\lceil \log_2(k) \rceil} - 1) - 1)\tau$. However, we know by hypothesis that $\lceil \log_2(n) \rceil + \lceil \log_2(k) \rceil = \lceil \log_2(kn) \rceil$ and that $k = 2^{p-1} - 1$; hence, we have $t = \lceil \log_2(kn) \rceil \beta + (kn-1)\tau$, and $f(k,n) \leq kn - 1$.

Consequently, the graph that we build by this method is able to gossip in the *linear-cost* model in minimum time. Hence, this is a linear gossip graph. Since this construction gives us graphs of the same size as in the *unit-cost* model, we directly get $G_{\beta,\tau}(kn) \leq kG_{\beta,\tau}(n) + \frac{n}{2} \cdot (G_{\beta,\tau}(2k) - k)$. ∎

Thanks to the previous method, we obtain, for the first time, as for the *unit-cost* model, infinitely many linear gossip graphs for which their number of edges does not exceed $(n/2) \cdot (\lceil \log_2(n) \rceil - 3)$ edges. The proof relies exactly on the same arguments as for Proposition 6, since the Knödel graph $W_{p-5,2k}$ is a gossip graph as well as a linear gossip graph, and gossiping can be achieved in

TABLE 2. Upper bounds for $G_{\beta,\tau(n)}$ (n even, $18 \leq n \leq 128$).

n	$G_{\beta,\tau}(n) \leq$	Formerly	Comments
18	27*	27*	[6]
20	30*	30*	[6]
22	44	44	$W_{4,22}$
24	36*	36*	[6]
26	52*	52*	$W_{4,26}$
28	56*	56*	[6]
30	60*	60*	[6]
32	80*	80*	[6]
34	68	85	$W_{4,34}$
36	72	90	$W_{4,36}$
38	76	95	$W_{4,38}$
40	80	100	$W_{4,40}$
42	84*	84*	[6]
44	88*	88*	[6]
46	115	115	$W_{5,46}$
48	96	96	2-way [24–24]
50	125	125	$W_{5,50}$
52	130	130	$W_{5,52}$
54	135	135	$W_{5,54}$
56	140	140	$W_{5,56}$
58	145*	145*	[6]
60	150*	150*	[6]
62	155*	155*	[6]
64	192*	192*	[6]
66	165	198	$W_{5,66}$
68	170	204	$W_{5,68}$
70	175	210	$W_{5,70}$
72	144*	144*	3-way
74	185	222	$W_{5,74}$
76	190	228	$W_{5,76}$
78	195	234	$W_{5,78}$
80	200	240	$W_{5,80}$
82	205	246	$W_{5,82}$
84	210	252	$W_{5,84}$
86	215	258	$W_{5,86}$
88	220	264	$W_{5,88}$
90	225	270	$W_{5,90}$
92	230	276	$W_{5,92}$
94	282	282	$W_{6,94}$
96	240	240	2-way [48–48]
98	294	294	$W_{6,98}$
100	300	300	$W_{6,100}$
102	306	306	$W_{6,102}$
104	312	312	$W_{6,104}$
106	318	318	$W_{6,106}$
108	324	324	$W_{6,108}$
110	330	330	$W_{6,110}$
112	336	336	$W_{6,112}$
114	342	342	$W_{6,114}$
116	348	348	$W_{6,116}$
118	354	354	$W_{6,118}$
120	360	360	$W_{6,120}$
122	366*	366*	[6]
124	372*	372*	[6]
126	378*	378*	[6]
128	448*	448*	[6]

both cases using the same gossip scheme [6, 8]. This is the purpose of the following observation:

Observation 2. $G_{\beta,\tau}(n') \leq (n'(p-3))/2$ *for all* $p \geq 7$ *and* $n' = 24 \cdot (2^{p-5} - 1)$.

Also, thanks to Theorem 4, it is possible to determine the exact value of $G_{\beta,\tau}(72)$.

Theorem 5. $G_{\beta,\tau}(72) = 144$.

Proof. The upper bound is given by Observation 2, where $p = 7$, that is, $k = 3$. Moreover, we know by Theorem 2.17 of [6] that a vertex of degree 3 can know only up to 66 pieces of information after seven rounds. Since $g_{72} = 7$, it follows that there is no vertex of degree less than or equal to 3 in an MLGG_{72}. Hence, $G_{\beta,\tau}(72) \geq 144$. Since the upper and lower bound coincide, we get the result. ∎

3.2. *Summary of the Results (Linear Cost)*

Table 2 presents the results given by the k-way split method and Observation 1 for even n with $18 \leq n \leq 128$. Note that for the values $n = 2^p, n = 2^p - 2, n = 2^p - 4$ and $n = 2^p - 6$ the result is optimal [6]. For $n = 2^p, G_{\beta,\tau}(n) = (pn)/2$, and for $n = 2^p - 2, n = 2^p - 4$, and $n = 2^p - 6, G_{\beta,\tau}(n) = ((p-1)n)/2$. The optimality for $G_{\beta,\tau}(n)$ is indicated by an asterisk (*).

The "Comments" column indicates how these bounds have been obtained, and the "Formerly" column is taken from [6].

4. CONCLUSIONS

In this paper, we presented a general compounding method which gives upper bounds for $G(kn)$ and $G_{\beta,\tau}(kn)$ for even n. Moreover, in the *unit-cost* model, it is possible, in some cases, to use variants of the general method, which are applicable for more (even) values. Thanks to these methods, some upper bounds can also be derived for $G(n)$ when n is odd, still in the *unit-cost* model. All these results, together with the ones of Proposition 8 and Observation 1, give the best-known upper bounds for $G(n)$ [respectively, $G_{\beta,\tau}(n)$], either matching or improving the upper bounds given in [5] and [6]. It is also interesting to note that these improvements can, in turn, be taken as entries for further values of upper bounds for $G(n)$ and that the recursion can obviously be applied several times. In a word, any improvement in the knowledge of $G(n)$ [respectively, $G_{\beta,\tau}(n)$], whether by our method or by any other, will help to improve the upper bounds on further values of $G(n)$ [respectively, $G_{\beta,\tau}(n)$].

We also proved that, for infinitely many n, there exists (linear) gossip graphs with $(n/2)\cdot(\lceil\log_2(n)\rceil - 3)$ [respectively, $(n/2)\cdot(\lceil\log_2(n)\rceil - 4)$ in the *unit-cost* model only] edges, something which was unknown before. The best result, formerly, was $(n/2)\cdot(\lceil\log_2(n)\rceil - 2)$, thanks to the Knödel graph $W_{\lceil\log_2(n)\rceil-2,n}$ (cf. Proposition 8 and Observation 1).

Although these methods have the same flavor as that of Farley's k-way split method concerning broadcasting, it is surprising that they had never been proved to be efficient for gossiping.

In addition to these general upper-bound results, we also gave some small improvements on certain lower bounds for $G(n)$, which were derived from a method used in [5].

Acknowledgments

The authors wish to thank the anonymous referees for their helpful comments and suggestions, which resulted in an improved version of the paper.

REFERENCES

[1] J.C. Bermond, P. Fraigniaud, and J.G. Peters, Antepenultimate broadcasting, Networks 26 (1995), 125–137.

[2] S.C. Chau and A.L. Liestman, Constructing minimal broadcast networks, J Combin Inf Syst Sci 10 (1985), 110–122.

[3] M.J. Dinneen, J.A. Ventura, M.C. Wilson, and G. Zakeri, Compound constructions of broadcast networks, DAMATH: Discr Appl Math Combin Oper Res Comput Sci (1999), 93.

[4] A. Farley, Minimal broadcast networks, Networks 9 (1979), 313–332.

[5] G. Fertin, A study of minimum gossip graphs, Discr Math 1–3(215) (2000), 33–57.

[6] 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, BC, 1994. Available at *http://fas.sfu.ca/pub/cs/techreports/1994/*.

[7] W. Knödel, New gossips and telephones, Discr Math 13 (1975), 95.

[8] R. Labahn, Some minimum gossip graphs, Networks 23 (1993), 333–341.

[9] H.-C. Pahlig, Approximative Algorithmen zur Konstruktion minimaler Gossipgraphen, Diplomarbeit, Universität Rostock–Fachbereich Mathematik, Feb. 1997.

[10] J.A. Ventura and M.X. Weng, A new method for constructing minimal broadcast networks, Networks 23 (1993), 481–497.

[11] M.X. Weng and J.A. Ventura, A doubling procedure for constructing minimal broadcast networks, Telecomm Syst 3 (1995), 259–293.